

i-MATH Winter School
DocCourse
Combinatorics and Geometry 2009
Discrete and Computational Geometry

Centre de Recerca Matemàtica
January to March 2009

Research Reports

Edited by

Marc Noy (Universitat Politècnica de Catalunya)
Julian Pfeifle (Universitat Politècnica de Catalunya)



© CRM

Centre de Recerca Matemàtica
Campus de Bellaterra, Edifici C
08193 Bellaterra (Barcelona)

First edition: July 2010

ISBN:

Legal deposit:

Contents

Combinatorial Convexity	
by Víctor Álvarez and Jeong Hyeon Park	7
What Does the Convex Hull of a Discrete Point Set Look Like?	
by Ina Voigt	19
f-vectors and Radon Partitions	
by Jeong Hyeon Park	27
Independence Complexes of Perfect Graphs	
by Ragnar Freij and Matthias Henze	39
The Last Delaunay Triangulation	
by Víctor Álvarez, Frederik von Heymann, and Canek Peláez	55
Lattice Paths and Lagrangian Matroids	
by Anna Gundert, Edward D. Kim, and Daria Schymura	63
Patterns in Ordered Trees	
by Ragnar Freij, Daria Schymura, and Lluís Vena	83
Prodsimplicial-neighborly Polytopes	
by Benjamin Matschke, Bernardo González Merino, Julian Pfeifle, and Vincent Pilaud	105
Realization of the Delaunay Triangulation	
by Víctor Álvarez and Marek Krčál	141
Recoloring Directed Graphs	
by Stefan Felsner, Clemens Huemer, and Maria Saumell	153
Sets with Small Neighborhood in the Integer Lattice	
by Aaron Dall, Frederik von Heymann, and Birgit Vogtenhuber	163

Part III

Research Reports

Combinatorial Convexity

Victor Alvarez, Jeong Hyeon Park

April 20, 2009

Abstract

Erdős-Szekeres theorem is one of classic results in combinatorial geometry. In this project we consider the colored version of the problem. Especially, we are interested in the number of empty monochromatic triangles and the existence of an empty monochromatic convex quadrilateral. We give some minor results and plausible ideas to solve the problems.

1 Introduction

The following problem is quite well known among combinatorial geometry problems : given n points in the plane that are in general position, are there k points in convex position when n is very large? This problem was solved by Erdős and Szekeres[7]. They proved the result by Ramsey numbers and it becomes one of the most fundamental and famous results in combinatorial geometry. This problem has many relatives. One of its famous relative proposed by Erdős is as follows : given n points in the plane that are in general position, are there k points in convex position whose convex hull is *empty* when n is very large? How many such k tuple of points are there? Those problems are investigated by many researchers and there are many results on the problem.

We are interested in the chromatic variant of the problem. The chromatic variants of the problem is as follows

Question 1. *We are given n points in the plane that are in general position and each point is colored with red or blue. Then, are there k points of the same color in convex position whose convex hull is empty? How many such monochromatic k tuple of points are there?*

Those problems were recently proposed but there are already some nontrivial results. In this project, we present some minor results related to the problems and plausible ideas which may be helpful.

2 Related Work

2.1 On Empty k -gons

The existence of an empty quadrilaterals is obvious by a classic observation : if there are five points in the plane, there is a convex quadrilateral among them. The existence of an empty convex pentagon is proved by Harboth[9]. On the contrary, there is an arbitrarily large point set which doesn't contain empty convex k -gons, $k \geq 7$. Such a point set is called




Figure 1: An example of an empty monochromatic quadrilateral




Figure 2: The Horton set

the Horton set, found by Horton[10]. The Horton sets are constructed in a recursive way : $H^+(n/2)$ are lie high above $H^-(n/2)$.

The existence of an empty convex hexagon was a wide open problem for many years. Recently, Gerken[8] proved the existence of an empty hexagon. He characterized a point set without an empty convex hexagon that layered convex polygons and proved that there can't be such point set when the number of point is large. Nicholas[13] also proved the argument in a different way and Valtr[16] simplified Gerken's proof.

How many empty convex k -gons are there? It is easy to show that there are quadratic number of empty triangles. Katchalski, Meir[12] proved that there are $\binom{n-1}{2}$ empty triangles. Bárány, Füredi[3] improved the bound to $n^2 - O(n \log n)$ which is the best lower bound so far known. The upper bound on the number is again quadratic. Katchalski, Meir[12] constructed a point set which have only $c \cdot n^2, c \leq 800$ empty triangles. Bárány, Füredi[3] asserted there is a point set with only $2n^2$ empty triangles by a probabilistic argument. Valtr[17] constructed a point set which gives $\leq 1.8n^2$ empty triangles. His construction is to place the Horton sets in four layers. His technique is improved several times : Dumitrescu[6], Bárány, Valtr[4]. So far the best known upper bound is $\leq 1.62n^2$. With those results on the lower bound and the upper bound, the number of empty triangles is now believed to be $(1 + c)n^2, c > 0$.

It is known that the number of empty convex k -gons is dependent on the number of empty triangles. Pinchasi, Radoičić, Sharir[14] derived the relation among the number of empty convex k -gons. They defined and proved moment sums on the number of empty convex k -gons and that gives the number of empty convex quadrilaterals is quadratic and the




Figure 3: Valtr[17]'s upper bound construction

number of empty convex pentagons is quadratic as well if the conjecture is true. We will briefly review their result.

2.2 On Empty Monochromatic k -gons

The existence of an empty monochromatic triangle is given by Devillers, Hurtado, K  ryoli, Seara[5] proving the number of empty monochromatic triangles is at least $\lceil n/4 \rceil - 2$. But there is a point set without an empty monochromatic pentagon. It is again given by Devillers, Hurtado, K  ryoli, Seara[5] coloring the Horton set. We don't know whether there is an empty monochromatic convex quadrilateral in a large point set. It is known that there is a point set of 36 points without an empty monochromatic convex quadrilateral[11]. And Devillers, Hurtado, K  ryoli, Seara[5] proved there is an empty monochromatic convex quadrilateral in the Horton set. Recently, Aichholzer, Hackl, Huemer, Hurtado, Vogtenhuber[2] showed there is an empty monochromatic quadrilateral(possibly non convex) in a large point set.

The lower bound on the number of empty monochromatic triangles is proven by Aichholzer, Fabila-Monroy, Flores-Pe  aloza, Hackl, Huemer, Urrutia[1]. They shows there are $\Omega(n^{5/4})$ empty monochromatic triangles. Their proof is based on the discrepancy lemma, there are at least $\Omega((n_r - n_b)^2)$ empty monochromatic triangles. Their technique is improved by Pach, T  th[15] to $\Omega(n^{4/3})$. The upper bound is still quadratic and it is conjectured that there are always $\Omega(n^2)$ empty monochromatic triangles.

3 Empty Monochromatic Triangles

3.1 Relation among the number of empty convex k -gons

Pinchasi, Radoi  i  , Sharir[14] derived the following inequalities with very long analysis.

Theorem 3.1.1. *Let $X_k(P)$ be the number of empty convex k -gons in the point set P . Then, the following inequalities hold*

$$X_4(P) \geq X_3(P) - \frac{n^2}{2} - O(n)$$

$$X_5(P) \geq X_3(P) - n^2 - O(n)$$

Here, we give a very simple proof of the first inequality.

Our proof is based on Bárány and Füredi[3]. First, let l be an arbitrary line which is very far from the point set P . Then, for each segment $p_i p_j$, consider points in the strip s_{ij} defined by l_i and l_j that are orthogonal to l . Then, if we count empty k -gons in each strip which are incident with p_i and p_j , they are counted only once.




Figure 4: Empty k -gons counted only once

Let t_{ij} be the number of points in s_{ij} that forms empty triangles with p_i, p_j . Then, $X_3(P) = \sum t_{ij}$. Let q_{ij} be the number of empty quadrilaterals in s_{ij} which are incident with p_i, p_j . Then, q_{ij} and t_{ij} have the following relation.

Observation 1. $q_{ij} \geq t_{ij} - 1$

Proof. Let $t_{ij}^+(t_{ij}^-)$ be the number empty triangles incident with p_i, p_j such that the other point is above(below) the segment $p_i p_j$. Then,

- (1) Two adjacent triangles form a convex quadrilateral
- (2) Two triangles in different sides form a convex quadrilateral




Figure 5: Two kinds of empty quadrilaterals

Thus, $q_{ij} \geq \max\{t_{ij} - 1, t_{ij}^+ \cdot t_{ij}^-\}$. \square

By the observation, we get the following relation between $X_3(P)$ and $X_4(P)$.

Theorem 3.1.2. $X_4(P) \geq X_3(P) - \binom{n}{2}$

Proof. $X_4(P) = \sum q_{ij} \geq \sum t_{ij} - 1 = X_3(P) - \binom{n}{2}$ \square

Unfortunately, we couldn't derive a relation between $X_3(P)$ and $X_5(P)$. But we believe there is a simpler proof.

3.2 Simple Upper Bound

Aichholzer, *etc* [1] showed that there is a bi-colored point set which only has $\frac{1}{4}T(n)$ empty monochromatic triangles where $T(n)$ is the minimum over the numbers of empty triangles of n points. Their construction is merely to duplicate the point set which has the the smallest number of empty triangles among the point sets with the same cardinality and move them in ϵ . Then, for every empty triangle there is only one empty monochromatic triangle. See the following figure.




Figure 6: Only one triangle is monochromatic and empty

Here, we give an probabilistic construction which gives a similar upper bound.

Theorem 3.2.1. *There is a bi-colored point set of cardinality n such that the number of monochromatic empty triangles is strictly less than $\frac{1}{4}T(n)$*

Proof. Let P be the point set of cardinality n which has the smallest number of empty triangles. Then, we color each vertex red or blue with probability $\frac{1}{2}$. The probability that an empty triangle remains an empty monochromatic triangle is $\frac{1}{4}$: the triangle could be red or blue with probability $\frac{1}{8}$, respectively.

Thus, by the following simple calculation, the expectation of the number $T_m(n)$ of empty monochromatic triangles is

$$E[T_m(n)] = \sum_{\text{triangle } t \text{ is empty}} Pr[t \text{ is monochromatic}] = \frac{1}{4}T(n)$$

And we may color the point set with one color which has $T(n)$ monochromatic empty triangles. That is bigger than the expectation number so




Figure 7: All the possible cases

there must be point set which has empty monochromatic triangles of cardinality strictly less than $\frac{1}{4}T(n)$. \square

3.3 Towards the Quadratic Lower Bound

Our method is based on Bárány and Füredi's construction, see [3], of the lower bound for empty triangles when the given set of points is colorless, which we briefly describe here.

Given a colorless set P of n points, let us consider the complete graph K_n^P defined by each pair of points of P . Now let us consider any straight line l such that $\text{Conv}(P) \cap l = \emptyset$. We will label the points of P with integers $1, \dots, n$ according to how they appear on their orthogonal projection on l from left to right. Note that this labelling has the property that for any two labels i, j , $1 \leq i < j \leq n$, the slab that is delimited by the orthogonal projection on l of $p_i, p_j \in P$ contains only points with labels strictly between i and j .

Now let us call an edge $e_{ij} = \overline{p_i p_j}$ of K_n^P *good* if and only if e_{ij} constitutes a side of at least two empty triangles of P where the third vertex p_k is such that $i < k < j$. That is, p_k is contained in the slab defined by the orthogonal projection on l of p_i, p_j .

In [3], Bárány and Füredi showed that K_n^P contains essentially $n^2 - O(n \log n)$ *good* edges and hence, that many empty triangles.

Having described the previous construction, we can now proceed to explain our idea.

From now on let us consider a bi-colored set of points $P = R \cup B$ of $2n$ points in general position, where $|R| = |B| = n$ and R and B will stand for red and blue respectively. Let us choose one chromatic class, say R , and let us label its points with Bárány *et al.*'s labelling. Our idea is based on the fact that most of the edges of K_n^R are *good* edges, without considering the blue chromatic class. That is, let $e_{ij} = \overline{r_i r_j} \in R$ be a *good* edge of K_n^R such that $i < j$. If at least one of the red triangles that have e_{ij} as one of its sides is empty of blue points, then we will charge such a triangle to e_{ij} and we will continue choosing another edge of K_n^R . Note that such a triangle will not be counted twice. However, it might be possible that both triangles are pierced by at least one point of B . In such a case, let us consider the slab S_{ij}^l formed by the orthogonal projection on l of r_i and r_j and let us suppose for one moment that in this slab we can find an empty blue triangle T^B , we will explain later on why this assumption makes some sense. Being T^B fully contained in S_{ij}^l , the only edges of K_n^R for which this triangle can be counted more than one time are the edges that have at least one endpoint outside S_{ij}^l , and that do not intersect T^B as we are not trying to count empty blue triangles that intersect our original edge e_{ij} . Now, such set of edges can be divided into three parts, the edges that lie strictly above e_{ij} and possibly containing as

one endpoint either r_i or r_j , called this set \mathcal{A}_{ij} . The edges that lie strictly below e_{ij} and again possibly containing as one endpoint either r_i or r_j , called this set \mathcal{B}_{ij} . Finally, the set of edges that intersect the interior of e_{ij} properly, called this set \mathcal{I}_{ij} .




Figure 8: "Walking" along red edges to find a blue empty triangle

Let us assume without loss of generality, that T^B lies above e_{ij} . For this blue triangle, crucial will be the edges in \mathcal{A}_{ij} and \mathcal{I}_{ij} only, as for the edges of \mathcal{B}_{ij} , if we do not first find one empty red triangle, we would find one empty blue triangle that intersects e_{ij} or that is outside S_{ij}^l , and hence we could not charge such blue triangle to e_{ij} .

Now, let us divide the analysis into two cases:

- $\mathcal{A}_{ij} \neq \emptyset$. Let $V(\mathcal{A}_{ij})$ be the set of vertices of the edges of \mathcal{A}_{ij} . Now, let $V(\mathcal{A}_{ij})_{left}$ and $V(\mathcal{A}_{ij})_{right}$ be the set of vertices of $V(\mathcal{A}_{ij})$ that lie to the left and outside of S_{ij}^l , and to the right and outside of S_{ij}^l respectively. Note that the former is on the side of r_i and the latter is on the side of r_j .

The idea is to consider all the triangles formed by the vertices r_i, r_j and $r_k \in V(\mathcal{A}_{ij})_{left}$. Note that by construction $k < i < j$. If at least one of these triangles is empty, then we are done and we can charge this triangle to e_{ij} and continue with another edge of K_n^R . If none of these triangles is empty, then there is one blue point b that along with two other vertices of T^B form another empty blue triangle that is intersected by all the edges $r_j r_k$, with $r_k \in V(\mathcal{A}_{ij})_{left}$, and therefore such triangle will not be charge to any of those edges. We can handle similarly the vertices of $V(\mathcal{A}_{ij})_{right}$ and as a final product we would get an empty blue triangle that is intersected by all the edges of \mathcal{A}_{ij} and therefore that will not be counted for any of those edges.

- $\mathcal{I}_{ij} \neq \emptyset$. This case is still not fully clear to us how to handle as here the things get more complicated since a combination of *good* and *bad* edges of K_n^R might lead us to the same blue triangle T^B and hence to an overcounting. We have some ideas on how to fix this problem

that we believe are correct, however, the details are quite intricate and we have not finished to analyze all the details, therefore we have decided to leave those details out of this note until we come up with the right proof.




Figure 9: How many times can a blue triangle be charged?

We still have to discuss why the assumption of the existence of T^B makes some sense. The idea behind this intuition is that as we said before, the number of *good* edges of K_n^R is essentially $n^2 - O(n \log n)$, that is to say, if e_{ij} is a *good* edge, most probably the edges that complete the empty red triangles (without taking the blue chromatic class into consideration) that have e_{ij} as a side are also *good* and so on, note that for n sufficiently large $O(n \log n)$ is negligible against n^2 . Therefore, for each one of this edges we have at least two empty red triangles, if none of these triangles is empty of blue triangles, then as we follow the chain of *good* edges inside S_{ij}^t we go collecting blue points that eventually will help us to form the desired empty blue triangle T^B .

Details are vague at this point but we keep on working and we are confident that this method could lead us to the desired quadratic bound.

4 Empty Monochromatic Quadrilaterals

4.1 Generalization of Odd Cycles

In this section, we give a new idea to find a convex quadrilateral. Consider a graph G and bicoloring of the vertex of G . When does G contain a monochromatic edge whose ends are of the same color for every coloring of G ? Probably you can easily answer the question: G has a monochromatic edge if and only if G contains an odd cycle. And it is well known that G doesn't have any odd cycle if and only if G is bipartite. As a result, if $|E(G)| \geq t_2(n)$ where $t_2(n)$ is the maximum number of edges a bipartite graph G can have, G contains an odd cycle.

Our idea is to generalize it to hypergraphs. Consider a bicoloring of a hypergraph G . Now the question is what kind of a hypergraph G contains an edge which contains four vertices of the same color for every coloring of G ? The following example shows one class of such a graph.

Example 4.1.1. Consider a vertex set $V = \{1, \dots, 6k+1\}$. Construct a hypergraph G on V : every consecutive 6 vertices $i, \dots, i+5 \pmod{6k+1}$ forms an edge. And consider any bi-coloring of G . To prevent all the edges from having four vertices of the same color, every edge must contain three

red vertices and three blue vertices. And it is easily seen that two edges, $(i, \dots, i+5)$, $(i+6, \dots, i+11)$ must have the same coloring. But it is not possible for every edge because the number of vertices is not divisible by 6. Thus, there must be an edge with four vertices of the same color.



Figure 10: "Odd Cycle"

This hypergraph can be realizable in the plane. An edge of six vertices represents a convex hexagon in the plane. Then, that represents a convex $6k+1$ -gon such that every consecutive 6 vertices form a empty hexagon. For any bi-coloring of this polygon, there must be an empty monochromatic quadrilateral.




Figure 11: Example of an "odd cycle" in the plane

However, this example is not very interesting. In this example, there is already an empty 7-gon : either the convex $6k+1$ -gon is empty or there is a point which forms an empty 7-gon with vertices $(i, \dots, i+5)$. Note that the existence of an empty convex 7-gon implies the existence of a convex quadrilateral. So, we ask the following question.

Question 2. *Is there a simple but nontrivial "odd cycle"?*

We believe that there must be such an "odd cycle" in a very large point set. And the existence of an empty hexagon is probably related to an "odd cycle". In other words, an empty hexagon plays an important role in an "odd cycle".

5 Concluding Remarks

As far we know, every result on the number of empty monochromatic triangles is based on the discrepancy lemma. However, we felt the lemma

is not powerful enough to give a better lower bound. It is because the statement of the lemma and its proof are very loose. Of course, one may improve the discrepancy lemma and that would give a slightly better lower bound. But we think the lower bound would not be drastically improved. So we think a proof not based on the discrepancy lemma is close the solution.

The upper bound on the number of monochromatic empty triangles is still believed to be quadratic. We also believe that conjecture but on the other hand, one might construct a point set with subquadratic number of empty monochromatic triangles. That might be achieved by coloring the point set with very small number of empty triangles or by randomized ϵ duplication of the point set.

The existence of empty monochromatic k -gons can be generalized in the following way. Instead of considering monochromatic k points in convex position, the goal is to find monochromatic k points whose convex hull doesn't contain other points. This generalization would be helpful : if we find such monochromatic 5 points that would imply the existence of an empty monochromatic quadrilateral. This problem is interesting by itself.

References

- [1] O. Aichholzer, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, J. Urrutia, Empty monochromatic triangles. Proc. 20th Canadian Conference on Computational Geometry, 20, Montreal, Quebec, Canada, (2008) 75-78.
- [2] O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, B. Vogtenhuber. Large bichromatic point sets admit empty monochromatic 4-gons. In Proc. 25th European Workshop on Computational Geometry EuroCG '09, Brussels, Belgium, (2009).
- [3] I. Bárány, Z. Füredi, Empty simplices in Euclidean space, Canad. Math. Bull. 30 (1987), 436-445.
- [4] I. Bárány, P. Valtr, Planar point sets with a small number of empty convex polygons, Stud. Sci. Math. Hung. 41 (2004), 243-266.
- [5] O. Devillers, F. Hurtado, G. Károlyi, C. Seara, Chromatic variants of the Erdős-Szekeres Theorem. Computational Geometry, Theory and Applications, 26(3), (2003), 193-208.
- [6] A. Dumitrescu, Planar point sets with few empty convex polygons, Studia Sci. Math. Hungar. 36 (2000).
- [7] P. Erdős, G. Szekeres, A combinatorial problem in geometry, Compositio Math. 2 (1935), 463-470.
- [8] T. Gerken, Empty convex hexagons in planar point sets. Discrete and Computational Geometry, 39 (2008) 239-272.
- [9] H. Harboth, Konvexe Fünfecke in ebenen Punktmengen. Elem. Math. 33 (1979), 116-118.
- [10] J. D. Horton, Sets with no empty convex 7-gons, Canad. Math. Bull. 26 (1983), 482-484.
- [11] C. Huemer, C. Seara, 36 two-colored points with no empty monochromatic convex fourgons. (Submitted)

- [12] M. Katchalski, A. Meir, On empty triangles determined by points in the plane, *Acta. Math. Hungar.* 51 (1988), 323-328.
- [13] C. M. Nicolás, The empty hexagon theorem. *Discrete and Computational Geometry*, 38 (2007) 389-397.
- [14] R. Pinchasi, R. Radoičić and M. Sharir, On empty convex polygons in a planar point set, *J. Combin. Theory Ser. A* 113 (2006), 385-419
- [15] J. Pach, G. Tóth, Monochromatic empty triangles in two-colored point sets , in: *Geometry, Games, Graphs and Education: the Joe Malkevitch Festschrift* (S. Garfunkel, R. Nath, eds.), COMAP, Bedford, MA, (2008), 195–198.
- [16] P. Valtr , On empty hexagons. in: J. E. Goodman, J. Pach, and R. Pollack, *Surveys on Discrete and Computational Geometry, Twenty Years Later*, AMS, (2008) 433-441.
- [17] P. Valtr , On the minimum number of empty polygons in planar point sets, *Studia Sci. Math. Hungar.* 30 (1995), 155-163.

What does the convex hull of a discrete point set look like?

Ina Voigt

Faculty of Mathematics, TU Dortmund University
Vogelpothsweg 87, D-44227 Dortmund, Germany
ina.voigt@tu-dortmund.de

March 13, 2009

Abstract – One certainly knows, what the convex hull of a finite point set looks like - it's a polytope. But what does it look like if the point set is discrete but infinite? In this article, the convex hull of an infinite discrete point set in \mathbb{E}^n will be described. It will be shown that even all extreme points of its closure are contained in the discrete set. Furthermore, one will see that boundary regions, which have an empty intersection with the convex hull, are unbounded and that in each boundary point exists at least one direction in which the boundary is linear.

Keywords: discrete set, convex hull, extreme points, boundary segments

MSC 2000 classification: 52A20, 32F27

1 Introduction and Notation

If one reviews the literature on convex geometry (cf. e.g. [4, 5]), one notices that the considered convex sets are nearly always assumed to be compact or at least to be closed. But as one gets non-closed convex sets very natural as convex hulls of infinite discrete point sets, it is worthwhile to investigate these sets in more detail.

To get an idea of the shape of the convex hull of an infinite discrete point set, we want to have a look on its boundary. The first thing one thereby notices is that it has not to be closed. So it is interesting to examine the extreme points of its closure (see Section 2). Furthermore, we will investigate what kind of boundary regions can occur and show that those boundary regions, which have an empty intersection with the convex hull, are unbounded (see Section 3). Surely, the convex hull of a discrete set cannot be arbitrarily round. We will show that in each boundary point exists at least one direction in which the boundary is linear (see Section 4). These results are one part of the Ph.D. thesis [7].

Let \mathbb{E}^n be the Euclidean space of dimension n and denote the Euclidean norm by $\|\cdot\|$ and the standard scalar product by $\langle \cdot, \cdot \rangle$. We denote the *interior* of a set $M \subset \mathbb{E}^n$ by M° and the *boundary* by ∂M . We assign a dimension to a set M by $\dim M := \dim(\text{aff } M)$, the dimension of the affine space generated by M . Let $K \subset \mathbb{E}^n$ be a convex set. Then we say that $x \in K$ is an *extreme point* of K if x is not representable as a non-trivial convex combination of points of $K \setminus \{x\}$. We denote the set of all extreme points of K by $\text{Ext}(K)$. If a point $y \in K$ is a face of K , we call y an *exposed point* of K , and we denote the set of all exposed points of K by $\text{Exp}(K)$. Exists only one supporting hyperplane of K in $z \in K$, we call z a *regular point*. The *convex hull* of a set M is denoted by $\text{conv}(M)$. It is well known that for any set $M \subset \mathbb{E}^n$ holds $\text{Ext}(\text{conv}(M)) \subset M$.




Figure 1: Examples: $\text{conv}(\mathcal{P}_1)$, $\text{conv}(\mathcal{P}_2)$, $\text{conv}(\mathcal{P}_3)$

Furthermore, we know that for any convex set $K \subset \mathbb{E}^n$ holds $\text{Exp}(K) \subset \text{Ext}(K)$ and we have $\text{Ext}(K) \subset \overline{\text{Exp}(K)}$, if K is closed and convex. For more details about convex theory see e.g. [1]. We define a *polyhedron* as the intersection of finitely many half spaces. A *polytope* is a bounded polyhedron. Equivalently, we can describe each polytope as the convex hull of a finite point set. A set $\mathcal{P} \subset \mathbb{E}^n$ is called *discrete* if the intersection of \mathcal{P} with every bounded set is finite. For the boundary points of $\text{conv}(\mathcal{P})$ that are contained in the discrete set \mathcal{P} , we want to use the notation $\partial\mathcal{P} := \mathcal{P} \cap \partial\text{conv}(\mathcal{P})$.

We want to start with some first observations. If the point set is finite, we know that its convex hull is a closed polytope. But for an infinite discrete point set the convex hull can be non-polyhedral, as one can easily see by considering the set $\mathcal{P}_1 := \{(z, z^2) \mid z \in \mathbb{Z}\} \subset \mathbb{E}^2$, see Figure 1 (left). Furthermore, the convex hull of an infinite discrete point set has not to be closed. It is even possible that no point of the point set \mathcal{P} lies on $\partial\text{conv}(\mathcal{P})$:

Example 1.1 (point set without boundary points) Let $\mathcal{P}_2 := \{(n, 1 - 1/n) \mid n \in \mathbb{N}\} \cup \{(-n, -1 + 1/n) \mid n \in \mathbb{N}\} \subset \mathbb{E}^2$. Then we get $\text{conv}(\mathcal{P}_2) = \{x \in \mathbb{E}^2 \mid \langle x, v_i \rangle < 1, i = 1, 2\}$ with $v_1 := (0, 1), v_2 := (0, -1)$. Therefore, $\text{conv}(\mathcal{P}_2)$ is not closed and $\partial\mathcal{P}_2 = \emptyset$, see Figure 1 (middle).

By combining the above examples, one can also easily find a discrete point set, whose convex hull is neither closed nor is the closure of its convex hull polyhedral. Let $\mathcal{P}_3 := \{(n, n^2) \mid n \in \mathbb{N}\} \cup \{(-n, 1 + 1/n) \mid n \in \mathbb{N}\} \subset \mathbb{E}^2$. Then $\text{conv}(\mathcal{P}_3)$ is non-polyhedral and $\partial\text{conv}(\mathcal{P}_3) \not\subset \text{conv}(\mathcal{P}_3)$, see Figure 1 (right).

Furthermore, the discreteness of the point set yields directly that if its convex hull is not closed, then it has to be unbounded and the point set has to be infinite. To sum up these results, we already have the following properties:

Property 1.1 (obvious properties) Let $\mathcal{P} \subset \mathbb{E}^n$ be a discrete point set.

- (1) If \mathcal{P} is infinite, then $\text{conv}(\mathcal{P})$ has not to be closed.
- (2) If \mathcal{P} is infinite, then $\overline{\text{conv}(\mathcal{P})}$ can be non-polyhedral.
- (3) If $\text{conv}(\mathcal{P})$ is not closed, then \mathcal{P} has to be infinite and $\text{conv}(\mathcal{P})$ is unbounded. \square




Figure 2: Example 2.1: $\{(0, 1)\} = \text{Ext}(\text{conv}(\mathcal{P}_4)) \neq \text{Ext}(\overline{\text{conv}(\mathcal{P}_4)}) = \emptyset$

2 All extreme points are in \mathcal{P}

Now, we want to have a look at the extreme points. One should notice that generally, there is no inclusion between the set of extreme points of a convex set and the set of extreme points of its closure. Obviously, for $x \in \text{Ext}(\overline{K})$ with $x \in K$ holds $x \in \text{Ext}(K)$. But in general, an extreme point of \overline{K} has not to be contained in K . Considering for example $K = B^\circ(x, r) := \{y \in \mathbb{E}^n \mid \|y - x\| < r\}$ the open ball with radius r and center x , we see that $\text{Ext}(\overline{K}) \not\subset \text{Ext}(K)$. The next example shows that there also exist convex sets with $\text{Ext}(K) \not\subset \text{Ext}(\overline{K})$:

Example 2.1 (extreme point of conv, but not of the closure) Consider the discrete set $\mathcal{P}_4 := (\mathbb{Z} \times \{0\}) \cup \{(0, 1)\}$. Then we see that $(0, 1)$ is an extreme point of the convex set $\text{conv}(\mathcal{P}_4)$, but $(0, 1) \notin \text{Ext}(\overline{\text{conv}(\mathcal{P}_4)})$, see Figure 2.

But we know that for every set M holds $\text{Ext}(\text{conv}(M)) \subset M$. In the following, we will show that for a discrete set \mathcal{P} also holds $\text{Ext}(\text{conv}(\mathcal{P})) \subset \mathcal{P}$.

Lemma 2.1 (minimal distance) Let $K \subset \mathbb{E}^n$ be a convex set and $P \subset \overline{K}$ be a polytope. Then there exists always an extreme point of P in which the minimal distance between P and the boundary ∂K is attained.

Proof. Because P and ∂K are both closed and because P is bounded, there exist points $x_0 \in \partial K$ and $y_0 \in P$ in which the minimum distance is attained. We now consider the translation K_0 of the set K by $x_0 - y_0$. It holds $P \subset \overline{K_0}$ and $P \cap \partial K_0 \neq \emptyset$. Thus ∂K_0 contains a face and therefore an extreme point of P . As all points of $\partial K_0 \cap P$ attain the minimum distance between P and ∂K , the assertion follows. \square

Lemma 2.2 (exposed points) Let $\mathcal{P} \subset \mathbb{E}^n$ be a discrete point set. Then $\text{Exp}(\overline{\text{conv}(\mathcal{P})}) \subset \mathcal{P}$.

Proof. We show that a point $x \in \overline{\text{conv}(\mathcal{P})} \setminus \mathcal{P}$ is never an exposed point of $\overline{\text{conv}(\mathcal{P})}$.

As it is known that $\text{Ext}(\text{conv}(\mathcal{P})) \subset \mathcal{P}$, it suffices to show that no point $x \in \partial \text{conv}(\mathcal{P}) \setminus \text{conv}(\mathcal{P})$ is exposed for $\text{conv}(\mathcal{P})$.

Let $x \in \partial \text{conv}(\mathcal{P}) \setminus \text{conv}(\mathcal{P})$. Then there exists a sequence $(x_i)_{i \in \mathbb{N}} \subset \text{conv}(\mathcal{P})$ converging against x . CARATHEODORY's Theorem (see [2, 3]) yields that we can find for each point x_i of this sequence finitely many points $p_1(i), \dots, p_{k(i)}(i) \in \mathcal{P}$ such that $x_i \in \text{conv}(p_1(i), \dots, p_{k(i)}(i)) =: P_i$. As $(x_i)_{i \in \mathbb{N}}$ converges against x , the sequence $(x_i)_{i \in \mathbb{N}}$ converges against a supporting hyperplane H of $\overline{\text{conv}(\mathcal{P})}$ in x . Therefore, also the minimal distance between the corresponding polytopes P_i and H tends to zero. By the Minimal-Distance Lemma 2.1 we know, that for each polytope P_i the minimal distance to H is attained in an extreme point

Theorem 3.1 (polytopal boundary) *Let $\mathcal{P} \subset \mathbb{E}^n$ be a discrete set and $T \in \mathcal{T}^c(\text{conv}(\mathcal{P}))$ be maximal and bounded. Then it exists a polytope P such that T equals the union of a collection of faces of P .*

Proof. As $T \subset \text{conv}(\mathcal{P})$, one can find for each $x \in T$ finitely many points $p_1(x), \dots, p_{k(x)}(x) \in \partial\mathcal{P}$ such that x is representable as a convex combination of these points. Because $\text{conv}(p_1(x), \dots, p_{k(x)}(x)) \subset \partial\text{conv}(\mathcal{P})$, and because T is maximal, it follows $\text{conv}(p_1(x), \dots, p_{k(x)}(x)) \subset T$. Thus we obtain $\text{conv}(T) = \text{conv}(T \cap \mathcal{P})$. Since T is bounded and \mathcal{P} is discrete, $\text{conv}(T \cap \mathcal{P})$ is a polytope and the assertion follows. \square

Lemma 3.1 (unbounded face) *Let $\mathcal{P} \subset \mathbb{E}^n$ be discrete, $\text{conv}(\mathcal{P})$ not closed and $t \in \partial\text{conv}(\mathcal{P}) \setminus \text{conv}(\mathcal{P})$. Let H_t be a supporting hyperplane of $\text{conv}(\mathcal{P})$ in t . Then $H_t \cap \partial\text{conv}(\mathcal{P})$ is unbounded.*

Proof. Suppose, $T_t := H_t \cap \partial\text{conv}(\mathcal{P})$ is bounded. As $T_t = H_t \cap \overline{\text{conv}(\mathcal{P})}$ is the intersection of two closed convex sets, T_t is closed and convex. Thus, T_t is a compact convex set and it holds $T_t = \text{conv}(\text{Ext}(T_t))$. Because of $\text{Ext}(T_t) = \text{Ext}(\overline{\text{conv}(\mathcal{P})} \cap H_t) = \text{Ext}(\overline{\text{conv}(\mathcal{P})}) \cap H_t$, we get $\text{Ext}(T_t) \subset \text{Ext}(\overline{\text{conv}(\mathcal{P})})$. By the Extreme-Points Theorem 2.1 we know $\text{Ext}(\overline{\text{conv}(\mathcal{P})}) \subset \mathcal{P}$, what yields $\text{Ext}(T_t) \subset \mathcal{P}$. Thus it follows $T_t \subset \text{conv}(\mathcal{P})$, what is a contradiction. \square

Theorem 3.2 (unbounded boundary) *Let $\mathcal{P} \subset \mathbb{E}^n$ be discrete and $\text{conv}(\mathcal{P})$ not closed. Then, every $T \in \mathcal{T}^\varnothing(\text{conv}(\mathcal{P}))$ is unbounded.*

Proof. Let $T \in \mathcal{T}^\varnothing(\text{conv}(\mathcal{P}))$. Suppose, T is bounded. Then, for every $t \in T$, the set $T \cap H_t$ is bounded, where H_t denotes a supporting hyperplane of $\text{conv}(\mathcal{P})$ in t . By the Unbounded-Face Lemma 3.1, $H_t \cap \partial\text{conv}(\mathcal{P})$ is unbounded and thus contains a ray. Hence, if a ray $S \subset H_t \cap \partial\text{conv}(\mathcal{P})$ has a non-empty intersection with T , it has to contain at least one point of $\text{conv}(\mathcal{P})$. If we denote the set of all those rays by $\mathcal{S}_T := \{S \subset H_t \cap \partial\text{conv}(\mathcal{P}) \mid S \text{ ray, } S \cap T \neq \emptyset\}$, it follows $T \cap H_t \subset \text{conv}(S \cap \text{conv}(\mathcal{P}) \mid S \in \mathcal{S}_T)$ and thus we get $T \cap H_t \subset \text{conv}(\mathcal{P})$, a contradiction. \square

4 Line segments

We want to show that the boundary of the convex hull of a discrete set is in each point in at least one direction linear, that is, in each boundary point exists a supporting hyperplane of $\text{conv}(\mathcal{P})$ which intersects $\partial\text{conv}(\mathcal{P})$ in more than one point.

Definition 4.1 (linear segment) *Let $K \subset \mathbb{E}^n$ be closed and convex and $\partial K \neq \emptyset$. We say that ∂K contains a non-trivial linear segment in $x \in \partial K$ (or ∂K is in $x \in K$ in at least one direction linear) if a supporting hyperplane H_x of K exists such that $\{x\} \subsetneq K \cap H_x$.*

If $x \in \text{conv}(\mathcal{P}) \setminus \mathcal{P}$, there exist points $p_1, \dots, p_k \in \partial\mathcal{P}$ such that $x \in \text{conv}(p_1, \dots, p_k)$, and the existence of a non-trivial linear boundary segment follows. For a point $x \in \partial\text{conv}(\mathcal{P}) \setminus \text{conv}(\mathcal{P})$, the Unbounded-Face Lemma 3.1 yields directly the assertion. Thus, it remains to investigate the points $p \in \partial\mathcal{P}$.

Lemma 4.1 (*sets with regular exposed points*) Let $K \subset \mathbb{E}^n$ be a closed convex set in which exists at least one regular, exposed point $p \in \partial K$. Then

- (1) $\dim K = n = \dim(\text{aff } \partial K)$,
- (2) ∂K is connected,
- (3) K is line-free.

Proof. (1) Let $\partial K \neq \emptyset$ and $p \in \partial K$ regular and exposed. Thus, it exists exactly one supporting hyperplane H_p of K in p and $H_p \cap K = \{p\}$ is a face of K . This implies $\{p\} \subsetneq \partial K$ and $\partial K \not\subset H_p$. Hence it follows $\dim(\text{aff } \partial K) = n$ and therefore $\dim K = n$.

(2) It is known that the boundary of a convex set is not connected if and only if it is the union of two parallel hyperplanes, that is, if the convex set is a stripe. Since K contains at least one exposed point, K cannot be a stripe. Thus, ∂K is connected.

(3) Assume, K contains a line L . Then it follows that H_p contains a line parallel to L meeting p . But this is a contradiction, since p is exposed. \square

Lemma 4.2 (*extreme points accumulate*) Let $K \subset \mathbb{E}^n$ be a closed convex set and $p \in \partial K$ regular and exposed. Then, for all $\varepsilon > 0$ it exists an extreme point of K in $B(p, \varepsilon) \cap \partial K \setminus \{p\}$.

Proof. Here, we only want to sketch the idea of the proof. For the exact processing see [7]. Consider a sequence of boundary points $(q_i)_{i \in \mathbb{N}} \subset \partial K$, converging against p . Then, for each q_i , we can find a polytope P_i which contains q_i and whose vertices are extreme points of K or are contained in extreme rays of K . One can show that the P_i 's can be chosen in such a way that each of them has at least one extreme point of K as a vertex. As the sequence $(q_i)_{i \in \mathbb{N}}$ converges against p , the distance of the corresponding polytopes P_i to the supporting hyperplane H_p of K in p converges against zero. By the Minimal-Distance Lemma 2.1 we know that for each P_i the minimal distance to H_p is attained in a vertex of P_i . One further shows that all P_i can be constructed in such a way that the following holds: every P_i has at least one vertex e_i in which the minimum distance to H_p is achieved and which is an extreme point of K . Thus we have found a sequence of extreme points $(e_i)_{i \in \mathbb{N}}$ of K converging against H_p . Because p is regular and exposed, the sequence $(e_i)_{i \in \mathbb{N}}$ has to converge against p . \square

One could remark that the Extreme-Point-Accumulation Lemma 4.2 is a converse statement to the theorem by STRASZEWICZ [6] that for a closed convex set, each extreme point is the limit of a sequence of exposed points. Now, we have in addition that each regular exposed point is the limit of a non-trivial sequence of extreme points if the convex set is closed.

Corollary 4.1 (*no regular exposed points*) Let $\mathcal{P} \subset \mathbb{E}^n$ be discrete. Then $\text{conv}(\mathcal{P})$ contains no regular exposed points. \square

Thus, if $p \in \partial \mathcal{P}$ is regular, it is not exposed and hence the supporting hyperplane of $\text{conv}(\mathcal{P})$ in p contains additional points of $\partial \text{conv}(\mathcal{P})$. That is, $\partial \text{conv}(\mathcal{P})$ contains a non-trivial linear segment in p . On the other hand, if a

point $q \in \partial\mathcal{P}$ is not regular and we assume $\text{conv}(\mathcal{P})$ to be full dimensional, it follows that at least one supporting hyperplane H_q of $\overline{\text{conv}(\mathcal{P})}$ has to fulfill $\{q\} \subsetneq H_q \cap \partial\text{conv}(\mathcal{P})$ (see [7] for the details). Hence, $\partial\text{conv}(\mathcal{P})$ is in at least one direction linear in q , and the whole assertion follows.

Theorem 4.1 (*linear boundary segments*) *Let $\mathcal{P} \subset \mathbb{E}^n$ be discrete. Then $\partial\text{conv}(\mathcal{P})$ is in each point in at least one direction linear.* \square

5 Conclusion

We have gained a geometrical perception of the shape of the convex hull of an infinite discrete point set: We have revealed the fact that not only all extreme points of the convex hull but also all extreme points of its closure are contained in \mathcal{P} . We have proven that bounded boundary regions which are contained in the convex hull look polyhedral, and that those boundary regions which have an empty intersection with the convex hull are unbounded. And finally, we have established that each boundary point of the convex hull is contained in a non-trivial linear boundary segment.

References

- [1] A. Barvinok, *A Course in Convexity*, Graduate Studies in Mathematics, vol. **54**, Amer. Math. Soc., Providence, 2002.
- [2] C. Carathéodory, Über den Variabilitätsbereich der Koeffizienten von Potenzreihen, die gegebene Werte nicht annehmen, Math. Ann. **64** (1907), 95–115.
- [3] C. Carathéodory, Über den Variabilitätsbereich der Fourier’schen Konstanten von positiven harmonischen Funktionen, Rend. Circ. Mat. Palermo **32** (1911), 193–217.
- [4] P.M. Gruber, J.M. Wills, *Handbook of convex geometry*, Vol. A, North-Holland Publishing Co., Amsterdam, 1993.
- [5] P.M. Gruber, J.M. Wills, *Handbook of convex geometry*, Vol. B, North-Holland Publishing Co., Amsterdam, 1993.
- [6] S. Straszewicz, Über exponierte Punkte abgeschlossener Punktmenge, Fund. Math. **24** (1935), 139–143.
- [7] I.K. Voigt, Voronoizellen diskreter Punktmenge, Ph.D. thesis, TU Dortmund University, Faculty of Mathematics, Dortmund, 2008.

f -VECTORS AND RADON PARTITIONS

Jeong Hyeon Park

April 20, 2009

Abstract

We consider a simplicial polytope whose vertices are in general position. We give a characterization of its faces in terms of Radon partitions, an alternative proof of Barnettes lower bound Theorem for special cases. Moreover, we give a conjecture on the f -vectors of simplicial k -neighborly polytopes.

1 Introduction

The f -vector $(f_0, f_1, \dots, f_{d-1})$ of a d -polytope P is the integer vector with f_i equal to the number of faces of dimension i . There are many results on the f -vector of a d -polytope. The upper bound theorem[8] says the f -vector of a d -polytope is maximal, coordinate-wise if and only if the polytope is neighborly. That is if all the $\lfloor d/2 \rfloor$ -tuples of vertices define simplicial faces. On the other hand, The lower bound theorem[1] says the f -vector of a simplicial d -polytope is minimal if and only if the polytope is a stacked polytope.

One can easily see that a k -tuple S of vertices doesn't define a $k - 1$ -face if there is a Radon partition $(Z_+, Z_-), S = Z_+$. Now one may ask the following question : What's the relation between Radon partitions and f -vectors? In this project, we tried to answer this question. The goals of the project are

- (1) Characterization of faces in terms of Radon partitions
- (2) To find another proof of the lower bound theorem

2 Related Work

2.1 Radon Theorem and Oriented Matroids

In this project, the following old theorem are the most basic theorem.

Theorem 2.1.1 (Radon's Theorem). *Given set of points S in R^d with $|S| \geq d + 2$, S can be partitioned into (Z_+, Z_-) such that $\text{conv}(Z_+) \cap \text{conv}(Z_-) \neq \emptyset$*

This elementary theorem can be easily proved by affine dependency. And it has a lot of generalized theorems and related results. Among them the most interesting result is every point set can be represented in a combinatorial way with oriented matroids. For the definition of oriented matroids, see Ziegler[14] and Bjorner and et al. [5]. For every point set, the following holds

- (1) Minimal affine dependencies form *circuits*

(2) Minimal affine functions form *cocircuits*
Oriented matroids on the point set in R^d produced a lot of results :
see Ziegler[15].

2.2 Upper Bound on f -vectors

What is the upper bound of f -vectors? The following is the definition of the cyclic polytope whose f -vector is the maximum.

Definition 2.2.1. *Cyclic Polytopes* The moment curve in R^d is defined by

$$x : R \leftarrow R^d, \quad t \mapsto x(t) := \begin{pmatrix} t \\ t^2 \\ \vdots \\ t^d \end{pmatrix} \in R^d$$

The cyclic polytope $C_d(t_1, \dots, t_n)$ is the convex hull

$$C_d(t_1, \dots, t_n) := \text{conv}\{\mathbf{x}(t_1), \mathbf{x}(t_2), \dots, \mathbf{x}(t_n)\}$$

Another way to define the cyclic polytope is by Gale's evenness condition. The cyclic polytope $C_d(n)$ is $\lfloor \frac{d}{2} \rfloor$ -neighborly that is any subset S of the vertices with $|S| \leq \frac{d}{2}$ forms a face.

McMullen[8] proved f -vectors of d -polytope are bounded above by the f -vector of the cyclic polytope.

Theorem 2.2.2 (Upper bound theorem). *If P is a d -polytope with $n = f_0$ vertices, then for every k it has at most as many k -faces as the corresponding cyclic polytope :*

$$f_{k-1}(P) \leq f_{k-1}(C_d(n))$$

Here equality for some k with $\lfloor \frac{d}{2} \rfloor \leq k \leq d$ implies that P is neighborly.

In the proof of the theorem, the shelling order and the h -vector play important roles. A shelling of a polytope P is a linear ordering F_1, F_2, \dots, F_s of the facets of P such that $\bigcup_{i=1}^k F_i$ is *connected* for $1 \leq k \leq s$ in some sense (For the details, see Ziegler[14]). The h -vectors of a polytope P is a vector $\mathbf{h}(P) = (h_0, h_1, \dots, h_d)$ such that h_i is the number of new $i-1$ -faces in shelling P . The \mathbf{h} -vector has the following properties.

Observation 2.2.3. *We consider the f -polynomial*

$$\mathbf{f}(x) := f_{d-1} + f_{d-2}x + \dots + f_0x^{d-1} + f_{-1}x^d = \sum_{i=0}^d f_{i-1}x^{d-i}$$

and the h -polynomial

$$\mathbf{h}(x) := h_d + h_{d-1}x + \dots + h_1x^{d-1} + h_0x^d = \sum_{i=0}^d h_i x^{d-i}$$

Then,

$$\mathbf{h}(x) = \mathbf{f}(x-1)$$

And the following theorem says \mathbf{h} -vectors are symmetric[6][?].

Theorem 2.2.4 (Dehn-Sommerville equations). *The h -vectors of the d -polytope P satisfies*

$$h_k = h_{d-k} \quad \text{for } k = 0, 1, \dots, d$$

Based on these two observation and theorem, McMullen showed the following lemma which implies the upper bound theorem.

Lemma 2.2.5. *Let P be a simplicial d -polytope on $f_0 = n$ vertices. Then for $0 \leq k \leq d$,*

$$h_k(P) \leq \binom{n - d - 1 + k}{k}$$

Equality holds for all k with $0 \leq k \leq l$ if and only if $l \leq \lfloor \frac{d}{2} \rfloor$ and P is l -neighborly.

2.3 Lower Bound on f -vectors

The lower bound theorem says the f -vector of a simplicial polytope is larger than the f -vector of a stacked polytope. The following are the definitions of stack operations and stacked polytopes.

Definition 2.3.1 (Stacking onto a facet). $P' = P \cup (F * p)$ where p is beyond F but beneath all other facets of P . In other words, $P \cap (F * p) = F$ and all proper faces of F are faces of P'




Figure 1: Staking operation

Definition 2.3.2 (Stacked polytopes). A stacked d -polytope on $d + n$ vertices arises from a d -simplex by $(n - 1)$ -times stacking onto a facet.




Figure 2: A staked polytope

After stacking a vertex p , the f -vector of P' is

$$f_0(P') = f_0(P) + 1$$

$$\begin{aligned} f_i(P') &= f_i(P) + f_{i-1}(F) \quad \text{for } 0 \leq i \leq d-2 \\ f_{d-1}(P') &= f_{d-1}(P) + f_{d-2}(F) - 1 \end{aligned}$$

Thus, you can easily show that the f -vector of a stacked polytope P is

$$\begin{aligned} f_i &\geq \binom{d}{i} f_0 - \binom{d+1}{i+1} i \quad \text{for all } 1 \leq i \leq d-2 \\ f_{d-1} &\geq (d-1)f_0 - (d+1)(d-2) \end{aligned}$$

Barnette[1] proved the following lower bound theorem.

Theorem 2.3.3 (Lower bound theorem). *For a simplicial d -dimensional polytope, the following inequalities hold*

$$\begin{aligned} f_k &\geq \binom{d}{k} f_0 - \binom{d+1}{k+1} k \quad \text{for all } 1 \leq k \leq d-2 \\ f_{d-1} &\geq (d-1)f_0 - (d+1)(d-2) \end{aligned}$$

The proof of the theorem consists of three parts. First, the following theorem says it is enough to prove the inequality for f_1 .

Theorem 2.3.4. *If the lower bound theorem is true for edges of simplicial polytopes, then it is true for faces of all dimensions.*

Then, Barnette showed it's enough to show $f_1 \geq df_0 - g(f_0)$ where g is a function on d and $f_1 \geq df_0 - 2\binom{d+1}{2}$.

The g -theorem generalizes the upper bound theorem and the lower bound theorem. The g -vectors are defined as follows

Definition 2.3.5. *The g -vector of d -polytope P , $g(P) := (g_0, g_1, \dots, g_{\lfloor \frac{d}{2} \rfloor})$ with $g_0 := h_0 = 1$, $g_k = h_k - h_{k-1}$ for $1 \leq k \leq \lfloor \frac{d}{2} \rfloor$*

Billera and Lee[3][4] and Stanley[12][13] showed following theorem.

Theorem 2.3.6 (The g -theorem). *A sequence $\mathbf{g} = (g_0, g_1, \dots, g_{\lfloor \frac{d}{2} \rfloor}) \in N_0^{\lfloor \frac{d}{2} \rfloor + 1}$ is the g -vector of a simplicial d -polytope if and only if it is an M -sequence where an M -sequence is a f -vector of a multicomplex.*

From this theorem, Björner[2] derived the "McMullen correspondence" which implies the upper bound theorem and the lower bound theorem.

We have more a generalized lower bound conjecture on f -vectors of a simplicial d -polytope proposed by McMullen and Walkup[9].

Conjecture 2.3.7. *Let P be a simplicial d -polytope, then for $k = 0, 1, \dots, \lfloor \frac{d}{2} \rfloor + 1$, $\sum_{j=-1}^k (-1)^{k-j} \binom{d-j}{d-j} f_j(P) \geq 0$ with equality for some k if and only if there is a triangulation without interior $(d-k-1)$ -faces.*

Note that the necessity is proved by g -theorem[13].

3 Characterization of Faces

3.1 Characterization of Faces in terms of Radon Partitions

From now on, we consider d -polytopes whose vertices are in general position. That is, for a subset S , $|S| \leq d+1$ of the point set, S spans a $\dim(|S|-1)$ -affine subspace. That implies the polytopes are simplicial.




Figure 3: Illustration of Lemma 3.1.2

Lemma 3.1.1. *A subset S of vertices of P defines a simplicial face if and only if there is no Radon partition (Z_+, Z_-) with $Z_+ \subset S$.*

Proof. \Rightarrow) Suppose that there is a Radon partition (Z_+, Z_-) with $Z_+ \subset S$. Consider a point set $S', S \subset S', |S'| = d$ which forms a facet. ($\text{conv}(S)$ is a face of $\text{conv}(S')$) Then, $Z_+ \subset S'$ and $Z_- \cap \text{aff}(S')^- \neq \emptyset$. (The latter follows from the definition of facets)

Without loss of generality, let $\text{aff}(S') := \{x | x_d = 0\}$. Then, for $z \in Z_-$, if $z \in \text{aff}(S')^+$, $z_d > 0$. Thus, the Radon partition can be simplified to (Z_+, Z'_-) , $Z_+, Z'_- \subseteq S'$. But this contradicts the assumption that the points are in general position.

\Leftarrow) Suppose that S is not a face. Then, $\text{aff}(S) \cap \text{conv}(P - S) \neq \emptyset$. Otherwise, there is a separating hyperplane between $\text{aff}(S)$ and $\text{conv}(P - S)$, which implies S is the vertex set of a face.

For the remaining part of the proof, we need the following lemma.

Lemma 3.1.2. *Let S, T be point sets in R^d with $|S| + |T| \geq d + 2$. If $\text{aff}(S) \cap \text{conv}(T) \neq \emptyset$, there is a subset $T' \subseteq T$, with $\text{aff}(S) \cap \text{conv}(T') \neq \emptyset$, $|S| + |T'| \leq d + 2$.*

Proof. Choose affinely independent points S' in $\text{aff}(S)$ such that $|S'| = |S|$ and their convex hull contains $\text{aff}(S) \cap \text{conv}(T)$ in its relative interior. By Kirchberger's theorem[10], there is a subset $T' \subseteq T$, $|S'| + |T'| \leq d + 2$ with $\text{conv}(S') \cap \text{conv}(T') \neq \emptyset$. Note that we can't take a proper subset of S' otherwise the intersection would be empty. Then, $\text{aff}(S) \cap \text{conv}(T') \neq \emptyset$. \square

By the lemma, there is a subset $T' \subseteq P - S$ with $\text{aff}(S) \cap \text{conv}(T') \neq \emptyset$, $|S| + |T'| = d + 2$. Here, the equality follows from the general position assumption. Then, $|\text{aff}(S) \cap \text{conv}(T')| = 1$ by the general position assumption. Let $p = \text{aff}(S) \cap \text{conv}(T')$. $\text{aff}(S)$ is homeomorphic to $R^{|S|-1}$. Thus, $S \cup \{p\}$ has a Radon partition on $\text{aff}(S)$ because

$|S \cup \{p\}| = \dim(\text{aff}(S)) + 2$. Let (Z_+, Z_-) be the Radon partition. Then, either Z_+ or Z_- is contained in S . Without loss of generality, Z_- be such a set. Then, $(Z_+, Z_- - \{p\} \cup \text{conv}(T'))$ is a Radon partition in R^d with $Z_+ \subseteq S$. This contradicts the assumption. \square

By Lemma 3.1.1, we get the following corollary.

Corollary 3.1.3. *P is k -neighborly if and only if for every Radon partition (Z_+, Z_-) , $\min(|Z_+|, |Z_-|) > k$.*

Proof. "if" part is obvious. "only if" part is true by Lemma 3.1.1. \square

4 Lower bound theorem and Radon Partitions

4.1 \overline{f} -vectors

We introduce the notion of \overline{f} -vectors. The following is the definition.

Definition 4.1.1. *Let P be a polytope in d -dimensional space. Then, \overline{f}_k is the number of $k+1$ -tuples of the vertices of P that don't define k -faces.*

You may ask what the relationship between f -vectors and \overline{f} -vectors is. For example, for k -neighborly polytopes, $f_i = \binom{f_0}{i+1}$, $0 \leq i \leq k-1$ while $\overline{f}_i = 0$, $0 \leq i \leq k-1$. Then, what about f_k and \overline{f}_k ? The following lemma tells you about the relationship.

Lemma 4.1.2. *For a k -neighborly polytope P , $f_k = \binom{f_0}{k+1} - \overline{f}_k$.*

Proof. By Lemma 3.1.1, a $k+1$ -tuple T of the vertices of P doesn't define a k -face if and only if there is a Radon partition (Z_+, Z_-) such that Z_+ or Z_- are the same as T . \square

Then, we can reformulate the lower bound theorem in terms of \overline{f} -vectors.

Theorem 4.1.3 (Lower Bound Theorem Revisited). *For a simplicial polytope P in d -dimensional space, Barnette's lower bound theorem is true if and only if*

$$\overline{f}_1 \leq \binom{f_0 - d}{2}$$

Proof. Proof is just by simple calculation

$$f_1 = \binom{f_0}{2} - \overline{f}_1 \geq \binom{f_0}{2} - \binom{f_0 - d}{2} = \frac{f_0(f_0 - 1)}{2} - \frac{(f_0 - d)(f_0 - d - 1)}{2} = d \cdot f_0 - \binom{d+1}{2}$$

And and by Lemma 2.3.4, this inequality implies the lower bound theorem. \square

Our problem is to prove the inequality in the language of Radon partitions. The inequality, $\overline{f}_1 \leq \binom{f_0 - d}{2}$, looks more elegant and simpler than the original inequality. Our goal is to prove the inequality with the Gale transform.




Figure 4: Pictorial summary of the Gale transform

4.2 Gale Transform

In this section we introduce the notion of the Gale transform. For details see Matoušek[7] and Ziegler[14]. The Gale transform is a mapping from n points in R^d to n vectors in R^{n-d-1} which transforms Radon partitions in R^d to linear functions on R^{n-d-1} . The following is the definition.

Definition 4.2.1 (Gale Transform). *Let $\{x_1, \dots, x_n\}$ be points in R^d . Let P be a matrix in $R^{n \times (d+1)}$ such that*

$$P = \begin{pmatrix} 1 & \dots & 1 \\ x_1 & \dots & x_n \end{pmatrix}$$

Let P' be the matrix in $R^{(n-d-1) \times n}$ which is orthogonal to P and each column is linearly independent. Then, the Gale transform of $\{x_1, \dots, x_n\}$ is $\{g_1, \dots, g_n\}$ such that each g_i is a column of P' .

For convenience, we denote $g_i = g(x_i)$.

The Gale transform has the following property which is useful in proving the theorem.

Property 4.2.1. *Let $P = \{x_1, \dots, x_n\}$ be points in R^d . Then, there is a Radon partition (Z_+, Z_-) , $Z_+, Z_- \subseteq P$ if and only if there is a vector $c \in R^{n-d-1}$ such that*

$$c \cdot v = \begin{cases} 0 & \text{if } v \in g(P \setminus (Z_+ \cup Z_-)) \\ > 0 & \text{if } v \in g(Z_+) \\ < 0 & \text{if } v \in g(Z_-) \end{cases}$$

Example 4.2.2. *The following figure is about the Gale transform from a convex pentagon in R^2 . For every Radon partition in the primal space, there is a vector c which satisfies the condition specified in Property. For example, $(12, 34)$ forms a Radon partition while there is a linear hyperplane pass through $g(5)$ which separates $g(1), g(2)$ from $g(3), g(4)$.*

4.3 Proof of the Lower bound theorem for polytopes with $d + 3$ vertices

We can easily prove the lower bound theorem for polytopes with $d + 3$ vertices with the Gale transform.

Theorem 4.3.1. *For a simplicial polytope P with $d + 3$ vertices, $\overline{f_1} \leq 3$.*




Figure 5: Example of the Gale transform

Proof. Consider a Gale diagram of P . It is enough to count the number of pairs $(g(p_i), g(p_j))$ that can be separated from the other vectors in the Gale diagram. Let's assume that there are already three pairs of vectors, $(g(p_1), g(p_2))$, $(g(p_3), g(p_4))$, $(g(p_5), g(p_6))$, separated from the other vectors. Then, we have three cases.

i) $p_i \neq p_j$ for every i, j

Label the vectors $g(p_{i_{1,1}}) \dots g(p_{i_{3,2}})$ in counterclockwise order if necessary. If there is a vector between $g(p_{i_{j,2}})$ and $g(p_{i_{j+1,1}})$, $(g(p_{i_{j,1}}), g(p_{i_{j,2}}))$ or $(g(p_{i_{j+1,1}}, g(p_{i_{j+1,2}}))$ can't be separated from the other vectors. If there is a vector between $g(p_{i_{j,1}})$ and $g(p_{i_{j,2}})$, $(g(p_{i_{1,1}}), g(p_{i_{1,2}}))$ can't be separated from the other vectors.




Figure 6: Illustration of Case i)

ii) Two pairs have a common vector and $d \geq 3$ Without loss of generality, let $(g(p_1), g(p_2)), (g(p_2), g(p_3))$ be such pairs. Then, the other pair, which can be separated from the other vectors, should have $g(p_1)$ or $g(p_3)$. Suppose not. Let $(g(p_4), g(p_5))$ be such a pair. Then, there can't be any vectors between $g(p_1)$ and $g(p_4)$, $g(p_3)$ and $g(p_5)$. That implies $d = 2$, which contradicts the assumption.

Similarly, there can't be another pair which can be separated.

iii) Four pairs among five vectors

This is the case of convex pentagons in the plane.

□




Figure 7: Illustration of Case ii)




Figure 8: Illustration of Case iii)

5 Conjecture on f -vectors

5.1 Stacked Polytope

5.2 Conjecture on f -vectors of k -neighborly polytopes

Recall the stacking operation. In stacking a vertex p , p doesn't kill any faces except one facet. And stacked polytopes have the smallest f -vectors. In this section, we give a conjecture on k -neighborly polytopes based on a generalized construction.

Imagine that the following operation is possible

1) After add a new vertex p to a 2-neighborly polytope, that is still 2-neighborly.

2) p doesn't kill any faces except $d - 1$ -faces, $d - 2$ -faces.

We don't know whether we can do this operation but it feels like a stacking operation. Based on this operation, we can calculate the partial part of f -vectors as follows.

$$\begin{aligned}
f_i &\geq \binom{d+1}{i+1} + \sum_{i=d+2}^n \binom{d-1}{i-1} i - \binom{d}{i} \cdot (i-1) \\
&= \binom{d-1}{i-1} \binom{n}{2} - \binom{i-1}{1} \binom{d}{i} n + \binom{i}{2} \binom{d+1}{i+1} \quad 2 \leq i \leq d-3
\end{aligned}$$




Figure 9: Generalized staking operation

The term in the summation represents the f -vector of the vertex figure of p . We may generalize it to k -neighborly polytopes that gives the following conjecture.

Theorem 5.2.1. *For a simplicial d -dimensional k -neighborly polytope,*

$$\begin{aligned}
f_i &\geq \binom{d-k+1}{i-k+1} f_{k-1} - \binom{i-k+1}{1} \binom{d-k+2}{i-k+2} f_{k-2} \dots \\
&\quad + (-1)^k \binom{i}{k} \binom{d+1}{i+1} \quad \text{for all } k \leq i \leq d-k-1
\end{aligned}$$

Note that this conjecture doesn't tell anything about neighborly polytopes because we can't guess the remaining part of the f -vector. And we are curious about how this conjecture is related to the generalized lower bound conjecture.

6 Conclusion

Our goal was to prove the lower bound theorem with oriented matroid axioms without considering any geometric flavors. But now we believe it's not possible. There may be nonrealizable oriented matroids. It leads to the following question.

Question 6.0.1. *A collection $\mathcal{V} \subseteq \{+, -, 0\}^n$ is the set of circuits of an oriented matroid M and M satisfies the following conditions*

- (1) *Every circuit has k nonzero components*
- (2) *There is no circuit with only one $+$ component.*

Among such oriented matroids M , is there a nonrealizable oriented matroid?

One might extend the proof to the cases with $d+4$ vertices by the affine Gale transform. In this case, a naive case analysis would be painful because of many cases. If you find a combinatorial behavior of the point

set in the affine Gale diagram that may be helpful with proving general cases.

We couldn't guess the remaining part of the f -vectors on k -neighborly polytopes because we don't know how many $d - 1$ -faces, \dots , $d - k$ -faces a new vertex kill. Probably an analysis of k -stacked polytopes would give an estimate. And we are curious about the relation between the generalized lower bound conjecture and our conjecture. Now we believe those two conjectures are equivalent.

References

- [1] D.W. Barnette , A proof of the lower bound conjecture for convex polytopes , Pac. J. Math. 46, (1973) 349-354.
- [2] A. Björner, Partial unimodality for f -vectors of simplicial polytopes and spheres, in Jerusalem Combinatorics 93 (eds. H. Barcelo and G. Kalai), Contemporary Math. Series, 178, Amer. Math. Soc., (1994) 45-54.
- [3] L. J. Billera , C. W. Lee , A proof of the sufficiency of McMullen's conditions for f -vectors of simplicial convex polytopes , J. Comb. Theory, Ser. A 31, (1981) 237-255.
- [4] L. J. Billera , C. W. Lee, Sufficiency of McMullens conditions for f -vectors of simplicial polytopes, Bull. Amer. Math. Soc. 2 , (1980) 181-185.
- [5] A. Björner , M. Las Vergnas, B. Sturmfels, N. White and G. M. Ziegler, Oriented Matroids, Cambridge Univ. Press, (1993).
- [6] M. Dehn , Die Eulersche Formel im Zusammenhang mit dem Inhalt in der Nicht-Euklidischen Geometrie , Math. Ann.61, (1905) 561-586.
- [7] J. Matoušek , Lectures on discrete geometry , volume 212 of Graduate Texts in Mathematics. , Springer-Verlag, New York, (2002).
- [8] P. McMullen , The maximum numbers of faces of a convex polytope. Mathematika 17, (1970) 179-184.
- [9] P. McMullen , D.W. Walkup , A generalized lower bound conjecture for simplicial polytopes , Mathematika 18, (1971) 264-273.
- [10] P. Kirchberger , Über Tschebyscheffsche Annäherungsmethoden, Math. Ann. 57 (1903), 509-540.
- [11] D. M. Y. Sommerville , The relations connecting the angle-sums and volume of a polytone in space ofn dimensions , Proc. Roy. Soc. London, Ser. A 115, (1927) 103-119.
- [12] R. P. Stanley , The upper-bound conjecture and Cohen-Macaulay rings , Stud. Appl. Math. 54, (1975) 135-142.
- [13] R. P. Stanley , The numbers of faces of a simplicial convex polytope. Adv. Math. 35, (1980) 236-238.
- [14] G. M. Ziegler , Lectures on Polytopes, Springer-Verlag, New York, (1994).
- [15] G. M. Ziegler , Oriented matroids today, Electron. J. of Combin. 3 (1996).

Independence complexes of perfect graphs

Ragnar Freij, Matthias Henze

Abstract

The motivation behind this project is Gil Kalai's conjecture from 1989 that a centrally symmetric d -polytope must have at least 3^d faces. We study the centrally symmetric Hansen polytopes, associated to perfect graphs. Among them, we find an infinite family of polytopes with $3^d + 16$ faces.

1 Introduction

The present work is a summary of the authors investigations on a problem proposed by Günter M. Ziegler during the *I-Math Winter School - DocCourse Combinatorics and Geometry 2009* at the Centre de Recerca Matemàtica (CRM), Barcelona. The problem is based on a work by Sanyal, Werner and Ziegler [10] and continues the quest for examples with properties related to validity questions of crucial conjectures in polytope theory.

Our main concern is about *centrally symmetric* polytopes P which have the origin as the center of symmetry, i.e., $P = -P$. For a d -polytope P the number of its i -dimensional faces is denoted by $f_i(P)$ and $f(P) = (f_0(P), \dots, f_{d-1}(P))$ is called the *f-vector* of P . We write $s(P) = 1 + \sum_{i=0}^{d-1} f_i(P)$ for the total count of non-empty faces (P itself is also considered as a face). Gil Kalai posed the following central problem about this quantity which is well-known as the “ 3^d -conjecture”.

Conjecture 1.1 (Kalai [8], 1989). *Every centrally symmetric d -polytope P has at least 3^d non-empty faces, i.e., $s(P) \geq 3^d$.*

Although, our main focus is on combinatorial aspects, occasionally we will also talk about relations to the Mahler-conjecture. In order to state it, we need the concept of the *polar* of a given convex body $K \subset \mathbb{R}^d$ which is defined as

$$K^\star = \{y \in \mathbb{R}^d \mid x^\top y \leq 1, \forall x \in K\}.$$

Conjecture 1.2 (Mahler [9], 1939). *For centrally symmetric d -dimensional convex bodies K the volume product $\mathbf{M}(K) = \text{vol}(K)\text{vol}(K^*)$ is minimized by the unit cube $C_d = [-1, 1]^d$, i.e., $\mathbf{M}(K) \geq \mathbf{M}(C_d) = \frac{4^d}{d!}$.*

The connection between these two problems is made by *Hanner polytopes*, since they attain equality in each case. They were first considered by Hanner [5] in 1977 and are recursively defined as follows.

Definition 1.3 (Hanner polytopes). *Any line segment is a Hanner polytope. A d -polytope P with $d > 1$ is a Hanner polytope if it can be written as the cartesian product of two Hanner polytopes or as the polar of a Hanner polytope.*

Note, that the polar operation of taking cartesian products of two convex bodies $K \subset \mathbb{R}^k$ and $L \subset \mathbb{R}^l$ is the so called *free-sum* defined by $K \oplus L = \text{conv}\{K \times \{0\}^l, \{0\}^k \times L\}$.

All centrally symmetric polytopes considered in the sequel are instances of so called *Hansen polytopes* (see Section 2) arising from perfect graphs. Recall, that a graph G is called *perfect*, if the chromatic number of each induced subgraph of G equals the size of its maximal clique. A fundamental result on perfect graphs is the *Strong perfect graph theorem* which was a long standing conjecture by Berge from 1960 and eventually proved by Chudnovsky, Robertson, Seymour and Thomas in 2006.

Theorem 1.4 ([2]). *A graph G is perfect if and only if it neither contains odd cycles of length at least five nor its complements as induced subgraphs.*

As our main result, we will give a family of perfect graphs whose Hansen polytopes have f-vector sum $3^d + 16$. Furthermore, we will give some evidence that the Mahler volume of these polytopes tends to the lower bound in Conjecture 1.2 if the dimension tends to infinity.

This work is organized as follows. In Section 2 we introduce the reader to Hansen polytopes which are the central objects in consideration. The work of Ziegler et al (see [10]) motivated us to detailly study self dual Hansen polytopes in Section 3. The subsequent part identifies many Hanner polytopes among Hansen polytopes. These are essential for our construction of the family with small f-vector sum which will be explained in Section 5.

Acknowledgements. We like to thank Günter M. Ziegler for showing us this beautiful problem and for valuable discussions.

This work has been done while the two authors enjoyed the hospitality of the Centre de Recerca Matemàtica in Barcelona.

We are also grateful to M. Joswig and E. Gawrilow for the development of polymake, without its help a big deal of this work would not have been possible.

2 Hansen polytopes

Our main object of study will be the so called Hansen polytopes, first defined by Hansen [6]. They are interesting because of their appealing combinatorial description, but also because they have in some senses proved to have "few faces". More precisely, certain Hansen polytopes occur in [10] as counterexamples to Kalai's so called B- and C-conjectures (see [8]), that would have been stronger versions of the 3^d -conjecture.

A natural pattern for constructing centrally symmetric $(d+1)$ -polytopes is the following. Start with any d -polytope P and take the *twisted prism*. This means, embed P in the hyperplane $x_0 = 1$ and $-P$ in the hyperplane $x_0 = -1$, and take the convex hull, i.e.,

$$\text{tprism}(P) = \text{conv} \{ \{1\} \times P, \{-1\} \times (-P) \}.$$

In this way, we get a centrally symmetric polytope, each of whose faces is the convex hull of a face in P and another one in $-P$. The Hansen polytopes are a special case of this construction. Recall, that a stable set (independent set) in a graph G is a set of nodes, no two of which are connected by an edge and a clique in G is a set of nodes where any two are connected by an edge.

Definition 2.1 (Hansen polytopes). *Let $G = (V, E)$ be a finite graph on $n = |V|$ nodes. The stable set polytope $\text{Stab}(G)$ of G (see [4, Section 9.2]) is the convex hull of indicator vectors of stable sets in G , i.e.,*

$$\text{Stab}(G) = \text{conv} \left\{ \sum_{i \in I} \mathbf{e}_i \mid I \subseteq V \text{ a stable set in } G \right\}.$$

The Hansen polytope of G is the twisted prism

$$\text{Hans}(G) = \text{tprism}(\text{Stab}(G))$$

over the stable set polytope.

Therefore, we obtain the vertex description

$$\text{Hans}(G) = \text{conv} \left\{ \pm(\mathbf{e}_0 + \sum_{i \in I} \mathbf{e}_i) \mid I \subseteq V \text{ a stable set in } G \right\}, \quad (1)$$

and the vertices of the Hansen polytope are in two-to-one correspondence with the stable sets in G .

For any graph, it is clear that a clique and a stable set cannot meet in more than one point. So for any point $x \in \text{Hans}(G)$ and any clique $C \subseteq V$ in G , we have $-1 \leq -x_0 + 2 \sum_{i \in C} x_i \leq 1$. It is easily seen that there are no redundancies among these inequalities. Moreover, they give a full inequality description of $\text{Hans}(G)$ if and only if G is perfect, as proven by Hansen [6].

In the rest of this paper, if not stated otherwise, we will assume that G is perfect, so we get the facet description

$$\text{Hans}(G) = \left\{ x \in \mathbb{R}^{n+1} \mid -1 \leq -x_0 + 2 \sum_{i \in C} x_i \leq 1, C \text{ a clique in } G \right\}. \quad (2)$$

In particular, we have a two-to-one correspondence between facets of $\text{Hans}(G)$ and cliques in G .

Examining the vertex description (1) and the facet description (2) of $\text{Hans}(G)$, we get a complete combinatorial description of the vertex-facet incidences. Indeed, a facet of a Hansen polytope corresponds to a clique C together with a sign $\epsilon \in \{-1, 1\}$, and a vertex corresponds to a stable set I together with a sign δ . The vertex (I, δ) is included in the facet (C, ϵ) if $\delta = \epsilon$ and $C \cap I \neq \emptyset$, or if $\delta \neq \epsilon$ and $C \cap I = \emptyset$.

It follows immediately that every facet of $\text{Hans}(G)$ contains exactly half of the vertices, so $\text{Hans}(G)$ will be a so called *weakly Hanner polytope*, meaning that it is the twisted prism over any of its facets.




Figure 1: The cube and the crosspolytope as Hansen polytopes.

Let \overline{G} denote the graph complement of G . Then a stable set in \overline{G} is just a clique in G , and a clique in \overline{G} is a stable set in G . Therefore, the vertices of $\text{Hans}(G)$ are in natural correspondence to the facets of $\text{Hans}(\overline{G})$ and vice versa. By the combinatorial description of the vertex-facet incidences, these carry over under this natural correspondence. Since the vertex-facet incidences determine the combinatorial type, we get a combinatorial equivalence

$\text{Hans}(G)^* \simeq \text{Hans}(\overline{G})$. In particular, if $G \cong \overline{G}$ is a self-complementary graph, then $\text{Hans}(G) \simeq \text{Hans}(G)^*$ will be self-dual.

As an example, consider the complete graph K_n . A set of nodes in K_n is stable if only if it has cardinality 0 or 1. Thus, $\text{Stab}(K_n)$ is the standard full dimensional n -simplex. Hence, the Hansen polytope as the twisted prism over this simplex is an $(n + 1)$ -dimensional crosspolytope $\text{Hans}(K_n) \simeq C_{n+1}^*$.

Dually, any set of nodes in the empty graph \overline{K}_n is stable. Thus, $\text{Stab}(\overline{K}_n)$ is the n -dimensional 0/1-cube, and $\text{Hans}(\overline{K}_n)$ is (affinely equivalent to) an $(n + 1)$ -cube. Figure 1 shows the Hansen polytopes of the complete and empty graph on 2 nodes.

3 Self dual Hansen polytopes

This section is devoted to the study of Hansen polytopes that come from a special class of graphs. Recall that a graph G is said to be *self-complementary*, if it is isomorphic to its complement, i.e., $G \cong \overline{G}$. Note, that such graphs only exist on $4t$ and $4t + 1$ nodes, since they must have exactly half the number of possible edges. For a survey on self-complementary graphs see Farrugia [3]. In the previous section we have seen that the Hansen polytopes of those graphs are combinatorially self dual. Besides the occurrence of these polytopes in [10] the motivation to study this special class comes from the fact, that in dimensions five and six they minimize the gap of the f-vector sum to 3^d among all Hansen polytopes that do not attain equality.

To make a start, we investigate all self-complementary graphs on 8 nodes. Alter [1] characterized these graphs and also gives a list of the 10 examples. It turns out that 9 of them are perfect and with the help of `polymake` [7] we analyzed the corresponding Hansen polytopes in terms of the 3^d - and the Mahler-conjecture. The result is, that the graph in Figure 2 minimizes both values in this particular class (for exact data we refer to Figure 3).




Figure 2: “Minimal” self-complementary graph on 8 nodes.

We have drawn the picture in such a way that one easily recognizes the two involved fourpaths and we propose the following generalization to graphs on $4t$ and $4t + 1$ nodes.

Definition 3.1 (Iterated fourpath construction). *Let G_1 be the path on four nodes and for all $t \in \mathbb{N}, t \geq 2$, define G_t recursively by the disjoint union of G_{t-1} and G_1 and additional edges between all nodes of G_{t-1} and the middle nodes of G_1 having degree two.*

Moreover, define G'_t as the graph arising from G_t joined with an additional node that is connected to exactly those nodes of G_t that have degree $\geq 2t$.

In Figure 3 the first five instances¹ of this construction are given. We abbreviate the f-vector of the Hansen polytope of a graph $G = (V, E)$ on $|V(G)| = d - 1$ nodes by $f(G) = f(\text{Hans}(G))$ and the f-vector sum by $s(G) = s(\text{Hans}(G)) = 1 + \sum_{i=1}^{d-1} f_i(G)$. Furthermore, we give rounded values for the quotients $sq(G) = \frac{s(G)}{3^d}$ and $mq(G) = \frac{\mathbf{M}(\text{Hans}(G))}{\mathbf{M}(C_d)}$.

The recursive nature of G_t and G'_t allows us to explicitly count the number of vertices (which equals the number of facets) of the Hansen polytopes $\text{Hans}(G_t)$ and $\text{Hans}(G'_t)$. Note, that they live in dimension $4t + 1$ and $4t + 2$, respectively.

Proposition 3.2. *For all $t \in \mathbb{N}$ we have*

$$f_0(G_t) = f_{4t}(G_t) = \frac{2}{3}(7 \cdot 4^t - 4)$$

and

$$f_0(G'_t) = f_{4t+1}(G'_t) = \frac{2}{3}(7 \cdot 4^t - 4) + 2^{2t+1}.$$

Proof. Independent sets in G_t and vertices of $\text{Hans}(G_t)$ are in one-to-two correspondence. To count this number, say $I(G_t)$, we proceed by induction on t . For $t = 1$ we deal with the fourpath which has exactly $I(G_1) = 8 = \frac{1}{3}(7 \cdot 4 - 4)$ independent sets. Now, let $t > 1$ and count the independent sets in G_t by containment of nodes from the last added fourpath. Since, its two middle nodes are adjacent to each node of G_{t-1} there are only four independent sets in G_t that contain one of these. Conversely, the two outer nodes are not connected to any node of G_{t-1} which means that we come up with four (any subset of the two) times the number of independent sets in G_{t-1} . Thus, by induction hypothesis we have

$$I(G_t) = 4 \cdot I(G_{t-1}) + 4 = 4 \cdot \frac{1}{3}(7 \cdot 4^{t-1} - 4) + 4 = \frac{1}{3}(7 \cdot 4^t - 4).$$

The additional node in G'_t has no connection to the independent set of size $2t$ that comes from all the outer nodes of the involved fourpaths, but it is

¹We thank Ewgenij Gawrilow from TU Berlin for computing $\text{Hans}(G_3)$.

connected to any other node. Therefore, for any subset of the $2t$ nodes we get one additional independent set in G'_t which explains the summand 2^{2t+1} in the second equality, since we have to multiply by two by the one-to-two correspondence. \square

Although, it is possible to similarly count the number of independent sets of any given size i in G_t and G'_t , we do not have a formula to count the face numbers f_i of the corresponding Hansen polytopes for $i \neq 0$. Nevertheless, by the results in low dimensions we suspect them to be “minimal” in the following sense.

Conjecture 3.3. *The polytopes $\text{Hans}(G_t)$ and $\text{Hans}(G'_t)$ have the minimal f -vector sum and the minimal Mahler-volume among all self dual Hansen polytopes in their respective dimensions.*

Remark 3.4. *The speciality of these graphs is also known to graph theorists. Rao proved (see [3, Chapter 2]) that the graphs G_t and G'_t have the following extremal properties.*

- *There is no other graph with the same degree sequence.*
- *They maximize the number of triangles among all self-complementary graphs of the same order.*

As can be read off of the data in Figure 3 the third instance G_3 of the iterated fourpath construction has a worse f -vector quotient $sq(G_3)$ than G_1 and G_2 . Therefore it is reasonable to neglect the supplementary condition of self-complementarity to continue the search for even closer examples.

But, before we pay detailed attention to other graphs, we like to point out that the above construction is just a very particular case of a more general one.

To this end we need the notion of a *split graph*, which is a graph whose nodes can be partitioned into an induced clique and an induced independent set. It is well-known that these graphs belong to the class of perfect graphs. Let F_1 be an arbitrary perfect self-complementary graph and F_2 be a self-complementary split graph. We now define a new graph $G = [F_1, F_2]$ as the disjoint union of F_1 and F_2 with additional edges between any node in F_1 and any node of the clique in F_2 . As special instances we find $G_t = [G_{t-1}, G_1]$ and $G'_t = [G'_{t-1}, G_1]$.

The following statement shows how the properties of the graphs F_1 and F_2 carry over to G in this construction.


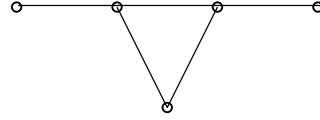
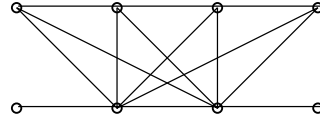
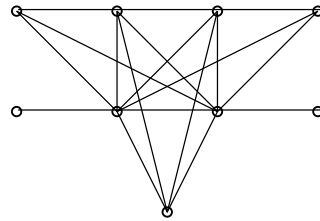
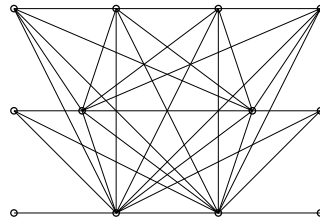

$f(G_1) = (16, 64, 98, 64, 16)$
$s(G_1) = 259 = 3^5 + 16$
$sq(G_1) \approx 1.06584, \quad mq(G_1) \approx 1.00833$

$f(G'_1) = (24, 116, 232, 232, 116, 24)$
$s(G'_1) = 745 = 3^6 + 16$
$sq(G'_1) \approx 1.02195, \quad mq(G'_1) \approx 1.00278$

$f(G_2) = (72, 596, 2216, 4652, 5922, 4652, 2216, 596, 72)$
$s(G_2) = 20955 = 3^9 + 1272$
$sq(G_2) \approx 1.064624, \quad mq(G_2) \approx 1.008667$

$s(G'_2) = 60361 = 3^{10} + 1312$
$sq(G'_2) \approx 1.022219, \quad mq(H(G'_2)) \approx 1.002977$

$s(G_3) = 1700611 = 3^{13} + 106288$
$sq(G_3) \approx 1.06666, \quad mq(G_3) \approx 1.0087256$

Figure 3: First examples of G_t and G'_t .

Proposition 3.5.

- (1) *If F_1 and F_2 are self-complementary and F_2 is a split graph, then $[F_1, F_2]$ is self-complementary.*
- (2) *If F_1 is perfect and F_2 is a split graph (thus also perfect), then $[F_1, F_2]$ is perfect.*

Proof. ad (1): To see that $G = [F_1, F_2]$ is self-complementary we only need to observe that the clique and the independent set in which F_2 splits up have to be of the same size when F_2 is self-complementary. The remainder follows immediately from the definition.

ad (2): For the perfectness of G we apply the strong perfect graph theorem (see Theorem 1.4) and show that there are no induced odd cycles of length at least five in G . This suffices because the complement of G is given by $\overline{G} = [\overline{F_1}, \overline{F_2}]$ and therefore has the same structure. Since, F_1 and F_2 were chosen to be perfect, we only have to consider induced subgraphs H with nodes in both of them. We observe that H can have at most two nodes in the split graph F_2 and that these have to be clique nodes, because otherwise we would get triangles in H . To complete an odd cycle the remaining nodes of H would induce a path of length at least three in F_1 . But then again we get triangles in H since the clique in F_2 is connected to any node in F_1 by construction. \square

Remark 3.6. *The definition of $G = [F_1, F_2]$ gives a recursive formula for the number of independent sets in G of any given size depending on those in F_1 and F_2 .*

4 Hanner graphs

Recall that one of our goals is to find Hansen d -polytopes with f-vector sum close to 3^d . Our approach is to start with polytopes with *exactly* 3^d non-empty faces, and then try to deform them in a way that does not change the face number too much. To this end, we make the following definitions.

Definition 4.1 (Hanner graph). *A graph G is a Hanner graph if $\text{Hans}(G)$ is a Hanner polytope.*

Definition 4.2 (recursive Hanner graph). *Denote by \circ the one node graph. The class of recursive Hanner graphs (RHG) is defined as follows.*

- 1. \circ is a RHG.

2. If G is a RHG, then \overline{G} and the disjoint union $G \sqcup \circ$ of G and \circ are RHGs.

The next lemma implies that Hansen polytopes of RHGs have exactly 3^d non-empty faces.

Lemma 4.3. *Any RHG is a Hanner graph.*

Proof. Clearly, the vertices of $\text{Hans}(\circ)$ are $\pm(1, 0)$ and $\pm(1, 1)$. So, $\text{Hans}(\circ)$ is affinely equivalent to a square, and hence a Hanner polytope. Hanner polytopes are closed under the polar operation and we have already seen that $\text{Hans}(\overline{G}) \simeq \text{Hans}(G)^*$. Thus, it suffices to show that $\text{Hans}(G \sqcup \circ)$ is a Hanner polytope whenever $\text{Hans}(G)$ is.

But a stable set in $G \sqcup \circ$ is just a stable set in G with or without the vertex of \circ added. So the vertices of $\text{Hans}(G \sqcup \circ)$ are

$$\left\{ \pm(\mathbf{e}_0 + \sum_{i \in I} \mathbf{e}_i), \pm(\mathbf{e}_0 + \sum_{i \in I} \mathbf{e}_i) + \mathbf{e}_{n+1} \mid I \text{ a stable set in } G \right\},$$

where \circ gets the label $n + 1$.

By the affine transformation $\mathbf{e}_0 \mapsto \mathbf{e}_0 - \mathbf{e}_{n+1}$, $\mathbf{e}_{n+1} \mapsto 2\mathbf{e}_{n+1}$, $\mathbf{e}_i \mapsto \mathbf{e}_i$ for $i = 1, \dots, n$, we get that $\text{Hans}(G \sqcup \circ) \simeq \text{Hans}(G) \times [-1, 1]$. This proves the lemma. \square

Consider the graph join $G * H$, i.e., the graph with all nodes and edges from G and H , and an edge between every node in G and every node in H . Note, that $G * H = \overline{\overline{G} \sqcup \overline{H}}$. Thus, an alternative way of defining recursive Hanner graphs is that they are closed under the operations $G \sqcup \circ$ and $G * \circ$.

It is trivial to see that neither of these operations destroy perfectness, so every RHG is perfect, and thus has a Hansen polytope with the facet description from (2).




Figure 4: The RHG corresponding to the sequence 10011.

We have a surjection from 01-sequences to RHGs, where a digit 1 corresponds to joining with \circ , and a digit 0 means taking the disjoint union

with \circ . Hence, sequences of length n corresponds to graphs on $n + 1$ nodes (the empty sequence corresponds to \circ). Moreover, every RHG either has an isolated node, or one that is connected to every other node, but not both. Deleting such a node gives a new RHG. This shows that the construction of a given RHG is unique, so the map from 01-sequences is bijective and so there are exactly 2^{n-1} RHGs on n nodes.

An example of a RHG is given in Figure 4, where the labels indicate in which order the nodes were added.

The corresponding Hansen polytopes are pairwise non-equivalent which means that we have at least 2^{d-2} Hansen d -polytopes that are Hanner polytopes. Since, for dimensions $d \geq 6$ there are more than these 2^{d-2} Hanner polytopes a natural question is the following.

Problem 4.4. *Are there any Hanner graphs that are not recursive Hanner graphs?*

From the duality between free-sums and cartesian products, it now follows that $\text{Hans}(G * \circ) \simeq \text{Hans}(G) \oplus I$. This is also easily seen geometrically, in Figure 5.




Figure 5: The polytopes $\text{Hans}(G)$, $\text{Hans}(G * \circ)$ and $\text{Hans}(G \sqcup \circ)$.

Finally, the following result on recursive Hanner graphs will be useful later.

Lemma 4.5. *If G is a RHG, then every induced subgraph of G is also a RHG.*

Proof. This is proven inductively over the number n of nodes in G . For $n = 1$, the statement is trivial. Suppose $n > 1$ and let $v \in G$ be the node added last in a recursive construction of G . Thus, $G \setminus \{v\}$ and all its induced subgraphs are RHGs. Now let H be any induced subgraph of G with $v \in H$. Then, v is connected to all or none of the other nodes of H , so either $H = (H \setminus \{v\}) \sqcup \{v\}$ or $H = (H \setminus \{v\}) * \{v\}$. Since, $H \setminus \{v\}$ is RHG by induction hypothesis, so is H and we are done. \square

5 Class of close examples

In the previous section we have seen that many Hanner polytopes are also Hansen polytopes. Here, we would like to investigate such Hansen polytopes that are almost Hanner polytopes in the sense that their graphs G are close to RHGs. This is formulated very vaguely, but the essence will be that there are only a few edges to remove from G in order to obtain a RHG. Our motivation is to somehow check if there are possibly counterexamples to the 3^d -conjecture in the “neighborhood” of Hanner polytopes and if not, to find classes of polytopes that exceed the conjectured lower bound very little.

Experiments that were done with extensive usage of `polymake` [7] did not come up with counterexamples but suggest that there could be classes of d -polytopes with f-vector sum $3^d + c$, where c is a constant independent on the dimension.

Based on these experiments we consider the following graphs. Given a RHG H_m on m nodes, we construct a graph $G[H_m]$ by adding an edge between any node of H_m and each of the two middle nodes of an adjoined fourpath (see Figure 6).




Figure 6: Constructing $G[H_m]$.

The corresponding Hansen polytope $\text{Hans}(G[H_m])$ lives in dimension $m + 5$ and is very close to a Hanner polytope, since removing the edge $\{v, w\}$ in $G[H_m]$ leads to the RHG $G'[H_m]$ on $m + 4$ nodes. For more convenience, we will deal for a moment with the complements of the graphs by considering the cliques as corresponding to vertices and the stable sets corresponding to facets of the Hansen polytope (recall that complementary graphs have combinatorially dual Hansen polytopes). In order to gain some geometric intuition for these polytopes we collect some nice properties that can easily be derived from the inequality description (2) of the involved Hansen polytopes.

Proposition 5.1. *Let H_m be a RHG on m nodes. Then*

$$i) \text{ Hans}(G[H_m]) = \text{conv} \{ \text{Hans}(G'[H_m]), \pm(\mathbf{e}_0 + \mathbf{e}_v + \mathbf{e}_w) \},$$

- ii) the vertices $\pm(\mathbf{e}_0 + \mathbf{e}_v + \mathbf{e}_w)$ see exactly two facets of $\text{Hans}(G'[H_m])$, i.e., the facets that correspond to the stable sets $\pm\{u, v, w\}$ and $\pm\{v, w\}$ of $G'[H_m]$,
- iii) any other facet of $\text{Hans}(G'[H_m])$ either contains the vertex $\mathbf{e}_0 + \mathbf{e}_v + \mathbf{e}_w$ in its affine hull or lies beyond.

Therefore, by considering $\text{Hans}(G[H_m])$ as in Proposition 5.1 i), we get exactly two new vertices and destroy exactly four facets. Furthermore, statement iii) suggests that most of the faces just get extended and thus the difference of the f-vector sums $s(G'[H_m])$ and $s(G[H_m])$ is supposed to be small. In fact, we have the following

Theorem 5.2. *Let H_m be a RHG on m nodes and let $d = m + 5$. Then*

$$s(G[H_m]) = 3^d + 16.$$

Proof. What we will show is, that there are constants c_1 and c_2 independent on m such that $s(G[H_m]) = c_1 \cdot 3^m + c_2$. The example of the fourpath (H_m is the empty graph) then will give the particular values of $c_1 = 3^5$ and $c_2 = 16$. We distinguish between the following two types of facets of $\text{Hans}(G[H_m])$, identified with their corresponding cliques in the graph.

- (1) $(X \cup C, \epsilon)$ for $X \subseteq \{x, v\}, \epsilon \in \{-1, 1\}$ and C a clique in H_m ,
- (2) (Y, ϵ) for $Y \in \{\{u\}, \{u, x\}, \{w\}, \{v, w\}\}$ and $\epsilon \in \{-1, 1\}$.

Each face of $\text{Hans}(G[H_m])$ can now be written as an intersection of some of these facets. Therefore, it suffices to show that there are $c'_1 \cdot 3^m + c'_2$ possible intersections and that we also counted $c''_1 \cdot 3^m + c''_2$ of them multiply. Again, the constants have to be independent on m .

To this end, consider the faces that can be written as intersections only of facets of type (1). A face F can be identified by a collection of signed cliques $(X \cup C, \epsilon)$ in $G[H_m]$. Let F^* be the face of $\text{Hans}(G_1)$ defined by the restrictions (X, ϵ) of these cliques to the adjoined fourpath G_1 . Similarly, let F_* be the face of $\text{Hans}(H_m)$ defined by their restrictions (C, ϵ) to H_m . We can choose the upper and the lower parts of the cliques independently. The upper part F^* will be defined by subsets of $\{x, v\}$, and the lower part F_* can be any face of $\text{Hans}(H_m)$. Since, H_m is a RHG, $\text{Hans}(H_m)$ has 3^{m+1} non-empty faces, so there are $3^{m+1} + 1$ choices for F_* . For each of these faces, F^* can be chosen in a constant number c of ways, not depending on m . This proves that the number of faces that are intersections between facets of type (1) is of the form $c \cdot 3^{m+1} + c$.

A similar argumentation can be done for faces that come as intersections of type (1) facets with type (2) facets. Moreover, the number of faces that arise only by intersections of type (2) facets does not depend on m at all, since for any H_m and any dimension there are exactly the eight of them mentioned above.

Finally, we have to consider the double counted faces. So fix a (possibly empty) collection \mathcal{F} of type (2) facets. We want to count the faces F that are double-counted by these. By this, we mean that there are collections \mathcal{D} and \mathcal{E} of type (1) facets, such that $F = \cap \mathcal{D} = (\cap \mathcal{F}) \cap (\cap \mathcal{E})$.

All vertices in H_m will have the same incidences with \mathcal{F} . With notations from before, we thus see that for any such \mathcal{F} we get some set of faces F^* that depends on \mathcal{F} , but we can choose F_* independently. Again, since $\text{Hans}(H_m)$ is Hanner, it follows that the number of faces that are double-counted by \mathcal{F} is of the form $c \cdot 3^m + d$. Therefore, the face number of $\text{Hans}(G[H_m])$ equals $c_1 \cdot 3^m + c_2$ for some constants c_1 and c_2 that do not depend on H_m as desired. \square

Remark 5.3. *In Section 4 we have seen that there are 2^{m-1} non-isomorphic RHGs on m nodes and therefore the above construction leads to an exponential number of examples with f-vector sum $3^d + 16$. In particular, there are 2^{m-2} of them, since the fourpath is self-complementary and therefore $G[H_m] \cong G[\bar{H}_m]$.*

The intuition that only very low and very high dimensional faces of $\text{Hans}(G[H_m])$ get affected when we move over to $\text{Hans}(G[H_m])$ gets further support by the following observation. The difference of 16 in the f-vector sum seems to have the same distribution over the dimensions. More precisely,

$$f(G[H_m]) - f(G'[H_m]) = (-4, -6, -2, 0, \dots, 0, 4, 12, 10, 2),$$

for any choice of H_m .

m	$mq(G^m)$
0	$= \frac{121}{120} \approx 1.00833$
1	$= \frac{361}{360} \approx 1.00278$
2	$= \frac{841}{840} \approx 1.00119$
3	$= \frac{1681}{1680} \approx 1.000595$
4	$= \frac{3025}{3024} \approx 1.000331$
5	$= \frac{5041}{5040} \approx 1.000198$

Figure 7: Mahler quotients of G^m in small dimensions.

Moreover, the Mahler-volume of $\text{Hans}(G[H_m])$ is very close to the conjectured lower bound and it also seems to be independent on the chosen RHG H_m as the data in Figure 7 indicates. Here, we denote with G^m any graph that arises from the above construction with attached RHG on m nodes. Based on these particular values we conjecture for the Mahler-volume of $\text{Hans}(G^m)$ that

$$mq(G^m) = \frac{(m+5)! + (m+1)!}{(m+5)!}$$

and therefore

$$\mathbf{M}(G^m) = \mathbf{M}(C_{m+5}) + \frac{4^{m+5}(m+1)!}{(m+5)!^2}.$$

As a final remark, we like to point out that further experiments suggest that there might be a generalization of Theorem 5.2. Replace the fourpath in the above construction by any split graph S_k on k nodes and for any given RHG H_m on m nodes consider the graph $G[H_m, S_k]$ which comes up as the disjointed union of H_m and S_k with additional edges connecting any node in H_m with any clique node in S_k . For example, we have $G[H_m] = G[H_m, G_1]$, where G_1 denotes the fourpath as before.

Conjecture 5.4. *Let S_k be a split graph on k nodes with $s(S_k) = 3^{k+1} + c(S_k)$ for some constant $c(S_k)$. Then, for any RHG H_m on m nodes*

$$s(G[H_m, S_k]) = 3^{m+k+1} + c(S_k),$$

i.e., the additive constant in the f-vector sum doesn't change under the operation of joining RHGs to the clique in S_k .




Figure 8: Another split graph S_5 .

For instance, we computed the f-vector sum of the Hansen polytopes for the graphs $G = [H_m, S_5]$ when $m \leq 4$. Here, S_5 is given by the split graph in Figure 8 and the constant is $c(S_5) = 64$.

References

- [1] R. Alter, *A characterization of self-complementary graphs of order 8*, Port. Math., **34** (3) (1975), 157-161.

- [2] M. Chudnovsky, N. Robertson, P.D. Seymour, R. Thomas, *The strong perfect graph theorem*, Ann. Math. **164** (2006), 51-229.
- [3] A. Farrugia, *Self-complementary graphs and generalisations: a comprehensive reference manual*, University of Malta, (1999), <http://www.alastairfarrugia.net/sc-graph/sc-graph-survey.pdf>.
- [4] M. Grötschel, L. Lovász, L. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, Algorithms and Combinatorics 2, Springer-Verlag, Berlin Heidelberg 1988, 2nd ed. 1993.
- [5] O. Hanner, *Intersections of translates of convex bodies*, Math. Scand. **4** (1956), 65-87.
- [6] A. B. Hansen, *On a certain class of polytopes associated with independence systems*, Math. Scand., **41** (1977), 225-241.
- [7] M. Joswig, E. Gawrilow, *Geometric reasoning with polymake*, <http://arxiv.org/abs/math.CO/0507273>.
- [8] G. Kalai, *The number of faces of centrally-symmetric polytopes*, Graphs and Combinatorics, **5** (1989), 389-391.
- [9] K. Mahler, *Ein Übertragungsprinzip für konvexe Körper*, Casopis Pest. Mat. Fys. **68** (1939), 93-102.
- [10] R. Sanyal, A. Werner, G. M. Ziegler, *On Kalai's conjectures concerning centrally symmetric polytopes*, Discrete Comput. Geom., **41** (2009), 183-198.

The Last Delaunay Triangulation

Víctor Álvarez, Frederik von Heymann, Canek Peláez

April 20, 2009

1 Introduction

During the DocCourse 2009 in Barcelona we were presented with the following problem: we know that the Delaunay graph of a set of n points (n pair) in the plane in general position always contains a perfect matching [2]. The proof, by Michael Dillencourt, relies heavily on the properties of the Euclidean metric, so the question remains open for other popular metrics.

Last year Ábrego et al. proved that for the metrics L_1 and L_∞ the Delaunay graph of points in the plane not only contains a perfect matching; the set of points always contains a Hamiltonian path [1].

The problem presented in the DocCourse was: Is it true that the Delaunay graph admits a perfect matching for all the L_p metrics?

Although interested in the problem, we faced a little difficulty with it: we were not able to even visualize how the Delaunay graph looked like in metrics different to L_1 , L_2 and L_∞ . So we first started to try and understand what happens to the Delaunay graph in arbitrary metrics L_p , and this led us to a different problem, which we called (by a suggestion of Ferran Hurtado) *The Last Delaunay Triangulation*.

2 Definitions

The Euclidean metric is defined by the usual distance function

$$|(x_1, y_1), (x_2, y_2)|_2 = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2} = (|x_2 - x_1|^2 + |y_2 - y_1|^2)^{\frac{1}{2}}.$$

If we define this distance function as L_2 , we can generalize it to a family of metrics L_p :

$$|(x_1, y_1), (x_2, y_2)|_p = (|x_2 - x_1|^p + |y_2 - y_1|^p)^{\frac{1}{p}}.$$

This family of metrics are sometimes called Minkowski metrics. We define L_∞ in the obvious way as

$$|(x_1, y_1), (x_2, y_2)|_\infty = \max(|x_2 - x_1|, |y_2 - y_1|).$$




Figure 1: Spheres in different L_p metrics.

We can think of how the different L_p metrics behave by looking at the unit sphere in each of them. For L_2 it's a circle, obviously, and it starts to look more and more like a diamond when p goes to 1, and more and more like a square when p goes to ∞ (Figure 1).

For $1 < p < \infty$ every point in a sphere has exactly one support line, which is not true for L_1 nor L_∞ because they have sharp corners; the spheres in L_p with $1 < p < \infty$ are a smooth curve in all of their points. Also, if two spheres in L_p intersect each other, they do it in one or two points (or they are the same ball); again, this is not true for L_1 nor L_∞ because they contain segments of lines (Figure 2).

Because of these properties, three non-collinear points determine exactly one sphere that passes through them in L_p , for $1 < p < \infty$. Once more, this is not true for neither L_1 nor L_∞ ; In fact, in those metrics there are many sets of three non-collinear points, for which there is no sphere passing through all of them.

Given a set P of n points in the plane in general position, we define the Delaunay graph $DG(P)$ of P as the graph with P as its vertex set, and where two points are connected by an edge if there is a ball (given a fixed metric) that covers them but no other point in P .

Obviously it is equivalent to consider empty spheres instead of balls, where a sphere S is called empty if there are no elements of P in the bounded region of S .

For L_2 the Delaunay graph is a triangulation, and in fact it's a triangulation for all L_p with $1 < p < \infty$. For L_1 and L_∞ it is not a triangulation, and even more, it's possible that it does not contains the convex hull of P (Figure 3).

Also we observe that for every edge in a Delaunay triangulation we find a




Figure 2: Differences between L_1 , L_∞ and L_p with $1 < p < \infty$.


triangle containing this edge, such that there is an empty sphere through the vertices of the triangle. Therefore the Delaunay triangulation can be expressed by the collection of empty spheres passing through three points of P .

3 Delaunay triangulations in L_p


Playing with different examples of sets of points and what happens when we increase the p in L_p starting from 2, we observed the following: A triangulation remains unchanged until one of the empty spheres that passes through three points of P suddenly touches a fourth point. In this moment technically we loose the general position because this sphere passes through four points, but we ignore that and keep growing p . What happens then is that we loose one of the edges of the triangulation, and we win the crossing edge that appears with the fourth point (Figure 4).

The interesting moment is when the sphere (actually two spheres converging into one) passes through these four points. Technically we have a crossing or an empty quadrangle there, but we take only one of the diagonal edges of the quadrangle so to get only triangulations for all $1 < p < \infty$. The two original spheres converge to a unique sphere that passes through the four points, and then we have our flip (and consequently two other spheres: we loose the other two because they now enclose other points of P). This keeps happening for other quadrangles in P as we grow p : we have flips in our Delaunay triangulations.

An interesting but erroneous idea that we had was that maybe when we loose an edge with a flip, that edge does not appear again in any other metric.



(a) Delaunay triangulation in L_2 , with Voronoi diagram in dashed lines.



(b) Non convex Delaunay graph in L_1 with the same set.

Figure 3: Delaunay graphs in the same set of points with different metrics.




Figure 4: A flip between two different Delaunay triangulations of the same point set.

Unfortunately this is not true: we have an example where we have a flip and then for a larger p we have another flip (back to the first configuration) in the same quadrangle (Figure 5).

But what does happen is that there comes a time when the spheres in our metric L_p are so close to a square that no matter how much we keep growing p , another flip is not possible.

And then what we have is The Last Delaunay Triangulation.




Figure 5: An edge that reappears when we grow p in L_p .




Figure 6: How the centres diverges or converges when p grows.

4 The Last Delaunay Triangulation

Instead of focusing on the empty spheres of our Delaunay triangulations, we are going to look at the centers of those spheres. As we grow p , the centers of those spheres start to move as curves; we have a flip every time these curves meet (simultaneously for the same p).

It is obvious that a curve described by a center does not intersect itself, and we haven't proved but we are convinced that the curves have the following behavior (up to symmetry): Take two of the three points that determine a sphere (without loss of generality let's assume that they do not lie on the same horizontal line), and consider the rectangle with these two points in its corners. If the third point is inside this rectangle, the curve described by the spheres passing through the three points will diverge asymptotically to a line with slope 1 or -1. If the third point is outside this rectangle, the curve will be bounded (Figure 6).

So when p is sufficiently large, the centers of the spheres will converge to a fixed point, or they will move away from the set in a stable way. There is no way in which they could intersect each other, and so the Delaunay triangulation will remain the same until p "reaches" ∞ .

That is The Last Delaunay Triangulation, and the name fits because in ∞ (as we already saw) there is the possibility that the Delaunay graph is not a triangulation. Although we have sketches of proofs for all of our claims, we will need to write them down carefully and formally before we present them.

5 Conjectures and future work

We believe that the Delaunay graph in L_∞ is a subgraph of The Last Delaunay Triangulation, and we are trying to prove it. This would give us an efficient

way to obtain The Last Delaunay Triangulation, because we will only need to calculate the Delaunay graph in L_∞ and then complete the triangulation (in the right way, of course).

As we saw, it's possible that a triangulation repeats itself, for growing p . How many times can this happen? How many times can an edge disappear and reappear later?

We also would like to know what is the maximum number of triangulations that a set of points can realize when we grow p starting in 2; we believe that this number of triangulations is polynomial in the number of points in the set, and that it will be n^2 or n^3 .

Much of the initial insight in this problem came after we programmed a little algorithm to calculate the circumcenter of three points in L_p for arbitrary p . This allowed us to actually see how a given Delaunay triangulation changes when p is growing; unfortunately our algorithm is not very fast, and we could only work with small point sets. Although not necessary for the main result and the conjectures we are investigating, it would be nice to come up with a better algorithm that would allow us to see how larger sets of points behave when growing the p .

Acknowledgments

We thank the organizers of the DocCourse and the CRM for supporting us and our research in the period Jan-Mar 2009. It was a wonderful experience and we hope this program can be continued for future generations of PhD-students and young researchers.

We are particularly grateful for the help and assistance of Ferran Hurtado.

References

- [1] B. Ábrego, E. Arkin, S. Fernández, F. Hurtado, M. Kano, J. Mitchell and J. Urrutia, *Matching points with squares*, Discrete & Computational Geometry, 41(1), pp. 77-95, January 2009.
- [2] M. B. Dillencourt, *A Non-Hamiltonian Delaunay Triangulation*, Information Processing Letters, 25(3), pp. 149-151, 1987.
- [3] M. B. Dillencourt, *Toughness and Delaunay Triangulations*, Discrete and Computational Geometry, 5(6), pp. 575-601, 1990.
- [4] D. T. Lee *Two-Dimensional Voronoi Diagrams in the L_p -Metric* Journal of the Association for Computing Machinery, 27(4), pp. 604-618, 1980.
- [5] Okabe A. Okabe, B. Boots and K. Sugihara, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Wiley, Chichester, UK (1992).

Lattice paths and Lagrangian matroids

Anna Gundert

Edward D. Kim*

Daria Schymura†

Abstract

Lagrangian matroids are a class of Coxeter matroids, an extension of ordinary matroids. In this paper, we investigate Lattice Path Lagrangian matroids (LPLMs), a family of Lagrangian matroids introduced by J. Bonin and A. de Mier. One definition for Lagrangian matroids involves ordinary matroids. Our main result states that for LPLMs the matroids arising from this construction are, as one would hope, lattice path matroids.

1 Introduction

Lagrangian matroids are a class of Coxeter matroids, an extension of ordinary matroids. While the precise definition of Coxeter matroids is rather intricate, many subclasses, such as Lagrangian matroids, of Coxeter matroids can be defined directly in terms of ordinary matroids or exchange axioms. In this paper, we will be studying the class of lattice path Lagrangian matroids (LPLMs) which was introduced by J. Bonin and A. de Mier. Coxeter matroids in general and also Lagrangian matroids are surveyed in [3].

Lagrangian matroids are families of transversals of the set $[n] \cup [n]^*$ satisfying a certain exchange axiom. Here, a transversal is a set $T \subset [n] \cup [n]^*$ with $|T \cap \{i, i^*\}| = 1$ for all $i \in [n]$. The precise definition of Lagrangian matroids will be presented in Section 2, along with the definition of LPLMs and some basic notation and terminology. We will see that it is possible to characterize Lagrangian matroids via sets of intersections with arbitrary transversals, which turn out to be ordinary matroids.

In Sections 3 to 5 we study the ordinary matroids that arise in this context when considering LPLMs. Our main result, Theorem 3.1, states that these are lattice path matroids, a class of ordinary matroids investigated in [1] and [2]. Section 3 explores the question and contains a proof that this property does not characterize LPLMs. Sections 4 and 5 contain two proofs of Theorem 3.1.

In the remaining sections, we explore other properties of lattice path Lagrangian matroids. One can define for Lagrangian matroids most of the usual matroid theoretic operations. In Section 6, we investigate for some of these operations whether they preserve lattice path Lagrangian matroids. Section 7 presents some open questions on LPLMs.

2 Basic Definitions

In this section, we give basic definitions and notation. First, we recall the definition of a matroid.

*Research partially supported by VIGRE NSF grant DMS-0135345

†Research partially supported by Deutsche Forschungsgemeinschaft Research Training Group *Methods for Discrete Structures*

2.1 (Ordinary) Matroids

Let E be a finite set and let \mathcal{B} be a non-empty family of subsets of E such that the following exchange property holds: If $B_1, B_2 \in \mathcal{B}$ and $x \in B_1 \setminus B_2$, then there exists a $y \in B_2 \setminus B_1$ such that $(B_1 \setminus \{x\}) \cup \{y\} \in \mathcal{B}$. Then, the pair $\mathcal{M} = (E, \mathcal{B})$ is called a *matroid*, each set $B \in \mathcal{B}$ is called a *basis*, and each subset $I \subseteq B \in \mathcal{B}$ is called an *independent set*. It is easy to see that all bases have the same size. In Section 3.1, an essential example of a matroid for this article will be presented: the class of lattice path matroids.

Two matroids $\mathcal{M}_1 = (E_1, \mathcal{B}_1)$ and $\mathcal{M}_2 = (E_2, \mathcal{B}_2)$ are *isomorphic* if there is a bijective map $\phi : E_1 \rightarrow E_2$ that maps bases of \mathcal{M}_1 to bases of \mathcal{M}_2 and vice versa: For all $B_1 \subseteq E_1 : B_1 \in \mathcal{B}_1 \Leftrightarrow \phi(B_1) \in \mathcal{B}_2$.

More about matroids can be found e.g. in [6].

2.2 Lagrangian matroids

Denote by $[n]^*$ the set $\{1^*, \dots, n^*\}$ and by S the set $[n] \cup [n]^*$. We consider the star operation as a function $*$: $S \rightarrow S$ by $(i^*)^* = i$. A subset $B \subset S$ is a *transversal* of S if, for all $i \in [n]$, B contains i or i^* , but not both. For two transversals of S , we say that i is a *divergence* if one of the transversals contains i and the other contains i^* . A set \mathcal{B} of transversals of S is the set of bases of a *Lagrangian matroid* if the following *symmetric exchange axiom* [3, Theorem 4.1.4] holds:

For all transversals $X, Y \in \mathcal{B}$ and each divergence i , there is a divergence j such that $X \Delta \{i, i^*, j, j^*\} \in \mathcal{B}$.

The symbol Δ denotes the symmetric difference. Note that the case $i = j$ is not excluded in the definition, so it is allowed just to replace i by i^* . Furthermore, note that a Lagrangian matroid is not an ordinary matroid. We will often simply refer to \mathcal{B} as a Lagrangian matroid.

Two Lagrangian matroids $\mathcal{B}_1, \mathcal{B}_2$ are *isomorphic* if there exists a map $\phi : [n] \cup [n]^* \rightarrow [n] \cup [n]^*$ with $\phi(i^*) = \phi(i)^*$ and $\phi(i) \in \{i, i^*\}$ for all $i \in [n] \cup [n]^*$, such that $B \in \mathcal{B}_1 \Leftrightarrow \phi(B) \in \mathcal{B}_2$. Such a map is determined by the transversal $T = \phi([n])$: Let $i \in [n]$. If $i \in T$ we have $\phi(i) = i$, otherwise $\phi(i) = i^*$. By definition, $\phi(i^*) = \phi(i)^*$. Every transversal determines such a map.

The following theorem gives a characterization of Lagrangian matroids in terms of ordinary matroids.

Theorem 2.1. [3, Theorem 4.1.1] *Let \mathcal{B} be a collection of transversals of S . For all transversals T of S define*

$$\mathcal{I}_T = \{I : I \subseteq B \cap T \text{ for some } B \in \mathcal{B}\}.$$

Then, \mathcal{B} is a Lagrangian matroid if and only if \mathcal{I}_T is the set of independent sets of an ordinary matroid on T .

Lagrangian matroids are e.g. studied in Chapters 3 and 4 of [3], in [5] and in [4].

2.3 Lattice Path Lagrangian Matroids

Fix $n \in \mathbb{N}$. An *n -path* S is any n -tuple $(S(1), \dots, S(n)) \in \{+1, -1\}^n$. The symbol $S(i)$ is used to denote the i^{th} element in the n -tuple S . Where there is no ambiguity, we simply call S a *path*. To every path S , we associate a *height function* $h_S : \{0, \dots, n\} \rightarrow \mathbb{Z}$ such that

$$h_S(i) = \begin{cases} 0 & \text{if } i = 0, \\ h_S(i-1) + S(i) & \text{if } i > 0. \end{cases}$$

We note that i and $h_S(i)$ have the same parity for every $i \in \{0, \dots, n\}$. Where convenient, we use \nearrow for $+1$ and \searrow for -1 , which is motivated by the shape of the height function h_S . An n -path S also corresponds to a transversal T of $[n] \cup [n]^*$ by $i \in T \Leftrightarrow S(i) = +1$ and $i^* \in T \Leftrightarrow S(i) = -1$.

If P and Q are two paths, we say that Q *stays above* P if $h_Q(i) \geq h_P(i)$ for every i . Let P and Q be paths such that Q stays above P . Then we say that a path X *stays between* P and Q if Q stays above X and X stays above P . We define the set $\mathcal{B}[P, Q]$ as the set of all paths X that stay between P and Q .

It is not hard to show that the set of transversals corresponding to $\mathcal{B}[P, Q]$ satisfies the symmetric exchange axiom. We call the Lagrangian matroid with set of bases $\mathcal{B}[P, Q]$ the *lattice path Lagrangian matroid given by P and Q* . (Where convenient, we abbreviate lattice path Lagrangian matroid by *LPLM*.) The representation of the bases of this Lagrangian matroid as lattice paths relies on the ordering $1 < 2 < \dots < n$. We want the class of lattice path Lagrangian matroids to be closed under reorderings and Lagrangian matroid isomorphisms: A *lattice path Lagrangian matroid* is a Lagrangian matroid \mathcal{B} for which an ordering of its ground set exists such that \mathcal{B} with this ordering is isomorphic to a Lagrangian matroid with set of bases $\mathcal{B}[P, Q]$. The family of lattice path Lagrangian matroids was introduced by J. Bonin and A. de Mier.

Figure 1 depicts a concrete example of $\mathcal{B}[P, Q]$ with $n = 11$ and $P = (1, 1, 1, -1, 1, -1, 1, -1, 1, -1, -1)$ and $Q = (-1, -1, 1, -1, 1, -1, 1, -1, 1, 1, 1)$.




Figure 1: An example lattice path Lagrangian matroid $\mathcal{B}[P, Q]$. The dashed and the dotted path depict the graphs of h_Q and h_P , respectively. Also depicted is the graph of h_B for a typical $B \in \mathcal{B}[P, Q]$.

Many of the proofs that follow require constructions of new paths from old paths, and we now fix our notation here.

Let S be an n -path. We define the selection operation as follows: Given integers $i_1, \dots, i_k \in [n]$, we use the notation $R = S(i_1, \dots, i_k)$ to define a k -path, where $R(j) := S(i_j)$. In particular, $S(1, \dots, n) = S$.

Let S be an n -path and R be an m -path. We define the concatenation operation as follows: We use the notation $V = S \circ R$ to describe the $(n + m)$ -path obtained by following S , then following R . Put concretely,

$$V(i) = \begin{cases} S(i) & \text{if } i \leq n, \\ R(i - n) & \text{if } i > n. \end{cases}$$

3 A characterization for LPLMs?

In Section 2.2 we saw a characterization of Lagrangian matroids via sets of intersections with arbitrary transversals. We now want to consider this for the case of lattice path Lagrangian matroids. What are the matroids appearing as \mathcal{I}_T for a LPLM \mathcal{B} ? We will see that these are matroids with a special structure known as lattice path matroids.

The other question we will address is whether this gives a characterization of LPLMs. More precisely, is every Lagrangian matroid \mathcal{B} such that for every transversal T the matroid \mathcal{I}_T is a lattice path matroid a LPLM? The answer to this is no, a counterexample will be presented in Section 3.3.

3.1 Lattice path matroids

In this section, we define lattice path (ordinary) matroids. A *North-East path* P (also abbreviated *NE-path*) is a path starting at $(0,0)$ that consists of North-steps $N = (1,0)$ and East-steps $E = (0,1)$. A North-East path from $(0,0)$ to (m,n) has length $m+n$. Let P and Q be NE-paths from $(0,0)$ to (m,n) such that Q stays above P , an example is shown in Figure 3.1.




Figure 2: The dotted path is an example for a path P , the dashed one for Q . The solid path is contained in the collection of all North-East paths between P and Q .

The collection \mathcal{B} of all North-East paths γ from $(0,0)$ to (m,n) between P and Q forms the set of bases of a matroid $\mathcal{M}[P,Q] = (E, \mathcal{B})$. The ground set E is $[m+n]$, and we identify each path γ with its set of North steps. A path γ corresponds to a subset $B_\gamma \subseteq [m+n]$ such that step i of γ is North if and only if $i \in B_\gamma$.

We show that \mathcal{B} fulfills the exchange property of ordinary matroids: Let $\gamma_1, \gamma_2 \in \mathcal{B}$, and let i be a North step with $i \in \gamma_1 \setminus \gamma_2$. If γ_1 is below γ_2 in step i , then there is a North step $j \in \gamma_2 \setminus \gamma_1$ with $j < i$ since both paths start at $(0,0)$. For $j = \max_k \{k : j \in \gamma_2 \setminus \gamma_1 \text{ and } k < i\}$, the path $\gamma_1 \setminus \{i\} \cup \{j\}$ stays between P and Q . If γ_1 is above γ_2 in step i , argue analogously with $j > i$. So \mathcal{M} is indeed a matroid. We call a matroid that is isomorphic to a matroid $\mathcal{M}[P,Q]$ of this type a *lattice path matroid*. Lattice path matroids are studied in [1] and [2].

3.2 Lattice path matroids from lattice path Lagrangian matroids

Let P and Q be n -paths such that Q stays above P . Let T be a transversal of length n . We know that the set $\mathcal{B}[P,Q]$ of all paths between P and Q is a Lagrangian matroid. Hence, by Theorem 2.1 the family of sets

$$\mathcal{I}_T = \mathcal{I}_T^{P,Q} = \{I : \text{there is a basis } B \in \mathcal{B}[P,Q] \text{ such that } I \subseteq B \cap T\}$$

is an ordinary matroid. Our main theoretical result on the structure of $\mathcal{I}_T^{P,Q}$ shows that it is actually a lattice path matroid:

Theorem 3.1. [Conjecture posed by A. de Mier] Let P , Q , and T be paths such that Q stays above P . The set $\mathcal{I}_T^{P,Q}$ is the collection of independent sets of a lattice path matroid.

A proof of this will be presented in Section 4. Following this, Section 5 contains a constructive proof, an algorithm that, given a transversal T , constructs the resulting lattice path matroid \mathcal{I}_T .

We will now fix some notation. The paths B such that $B \cap T$ gives a basis of the matroid $\mathcal{I}_T^{P,Q}$ are those paths $B \in \mathcal{B}[P,Q]$ that agree with T in as many steps as possible while staying between P and Q . They are called *maximal T -agreement paths*. Where there is no ambiguity, we say T is a *maximal path*.

and

$$h_{T \triangle \{i, i^*\}}(k) = \begin{cases} h_T(k) & \text{if } 1 \leq k \leq i-1, \\ h_T(k) + 2 & \text{if } i \leq k \leq n. \end{cases}$$

But $T \triangle \{i, i^*\}$ is not contained in $\phi(\mathcal{B})$. □

Lemma 3.3. *For every transversal T of $[n] \cup [n]^*$ the set $\mathcal{I}_T(\mathcal{B}) = \{I \mid I \subseteq B \cap T \text{ for some } B \in \mathcal{B}\}$ is the collection of independent sets of a lattice path matroid.*

Proof. For this proof, we set $B_{i,j} = \{1, \dots, i-1, i^*, i+1, \dots, j-1, j^*, j+1, \dots, n\}$ for $i, j \in [n], i \neq j$ such that $\mathcal{B} = \{\{1, 2, \dots, n\}\} \cup \{B_{i,j} \mid i, j \in [n], i \neq j\}$.

Let T be a transversal of $[n] \cup [n]^*$. We set $k = |T \cap [n]^*|$ and consider different values of k . If $k = 0$ or $k = 2$ then $T \in \mathcal{B}$ and hence the set $\mathcal{I}_T(\mathcal{B})$ consists of all subsets of T . This is a lattice path matroid defined by just one path, the one corresponding to T .

If $k = 1$ then T has exactly one starred element, say $i^* \in T$. The bases of $\mathcal{I}_T(\mathcal{B})$ are determined by those $B \in \mathcal{B}$ that have the largest possible intersection with T among all elements of \mathcal{B} . We have $T \cap \{1, 2, \dots, n\} = T \setminus \{i^*\}$ and $T \cap B_{i,j} = T \setminus \{j\}$ for $j \neq i$. Because $T \notin \mathcal{B}$, these are thus the transversals defining the bases of $\mathcal{I}_T(\mathcal{B})$. Hence, $\mathcal{I}_T(\mathcal{B})$ has bases $\{T \setminus \{j\} \mid j \in T\}$ and is the lattice path matroid with boundary paths (N, N, \dots, N, E) and (E, N, N, \dots, N) .

Now consider the case where $k \geq 3$, say $T \cap [n]^* = \{i_1^*, i_2^*, \dots, i_k^*\}$. Again, the bases of $\mathcal{I}_T(\mathcal{B})$ are determined by those $B \in \mathcal{B}$ that have the largest possible intersection with T among all elements of \mathcal{B} . These are the transversals B_{j_1, j_2} where $j_1, j_2 \in \{i_1, i_2, \dots, i_k\}$. They have $n - k + 2$ agreements with T , while all other elements of \mathcal{B} have $n - k$ or $n - k - 2$. Thus, the set of bases of $\mathcal{I}_T(\mathcal{B})$ is $\{([n] \setminus \{i_1, i_2, \dots, i_k\}) \cup \{j_1^*, j_2^*\} \mid j_1, j_2 \in \{i_1, i_2, \dots, i_k\}\}$. The set $\{\{k+1, k+2, \dots, n\} \cup \{i, j\} \mid 1 \leq i < j \leq k\}$ describes an isomorphic matroid. This is clearly a lattice path matroid: The direct sum of the family of all paths with 2 North steps and $k-2$ East steps and of the set containing one path just consisting $n-k$ North steps. □

4 Proof of the main theorem

In this section we give a first proof for Theorem 3.1.

The following observation is simple but crucial: If two maximal paths S and R meet at step i , i.e. $h_S(i) = h_R(i)$, then neither $V = S(1, \dots, i) \circ R(i+1, \dots, n)$ nor $W = R(1, \dots, i) \circ S(i+1, \dots, n)$ can have more agreements with T than S or R . Hence, S and R have the same number of agreements on steps 1 to i (and also on steps i to n) and V and W are also maximal. (See Figure 4.)




Figure 4: The solid path V and the dashed path W are also maximal.

From this it follows immediately that there must be a lowest maximal path P^{\max} and a highest maximal path Q^{\max} . Since the bases of $\mathcal{I}_T^{P,Q}$ are determined by maximal paths, we have $\mathcal{I}_T^{P^{\max},Q^{\max}} = \mathcal{I}_T^{P,Q}$. Thus, we only have to consider paths between P^{\max} and Q^{\max} .

Before we can come to the proof of Theorem 3.1, we will establish several lemmata that reveal the structure of the set of maximal paths. Lemma 4.1 states that, while P and Q can be chosen with arbitrary endpoints, P^{\max} and Q^{\max} will always meet in the last step. Therefore, all maximal paths will always meet at least in the last step.

Lemma 4.1. *All maximal paths end in the same point, i.e. $h_{P^{\max}}(n) = h_{Q^{\max}}(n)$. In particular, if S and R are two maximal paths, then $h_S(n) = h_R(n)$.*

Proof. Suppose, for a contradiction, that S and R are two maximal paths such that $h_S(n) \neq h_R(n)$. Without loss of generality, assume that $h_S(n) > h_R(n)$. We can assume that S and R do not meet after the first step, otherwise we just consider the interval after their last intersection. So we have $h_S(i) \geq h_R(i) + 2$ for all $i \geq 1$.

The steps in which S and R agree don't alter the distance of the two paths. So, since $h_R(n) < h_S(n)$, among the disagreeing steps there have to be more steps in which R goes down and S goes up. Because S and R have the same number of agreements with T , R cannot agree with T in every of these down steps. Hence, there exists a step j with $R(j) = \searrow$ and $T(j) = \nearrow$. Now, consider the following path (See Figure 5):

$$W := R(1, \dots, j-1) \circ \nearrow \circ R(j+1, \dots, n).$$

The path W agrees with T one more time than R and, since $h_S(n) \geq h_R(n) + 2$, it stays between S and R




Figure 5: W is the dashed path.

because of

$$h_W(i) = \begin{cases} h_R(i) & 1 \leq i \leq j-1 \\ h_R(i) + 2 & j \leq i \leq n. \end{cases}$$

However, a path W of this kind cannot exist, by maximality of the path R . □

Consider two maximal paths S and R such that S is above R at a step $i \in [n]$. In the following, we will call a height $j \in [n]$ with $h_R(i) \leq j \leq h_S(i)$ *feasible* if $j = h_R(i) + 2k$ for some $k \geq 0$. This means feasible heights are the possible heights at the i -th step for maximal paths between R and S .

There are two kinds of steps i : Either all maximal paths are parallel, i.e. they all agree at i , or there are two maximal paths that differ at i . The next lemmata explore the behaviour of the maximal paths in these two types of steps. First, we will see that in a parallel step all feasible heights can be reached by a maximal path. Figure 6 tries to illustrate this - feasible heights that can be reached by maximal paths are highlighted.

$T(j) = T(1) = \nearrow$. Let us first look at the case $S(j) = \searrow, R(j) = \nearrow$. Consider the path

$$W = \nearrow \circ R(2, \dots, j-1) \circ \searrow \circ R(j+1, \dots, n).$$

(See Figure 7.) It lies between R and S because

$$h_W(i) = \begin{cases} h_R(i) + 2 & 1 \leq i \leq j-1, \\ h_R(i) & j \leq i \leq n. \end{cases}$$

The paths W and R agree at every step except for 1 and j . If $T(j) = \searrow$, W would agree with T at step 1 and j while R would not. So W would have two more agreements than R which is not possible.




Figure 7: The path W

The second case, $S(j) = \nearrow, R(j) = \searrow$, is treated similarly: Because $S(n) = \searrow$ and $R(n) = \nearrow$, the first case shows that $T(n) = \nearrow$. Using this, we can proceed as above, now with the path

$$W = S(1, \dots, j-1) \circ \searrow \circ S(j+1, \dots, n) \circ \nearrow.$$

□

Now, we can explore the behaviour of T in the non-parallel steps. As mentioned before, this lemma will show that every feasible height that can be reached by maximal paths from above and from below actually is the height of a maximal path passing from above and also of one passing from below.

Lemma 4.4 (“Zig-zag steps”). *Suppose that not all maximal paths agree at step i . Then for every feasible j with $h_{P^{\max}}(i) \leq j \leq h_{Q^{\max}}(i)$ there is a maximal path B with $h_B(i) = j$. Moreover, if $h_{P^{\max}}(i-1) \leq j-1 < j+1 \leq h_{Q^{\max}}(i-1)$, there are maximal paths V and W such that $h_W(i) = h_V(i) = j$, $h_V(i-1) = j+1$ and $h_W(i-1) = j-1$. (See Figure 8.)*

Proof. We will prove the claim for all feasible j with $h_R(i) \leq j \leq h_S(i)$ for two arbitrary maximal paths S and R with $S(i) \neq R(i)$ and such that S is above R at i or at $i-1$: $h_S(i) > h_R(i)$ or $h_S(i-1) > h_R(i-1)$. Why does this suffice to prove the lemma?

If $P^{\max}(i) \neq Q^{\max}(i)$, we can simply choose $S = P^{\max}$ and $R = Q^{\max}$. If $P^{\max}(i) = Q^{\max}(i)$, by assumption there has to be some path B with $B(i) \neq P^{\max}(i)$. Then either $h_{P^{\max}}(i) < h_B(i) < h_{Q^{\max}}(i)$ and we can apply the above claim twice, or B meets one of the maximal boundary paths and we can apply it to B and the other boundary path.

So, let S and R be two maximal paths with $S(i) \neq R(i)$ and such that $h_S(i) > h_R(i)$ or $h_S(i-1) > h_R(i-1)$. We will proceed by induction on the minimal distance of S and R at steps i and $i-1$:

$$d(S, R) = \min\{h_S(i) - h_R(i), h_S(i-1) - h_R(i-1)\}.$$




Figure 10: The paths V and W

and

$$h_W(j) = \begin{cases} h_R(j) & 1 \leq j \leq i-1, j = n \\ h_R(j) + 2 & i \leq j \leq n-1. \end{cases}$$

Thus, V and W lie between S and R and are indeed the paths we needed.

For the other case, we can use the fact that $S(1) = \nearrow, R(1) = \searrow$ and proceed with a similar argument on the interval $[1, i]$. \square

Lemma 4.5. *Let $1 \leq i_0 < i_1 \leq n$ such that P^{\max} and Q^{\max} do not meet in any step $i_0 \leq i \leq i_1 - 1$. Then the following two statements are equivalent:*

1. *In the interval $\{i_0, \dots, i_1\}$ there are no steps where all maximal paths agree.*
2. *T is constant on $\{i_0, \dots, i_1\}$.*

Proof. We can w.l.o.g. assume that $i_0 = 1$ and $i_1 = n$, i.e. that P^{\max} and Q^{\max} only meet at the last step.

First, suppose that there are no steps in the interval $\{1, \dots, n\}$ in which all maximal paths agree. To show that T is constant we will prove that $T(i) = T(i-1)$ for every $i \in \{2, \dots, n\}$.

So, let $i \geq 2$. Because P^{\max} and Q^{\max} cannot meet in step $i-1$, we have $h_{Q^{\max}}(i-1) \geq h_{P^{\max}}(i-1) + 2$. Thus, there is a feasible height j such that $h_{P^{\max}}(i-1) \leq j < j+2 \leq h_{Q^{\max}}(i-1)$.

By Lemma 4.4 we can find maximal paths S and R that surround the square $(i-2, j+1), (i-1, j), (i, j+1), (i-1, j+2)$, i.e. $h_S(i-2) = h_R(i-2) = h_S(i) = h_R(i) = j+1$, $h_S(i-1) = j+2$ and $h_R(i-1) = j$. (See Figure 11.)




Figure 11: Paths surrounding a square

These paths disagree at steps $i-1$ and i , thus we have $T(i) = T(i-1)$ by Lemma 4.3.

Now, suppose that T is constant (w.l.o.g. $T = \nearrow(1, \dots, n)$) and let again $1 < i < n$. We have to show that there are maximal paths which disagree at step i .

Because of the Lemmas 4.2 and 4.4 we know that there are two maximal paths S and R agreeing with T at step i with $h_S(i) = h_R(i) + 2$. Until the i -th step, S has one more up step, i.e. one more agreement than R . Thus, the path $S(1, \dots, i-1) \circ \searrow \circ R(i+1, \dots, n)$ has the same number of agreements with T until the i -th step as R . Therefore it is a maximal path that lies between P^{\max} and Q^{\max} and disagrees with S and R at step i . \square

With one more remark we have all the tools needed to prove Theorem 3.1: The proof will use some matroid operations, Section 6 contains the definitions of these. We will also use that lattice path matroids are preserved under taking the direct sum which is easy to see.

Proof of Theorem 3.1. It suffices to prove the theorem for paths P and Q such that P^{\max} and Q^{\max} do not meet: Suppose they meet at step i . Let $I_1 = \{1, \dots, i\}$, $I_2 = \{i+1, \dots, n\}$ and $W' = W(I_1)$, $W'' = W(I_2)$ for $W = P, Q, T$. Then every maximal path splits into a maximal path w.r.t. P', Q' and T' and a maximal path w.r.t. P'', Q'' and T'' . Thus, $\mathcal{I}_T^{P,Q} = \mathcal{I}_{T'}^{P',Q'} \oplus \mathcal{I}_{T''}^{P'',Q''}$. Since lattice path matroids are preserved under taking direct sums, it suffices to prove the statement for both summands.

So, let us assume that P^{\max} and Q^{\max} do not meet. Lemma 4.2 and Lemma 4.4 show that there are only two kinds of steps i , parallel steps and zig-zag steps. It actually suffices to prove the theorem for the case where there are no parallel steps: Suppose step i is a step where all maximal paths agree. Then by Lemma 4.2 they all agree with T at i which means that i is contained in every basis of $\mathcal{I}_T^{P,Q}$. Now we take the i -th step out of all paths we are considering, i.e. we look at

$$\begin{aligned} P' &:= P^{\max}(1, \dots, i-1) \circ P^{\max}(i+1, \dots, n), \\ Q' &:= Q^{\max}(1, \dots, i-1) \circ Q^{\max}(i+1, \dots, n) \text{ and} \\ T' &:= T(1, \dots, i-1) \circ T(i+1, \dots, n). \end{aligned}$$

Then the resulting matroid $\mathcal{I}_{T'}^{P',Q'}$ is the restriction of $\mathcal{I}_T^{P,Q}$ to $[n] \setminus \{i\}$, $\mathcal{I}_{T'}^{P',Q'} = \mathcal{I}_T^{P,Q}|([n] \setminus \{i\})$. Because every basis contains i , we regain the original matroid by the following direct sum: $\mathcal{I}_T^{P,Q} = \mathcal{I}_{T'}^{P',Q'} \oplus \{\{i\}\}$. Lattice path matroids are preserved under taking direct sums, and $\{\{i\}\}$ is the lattice path matroid where both defining paths are simply one North step. So it suffices to show that $\mathcal{I}_{T'}^{P',Q'}$ is a lattice path matroid.

Thus, we can also assume that P, Q and T are such that in every step there are two maximal paths that disagree. We need to find two North-East-paths V and W starting at $(0, 0)$ and ending at the same point such that the bases of $\mathcal{I}_T^{P,Q}$ are the North-steps of paths between V and W .

By Lemma 4.4 all possible paths between P^{\max} and Q^{\max} are maximal paths. Thus,

$$\mathcal{I}_T^{P,Q} = \{I \mid I \subseteq B \cap T, B \in \mathcal{B}(P^{\max}, Q^{\max})\}.$$

By Lemma 4.5, T is furthermore constant. This means that the bases of our matroid $\mathcal{I}_T^{P,Q}$ are either all sets of up-steps of paths between P^{\max} and Q^{\max} , or all sets of down-steps of such paths.

If $T = \searrow(1, \dots, n)$, considering P and Q reflected on the x -axis and $T' = \nearrow(1, \dots, n)$ leads to the same matroid. Thus, we can assume that $T = \nearrow(1, \dots, n)$.

If we use the following function to map up-down-paths to NE-paths

$$\varphi(W)(i) = \begin{cases} N & \text{if } W(i) = \nearrow, \\ E & \text{if } W(i) = \searrow, \end{cases}$$

the paths $V = \varphi(P^{\max})$ and $W = \varphi(Q^{\max})$ fulfill our conditions:

- Since P^{\max} and Q^{\max} end in the same point, also V and W do.
- The set $\mathcal{B}(P^{\max}, Q^{\max})$ is mapped bijectively to $\mathcal{B}(V, W)$.
- Because $T = \nearrow (1, \dots, n)$, the bases of $\mathcal{I}_T^{P, Q}$ correspond to up-steps of paths between P^{\max} and Q^{\max} , thus to North-steps of paths between V and W .

Hence, $\mathcal{I}_T^{P, Q}$ is the lattice path matroid $\mathcal{M}[V, W]$. \square

This concludes the proof of Theorem 3.1. The next section will present a second, more constructive proof. We will see an inductive algorithm that constructs the boundary paths of the lattice path matroid $\mathcal{I}_T^{P, Q}$.

5 An algorithm to construct the lattice path matroid

Let $\mathcal{B}[P, Q]$ be a lattice path Lagrangian matroid of n steps and let T be a traversal of length n . In this section, we describe an algorithm that constructs the lattice path matroid $\mathcal{I}_T^{P, Q}$ by explicitly defining the boundary North-East paths of Section 3.1. This algorithm provides a second proof of Theorem 3.1.

The matroid $\mathcal{I}_T^{P, Q}$ is a lattice path matroid in the sense that it is such a matroid only up to isomorphism. That is, the first step of the lattice path matroid may not record the first step of the maximal paths in the lattice path Lagrangian matroid. It will be necessary to record which step in the lattice path matroid records which step in the lattice path Lagrangian matroid by defining an explicit permutation $\phi \in S_n$. We will use the convention $\phi(i) = j$ to mean the i -th step in the lattice path matroid records the j -th step in the lattice path Lagrangian matroid.

We approach the problem by induction on n . For the base case, $n = 1$, we consider the following cases:

1. If (T, P, Q) is one of the triples $(\nearrow, \nearrow, \nearrow)$, $(\nearrow, \searrow, \nearrow)$, $(\searrow, \searrow, \nearrow)$, or $(\searrow, \searrow, \searrow)$, then T agrees with at least one of P or Q . In particular, the one-element path B that agrees with T is in $\mathcal{B}[P, Q]$. In this case, the associated lattice path matroid consists of the unique path from $(0, 0)$ to $(0, 1)$.
2. The two cases that (T, P, Q) is one of the triples $(\nearrow, \searrow, \searrow)$ or $(\searrow, \nearrow, \nearrow)$ describe the situations where $T \neq P = Q$. In particular, the only path B in $\mathcal{B}[P, Q]$ does not agree with T and the associated lattice path matroid consists only of the path from $(0, 0)$ to $(1, 0)$.

The remaining two cases (where $Q = \searrow$ and $P = \nearrow$) do not satisfy the definition of lattice path Lagrangian matroids. In all the base case examples, ϕ is of course the identity map in S_1 .

For the induction step, we suppose that we are given appropriate paths P and Q of length n , and a transversal T of length n . By ignoring the n^{th} step, we have paths P' and Q' of length $n - 1$, and a transversal T' of length $n - 1$ which determine an already-constructed lattice path matroid L' and matroid isomorphism ϕ' , by the induction hypothesis. The algorithm describes the construction of the lattice path matroid L and matroid isomorphism ϕ corresponding to P , Q , and T by making appropriate modifications to L' and ϕ' .

We may assume, without loss of generality, that $T(n) = \nearrow$. If this is not the case, then we can “negate” each of P , Q , and T , and note that $\mathcal{I}_{-T'}^{-P', -Q'} = \mathcal{I}_{T'}^{P', Q'}$.

Recall Lemma 4.1, which states that all maximal paths must end at the same point. We distinguish two possibilities that can occur when moving from the triple (P', Q', T') to (P, Q, T) .

1. In the first case, every maximal T' -agreement path B' for (P', Q', T') ends on h_Q (that is $h_{B'}(n) = h_Q(n)$ for each maximal T' -agreement path B') and $Q(n) = -1$. In this case, every path B' must go down on the n^{th} step.

2. Otherwise, paths B' may go up or down on the last step, but for maximally-agreeing paths, we extend B' to a path B that goes up on the n^{th} step.

In the second case, we say that the triple (P, Q, T) is *extendible*.

Lemma 5.1. *If the triple (P, Q, T) is extendible, then $\text{rank}(\mathcal{I}_T^{P, Q}) = 1 + \text{rank}(\mathcal{I}_{T'}^{P', Q'})$.*

Proof. Suppose (P, Q, T) is an extendible triple. Let B be any maximal T' -agreement path. Suppose that B agrees with T' exactly r times. Define B^+ to follow B for the first $n - 1$ steps and define $B_n = +1$. Since (P, Q, T) is extendible, B^+ agrees with T exactly $r + 1$ times, so the rank increases by at least one.

Suppose, for a contradiction, that the rank increases by more than one. Then, there is a T -agreement path C of length n with s agreements and $s \geq r + 2$. From C , we obtain a path C' of length $n - 1$ by ignoring the final step. Then C' agrees with T at least $s - 1 \geq r + 1$ times. This contradicts the construction of B , which agreed with T exactly r times and agreed maximally with T in $n - 1$ steps. \square

Lemma 5.2. *If the triple (P, Q, T) is not extendible, then $\text{rank}(\mathcal{I}_T^{P, Q}) = \text{rank}(\mathcal{I}_{T'}^{P', Q'})$.*

Proof. Let $r = \text{rank}(\mathcal{I}_{T'}^{P', Q'})$. Given any maximal T' -agreement path B' , define $B = B' \circ \searrow$. Then, B is path of length n that agrees with T exactly r times, thus $r \leq \text{rank}(\mathcal{I}_T^{P, Q})$.

Now, suppose for a contradiction that $r + 1 \leq s$, where $s = \text{rank}(\mathcal{I}_T^{P, Q})$. Then, there is a maximal T -path B which agrees with T exactly s times. Since the triple (P, Q, T) was not extendible, the path B did not agree with T in the final step. The path $B' = B(1, \dots, n - 1)$ which ignores the last step still agrees with T' at least $s \geq r + 1$ times, which contradicts the definition of r . \square

The ordered listing $\phi(1), \phi(2), \dots, \phi(n)$ of the image of the permutation ϕ will consist of three distinct sections. Using the terminology of Section 4, the image first contains all of the steps that are permanently parallel steps. Then it contains all of the steps that are permanently zig-zag steps. All remaining steps i (where it is not yet decidable if i is a parallel step or zig-zag step) are last in the image of ϕ . In each segment of the image, the values are in increasing order. For example, if $\phi(1), \phi(2), \dots, \phi(10)$ is, in order, 2, 3, 5, 7, 1, 4, 6, 8, 9, 10, then steps 2, 3, 5, and 7 of the lattice path Lagrangian matroid are known to be parallel steps. And, for instance, steps 1, 4, 6, and 8 may record steps known to be zig-zag steps while it is not known what kind of steps the ultimate three are.

The algorithm treats the extendible and non-extendible cases differently. The case that the triple (P, Q, T) is extendible is easily handled. We describe the modifications to the lattice path matroid L' and to ϕ' :

1. Given L' which ends at (m, n) , define L to be the lattice path matroid ending at $(m + 1, n)$. The boundary paths P and Q defined for the lattice path matroid L will go North on the last step.
2. The last step of the lattice path matroid will record the last step of the lattice path Lagrangian matroid. That is, we extend ϕ' to ϕ by setting $\phi(n) = n$.

The algorithm correctly handles the case that the triple (P, Q, T) is extendible as every maximal T' -path B' gives a maximal T -path $B' \circ \nearrow$ and every maximal T -path B restricts to a maximal T' -path $B(1, \dots, n - 1)$. In particular, all maximal T -paths coincide on the n -th step.

The case where the triple (P, Q, T) is not extendible is more complicated. By the proof Lemma 5.2, every path of length n of the form $B' \circ \searrow$ is a maximal T -path, where B' is a maximal T' -path. Let P^{\max} denote the lowest maximal path of (P', Q', T') on $n - 1$ elements. We identify a special step k :

$$k = \max_{i=1, \dots, n-1} \{i \mid P^{\max}(1, \dots, i-1) \circ P^{\max}(i)^* \circ P^{\max}(i+1, \dots, n-1) \circ \nearrow \in \mathcal{B}(P, Q)\}. \quad (1)$$

If the set described in (1) is empty, then no modification to any step of the lowest T -path P^{\max} is in $\mathcal{B}[P, Q]$. If the set described in (1) is non-empty, then k is the latest (or right-most) step of P^{\max} which can be modified to still give a path in $\mathcal{B}[P, Q]$. Let R be the path

$$R = P^{\max}(1, \dots, k-1) \circ P^{\max}(k)^* \circ P^{\max}(k+1, \dots, n-1) \circ \nearrow. \quad (2)$$

Note that the path R agrees with T in the n -th step. By Lemma 5.2, it follows that $R(k) = \searrow$, $T(k) = \nearrow$, and $P^{\max}(k) = \nearrow$. Then, the path R is a maximal T -path.




Figure 12: All new maximal T -paths are found by modifying the i -th step of P^{\max} for $k \leq i < n$ and $P^{\max}(i) = \nearrow$.

For all i such that $k \leq i < n$, the following must be true:

1. If $P^{\max}(i) = \nearrow$, then $T(i) = \nearrow$ (by Lemma 5.2) and $P^{\max}(1, \dots, i-1) \circ P^{\max}(i)^* \circ P^{\max}(i+1, \dots, n-1) \circ \nearrow$ is a maximal T -path which disagrees with T in the i -th step.
2. If $P^{\max}(i) = \searrow$, then $T(i) = \searrow$ (by Lemma 5.2) and $P^{\max}(1, \dots, i-1) \circ P^{\max}(i)^* \circ P^{\max}(i+1, \dots, n-1) \circ \nearrow$ disagrees with T two more times than P^{\max} does.

Thus, the steps i such that $k \leq i < n$ and $P^{\max}(i) = \nearrow$ are exactly all of the steps where the lowest path $P^{\max} \circ \searrow$ can be modified to give a maximal T -path. In Figure 12, the shaded region shows the modifications of P^{\max} at step i for $k \leq i < n$. By Lemma 5.2, the set

$$\{i \mid R(1, \dots, i-1) \circ R(i)^* \circ R(i+1, \dots, n-1) \circ \nearrow \in \mathcal{B}(P, Q)\} \quad (3)$$

must be empty. This implies that there is no maximal T -path below R , and R is the new lowest maximal T -path. Further it implies that all new maximal T -paths must fall within the shaded region in Figure 12.

Since the set described in (3) is empty, the path on the lattice path matroid L' encoding P^{\max} must be on the right edge of L' starting at step k . Let P' be the lower boundary path of L' . We modify P' at the step recording the k -th step of the lattice path Lagrangian matroid by inserting an East-step. The permutation ϕ is reordered so that any steps i which are known to permanently be parallel steps or zig-zig steps appear in the first or second segments of the ordered image of ϕ , respectively.

6 Matroid operations on LPLMs

Unlike Coxeter matroids in general, one can define for Lagrangian matroids most of the usual matroid theoretic operations, like duality, contraction (and deletion, by combining them), direct sums and truncations. In this

section, we investigate for some of these operations whether they preserve lattice path Lagrangian matroids. While lattice path (ordinary) matroids are preserved under all operations presented here [1, 2], we will see that the direct sum of two LPLMs is not always a Lattice path Lagrangian matroid.

For each operation, we will first recapitulate its definition for matroids before adapting it to Lagrangian matroids. For the matroid operations, we follow [6], the definitions for the constructions on Lagrangian matroids can be found in [3, Section 3.9].

6.1 Duality

For a matroid \mathcal{M} with ground set E and set of bases \mathcal{B} , the *dual matroid* \mathcal{M}^* also has E as ground set. Its bases are the complements of the bases of \mathcal{M} : $\mathcal{B}^* = \{E \setminus B \mid B \in \mathcal{B}\}$.

This concept directly translates to Lagrangian matroids: Let \mathcal{B} be the set of bases of a Lagrangian matroid \mathcal{M} on the set $S = [n] \cup [n]^*$. Then $\mathcal{B}^* = \{S \setminus B \mid B \in \mathcal{B}\}$ is again a set of bases of a Lagrangian matroid, the *dual Lagrangian matroid* \mathcal{M}^* of \mathcal{M} . Using the operation $*$ we can write $S \setminus B = B^*$ which even more clearly shows that the elements of \mathcal{B}^* are also transversals. It is easy to see that \mathcal{B}^* fulfills the symmetric exchange axiom.

For lattice path Lagrangian matroids taking the dual corresponds to reflecting all paths with respect to the x -axis: For a path B the dual path B^* is the path taking the exact opposite steps. Thus, $(\mathcal{B}[P, Q])^* = \mathcal{B}[P^*, Q^*]$ and LPLMs are preserved by duality.

6.2 Deletion and contraction

For deletion and contraction the parallels between matroids and Lagrangian matroids are not as direct. We will nevertheless first consider the definition for ordinary matroids.

Let \mathcal{M} be a matroid with ground set E and independent sets \mathcal{I} , and let $X \subseteq E$. Then the *restriction* $\mathcal{M}|X$ of \mathcal{M} to X is the matroid on the ground set X with independent sets $\{I \subseteq X \mid I \in \mathcal{I}\}$. The *deletion* $\mathcal{M} \setminus X$ of X from \mathcal{M} is the restriction of \mathcal{M} to $E \setminus X$: $\mathcal{M} \setminus X = \mathcal{M}|(E \setminus X)$. Its independent sets are $\{I \subseteq E \setminus X \mid I \in \mathcal{I}\}$. The operation of contraction is the dual of the operation of deletion: The *contraction* \mathcal{M}/X of X from \mathcal{M} is $\mathcal{M}/X = (\mathcal{M}^* \setminus X)^*$.

For Lagrangian matroids we will keep this concept of duality between deletion and contraction. Let \mathcal{B} be the set of bases of a Lagrangian matroid \mathcal{M} on the set $S = [n] \cup [n]^*$. In [3, Section 3.9] the contraction of one element $i \in S$ is defined as $\mathcal{B}/\{i\} = \{B \setminus \{i\} \mid B \in \mathcal{B}, i \in B\}$. With $\mathcal{B}/\{x_1, x_2, \dots, x_k\} = (\dots((\mathcal{B}/\{x_1\})/\{x_2\})\dots)/\{x_k\}$ this yields the following definitions for $X \subseteq S$:

If X is a subtransversal of S , i.e. $|X \cap \{i, i^*\}| \leq 1$ for all $i \in [n]$, the *contraction* \mathcal{M}/X of X from \mathcal{M} has the set $\mathcal{B}/X = \{B \setminus X \mid B \in \mathcal{B}, X \subset B\}$ as set of bases. The *deletion* $\mathcal{M} \setminus X$ of X from \mathcal{M} is defined as the dual concept and therefore has bases $\mathcal{B} \setminus X = (\mathcal{B}^*/X)^* = \{B \setminus X^* \mid B \in \mathcal{B}, X^* \subset B\}$.

If there is $i \in [n]$ such that $i, i^* \in X$, i.e. if X is not a subtransversal of S , the deletion and the contraction of X from \mathcal{M} are both empty: $\mathcal{B} \setminus X = \mathcal{B}/X = \emptyset$. With this definition we have the following relations:

$$\mathcal{B} \setminus X = (\mathcal{B}^*/X)^* = \mathcal{B}/X^* \text{ and } \mathcal{B}/X = (\mathcal{B}^* \setminus X)^* = \mathcal{B} \setminus X^*$$

One can also see that $\mathcal{B} \setminus \{x_1, x_2, \dots, x_k\} = (\dots((\mathcal{B} \setminus \{x_1\}) \setminus \{x_2\}) \dots) \setminus \{x_k\}$.

Because of these observations, for lattice path Lagrangian matroids it suffices to consider the deletion of one element. One can see that this preserves LPLMs:

Let P and Q be n -paths such that Q stays above P and let $i \in [n] \cup [n]^*$. We want to consider the Lagrangian matroid

$$\mathcal{B}[P, Q] \setminus \{i\} = \{B \setminus \{i^*\} \mid B \in \mathcal{B}[P, Q], i^* \in B\}.$$

We can w.l.o.g. assume that $i \in [n]$. Then $\mathcal{B}[P, Q] \setminus \{i\}$ is constructed by taking all paths $B \in \mathcal{B}[P, Q]$ that go down at the i -th step and deleting their i -th step. If we let P' and Q' be the lowest and the topmost of the paths in $\mathcal{B}[P, Q] \setminus \{i\}$, we want to see that $\mathcal{B}[P, Q] \setminus \{i\} = \mathcal{B}[P', Q']$.

All paths in $\mathcal{B}[P, Q] \setminus \{i\}$ lie between P' and Q' . Are also all paths in $\mathcal{B}[P', Q']$ in $\mathcal{B}[P, Q] \setminus \{i\}$? Let $B' \in \mathcal{B}[P', Q']$. Consider the paths

$$\begin{aligned} B &= B'(1, \dots, i-1) \circ \searrow \circ B'(i+1, \dots, n), \\ P'' &= P'(1, \dots, i-1) \circ \searrow \circ P'(i+1, \dots, n), \\ Q'' &= Q'(1, \dots, i-1) \circ \searrow \circ Q'(i+1, \dots, n). \end{aligned}$$

Because B' is between P' and Q' , the path B lies between P'' and Q'' . The paths P' and Q' were constructed from P'' and Q'' . Thus, $P'', Q'' \in \mathcal{B}[P, Q]$. Therefore, we have $B \in \mathcal{B}[P, Q]$ and $B' \in \mathcal{B}[P, Q] \setminus \{i\}$.

So we see that lattice path Lagrangian matroids are preserved under contraction and deletion.

6.3 Direct sums

Let \mathcal{M}_1 and \mathcal{M}_2 be matroids with ground sets E_1 and E_2 and sets of bases \mathcal{B}_1 and \mathcal{B}_2 . The *direct sum* of \mathcal{M}_1 and \mathcal{M}_2 is the matroid $\mathcal{M}_1 \oplus \mathcal{M}_2$ on the ground set $E_1 \dot{\cup} E_2$, the disjoint union of E_1 and E_2 , with bases $\{B_1 \dot{\cup} B_2 \mid B_1 \in \mathcal{B}_1, B_2 \in \mathcal{B}_2\}$.

For Lagrangian matroids we can directly translate this concept, using shifting for the disjoint union: Let \mathcal{B}_1 and \mathcal{B}_2 be the sets of bases of two Lagrangian matroids \mathcal{M}_1 and \mathcal{M}_2 on the ground sets $[n_1] \cup [n_1]^*$ and $[n_2] \cup [n_2]^*$. The *direct sum* of \mathcal{M}_1 and \mathcal{M}_2 has the ground set $[n_1 + n_2] \cup [n_1 + n_2]^*$. Its set of bases is $\mathcal{B}_1 \oplus \mathcal{B}_2 = \{B_1 \cup (B_2 + n_1) \mid B_1 \in \mathcal{B}_1, B_2 \in \mathcal{B}_2\}$.

Lattice path Lagrangian matroids are **not** preserved under taking the direct sum. This is demonstrated by the following counterexample:

Let $\mathcal{B}_1 = \mathcal{B}_2 = \{\{1, 2\}, \{1, 2^*\}, \{1^*, 2\}\}$. Then

$$\mathcal{B}_1 \oplus \mathcal{B}_2 = \{1234, 1234^*, 123^*4, 12^*34, 1^*234, 12^*34^*, 12^*3^*4, 1^*23^*4, 1^*234^*\}$$

where $abcd$ stands for $\{a, b, c, d\}$.

Considering all possibilities, one can see that there are only four choices of 4-paths P and Q such that the corresponding LPLM $\mathcal{B}[P, Q]$ has nine elements. None of these LPLMs contains 1234 or $1^*2^*3^*4^*$. They all have $h_Q(4) = h_P(4) + 2$.

Reordering the elements of [4] doesn't change the facts that $\mathcal{B}_1 \oplus \mathcal{B}_2$ contains the transversal without any stars and that it has nine elements. Let ϕ be a Lagrangian matroid isomorphism corresponding to a transversal T . The image $\phi(\mathcal{B}_1 \oplus \mathcal{B}_2)$ contains T and every transversal $T' = T \triangle \{i, i^*\}$ that differs from T in exactly one step. If T has no or exactly one starred element, the transversal 1234 is among these and for T with three or four starred elements the transversal $1^*2^*3^*4^*$. For T with exactly two starred elements $\phi(\mathcal{B}_1 \oplus \mathcal{B}_2)$ has to contain at least one transversal with exactly one, one with exactly two and also one with three starred elements. These correspond to paths B_1, B_2, B_3 with $h_{B_1}(4) = h_{B_2}(4) + 2 = h_{B_3}(4) + 4$.

Thus, there is no isomorphism ϕ for which $\phi(\mathcal{B}_1 \oplus \mathcal{B}_2)$ can be reordered to be isomorphic to any of the four LPLMs with nine elements.

This shows that taking the direct sum does not preserve Lattice path Lagrangian matroids.

7 Additional topics

In this section, we explore further questions that might be interesting to study.

7.1 Other classes of Lagrangian matroids

There are several other classes of Lagrangian matroids. This section introduces some coming from matrices, graphs and maps. The relationship between LPLMs and these classes haven't been studied so far.

7.1.1 Representable Lagrangian matroids

Let A and B be two $n \times n$ -matrices such that $AB^t = BA^t$ and consider $C = (A|B)$ where A is indexed by the elements of $[n]$, B by those of $[n]^*$. Then

$$\mathcal{B}_C = \{X \subseteq [n] \cup [n]^* \mid X \text{ transversal, the columns corresponding to elements of } X \text{ are linearly independent}\}$$

is a Lagrangian matroid. Lagrangian matroids of this type are called *representable*. A description can be found in [3, Section 3.4].

7.1.2 Eulerian Lagrangian matroids

Eulerian Lagrangian matroids are introduced in [5].

Let G be a connected, simple graph where all vertices have degree 4 or 2. Assume that there are n vertices $1, \dots, n$ of degree 4. A *transition* is a set of two distinct edges that share a vertex. A *bitransition* is a set of two disjoint transitions at the same vertex. At every vertex of degree 4 there are 3 distinct bitransitions, we choose two of them. For vertex i we identify one of the chosen bitransitions with $i \in [n]$, the other with i^* . A transversal of the set $[n] \cup [n]^*$ thus corresponds to a choice of one bitransition at every vertex of degree 4.

We refer to the following operation as *splitting* a bitransition $b = \{t, t'\}$ at vertex v : Delete v from G and add two new vertices v and v' such that v is connected to the edges in t and v' to the ones in t' . The resulting graph is denoted by $G|b$. For a set $B = \{b_1, b_2, \dots, b_k\}$ of bitransitions with pairwise distinct ends we write $G|B$ for $(\dots((G|b_1)|b_2)|\dots|b_k)$. We call such a set *eulerian* if

1. $G|B$ is connected, and
2. B contains one bitransition with end v for every vertex v of degree 4.

Then the collection of sets

$$\mathcal{B} = \{B \subset [n] \cup [n]^* \mid \text{The set of bitransitions corresponding to } B \text{ is eulerian.}\}$$

is a Lagrangian matroid.

7.1.3 Lagrangian matroids from maps on orientable compact surfaces

This class of Lagrangian matroids is described in [3, Section 4.3]. Let $G = (V, E)$ be a graph that is embedded on an oriented connected compact surface S . Then $\mathcal{M} := (G, S)$ is a *map* if the connected components of $S \setminus G$ are open 2-cells, called the *faces* of \mathcal{M} .

For a map $\mathcal{M} = (G, S)$, its dual map is $\mathcal{M}^* = (G^*, S)$ where G^* is defined as follows: For every face F of \mathcal{M} we choose a point v_F inside F . Then the covertices, the vertices of G^* are $\{v_F \mid F \text{ face of } \mathcal{M}\}$. For every edge e of G we let $e^* = \{v_{F_1}, v_{F_2}\}$ where F_1 and F_2 are the faces adjacent to e . We choose the embedding of e^* such that $e^* \cap f = \emptyset$ for every $f \in E$, $f \neq e$ and $|e^* \cap e| = 1$. These are called the coedges.

We have $|E| = |E^*|$ and index the edges of G by $[n]$, the coedges by $[n]^*$. A *basis* of \mathcal{M} is a set $B \subset E \cup E^*$ such that B corresponds to a transversal of $[n] \cup [n]^*$ and $S \setminus \bigcup_{e \in B} e$ is connected. Then the collection of sets

$$\mathcal{B}(\mathcal{M}) = \{B \subset E \cup E^* \mid B \text{ is a basis of } \mathcal{M}\}.$$

is a Lagrangian matroid.

7.2 Lagrangian matroid polytopes

Every Coxeter matroid can be represented by a polytope, the matroid polytope. The polytopes coming from LPLMs might have a special form. Let $\mathcal{B}(P, Q)$ be a lattice path Lagrangian matroid of length n . We associate to each transversal B in $\mathcal{B}(P, Q)$ its *characteristic vector* $\psi_B \in \mathbb{R}^n$ with

$$\psi_B(i) = \begin{cases} +1 & B(i) = \nearrow, \\ -1 & B(i) = \searrow. \end{cases} \quad (4)$$

We note that the definition extends to Lagrangian matroids and to transversal matroids in general. If B is a (partial) transversal in a transversal matroid \mathcal{B} on $S = [n] \cup [n]^*$ then its characteristic vector $\psi_B \in \mathbb{R}^n$ is

$$\psi_B(i) = \begin{cases} +1 & i \in B, \\ -1 & i^* \in B, \\ 0 & i \notin B \text{ and } i^* \notin B. \end{cases}$$

The matroid polytope \mathcal{P} of $\mathcal{B}(P, Q)$ is the convex hull of the characteristic vectors of the bases of $\mathcal{B}(P, Q)$. That is,

$$\mathcal{P} = \text{conv}\{\psi_B \in \mathbb{R}^n \mid B \text{ is a basis of } \mathcal{B}(P, Q)\}.$$

It is known that the length of edges of \mathcal{P} are either 2 or $2\sqrt{2}$ and that \mathcal{P} is a ± 1 -polytope. Beyond a simple combinatorial interpretation of the edges of different lengths, nothing else is known about the matroid polytope \mathcal{P} of LPLMs. In particular, it is not known if the faces of the polytopes coming from LPLMs have a special form or if the f -vector of these polytopes have special forms.

Acknowledgments

The authors are grateful to Anna de Mier and the Centre de Recerca Matemàtica for the opportunity to work on this project.

References

- [1] J. Bonin, A. de Mier. Lattice path matroids: structural properties. *European J. Combin.*, **27** (2006), 701–738.
- [2] J. Bonin, A. de Mier, M. Noy. Lattice path matroids: enumerative aspects and Tutte polynomials. *J. Combin. Theory Ser. A*, **104** (2003), 63–94.
- [3] A. V. Borovik, I. M. Gel’fand, N. White. *Coxeter matroids*. Progress in Mathematics. 216. Birkhäuser Boston, Inc. Boston, MA. 2003.
- [4] A. V. Borovik, I. M. Gel’fand, N. White. Symplectic Matroids. *J. of Algebraic Combinatorics*, **8** (1998), 235–252.
- [5] A. Bouchet. Greedy algorithm and symmetric matroids. *Math. Program.*, **38** (1987), 147–159.
- [6] J. G. Oxley. *Matroid theory*. Oxford Science Publications. Oxford University Press. New York, 1992.

Anna Gundert
Institut für Mathematik. Freie Universität, Berlin 14195, Germany.
gundert@math.fu-berlin.de

Edward D. Kim

Department of Mathematics. University of California, Davis, CA 95616, USA
ekim@math.ucdavis.edu

Daria Schymura

Institut für Informatik. Freie Universität, Berlin 14195, Germany
schymura@mi.fu-berlin.de

Patterns in Ordered Trees

Ragnar Freij, Daria Schymura, Lluís Vena

April 20, 2009

Abstract

We consider trees whose vertices are numbered $1, \dots, n$. A nesting in such a tree is a pair of edges (uv, xy) with $u < x < y < v$, and a crossing is a pair (uv, xy) with $u < x < v < y$. Our main result is a bijective proof that there are more trees that avoid nestings than trees that avoid crossings.

Furthermore, we show that the number of non-nesting paths is between $1/n^{3/2}2^n$ and $2n^23^n$. We also give exact enumerations for the numbers of non-crossing, respectively non-nesting, embeddings of some particular isomorphism classes of trees, showing that there are trees that have exponentially many non-crossing embeddings, but only linearly many non-nesting embeddings.

1 Introduction

We consider labelled graphs on the linearly ordered vertex set $V = [n]$. A *crossing* is a pair of disjoint edges (uv, xy) with $u < x < v < y$, and a *nesting* is a pair of disjoint edges (uv, xy) with $u < x < y < v$. If we embed the vertices on a horizontal line and draw the edges as arcs above the line, these two notions have an obvious geometric meaning. We call graphs that avoid crossings *non-crossing* and graphs that avoid nestings *non-nesting*. It has been shown in [2] that the numbers of non-crossing and non-nesting graphs on n vertices are equal. They equal $2^n s_{n-1}$ where s_n is the n th Schröder number, which grows approximately like $(3 + 2\sqrt{2})^n$.

Anna de Mier and Marc Noy conjectured that there are at least as many non-nesting trees on n vertices as non-crossing trees. Our main contribution is the proof of this conjecture, which will be presented in Sections 2 and 3. In Section 2, we will present an algorithm that allows us to generate all non-nesting trees. We will also define two strict subclasses of the set of non-nesting trees and give algorithms that generate all trees in these subclasses.

In Section 3, we prove that both subclasses are in bijective correspondence with certain lattice paths, namely North/East lattice paths from $(0, 0)$ to $(n, 2n)$ that stay weakly below the line $y = 2x$. The number of these lattice paths happens to be the number of non-crossing trees on n vertices, which proves our main result. The number of non-crossing trees is well-known (see [1]). It equals $\frac{1}{2n+1} \binom{3n}{n}$, which is roughly $(27/4)^n$ by Stirling's estimate. The number of non-nesting trees seems difficult to compute. Our result and the formula for the number of non-nesting graphs imply the following bounds on the number of non-nesting trees on n vertices t_n

$$(27/4)^n < t_n < (6 + 4\sqrt{2})^n.$$

Since we now know that there are more non-nesting than non-crossing trees, we will discuss two refinements of this inequality in Sections 4 and 5. In Section 4, we will associate with every tree G a Dyck path $D = D(G)$. This partitions the class of non-nesting trees and the class of non-crossing trees into sets of trees that are associated with the same Dyck path D . Let $NN(D)$ be the size of the set of non-crossing trees that are associated with D , and, analogously, let $NC(D)$ be the number of non-crossing trees associated with D . We study the questions for which Dyck paths $NN(D) \geq NC(D)$ holds. We show that this inequality is not true in general and give examples for both cases.

In Section 5, we study the question how many non-nesting and non-crossing embeddings particular unlabelled trees have. We will give an example of a class of trees that have exponentially many non-crossing embeddings, but only linearly many non-nesting embeddings. We then restrict our attention to paths. It is easy to compute that the number of non-crossing paths on n vertices equals $n2^{n-3}$. Again, the number of non-nesting paths p_n seems difficult to compute. We show that up to polynomial factors

$$2^n < p_n < 3^n.$$

Let f_n be the number of non-crossing paths on n vertices and F_n be the n th Fibonacci number. We conjecture that

$$f_n \leq p_n \leq F_{2n}$$

with strict lower bound for $n \geq 5$ and strict upper bound for $n \geq 7$.

2 Constructive algorithms for trees

In Subsection 2.1 we present some observations that will allow us to build, iteratively, all non-nesting trees. Subsection 2.2 is devoted to an iterative

algorithm to build the class of non-crossing trees and Subsection 2.3 presents two strict subclasses of non-nesting trees.

2.1 Observations and construction of the non-nesting trees

The non-nesting (and the non-crossing) property is hereditary: if G is a non-nesting graph, then so is $G \setminus e$, where e is an edge of G . The class of non-nesting trees is not hereditary, but the class of subgraphs of non-nesting trees is hereditary.

The *right neighbourhood* of the vertex i , R_i , in a labelled graph is the set of edges $\{i, j\}$ with $j > i$. Notice that any labelled graph is completely characterized by its sequence of right neighbourhoods (R_1, \dots, R_{n-1}) , which is a partition of the edge set. We can also identify $R_i = \{\{i, j_1\}, \dots, \{i, j_k\}\}$ with the set $\{j_1, \dots, j_k\}$.

Let T be a non-nesting tree on n vertices. We say that t_i is the i -th *embryonal tree* of T if it is formed by the union of the right neighbourhoods of the vertices $\{1, \dots, i\}$. Thus, the embryonal tree t_i is built by adding the edges of R_i , the right neighbourhood of i , to t_{i-1} . As $\{R_i\}_{i \in [1, n-1]}$ is an edge partition of T , $t_{n-1} = T$ and the union $t_{i-1} \cup \{R_i\} = t_i$ is disjoint.

Let $NNT(i, j)$ be the class of all the embryonal trees t_i , that arise from some non-nesting tree T , where j is the last vertex reached in each embryonal tree in the class. If $j = \max_{k \in [1, i]} \{\text{right neighbourhood of } k \text{ in } t_i\}$ then $t_i \in NNT(i, j)$. Notice that $i < j$.

Let t_i and t_{i+1} be the i -th and the $i+1$ -th embryonal tree of T respectively. Assume $t_i \in NNT(i, j)$. Let R_{i+1} be the right neighbourhood of $i+1$ in T . Let $\mathcal{P}([j_1, j_2])$ denote all the possible subsets from $[j_1, j_2]$. The relation between $i+1$ and j in t_i will determine some restrictions on R_{i+1} .

- $R_{i+1} \in \mathcal{P}([j, n])$. If $\min_{R_{i+1}} = \min(R_{i+1}) < j$, then the edge $\{i+1, \min_{R_{i+1}}\}$ is in T , but so is an edge $\{k, j\}$ for some $k \leq i$. Thus the edges $\{i+1, \min_{R_{i+1}}\}$ and $\{k, j\}$ form a nesting, contradicting the non-nesting property of T . Moreover, the condition $R_{i+1} \in \mathcal{P}([j, n])$ characterize the non-nesting property. As all the edges in t_i have either a smaller than i left-vertex or a smaller than j right-vertex, the new edges do not create a nesting. All the edges in R_{i+1} have the same left vertex, so they do not create a nesting by themselves.
- If $i+1$ and j belong to the same connected component in t_i , then $R_{i+1} \in \mathcal{P}([j+1, n])$. If $j \in R_{i+1}$ then $i+1$ and j would be connected by two paths in T . One path is formed by the edge $\{i+1, j\}$ and

the other path, is deduced in t_i , as $i + 1$ and j belong to the same connected component. As the edge $\{i + 1, j\}$ is not in t_i , the two paths are disjoint, so they close a cycle in T . This contradicts the fact that T is a tree.

- If $i + 1 < n$ is the last element from its connected component in t_i (the vertex with largest label in its connected component) then $R_{i+1} \neq \emptyset$. Assume $R_{i+1} = \emptyset$. The edges from the right neighbourhoods $\{R_k\}_{k \geq i+2}$, will have both vertices larger than any label from the connected component where $i + 1$ lies on. Thus T will have at least two connected components: one containing n and the other containing $i + 1$. This contradicts the fact that T is a tree.

Thus, we have 4 possible cases where t_i might fit. Each configuration restricts the possible right neighbourhood R_{i+1} :

- (C1) $i + 1$ is connected with j and $i + 1$ is the last element in its connected component (this means that $i + 1 = j$). $R_{i+1} \in \mathcal{P}([j + 1, n]) \setminus \emptyset$.
- (C2) $i + 1$ is not connected with j and $i + 1$ is the last element in its connected component. $R_{i+1} \in \mathcal{P}([j, n]) \setminus \emptyset$.
- (C3) $i + 1$ is connected with j and $i + 1$ is not the last element in its connected component. $R_{i+1} \in \mathcal{P}([j + 1, n])$.
- (C4) $i + 1$ is not connected with j and $i + 1$ is not the last element in its connected component. $R_{i+1} \in \mathcal{P}([j, n])$.

Notice that any choice R_1, \dots, R_{n-1} , where R_i has been chosen appropriately (depending on whether $t_{i-1} = \bigcup_{k=1}^{i-1} R_k$ has fit in C1, C2, C3 or C4) generates a non-nesting tree. As each tree can be decomposed in its right neighbourhoods $\{R_i\}$, by taking all the possible choices for R_i each time, we obtain all non-nesting trees.

2.2 Construction of the non-crossing trees

Let T_c be a non-crossing tree. Define $\{t_1^c, \dots, t_{n-2}^c, t_{n-1}^c = T_c\}$ as the embryonal trees of T_c . Like in the non-nesting case, t_i^c is built from t_{i-1}^c by adding the edges of the right neighbourhood of i , R_i , to t_{i-1}^c . So, $t_i^c = \bigcup_{k=1}^i R_k$, where R_k are the edges $\{k, j\} \in T_c$ with $j > k$.

Let $NCT(i, j)$ be the embryonal trees t_i^c , for some non-crossing tree T^c . In this setting, the variable j is the first vertex, from $i + 2$ on ($i + 2$ included),

with an edge incoming from the left. If there is no such vertex, then the embryonal tree belongs to $NCT(i, n)$.

We can proceed by an analogous reasoning as for the non-nesting trees and observe that, for each $t_i^c \in NCT(i, j)$, the right neighbourhood of a non-crossing tree with the i -th embryonal tree t_i^c , R_{i+1} , should be a set in $\mathcal{P}([i+2, j])$, otherwise we would create a crossing. Also:

- if $i+1$ is connected with j then $R_{i+1} \in \mathcal{P}([i+2, j-1])$.
- if $i+1$ is the last element in its connected component then $R_{i+1} \in \mathcal{P}([i+2, j]) \setminus \emptyset$.

As in the non-nesting case we may write down the four cases depending on whether $i+1$ is connected to j and if $i+1$ is the last element in its connected component.

However, there are no non-crossing embryonal trees fulfilling the fourth case. Namely, in no $t_i^c \in NCT(i, j)$ we have that $i+1$ is not connected with j and $i+1$ is not the last element in its connected component. Thus, there are just 3 possible cases for each embryonal tree in $NCT(i, j)$.

Let us show the non-existence of the fourth case for the non-crossing trees. Let $j^* \neq j$ be the last element in the component of $i+1$. Let r_j be the vertex connected to j with the largest index. By the definition of $NCT(i, j)$ we have $j^* > j$.

We must have a path from $i+1$ to j^* . However, this path must either contain r_j as one of its vertex or cross the edge $\{r_j, j\}$, as the edge $\{r_j, j\}$ closes a region with $i+1$ inside. In both cases we contradict either the assumption about the connection between $i+1$ and j or the non-crossing property of t_i^c .

2.3 Non-nesting trees strict subclasses

With the observations in Section 2.2, it is natural to consider two (strict) subclasses of non-nesting trees. Assume t_i is a non-nesting embryonal tree belonging to $NNT(i, j)$ in which we may apply (C4).

Instead of trying all cases that (C4) allows, we might restrict ourselves to consider t_i as if it was a (C2) or a (C3). In the constructive algorithm, whenever we find a (C4) embryonal tree, we may construct all possible embryonal trees t_{i+1} out of t_i as if t_i was a (C2), this subclass will be called “avoid the \emptyset ”. If we treat t_i in (C4) as if it was a (C3) we call this subclass “do not connect $i+1$ and j ”.

If $n \geq 5$ both subclasses are strictly contained in the class of the non-nesting trees with n vertices.

3 Bijections to lattice paths

This section shows two bijections between the class of non-crossing trees and two subclasses of non-nesting trees. This allows us to show that there are more non-nesting trees than non-crossing trees (Corollary 3.10). The bijections are done via lattice paths (Theorem 3.1, Theorem 3.6).

3.1 First bijection

Let us present the first bijection.

Theorem 3.1. *The subclass of non-nesting trees on $n+1$ vertices “avoid the \emptyset ” is in bijection with P_n ; the lattice paths from $(0,0)$ to $(n,2n)$, with n East steps and $2n$ North steps, and weakly below the line $y = 2x$.*

Proof. Let $p \in P_n$ be a lattice path and let T be a tree in the “avoid the \emptyset ” subclass.

To build T we will move along the path p , starting from $(0,0)$ and ending at $(n,2n)$. We need to define $3n$ steps corresponding to the n East steps and the $2n$ North steps from the lattice path p .

Let T_0, T_1, \dots, T_{3n} be the intermediate steps. Each T_i consists of a 3-tuple (t_i, u_i, l_i) where t_i is a subset of edges from some non-nesting tree (they will be embryonal trees) and u_i and l_i are two positive integers with $u_i, l_i \in [1, n+1]$ and $u_i \leq l_i$, for all $i \in [1, 3n]$. The initial condition is $T_0 = (\emptyset, 1, 1)$ and the final is $T_{3n} = (T, n+1, n+1)$.

To go from $T_i = (t_i, u_i, l_i)$ to $T_{i+1} = (t_{i+1}, u_{i+1}, l_{i+1})$ we look at the following step $s_{i+1} \in \{\text{North}, \text{East}\}$ in the path p .

- If s_{i+1} is an East step then $t_{i+1} = t_i$, $u_{i+1} = u_i$ and $l_{i+1} = l_i + 1$. We move the l (last) index to the right and leave everything else untouched.
- If s_{i+1} is a North step we have two subcases:
 - If we can connect u_i and l_i then we connect them. We insert the edge $\{u_i, l_i\}$ if it is not in t_i and if the edge $\{u_i, l_i\}$ does not induce a cycle. So $t_{i+1} = t_i \cup \{u_i, l_i\}$, $u_{i+1} = u_i$ and $l_{i+1} = l_i$.
 - In the case we cannot connect u_i and l_i , either because we lose the acyclic condition for the tree or the simple graph condition, we push u_i one step to the right. So, $t_{i+1} = t_i$, $u_{i+1} = u_i + 1$ and $l_{i+1} = l_i$.

Before proceeding further, let us prove a technical proposition:

Proposition 3.2. *Let $T_i = (t_i, u_i, l_i)$ be a configuration coming from some path p at the step i . Let $\{c_1, \dots, c_k\}$ be the connected components of t_i with some vertex in the set $[1, u_i]$. Then we have:*

- (i) $u_i \leq l_i$
- (ii) *The last vertex (vertex with largest label) in each c_j is in the interval $[u_i, l_i]$.*
- (iii) *If $u_i = l_i$ then t_i is a non-nested tree on u_i vertices and $2|\{\text{East steps seen}\}| = |\{\text{North steps seen}\}|$.*
- (iv) t_i does not have a nesting.
- (v) *If at step i , $2|\{\text{East steps seen}\}| = |\{\text{North steps seen}\}|$, then we have $u_i = l_i$.*

Proof of Proposition 3.2. We proceed by induction to prove the first four parts. Part (v) will be proved at the end. Since the path is always below the line $y = 2x$ the first step is always an East step, thus if $i = 1$ the result holds (it also holds if $i = 0$). Assume the induction hypothesis: the result holds for $1, \dots, i$ and let us proof it for $i + 1$.

Let s_{i+1} be the $i + 1$ -th step in the path p . If s_{i+1} is an East step:

- (i) $u_{i+1} = u_i \leq l_i \leq l_{i+1} = l_i + 1$.
- (ii) Since $u_{i+1} = u_i$ and no edge has been added, all the connected components $\{c_i\}$ are the same, also, $[u_i, l_i] \subset [u_{i+1}, l_{i+1}]$.
- (iii) We cannot have this case.
- (iv) As we have not inserted any edge, we have not destroyed the non-nesting property.

Let us assume s_{i+1} is a North step. We have two subcases, namely $u_i < l_i$ or $u_i = l_i$. Assume the latter, $u_i = l_i$. In this case, by induction hypothesis we already have a tree on u_i vertices and $2|\{\text{East steps seen}\}| = |\{\text{North steps seen}\}|$. This means that we are on the line $y = 2x$ and a North step in this situation is impossible.

Assume now $u_i < l_i$. Suppose we are able to insert an edge without closing a cycle. Then:

- (i) $u_{i+1} = u_i < l_i = l_{i+1}$.
- (ii) As we have just inserted another edge $\{u_i, l_i\}$, this second part also holds.

- (iii) This part is not applicable.
- (iv) The edge we have inserted does not create a nesting: there are no edges “below” $\{u_i, l_i\} = \{u_{i+1}, l_{i+1}\}$ as no edges have a left vertex larger than $u_i = u_{i+1}$. Also there are no edges “above” as no edge have a right vertex larger than $l_i = l_{i+1}$.

Suppose that the edge between u_i and l_i close a cycle (it is not allowed to be put in the embryonal trees) or we have already the edge $\{u_i, l_i\}$. Thus, we should update u_i to $u_i + 1$.

- (i) $u_{i+1} = u_i + 1 \leq l_i = l_{i+1}$.
- (ii) By the induction hypothesis all connected components in t_i have a vertex in $[u_i, l_i]$. Let c_{u_i} be the connected component where the vertex u_i lies. If we cannot place an edge between u_i and l_i , the connected component c_{u_i} should contain also the vertex l_i . Thus, $[u_i + 1, l_i] = [u_{i+1}, l_{i+1}]$ contains a vertex for each connected component.
- (iii) In particular, if $u_i + 1 = u_{i+1} = l_{i+1} = l_i$ the reasoning for ii shows that we have a tree as it is clear, by the algorithm, that we do not have any cycles. Since the North steps either insert an edge or add one to u and the East steps add one to l , we obtain the equality $y = 2x$ (where x are the East and y are the North steps). If $u_{i+1} < l_{i+1} = l_i$ this part is not applicable.
- (iv) As we have not inserted an edge $t_{i+1} = t_i$, it remains non-nesting.

Let us prove (v). Assume the contrary: at step i , $u_i < l_i$ (by (i) $u_i \leq l_i$). As t_i is a simple graph, acyclic and all the edges connect indices in $[1, l_i]$, it has at most $l_i - 1$ edges. As the number of North steps is $u_i - 1 + \#\{\text{edges in } t_i\}$ and the number of East steps is $l_i - 1$, we obtain that $2|\{\text{East steps seen}\}| > |\{\text{North steps seen}\}|$. \square

Let us present the following claims that prove the injection.

Claim 3.3. *This procedure produces non-nesting trees.*

Claim 3.4. *The procedure is injective: two different lattice paths produce two different trees.*

Claim 3.5. *All trees belong to the “avoid the \emptyset ” subclass.*

Proof of Claim 3.3. The claim follows from the third part of Proposition 3.2. \square

Proof of Claim 3.4. Let p_1 and p_2 be two paths and let $i + 1$ be the first step in which they differ. Let $T_i = (t_i, u_i, l_i)$ be the configuration at step i .

Suppose p_1 takes a North step where p_2 takes an East step. If the edge $\{u_i, l_i\}$ was in t_i or its presence would close a cycle in p_1 , then $u_{i+1} = u_i + 1$. In this case the degree of u_i will remain as it was in t_i at the final tree. However, in p_2 , as at some point we will have a North step (a North step is the final step), u_i will get at least one more edge attached to it.

If we are allowed to insert the edge $\{u_i, l_i\}$, then p_1 will have such an edge while p_2 will not have it. This makes both (labelled) trees different. \square

Proof of Claim 3.5. Assume that, for some path p , we obtain a configuration $T_i = (t_i, u_i, l_i)$ where the t_i might be classified as (C4). This means that the connected component where u_i belongs, c_{u_i} , is not connected to l_i . Also, u_i is not the last element in c_{u_i} .

In this case we never choose the empty set: whenever we find a North step we would connect u_i to the corresponding vertex l_j . As the final step is a North step and we have $u_i + 1 < l_i$ (meaning that i is not the last step by Proposition 3.2), there remains a North step. \square

It is left to see that any non-nesting tree in the “avoid the \emptyset ” class can be associated with a path. This follows from the algorithm for constructing all non-nesting trees restricted to this subclass.

The iterative algorithm presented in Section 2.2 consists of n steps. In the i -th one, we choose a subset deciding where i should be connected to. Assume a path is built until the i -th step of the algorithm and we should perform the $i + 1$ -th one. Let $T_j = (t_j, u_j, l_j)$ be the configuration from the path where $u_j = i + 1$ and l_j is the last element reached in the embryonal tree t_j .

Let S be the subset where $i + 1$ should be connected to. If S is the empty set, then we are in the case (C3) (as the (C4) case is treated as (C2)) and we are not able to connect $u_j = i + 1$ with l_j . Thus, the next step in p would be a North step and $u_{j+1} = u_j + 1$.

So, assume S is non-empty. Let $B(S) \subset \{0, 1\} \times [l_j, n]$ be the binary representation of the set. It has a 1 in the k -th position if k is in S . Let $L(S)$ be the position of the rightmost one in $B(S)$ (the largest element in S). We will move an index o from the position l_j to the position $L(S)$ and:

- If $B(S)$ contains a 1 in the o -th position put a North step in p , this will mean an edge between u_j and $l_{j'}$. Update T_i accordingly.
- If $B(S)$ contains a 0 in the o -th position put an East step. This means that we do not put any edge between u_j and $l_{j'}$ but there are some more edges to put in u_j so we move $l_{j'}$ to $l_{j'+1} = l_{j'} + 1$.

After $L(S)$ is reached and we have placed the corresponding North step (and an edge $\{u_j, l_{j'}\} = \{u_j, L(S)\}$), put another North step to move u_j to $u_j + 1$ as the edge $\{u_j, l_{j'}\} = \{u_j, L(S)\}$ is already there (so, we cannot put that edge two times).

The number of East steps should always be below n as we would not surpass the $n + 1$ vertex.

Also, twice the number of East steps is larger than number of North steps. An embryonal tree is a subgraph of a tree on $l_{j'}$ vertices (for some suitable last vertex reached $l_{j'}$), thus the number of edges is below $l_{j'} - 1$.

Moreover, the vertex being examined, $i + 1$, is such that $i \leq l_{j'} - 1$. As $l_{j'} - 1$ is the number of East steps and $i + \#\{\text{edges in } t_i\}$ corresponds to the number of North steps, we obtain the desired inequality $y \leq 2x$. \square

Figure 3.1 depicts an example of the bijection.




Figure 1: Correspondence between a path and a non-nesting tree on 6 vertices.

3.2 Second bijection

Let us state a second theorem concerning the “do not connect u and l ” subclass and the lattice paths.

Theorem 3.6. *The subclass of non-nesting trees on $n + 1$ vertices “do not connect u and l ” is in bijection with P_n : the lattice paths from $(0, 0)$ to $(n, 2n)$, with n East steps and $2n$ North steps, and weakly below the diagonal $x = 2y$.*

Proof. The bijection follows the same lines as Theorem 3.1. However, instead of looking at the densest acyclic graph, we will look at the sparsest connected one, with an slight modification:

- If s is an East step then $t_{i+1} = t_i$, $u_{i+1} = u_i$ and $l_{i+1} = l_i + 1$. We move the l index one vertex to the right.

- If s is a North step we have two subcases:
 - Insert the edge $\{u_i, l_i\}$ in t_i :
 - (a) If the vertex l_i has no incoming edge (or, more generally, total degree equals zero).
 - (b) If u_i is the last element in its connected component.
 So, in these cases, $t_{i+1} = t_i \cup \{u_i, l_i\}$, $u_{i+1} = u_i$ and $l_{i+1} = l_i$.
 - Otherwise, move u_i one step to the right. So $u_{i+1} = u_i + 1$, leave $t_{i+1} = t_i$ and $l_{i+1} = l_i$.

Let us proof Proposition 3.2 for this case.

Proof. As before, the induction works for the first case as we have an East step. Assume the statement for up to i steps.

- If the $i + 1$ -th step is an East step, then the result holds as we have just grown the interval.
- If the $i + 1$ -th step is to the North and we put an edge, then the result also holds: we place edges between disconnected components so we do not produce a cycle and the new edges does not induce a nesting.
- If it is a North step, but have not put an edge, then the connected component of u_i has a member in $[u_i + 1, l_i]$. The result holds by moving u_i one step to the right.

If $u_i = l_i$ we should obtain a connected graph as we do not move u_i unless it is connected with a larger vertex in $[u_i + 1, l_i]$. As the exceptional rule about connecting the vertices of degree zero l_i with u_i do not create a cycle, we will obtain an acyclic connected graph, thus it will have $u_i - 1 = l_i - 1$ edges. As we are always below the line $y = 2x$ we obtain that $u_i \leq l_i$, with $u_i = l_i$ if and only if $2|\{\text{East steps seen}\}| > \{\text{North steps seen}\}$. \square

Claim 3.7. *This procedure produces non-nesting trees.*

Proof of Claim 3.7. The claim follows from the third part of Proposition 3.2. \square

Claim 3.8. *The procedure is injective: two different lattice paths produce two different trees.*

Proof of Claim 3.8. Let s_i be the first step where p_1 and p_2 differ. We have $i < 3n$. Assume s_i in p_1 is an East step.

Assume that l_{i-1} has incoming degree. As the last step of p_1 will be a North step, u_{i-1} will have a larger degree in the tree from p_1 than in the tree from p_2 (for which $u_i = u_{i-1} + 1$).

If l_{i-1} has no incoming degree, then p_2 produce a tree with the edge $\{u_{i-1}, l_{i-1}\}$, but the tree from p_1 will not have this edge. \square

Claim 3.9. *All trees belong to the “do not connect u and l ” subclass.*

Proof of Claim 3.9. If the path leads to a embryonal tree, t_i , of the type (C4), and the following is a North step, we will always move u_i to $u_{i+1} = u_i + 1$ as l_i will have incoming degree and u_i will not be the larger element in its connected component. \square

The inverse function comes from the same reasoning: this application mimics the algorithm to produce non-nesting trees from the subclass of “do not connect u and l ”. \square

3.3 Paths and non-crossing trees

In [4] it is shown that P_n , the number of paths with n East steps and $2n$ North steps, from $(0, 0)$ to $(n, 2n)$ and strictly below $y = 2x + 1$ (so weakly below $y=2x$) is

$$\frac{1}{2n+1} \binom{3n}{n}.$$

In [1] it is proved that the number of non-crossing trees with $n+1$ vertices is $\frac{1}{2n+1} \binom{3n}{n}$. We observe that we need five vertices to have a C4 embryonal tree; the tree in Figure 3.1 without the sixth vertex is the example with less vertices. Thus, from Theorem 3.1 we obtain:

Corollary 3.10. *There are more non-nesting trees than non-crossing trees on n vertices. The inequality is strict if $n \geq 5$.*

4 Connections to Dyck paths

4.1 Degree Sequences

In this section, we will again associate certain lattice paths to our pattern-avoiding graphs. The starting point is observing that if we constrain a graph to be non-crossing or non-nesting, then it is uniquely determined by some very local behaviour.

For a graph G on the vertex set $[n]$ and for any i , define r_i to be the number of edges (i, j) in G with $i < j$. Similarly, l_i is the number of edges (i, j) in G with $j < i$. So r_i and l_i count the number of rightbound, respectively leftbound, edges from i .

Note that if we know l_i and r_i for every i , and we also know that G is non-nesting, then we can reconstruct G completely. Because we can iterate through the nodes $1, \dots, n$, starting r_i new edges at node i , and finishing l_i . To avoid nestings, we must always choose to first finish the edges that were started first.

Similarly, if we know that G is non-crossing, then anytime we finish an edge, we must choose the last one that is started but not yet finished. So in both the non-nesting and the non-crossing case, we can reconstruct our graphs from left to right, from the information in the (l_i, r_i) sequences. This also gives a bijection between non-nesting multigraphs and non-crossing multigraphs, via the (l_i, r_i) sequences, illustrated in Figure 2. However, this bijection does not respect natural graph properties, such as connectedness, and in particular does not restrict to trees. As Figure 2 shows, it does not even map simple graphs to simple graphs.




Figure 2: A non-crossing graph and a non-nesting graph with the same sequence: $(0, 2), (0, 1), (1, 0), (2, 1), (1, 0)$.

The bijection does, however, respect the multiset of vertex degrees, since these are encoded as $\{l_i + r_i \mid i \in [n]\}$. This proves the following lemma.

Lemma 4.1. *Let \mathcal{A} be a collection of multisets of natural numbers. Let $NN_{\mathcal{A}}$ be the number of non-nesting graphs whose degree sequence is in \mathcal{A} , and let $NC_{\mathcal{A}}$ be the number of non-crossing graphs whose degree sequence is in \mathcal{A} .*

Then $NN_{\mathcal{A}} = NC_{\mathcal{A}}$.

When \mathcal{A} contains *all* n -element multiset of naturals, we have just seen again that the number of non-nesting multigraphs equals the number of non-crossing multigraphs on n nodes. Other interesting special cases are when \mathcal{A} is the collection of n -element multisets that sum to m , where we get the following corollary.

Corollary 4.2. *For any n and m , the number of non-nesting multigraphs on node set $[n]$ with m edges, equals the number of non-crossing multigraphs on node set $[n]$ with m edges.*

Yet another corollary will be of particular interest.

Corollary 4.3. *There is a bijection between the sets of non-nesting and of non-crossing perfect matchings on $[2n]$.*

Proof. A perfect matching is just a graph where every vertex has degree 1. Thus, the claim follows from Lemma 4.1. \square

4.2 Dyck paths

To further exploit the degree sequence idea, we must understand what sequences can show up as (l_i, r_i) sequences. To start with, we clearly must have $\sum_i l_i = \sum_i r_i$, since each of these sums counts the number of edges in the graph. Moreover, when looking at the graph from left to right, we can never have finished more edges than we have started. In terms of (l_i, r_i) , this means that $\sum_{i \leq j} l_i \leq \sum_{i \leq j-1} r_i$ for each j . These two restrictions lead us to the notion of Dyck paths.

A Dyck path of length $2n$ is a lattice path from $(0, 0)$ to $(2n, 0)$ with steps $(1, 1)$ and $(1, -1)$, that has no point below the x -axis. An example is given in Figure 3. Dyck paths are counted by the well-known Catalan numbers $\frac{1}{n+1} \binom{2n}{n}$, and are thus in bijection to several interesting combinatorial structures, see [5], pages 219–229.

Relatedly to the sequences (l_i, r_i) , we can associate a Dyck path $D(G)$ to any graph G , in the following way. Start the path at the origin. Going from left to right in the graph, take l_i steps in direction $(1, -1)$, then take r_i steps in direction $(1, 1)$, and continue to the next vertex.

Since $\sum_i l_i = \sum_i r_i = m$ is the number of edges in G , the path will go from $(0, 0)$ to $(2m, 0)$. Since $\sum_{i \leq j} l_i \leq \sum_{i \leq j-1} r_i$ for each j , the path will never go below the x -axis, and will thus be a Dyck path. As an illustration, Figure 3 shows the Dyck path corresponding to the two graphs in Figure 2.

If G is a perfect matching, then every node in G gives exactly one step in $D(G)$. Thus the mappings from non-nesting and non-crossing multigraphs to Dyck paths are both bijective on the set of non-nesting matchings, as well as on non-crossing matchings. This proves that there are exactly $\frac{1}{n+1} \binom{2n}{n}$ non-crossing perfect matchings, and equally many non-nesting ones.

For a Dyck path D , let $NN(D)$ and $NC(D)$ be the number of non-nesting, respectively non-crossing trees G such that $D(G) = D$. We know from Corollary 3.10 that $\sum_D NN(D) \geq \sum_D NC(D)$, so it is reasonable to




Figure 3: A Dyck path with the corresponding non-nesting and non-crossing matchings.

ask whether $NN(D) \geq NC(D)$ holds for any D , and if not, for which Dyck paths D it does hold.

We can read $NN(D)$ and $NC(D)$ from the Dyck path D of length $2m$ as follows. Consider the corresponding non-nesting, respectively non-crossing matching, and then contract consecutive vertices until there are $m+1$ vertices left. The contractions are restricted by declaring that when an up-step is followed by a down-step in D , these two nodes must not be contracted.

Now, $NN(D)$ and $NC(D)$ count how many of these contracted graphs are connected, in the non-nesting and the non-crossing case, respectively. Note that since the contracted graphs will have m edges and $m+1$ nodes, they are trees if and only if they are connected.

For example, if D is the Dyck path in Figure 3, we get $NN(D) = NC(D) = 4$.

Some particular cases are easy to solve. Notably, if $D = EF$ is the concatenation of two Dyck paths E and F , and thus has a return to the x -axis, then $NN(D) = NN(E) \cdot NN(F)$ and $NC(D) = NC(E) \cdot NC(F)$. Indeed, to connect the graph we have to contract the last node of E with the first node of F , and then we need to connect each of the two parts of the graph.

Another induction that goes through smoothly is the following: Suppose D starts with $j-1$ up-steps immediately followed by a down-step. Let D' be the Dyck path obtained from D by adding one up-step and one down-step immediately after the $j-1$ first steps. Then a non-crossing tree corresponding to D' is obtained from one corresponding to D , by contracting either $(j-1, j)$ or $(j+1, j+2)$. So $NC(D') = 2 \cdot NC(D)$. Similarly, a non-nesting tree on D gives one on D' by contracting either $(1, 2)$ or $(j+1, j+2)$. This shows that $NN(D') = 2 \cdot NN(D)$.

In particular, these inductions show that when searching for small examples where $NN(D) \neq NC(D)$, we need only look among Dyck paths with no returns to the x -axis. Also, we can assume that the first slope of down-steps only is one step long, and by symmetry that the same thing holds for the

last slope of up-steps.

A minimal example of a Dyck path where $NC(D) < NN(D)$ has length 8, and is depicted in Figure 4.




Figure 4: A Dyck path D with $NC(D) = 6 < 7 = NN(D)$.

Not surprisingly, we have to work slightly harder to find an example with $NN(D) < NC(D)$. Figure 5 shows a minimal counterexample, with length 14.




Figure 5: A Dyck path D with $NN(D) = 24 < 27 = NC(D)$.

The number of non-crossing trees on a given Dyck path is rather easy to calculate, as follows. Consider a non-crossing matching M on $2m$ vertices. This is, by the earlier discussion, equivalent with a Dyck path of length $2m$. A *cycle* of length k in M is a sequence of arcs

$$\{a_1 + 1, a_2\}, \{a_2 + 1, a_3\}, \dots, \{a_{k-1} + 1, a_k\}, \{a_k + 1, a_1\}.$$

Note that every arc of length greater than one, defines one cycle in which it is the longest arc $\{a_i + 1, a_1\}$.

To contract the matching to a tree, we need to connect all the arcs in the cycle to each other. But we must also avoid to build a cycle (in the ordinary graph sense) in the contracted graph. Hence, for a cycle

$$\{a_1 + 1, a_2\}, \{a_2 + 1, a_3\}, \dots, \{a_{k-1} + 1, a_k\}, \{a_k + 1, a_1\},$$

we must contract all but one of the pairs $(a_i, a_i + 1)$. For each k -cycle in the matching, this gives us k choices. Hence, on the non-crossing matching M , we have $\prod_c \ell(c)$ trees, where the product is over all cycles in M and $\ell(c)$ is the length of the cycle.

Equivalently, we could use Dyck path terminology. Suppose the Dyck path D passes through the points (j, i) and (k, i) , and stays strictly above the i level between these points. This means that there is an arc $\{j + 1, k\}$ in the corresponding non-crossing matching. The cycle defined by this arc, will consist of arcs that go between the returns to level $i + 1$, between the two endpoints $j + 1$ and k . So the length of the cycle will be one plus the number of returns to level $i + 1$, between two consecutive returns to i . To get the total number of NC-trees on D , we have to multiply these lengths over all pairs of consecutive returns to the same level.

Using this, we can for example read directly from the pictures that there are $27 = 3 \cdot 3 \cdot 3$ NC-trees on the path in Figure 5. We also see that there are $6 = 3 \cdot 2$ NC-trees on the path in Figure 4.

However, there seems to be no similar easy way to calculate the number of non-nesting trees on a given Dyck path.

5 Embeddings of particular graphs

5.1 Not every unlabelled tree has more non-nesting than non-crossing embeddings

In Section 3, we showed that the number of non-nesting trees on n vertices is at least as large as the number of non-crossing trees on n vertices. Let us refine this question. Is this inequality perhaps true for every tree? Is it true that, for every tree G , the number of non-nesting embeddings of G is at least as large as the number of non-crossing embeddings of G ? The answer to this question is negative in a strong sense. We will give an example of a class of trees that have exponentially many non-crossing embeddings, but only linearly many non-nesting embeddings.

Let $Sp(n, m)$ be the set of unlabelled trees on n vertices that have one vertex of degree m and otherwise only vertices of degree at most 2. We call these trees *spiders* on n vertices with m legs. The vertex of degree m is called the *body* of the spider.




Figure 6: A tree from $Sp(14, 7)$.

Lemma 5.1. *For each $G \in Sp(n, m)$, the number of non-crossing embeddings equals $\frac{nm!}{\prod_i m_i!} 2^{n-m-1}$ where m_i is the number of legs of length i of G .*

Proof. Embed the spider on S^1 . First, embed the body. Then, choose one of the $\frac{m!}{\prod_i m_i!}$ orderings of the spider's legs and embed the neighbors of the body. For each remaining vertex there are 2 possibilities; it can be embedded to the left or to the right of its predecessor.




Figure 7: A spider embedded on S^1 .

There are n possibilities to cut the sphere, and each gives a different embedding of the spider. \square

Corollary 5.2. *If G is a spider on n vertices with m legs that have all the same length, then G has $n2^{n-m-1}$ non-crossing embeddings.*

The case $m = 2$ gives the following corollary.

Corollary 5.3. *A path on n vertices has $n2^{n-3}$ non-crossing embeddings.*

Let n be odd and $m = \lfloor \frac{n}{2} \rfloor$. Then $Sp(n, m)$ consists of one spider $Sp_{n,2}$ whose legs all have length 2 and that has $n2^{\lfloor \frac{n}{2} \rfloor}$ non-crossing embeddings.

Proposition 5.4. *For $n \geq 7$ and odd, the number of non-nesting embeddings of $Sp_{n,2}$ equals $33\lfloor \frac{n}{2} \rfloor - 49$.*

Proof. When we embed $Sp_{n,2}$, some of the neighbors of the body go to the right and some go to the left to the body. There are $\lfloor \frac{n}{2} \rfloor + 1$ possibilities to distribute the neighbors of the body to the left and right. Let j be the number of neighbors of the body embedded to the left. For $j = 0$ and $j = \lfloor \frac{n}{2} \rfloor$, there are 5 possibilities to embed the remaining arcs. Let us look w.l.o.g. at $j = 0$. For all, but 2, neighbors of the body, the embedding of the remaining arc is unique. For the first and last neighbor of the body, there are 2 and 3 possibilities to embed the remaining arc, as is shown in Figure 5.1. One of the combinations is forbidden because it causes a nesting.




Figure 8: The solid arcs are determined through $j = 0$ and non-nestingness.

For $j = 1$ and $j = \lfloor \frac{n}{2} \rfloor - 1$, there are 20 possibilities to embed the remaining arcs. A careful count shows that for each $2 \leq j \leq \lfloor \frac{n}{2} \rfloor - 2$, there are 33 embeddings of $Sp_{n,2}$. Summing over all j s gives the equation. \square

5.2 Bounds on the number of non-nesting paths

We have seen that, in general, it is wrong that there are more non-nesting embeddings of a tree than non-crossing embeddings. But maybe this is true for the simplest trees, namely for stars and paths. For stars $K_{1,n}$, all embeddings are non-crossing and non-nesting, so their number is equal.

We know the number of non-crossing embeddings of a path on n vertices, which is $n2^{n-3}$. In the next section, we will show that there are at least $\Omega(n^{-3/2}2^n)$ non-nesting embeddings of a path.

We will now use the same basic idea as in Section 4, but a different bijection, to get some insight in the asymptotic behaviour of the number of non-nesting paths. In fact, we will count non-nesting paths with some restrictions. The first, and least essential restriction, is that we only look at paths whose endpoints are also endpoints of the interval to which we embed them, i.e. 1 and n . More essentially, we require the paths to be alternating, so every node has either both outgoing arcs to the left, or both to the right. In the words of Section 4, this means that for any i , (l_i, r_i) must be either $(0, 2)$ or $(2, 0)$.

With these restrictions, a non-nesting path is determined by deciding, for each $1 < i < n$, whether i should have rightbound or leftbound arcs. These choices can be made independently, but still with the restriction that $\sum_i l_i = \sum_i r_i$ and $\sum_{i \leq j-1} l_i \leq \sum_{i \leq j} r_i$ for each j . We thus get a Dyck path of length $2(n-2)$, where the i th step is an up-step if $(l_i, r_i) = (0, 2)$, and a down-step if $(l_i, r_i) = (2, 0)$. It follows that there are at least $\frac{1}{k} \binom{2k-2}{k-1} = \Omega(k^{-3/2}2^{2k})$ non-nesting paths of even length $2k$. Moreover, there are obviously at least as many non-nesting paths on $2k+1$ nodes, which we see by simply adding an edge $(2k, 2k+1)$ at the end. This shows the following proposition.

Proposition 5.5. *There are $\Omega(n^{-3/2}2^n)$ non-nesting paths of length n .*

This is probably a rather big underestimate. It could be expected that when we delete the restrictions on where the endpoints should be placed, this should multiply the number of paths with a roughly a factor n^2 . More notably, the request that the paths should be alternating is a huge restriction, that should influence the exponent, and thus the growth rate, significantly. Note that the fact that the growth rate is at least 2, does not prove that there are more non-nesting paths than non-crossing paths, though we strongly believe this to be true.

Now we will show that the number of non-nesting paths has a growth rate that is bounded from above by 3. Thus, there is still a significant gap between the lower and the upper bound.

Proposition 5.6. *There are at most $2n(n-1)3^n$ non-nesting embeddings of a path on n vertices.*

Proof. Look at the situation at a vertex v . If v is a leaf, the incoming edge can go to the left or to the right. There are $\binom{n}{2}$ possibilities to choose the leaves, and for each choice, there are at most 4 different situations how the incoming edges look like. The remaining $n-2$ vertices have degree 2. At each vertex both incoming edges can go to the left, to the right, or one goes to the left and the other one to the right. So at each vertex, there are at most 3 configurations possible. Once we fixed the situation at the vertices, there is at most one non-nesting path that fulfills these constraints. \square

Disregarding polynomial factors, we now know that the number of non-nesting paths is between 2^n and 3^n . We implemented a counting algorithm and computed the exact values up to 9. In the next table, the number of non-crossing and non-nesting paths, as well as every other Fibonacci number F_{2n} are compared.

n	nc paths	nn paths	F_{2n}
2	1	1	1
3	3	3	3
4	8	8	8
5	20	21	21
6	48	55	55
7	112	142	144
8	256	369	377
9	576	950	987

Conjecture 5.7. *The number of non-nesting paths on n vertices is at most the $2n$ th Fibonacci number F_{2n} , which is $\Theta((\frac{1+\sqrt{5}}{2})^n)$, that is roughly 2.62^n .*

6 Open problems

Although it seems complicated, it remains to count the number of non-nesting embeddings of trees or, more generally, give better asymptotical bounds.

It remains to see whether Conjecture 5.7 holds or, more generally, to count or narrow the asymptotical bounds on the number of non-nesting embeddings of paths in n vertices.

We have shown that, for some classes of graphs, the number of non-nesting graphs is larger than the for the non-crossing graph and viceversa. The search for other equalities or inequalities is wide open.

Acknowledgements

This paper was prepared during the DocCourse *Combinatorics and Geometry 2009* at CRM, Barcelona. The authors thank the organizers of this event, and especially Anna de Mier and Marc Noy for proposing and helpfully supervising the project.

References

- [1] P. Flajolet, M. Noy, *Analytic Combinatorics of Non-crossing Configurations*, Discrete Math. **204** (1999), 203–229.
- [2] Martin Klazar, *Counting pattern-free set partitions II: noncrossing and other hypergraphs*, Electron. J. Comb. **7** (2000), R34.
- [3] Anna de Mier, *k-noncrossing and k-nonnesting graphs and fillings of Ferrers diagrams*, Combinatorica **27** (2007), 699–720.
- [4] Gopal Mohanty, *Lattice Path Counting and Applications*, Academic Press, New York 1979.
- [5] Richard Stanley, *Enumerative Combinatorics*, vol.2, Cambridge University Press, New York 1999.

PRODSIMPLICIAL-NEIGHBORLY POLYTOPES

BENJAMIN MATSCHKE, BERNARDO GONZÁLEZ MERINO,
JULIAN PFEIFLE, AND VINCENT PILAUD

ABSTRACT. Given integers $n_1, \dots, n_r \geq 1$ and $k \geq 0$, we are interested in polytopes whose proper faces are combinatorially equivalent to faces of the product of simplices $\Delta_{n_1} \times \dots \times \Delta_{n_r}$, and whose k -skeleton is combinatorially equivalent to that of $\Delta_{n_1} \times \dots \times \Delta_{n_r}$. We prove in particular that such a polytope exists in dimension $2k + r + 1$. Using topological obstructions we see that this dimension is minimal if we require moreover our polytope to be obtained as a projection of a polytope combinatorially equivalent to $\Delta_{n_1} \times \dots \times \Delta_{n_r}$ and if the n_i 's are large compared to k . If some n_i 's are small we give even lower-dimensional examples using projections of deformed products.

Furthermore the last construction generalizes to products of arbitrary simple polytopes.

CONTENTS

1. Introduction	2
1.1. Examples of polytopes from graphs...	2
1.2. Definitions	2
1.3. Outline and main results	3
2. Construction of PPSN polytopes	4
2.1. Products of cyclic polytopes	5
2.2. Minkowski sums of cyclic polytopes	5
2.3. Reflections of cyclic polytopes	8
2.4. Projections of deformed products of polytopes	10
3. Topological obstruction	16
3.1. Preserved faces, Gale transform and topological obstruction	16
3.2. Preserving the k -skeleton of a product of simplices	18
3.3. Colorings of $\text{KG}(\mathcal{Z})$	19
Acknowledgments	25
References	25
0. Appendix: background material	26
0.1. Products of polytopes	26
0.2. Cyclic polytopes	27
0.3. Gale duality	29
0.4. Projections of polytopes and strictly preserved faces	31
0.5. Kneser graphs	32
0.6. Embeddability of simplicial complexes	33

1. INTRODUCTION

1.1. Examples of polytopes from graphs... During his course untitled *Convex Polytopes: Examples and Conjectures* given for the *Doc-Course on Combinatorics and Geometry 2009*, Günter Ziegler mentioned two problems:

- (1) the first one is our (relative) “lack of examples” on polytopes: according to him, knowing more families of polytopes would be profitable for our intuition on the subject (one can think in particular of disproving conjectures);
- (2) the second one is the question of the polytopality of a given graph: for example, one of the exercises asked whether the product of two Petersen graphs is or not the graph of a polytope (and what may be the dimension of this polytope).

This project is related with these two points, since it gives example of polytopes answering the following question: given an initial polytope, what is the minimal dimension of a polytope with the same graph (or more generally, with the same k -skeleton, $k \geq 0$)? For example, for any $k \geq 0$, the $(2k+2)$ -dimensional cyclic polytope with $n+1$ vertices has the same k -skeleton as the n -dimensional simplex.

In this project, we only tackle this problem for products of simplices, and deal with the following question: given $n_1, \dots, n_r \geq 1$ and $k \geq 0$, what is the minimal dimension of a polytope whose k -skeleton is combinatorially equivalent to that of the product of simplices $\Delta_{n_1} \times \dots \times \Delta_{n_r}$?

1.2. Definitions. Let Δ_n denote the n -dimensional simplex. For any tuple $\underline{n} = (n_1, \dots, n_r)$ of integers, we denote by $\Delta_{\underline{n}}$ the product of simplices $\Delta_{n_1} \times \dots \times \Delta_{n_r}$. The combinatorial structure of this polytope is that of a product (see Appendix 0.1): its dimension is $\sum n_i$ and its non-empty faces are obtained as products of non-empty faces of the simplices $\Delta_{n_1}, \dots, \Delta_{n_r}$. For example, Fig. 1 represents the graphs of $\Delta_i \times \Delta_6$, for $i = 0, 1, 2, 3$. We are interested in polytopes with the same “initial” structure as such products:

Definition 1.1. Let $k \geq 0$ and $\underline{n} = (n_1, \dots, n_r)$, with $n_i \geq 1$. A polytope is (k, \underline{n}) -prodsimplicial-neighborly – or (k, \underline{n}) -PSN for short – if

- (i) its proper faces are combinatorially equivalent to faces of $\Delta_{\underline{n}}$, and
- (ii) its k -skeleton is combinatorially equivalent to that of $\Delta_{\underline{n}}$.

This definition is essentially motivated by two particular classes of PSN polytopes:




FIGURE 1. The graphs of the products $\Delta_{(i,6)} = \Delta_i \times \Delta_6$, for $i = 0, 1, 2, 3$.

- (1) *neighborly*¹ *simplicial* polytopes arise when $r = 1$; for example, the cyclic polytope $C_d(n + 1)$ is simplicial and has the same $(\lfloor \frac{d}{2} \rfloor - 1)$ -skeleton as the simplex Δ_n (see Appendix 0.2).
- (2) *neighborly cubical* polytopes arise when $\underline{n} = (1, 1, \dots, 1)$ (see [JS07], [SZ09]).

Obviously, the product $\Delta_{\underline{n}}$ itself is a (k, \underline{n}) -PSN polytope of dimension $\sum n_i$. We are naturally interested in finding (k, \underline{n}) -PSN polytopes in smaller dimensions. We denote $\delta(k, \underline{n})$ the smallest possible dimension that a (k, \underline{n}) -PSN polytope may have.

Certain particular examples of such polytopes will be obtained as projections into a smaller subspace of the product $\Delta_{\underline{n}}$ itself, or of a combinatorially equivalent polytope. For example, the cyclic polytope $C_d(n + 1)$ (as any polytope with $n + 1$ vertices) can be seen as a projection of the simplex Δ_n into \mathbb{R}^d . We say that a (k, \underline{n}) -PSN polytope which is a projection of a polytope combinatorially equivalent to $\Delta_{\underline{n}}$ is a (k, \underline{n}) -*projected-prodsimplicial-neighborly* polytope – or (k, \underline{n}) -*PPSN* polytope for short. We denote $\delta_{pr}(k, \underline{n})$ the smallest possible dimension of a (k, \underline{n}) -PPSN polytope.

1.3. Outline and main results. In the second section of this paper, we present various constructions of PSN polytopes. With a well-chosen Minkowski sum of cyclic polytopes we obtain in particular the following theorem:

¹In the literature, a polytope is *k-neighborly* if any subset of at most k of its vertices forms a face. Observe that such a polytope is $(k - 1, n)$ -PSN with our notation, as long as it is simplicial.

Theorem 1.2. *For any $k \in \mathbb{N}$ and $\underline{n} = (n_1, \dots, n_r)$, there exists a (k, \underline{n}) -PPSN polytope of dimension $2k + r + 1$. In other words:*

$$\delta(k, \underline{n}) \leq \delta_{pr}(k, \underline{n}) \leq 2k + r + 1.$$

For $r = 1$ we recover simplicial neighborly polytopes. We then discuss improvements of this bound when some n_i are small compared to k , using projections of deformed products of polytopes. If all $n_i = 1$ we recover the cubical neighborly polytopes of [SZ09]. This construction is not only applicable to products of simplices but also to products of arbitrary simple polytopes.

In our third section, we use a topological obstruction method to prove the following lower bounds on $\delta_{pr}(k, \underline{n})$:

Theorem 1.3. *Let $k \in \mathbb{N}$, $\underline{n} = (n_1, \dots, n_r)$, and $R = \{i \in [r] \mid n_i \geq 2\}$. If $k \leq \sum_{i \in R} \lfloor \frac{n_i - 2}{2} \rfloor$, then the minimal dimension of a (k, \underline{n}) -PPSN polytope satisfies:*

$$\delta_{pr}(k, \underline{n}) \geq 2k + |R| + 1.$$

In particular, we obtain:

Corollary 1.4. *When all $n_i \geq 2$ and if $k \leq \sum_{i \in [r]} \lfloor \frac{n_i - 2}{2} \rfloor$, the minimal dimension is exactly*

$$\delta_{pr}(k, \underline{n}) = 2k + r + 1.$$

2. CONSTRUCTION OF PPSN POLYTOPES

In this section, we give upper bounds on $\delta_{pr}(k, \underline{n})$ by constructing examples of polytopes with the prescribed skeleton whose dimension is as small as we can. We present four different constructions:

- (1) the first (elementary) construction is just a product of cyclic polytopes, but already proves that $\delta_{pr}(k, \underline{n})$ is at most $(2k + 2)r$ (and thus, does not depend on how big are the n_i 's).
- (2) the second construction is a well-chosen Minkowski sum of cyclic polytopes, and provides a (k, \underline{n}) -PPSN polytope of dimension $2k + r + 1$. When k is small compared to the n_i 's, we will later prove that this dimension is the best possible dimension of a (k, \underline{n}) -PPSN polytope (see Section 3).
- (3) the third one is a special construction of a $(2k + 2)$ -dimensional $(k, (1, n))$ -PPSN polytope, that only uses elementary combinatorial properties of the face description of the cyclic polytope.
- (4) finally, our last construction uses projections of deformed products of polytopes, and deals more generally with the case when some of the n_i 's are small.

2.1. Products of cyclic polytopes. As mentioned previously, the cyclic polytope $C_d(n+1)$ is simplicial and has the same $(\lfloor \frac{d}{2} \rfloor - 1)$ -skeleton as the simplex Δ_n (see Appendix 0.2). Thus, we can obtain PSN polytopes using products of cyclic polytopes.

Given the integer k and the tuple $\underline{n} = (n_1, \dots, n_r)$, we denote $I = \{i \in [r] \mid n_i > 2k + 2\}$. Then the product

$$P = \prod_{i \in I} C_{2k+2}(n_i + 1) \times \prod_{i \notin I} \Delta_{n_i}$$

is a (k, \underline{n}) -PPSN polytope of dimension $(2k + 2)|I| + \sum_{i \notin I} n_i$ (which is smaller than $\sum n_i$ when I is not empty). Consequently,

$$\delta(k, \underline{n}) \leq \delta_{pr}(k, \underline{n}) \leq (2k + 2)|I| + \sum_{i \notin I} n_i.$$

2.2. Minkowski sums of cyclic polytopes. Our second examples will be Minkowski sums of cyclic polytopes. Despite their simplicity they turn out to be the best possible PSN polytopes for large n_i 's.

We will choose for all $i \in [r]$ an $(n_i + 1)$ -set $I_i \subset \mathbb{R}$ such that the intervals $\text{conv}(I_i)$ are pairwise disjoint. Then $I_1 \times \dots \times I_r$ indexes the vertex set of $\Delta_{\underline{n}} = \Delta_{n_1} \times \dots \times \Delta_{n_r}$. For all $(a_1, \dots, a_r) \in I_1 \times \dots \times I_r$ define the following points in \mathbb{R}^{2k+r+1} :

$$v_{a_1, \dots, a_r} = \left(a_1, \dots, a_r, \sum_{i \in [r]} a_i^2, \dots, \sum_{i \in [r]} a_i^{2k+2} \right)^T.$$

Let P be the convex hull of these $\prod (n_i + 1)$ vertices.

Proposition 2.1. *The above index sets I_1, \dots, I_r can be chosen such that the resulting $(\leq 2k + r + 1)$ -dimensional polytope P has combinatorially the same k -skeleton as $\Delta_{\underline{n}}$.*

Observe that this proposition proves Theorem 1.2: $\delta_{pr}(k, \underline{n}) \leq 2k + r + 1$. We will see in Section 3 that this bound is sharp when k is small compared to the n_i 's.

Proof of Proposition 2.1. Let F be a k -face of $\Delta_{\underline{n}}$ whose vertex set is $A_1 \times \dots \times A_r \subset I_1 \times \dots \times I_r$, each A_i being non-empty. We will construct a *monic* polynomial (that is, a polynomial with leading coefficient equal to 1)

$$f(t) = f_{A_1, \dots, A_r}(t) = \sum_{j=0}^{2k+2} c_j t^j$$

that has, for all $i \in [r]$, the form

$$f(t) = Q_i(t) \prod_{a \in A_i} (t - a)^2 + s_i t + r_i$$

with positive polynomials $Q_i(t)$ and certain reals r_i and s_i . From the coefficients of this polynomial, we build the vector

$$n_F = (s_1 - c_1, \dots, s_r - c_1, -c_2, -c_3, \dots, -c_{2k+2}) \in \mathbb{R}^{2k+r+1}.$$

To prove that n_F is a face-defining normal vector for F , take an arbitrary $(a_1, \dots, a_r) \in I_1 \times \dots \times I_r$, and consider the following inequality for the inner product:

$$\begin{aligned} n_F \cdot v_{a_1, \dots, a_r} &= \sum_{i \in [r]} \left(s_i a_i - \left(\sum_{j=1}^{2k+2} c_j a_i^j \right) \right) \\ &= \sum_{i \in [r]} (s_i a_i + c_0 - f(a_i)) \\ &= \sum_{i \in [r]} \left(c_0 - Q_i(a_i) \prod_{a \in A_i} (a_i - a)^2 - r_i \right) \\ &\leq r c_0 - \sum_{i \in [r]} r_i. \end{aligned}$$

Equality holds if and only if $(a_1, \dots, a_r) \in A_1 \times \dots \times A_r$. Given the existence of a polynomial f with the claimed properties, this proves that $A_1 \times \dots \times A_r$ indexes all v_{a_1, \dots, a_r} 's that lie on a face F' in P , and they of course span F' by definition of P . To prove that F' is combinatorially equivalent to F it suffices to show that each $v_{a_1, \dots, a_r} \in F'$ is in fact a vertex of P , since P is a projection of $\Delta_{\underline{n}}$. This can be shown with the normal vector $(2a_1, \dots, 2a_r, -1, 0, \dots, 0)$, using the same calculation as before.

We still need to show how to construct our polynomial $f(t)$. Finding $f(t)$ is equivalent to the task of finding polynomials $Q_i(t)$ such that

- (1) $Q_i(t)$ is monic of degree $2k + 2 - 2|A_i|$.
- (2) The r polynomials $f_i(t) = Q_i(t) \cdot \prod_{a \in A_i} (t - a)^2$ are equal up to possibly the coefficients in front of t^0 and t^1 .
- (3) $Q_i(t) > 0$ for all $t \in \mathbb{R}$.

The first two items form a linear equation system on the coefficients of the $Q_i(t)$'s, which has the same number of equations as variables. We will show that it has a unique solution if one chooses the index sets I_i in a good way (the third item will be dealt with in the end). To

do this, choose pairwise distinct reals $\bar{a}_1, \dots, \bar{a}_r \in \mathbb{R}$ and look at the similar equation system:

- (1) $\bar{Q}_i(t)$ are monic polynomials of degree $2k + 2 - 2|A_i|$.
- (2) The r polynomials $\bar{f}_i(t) = \bar{Q}_i(t) \cdot (t - \bar{a}_i)^{2|A_i|}$ are equal up to possibly the coefficients in front of t^0 and t^1 .

The first equation system moves into the second when we deform the points of the sets A_i continuously to \bar{a}_i , respectively. If we show that the second equation system has a unique solution then so has the first equation system as long as we have chosen the sets I_i close enough to the \bar{a}_i 's, by continuity of the determinant (note that in the end this closeness condition will be required for all k -faces, which are only finitely many, hence we can fulfill that).

Note that a polynomial $\bar{f}_i(t)$ of degree $2k + 2$ has the form

$$(1) \quad \bar{Q}_i(t) \cdot (t - \bar{a}_i)^{2|A_i|} + s_i t + r_i$$

for a monic polynomial \bar{Q}_i and some reals s_i and r_i , if and only if $\bar{f}_i''(t)$ has the form

$$(2) \quad R_i(t) \cdot (t - \bar{a}_i)^{2(|A_i|-1)}$$

for a polynomial $R_i(t)$ with leading coefficient $(2k + 2)(2k + 1)$. The backward direction can be settled by assuming without loss of generality $\bar{a}_i = 0$ (otherwise just make a variable shift $(t - \bar{a}_i) \mapsto t$) and then integrating (2) twice with integration constants zero to obtain (1).

Therefore the second equation system is equivalent to the following third one:

- (1) $R_i(t)$ are polynomials of degree $2k - 2(|A_i| - 1)$ with leading coefficient $(2k + 2)(2k + 1)$.
- (2) The r polynomials $g_i(t) = R_i(t) \cdot (t - \bar{a}_i)^{2(|A_i|-1)}$ all equal the same polynomial, say $g(t)$.

Since $\sum_i 2(|A_i| - 1) = 2k$, it has the unique solution

$$R_i(t) = (2k + 2)(2k + 1) \prod_{j \neq i} (t - \bar{a}_j)^{2(|A_j|-1)},$$

with

$$g(t) = (2k + 2)(2k + 1) \prod_{j \in [r]} (t - \bar{a}_j)^{2(|A_j|-1)}.$$

Therefore the second system also has a unique solution, where the $\bar{f}_i(t)$ are obtained by integrating $g_i(t)$ twice with some specific integration constants. For a fixed i we can again assume $\bar{a}_i = 0$. Then both integration constants have been zero for this i , hence $\bar{f}_i(0) = 0$ and $\bar{f}_i'(0) = 0$. Since g_i is non-negative and zero only at isolated points, \bar{f}_i

is strictly convex, hence non-negative and zero only at $t = 0$. Therefore $\bar{Q}_i(t)$ is positive for $t \neq 0$. Since we chose $\bar{a}_i = 0$, we can quickly compute the correspondence between the coefficients of $\bar{Q}_i(t) = \sum_j \bar{q}_{i,j} t^j$ and of $R_i(t) = \sum_j r_{i,j} t^j$:

$$r_{i,j} = (2|A_i|(2|A_i| - 1) + 4j|A_i| + j(j - 1))\bar{q}_{i,j}.$$

In particular

$$\bar{Q}_i(0) = \bar{q}_{i,0} = \frac{r_{i,0}}{2|A_i|(2|A_i| - 1)} = \frac{R_i(0)}{2|A_i|(2|A_i| - 1)} > 0,$$

therefore $\bar{Q}_i(t)$ is everywhere positive, hence so is $Q_i(t)$ if one chooses I_i possibly even closer to \bar{a}_i , since the solutions of linear equation systems move continuously when one deforms the entries of the equation system by a homotopy (as long as the determinant stays non-zero), since the determinant and taking the adjoint matrix are continuous maps. The positivity of $Q_i(t)$ finishes the proof. \square

Remark 2.2. We can use the same idea to obtain another PPSN polytope (the result is weaker, but the proof is much simpler and avoids all technicalities). Let $I_i \subset \mathbb{R}$ be pairwise disjoint $(n_i + 1)$ -sets as before. For all $(a_1, \dots, a_r) \in I_1 \times \dots \times I_r$ define points

$$w_{a_1, \dots, a_r} = \left(\sum_{i \in [r]} a_i, \sum_{i \in [r]} a_i^2, \dots, \sum_{i \in [r]} a_i^{2k+2r} \right)^T \in \mathbb{R}^{2k+2r}.$$

The convex hull P of these points is a (k, \underline{n}) -PPSN polytope, without further restrictions on the I_i 's.

Sketch of proof. Let $A_1 \times \dots \times A_r \subset I_1 \times \dots \times I_r$ define a k -face of $\Delta_{\underline{n}}$. To prove that the corresponding points in P defines a face, show that the coefficient vector $\prod_{i \in [r]} \prod_{a \in A_i} (t - a)^2$ without the coefficient of t^0 – as a normal vector – defines this face. \square

2.3. Reflections of cyclic polytopes. Our third example deals with the special case of the product $\Delta_1 \times \Delta_n$ of a segment by a simplex, for which the previous construction is not optimal. Even if it is only a particular case of a projection of a deformed product of $\Delta_1 \times \Delta_n$ (see Section 2.4), we present it separately since we have an explicit description and a nice combinatorial proof of it.

We consider the cyclic polytope $C_{2k+2}(n+1)$ with $n+1$ vertices in dimension $2k+2$ (see Appendix 0.2). We obtain a $(k, (1, n))$ -PSN polytope as the convex hull of the union of this polytope with another copy of it, reflected with respect to a well-chosen hyperplane:

Proposition 2.3. *For any $\lambda \in \mathbb{R}$ sufficiently large, the polytope P defined as*

$$\text{conv}(\{(t_i, \dots, t_i^{2k+2})^T \mid i \in [n+1]\} \cup \{(t_i, \dots, t_i^{2k+1}, \lambda - t_i^{2k+2})^T \mid i \in [n+1]\})$$

is a $(k, (1, n))$ -PSN polytope of dimension $2k + 2$.

Proof. The polytope P is obtained as the convex hull of two copies of the cyclic polytope $C_{2k+2}(n+1)$: the first one $Q = \text{conv}(\{\mu_{2k+2}(t_i) \mid i \in [n+1]\})$ lies on the moment curve, while the second one is obtained as a reflexion of Q with respect to an hyperplane far enough along and orthogonal to the last coordinate vector u_{2k+2} . During this process,

- (1) we destroy all the faces of Q only contained in *upper* facets of Q (those whose supporting hyperplane lies above the whole cyclic polytope while looking at the last coordinate; see Appendix 0.2);
- (2) we create prisms over faces of Q that lie in at least one upper and one lower facet of Q . In other words (see Appendix 0.4), we create prisms over the faces of Q strictly preserved under the orthogonal projection $\pi : \mathbb{R}^{2k+2} \rightarrow \mathbb{R}^{2k+1}$ with kernel $\mathbb{R}u_{2k+2}$.

The projected polytope $\pi(Q) = \text{conv}\{\pi(\mu_{2k+2}(t_i)) \mid i \in [n+1]\}$ is nothing else but the cyclic polytope $C_{2k+1}(n+1)$. This polytope being k -neighborly, any $(\leq k-1)$ -face of Q is strictly preserved by π , and thus, we take a prism over all $(\leq k-1)$ -faces of Q .

Thus, in order to prove that the k -skeleton of P is that of $\Delta_1 \times \Delta_n$, it only remains to prove that any k -face of Q still remains in P . It is obviously the case if this k -face is also a k -face of $C_{2k+1}(n+1)$, and follows from the following combinatorial lemma otherwise. \square

Lemma 2.4. *A k -face of $C_{2k+2}(n+1)$ which is not a k -face of $C_{2k+1}(n+1)$ is only contained in lower facets of $C_{2k+2}(n+1)$.*

Proof. Let $F \subset [n+1]$ be a k -face of $C_{2k+2}(n+1)$ contained in at least one upper facet of $C_{2k+2}(n+1)$. Then, there exists $G \subset [n+1] \setminus F$ such that $|F \cup G| = 2k+2$, all inner blocks of $F \cup G$ have even size, and the final block of $F \cup G$ has odd size (in other words, G completes F into an upper facet). Then,

- (1) if $n+1 \in G$, then $G' = G \setminus \{n+1\}$ is such that $|F \cup G'| = 2k+1$ and all inner blocks of $F \cup G'$ have even size.
- (2) otherwise, $n+1 \in F$, and $F' = F \setminus \{n+1\}$ has only k elements and thus, is a face of $C_{2k+1}(n)$, that we can complete into a face with $G' \subset [n] \setminus F$.

In both cases, G' completes F into a facet of $C_{2k+1}(n+1)$, which implies that F is a k -face of $C_{2k+1}(n+1)$. \square

2.4. Projections of deformed products of polytopes. Our next examples will be projections of deformed products of simplices. They give better results than the construction of Section 2.2 when some of the n_i 's are small compared to k , and in particular when some of the simplices are segments (as in the previous construction).

We first give a short description of deformed products of polytopes, which have been invented in [AZ99] and used in [Zie04] and [SZ09], and then apply this construction to products of simplices.

Deformed products. Let P_1, \dots, P_r be *simple* polytopes given by $P_i = \{x \in \mathbb{R}^{n_i} \mid A_i x \leq b_i\}$ (where $A_i \in \mathbb{R}^{n_i \times n_i}$ and $b_i \in \mathbb{R}^{n_i}$). The product $P = P_1 \times \dots \times P_r$ is then given by the inequality system

$$\begin{pmatrix} A_1 & & \\ & \ddots & \\ & & A_r \end{pmatrix} x \leq \begin{pmatrix} b_1 \\ \vdots \\ b_r \end{pmatrix}$$

with $n = \sum n_i$ inequalities. The left hand side matrix shall be denoted by A . In [AZ99] it is proved that one can make A to a matrix A' by changing the zero entries below the diagonal blocks in A *arbitrarily* such that for a certain new right hand side b' the new inequality system $A'x \leq b'$ defines a polytope P' that is combinatorially equivalent to P . This equivalence is given by letting the facet of P corresponding to the i 'th inequality of $Ax \leq b$ correspond to the facet of P' corresponding to the i 'th inequality of $A'x \leq b'$.

Projections of deformed products of simplices. We want to construct a d -dimensional polytope that has combinatorially the same k -skeleton as $\Delta_{\underline{n}} = \Delta_{n_1} \times \dots \times \Delta_{n_r}$, such that d is as small as possible for fixed k , or k is as large as possible for fixed d . The former formulation seems to be the more natural question, whereas the results for the latter one are easier to write down. Hence for the following propositions, think of the n_i 's and d as being given, and ask for the largest k . The results fit into two propositions, since we have to do a case distinction on the number of segments $s = |\{i \in [r] \mid n_i = 1\}|$ in the product.

Proposition 2.5. *Suppose that $s \leq d$, and assume $1 = n_1 = \dots = n_s < n_{s+1} \leq \dots \leq n_r$. Let $t \in \{0, \dots, r - s\}$ be maximal such that $\sum_{i=1}^{s+t} n_i \leq d$. If further $d \geq 2k + 2(r - s) - t$, then there exists a d -dimensional (k, \underline{n}) -PPSN polytope.*

Proposition 2.6. *Suppose that $s \geq d + 1$ and $d \geq 2k + 2(r - s) + 2$. Then there is a d -dimensional (k, \underline{n}) -PPSN polytope.*

Proof of Proposition 2.5. The simplex Δ_{n_i} will for us be the set $\{x \in \mathbb{R}^{n_i} \mid A_i x \leq b_i\}$, where A_i is the matrix

$$A_i = \begin{pmatrix} 1 & & \\ & \ddots & \\ & & 1 \\ -1 & \dots & -1 \end{pmatrix}$$

and b_i is a suitable right hand side. The matrix of the inequality system defining $\Delta_n = \Delta_{n_1} \times \dots \times \Delta_{n_r}$ is a block matrix with blocks A_i . We deform Δ_n to a combinatorially equivalent polytope Δ'_n by changing A , first in the block A_{s+t+1} and then below the diagonal blocks (the right hand side of the inequality system has to be changed accordingly). We obtain the following matrix A' :

$$\left(\begin{array}{c|c} \begin{array}{c} \boxed{A_r} \\ \vdots \\ \boxed{A_{s+t+2}} \end{array} & \begin{array}{c} \\ \\ 1 \\ \vdots \\ 1 \end{array} \\ \hline \begin{array}{c} g_{d+t}^T \\ \vdots \\ g_{j+1}^T \\ g_{d+t+1}^T \\ g_j^T \\ g_{j-1}^T \\ \vdots \\ g_{s+2}^T \\ g_{s+1}^T \\ g_s^T \\ g_s^T \\ \vdots \\ g_1^T \\ g_1^T \end{array} & \begin{array}{c} M \\ \vdots \\ M \\ -1 \quad \dots \quad -1 \\ \\ \boxed{A_{s+t}} \\ \vdots \\ \boxed{A_{s+1}} \\ \\ \begin{array}{c} +1 \\ -1 \end{array} \\ \vdots \\ \begin{array}{c} +1 \\ -1 \end{array} \end{array} \end{array} \right)$$

The vertical line in A' separates the first $n-d$ from the last d columns. It possibly slices the block A_{s+t+1} , which is the block with the smallest index that contains elements of the left side. The left part of A' is denoted by A_L .

The g_i 's are chosen such that $G = \{u_1, \dots, u_{n-d}, g_1, \dots, g_{d+t+1}\}$ is a Gale dual (see Appendix 0.3) of the cyclic polytope $C_{d+t}(n+t+1)$ with the further requirement that the vertices corresponding to g_1, \dots, g_{d+t} lie in a facet. This makes $\{u_1, \dots, u_{n-d}, g_{d+t+1}\}$ positively spanning. Hence g_{d+t+1} has only negative entries. Therefore the rows of the new block A_{s+t+1} are still normal vectors of a simplex as long as we choose $M > 0$ big enough with respect to the g_i 's.

The cyclic polytope $C_{d+t}(n+t+1)$ is $\lfloor \frac{d+t}{2} \rfloor$ -neighborly. Therefore, every subset of $n+t+1 - \lfloor \frac{d+t}{2} \rfloor = n-d+1 + \lfloor \frac{d+t+1}{2} \rfloor$ vectors of G is positively spanning.

Now, a k -face F of Δ'_n is an intersection of $n-k$ facets of Δ'_n . These $n-k$ facets have the property that they do not contain all the facets corresponding to some factor Δ_{n_i} . The $n-k$ facets determine $n-k$ *pairwise unequal* row vectors in A_L . Except for possibly $r-s-t-1$ rows of the form $(0, \dots, 0, -1, \dots, -1, 0, \dots, 0)$, all of them are vectors of G . Therefore, if we know that

$$(3) \quad n-k \geq (n-d+1 + \lfloor \frac{d+t+1}{2} \rfloor) + (r-s-t-1),$$

then the $n-k$ rows are positively spanning and the face F is strictly preserved under the projection π to the last d coordinates (see Appendix 0.4). But (3) is exactly the assumption $d \geq 2k + 2(r-s) - t$. Therefore $\pi(\Delta'_n)$ is a d -polytope with the desired combinatorial k -skeleton. \square

Proof of Proposition 2.6. The situation here is similar to the previous proof. At first we scale the A_i 's for $i \in \{d+1, \dots, s\}$ by a small ε , which changes Δ_n only by an affine transformation. Then we deform Δ_n by changing A below the diagonal to the following matrix A' :

$$\begin{pmatrix}
 & \overset{n-d}{\begin{matrix} A_r \\ \vdots \\ A_{s+1} \\ \begin{matrix} +\varepsilon \\ -\varepsilon \end{matrix} \\ 1 \quad \begin{matrix} +\varepsilon \\ -\varepsilon \end{matrix} \\ 1 \quad 1 \\ \vdots \\ 1 \end{matrix}} & \overset{d}{\begin{matrix} \\ \\ \\ \\ \\ \\ \\ \\ \\ \end{matrix}} \\
 \hline
 \begin{matrix} g_{d-1}^T \\ g_{d-1}^T \\ \vdots \\ g_1^T \\ g_1^T \end{matrix} & \begin{matrix} \begin{matrix} +\varepsilon \\ -\varepsilon \end{matrix} \\ 1 \quad +1 \\ 1 \quad -1 \\ \begin{matrix} +1 \\ -1 \end{matrix} \\ \vdots \\ \begin{matrix} +1 \\ -1 \end{matrix} \end{matrix}
 \end{pmatrix}$$

where the g_i 's are chosen such that $G = \{u_1, \dots, u_{n-d}, g_1, \dots, g_{d-1}\} \subset \mathbb{R}^{n-d}$ is a Gale dual of the cyclic polytope $C_{d-2}(n-1)$, which is $\lfloor \frac{d-2}{2} \rfloor$ -neighborly. Therefore, each choice of $n-1 - \lfloor \frac{d-2}{2} \rfloor = n - \lfloor \frac{d}{2} \rfloor$ vectors of G are positively spanning. This property of G does not change if we perturb the vectors by a small ε .

Now, a k -face F of Δ'_n is an intersection of $n-k$ facets of Δ'_n . These $n-k$ facets have the property that they do not contain all the facets corresponding to some factor Δ_{n_i} . The $n-k$ facets determine $n-k$ row vectors in A_L . Except for possibly the $r-s$ rows $(0, \dots, 0, -1, \dots, -1, 0, \dots, 0)$ of A_L and *one* of the special rows $(0, \dots, \pm\varepsilon, \dots, 0)$ of A_L , all the $n-k$ rows are up to an ε -entry elements of G and they are pairwise distinct. The assumption $d \geq 2k+2(r-s)+2$ is equivalent to $n-k \geq (n - \lfloor \frac{d}{2} \rfloor) + r-s + 1$. Therefore the $n-k$ rows for F are positively spanning, even if we remove the possible row $(0, \dots, \pm\varepsilon, \dots, 0)$. Hence F is strictly preserved under π , the projection to the last d coordinates. Therefore $\pi(\Delta'_n)$ is a d -polytope with the desired combinatorial k -skeleton. \square

Finally, we reformulate Propositions 2.5 and 2.6 to the following theorem to express what is the smallest d such that for given n_i 's and k we can construct a desired d -dimensional polytope.

Theorem 2.7. *Let $1 \leq n_1 \leq \dots \leq n_r$, $\underline{n} = (n_1, \dots, n_r)$, $k \geq 0$ and $s = |\{i \mid n_i = 1\}|$.*

- *If $3s \leq 2k + 2r$, let $t \in \{0, \dots, r - s\}$ be maximal such that*

$$\sum_{i=s+1}^{s+t} (n_i + 1) + 3s \leq 2k + 2r.$$

Then $\delta_{pr}(k, \underline{n}) \leq 2k + 2(r - s) - t$.

- *If $3s = 2k + 2r + 1$, then $\delta_{pr}(k, \underline{n}) \leq 2k + 2(r - s) + 1$.*
- *If $3s \geq 2k + 2r + 2$, then $\delta_{pr}(k, \underline{n}) \leq 2k + 2(r - s) + 2$.*

Proof. If $3s \leq 2k + 2r + 2$ then apply Proposition 2.5, otherwise apply Proposition 2.6. \square

An extension to products of arbitrary simple polytopes. A very similar construction gives us examples of d -dimensional polytopes that are on a high skeleton combinatorially equivalent to an arbitrary product $P = P_1 \times \dots \times P_r$ of simple polytopes.

Theorem 2.8. *Let P_1, \dots, P_r be simple polytopes. Let $n_i = \dim P_i$ be their dimension, $n = \sum_i n_i$ and $f_i = f_{n_i-1}(P_i)$ their number of facets. Let $\chi_i^\Delta = \chi((P_i^\Delta)_1)$ be the chromatic number of the graph of the polar polytope of P_i . Let q be maximal with $\sum_{i=1}^q n_i \leq d$. Define $h = \sum_{i=1}^q n_i - d$, and let $\ell = h + \sum_{i=1}^q \chi_i^\Delta$. Then for*

$$k = d - \left\lfloor \frac{\ell}{2} \right\rfloor - 1 - \sum_{i=q+1}^r (f_i - n_i),$$

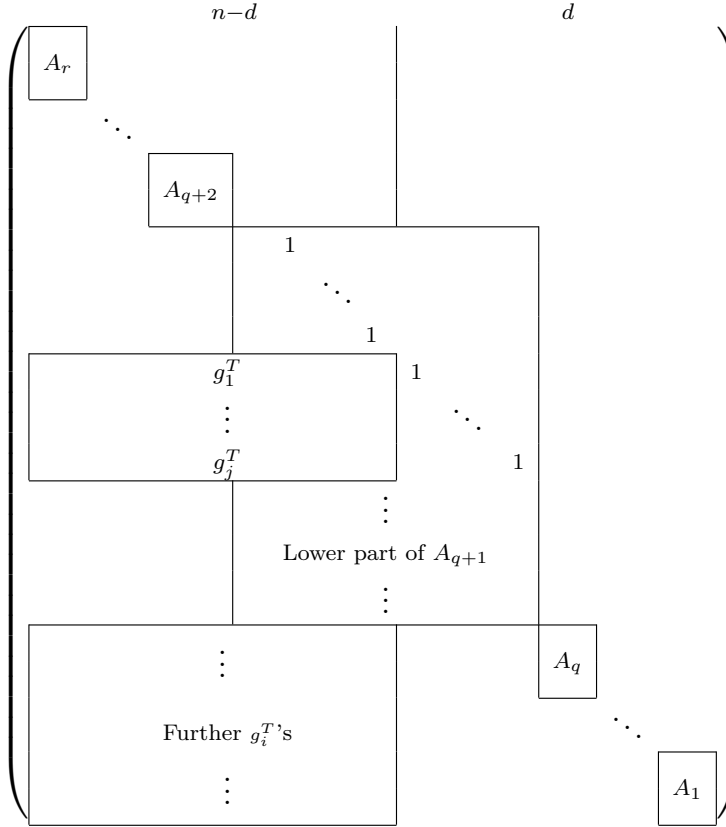
there is a d -dimensional polytope whose k -skeleton is combinatorially equivalent to the one of $P = P_1 \times \dots \times P_r$.

Remarks 2.9. (1) Heuristically (if r is much larger than the n_i 's), one should sort the factors P_1, \dots, P_r in descending order by $\frac{2f_i - \chi_i^\Delta}{n_i}$ to make k large.

- (2) The construction in the proof only needs P_1, \dots, P_{q+1} to be simple. In fact, using the next remark we only need to assume P_1, \dots, P_q to be simple, since then $h = 0$.
- (3) If $h > 0$, then one can add h intervals to the factors P_1, \dots, P_r : $\tilde{P} = (\Delta^1)^h \times P$, $\tilde{h} = 0$, $\tilde{r} = r + h$, $\tilde{\chi}_i^\Delta = 1$ for $i \in \{1, \dots, h\}$, $\tilde{\ell} = \ell$, $\tilde{q} = q + h$, $\tilde{k} = k$. Thus, we can even find a d -dimensional polytope whose k -skeleton is equivalent to the one of \tilde{P} .
- (4) If some of the factors P_1, \dots, P_r are segments or polygons with an even number of vertices, then we can sometimes strengthen the result in the same way as done in Proposition 2.6 for the

segments or in [Zie04] and [SZ09, Section 4] for polygons. That is, assume the corresponding A_i blocks to be (diagonally) adjacent, scale then down by a small ε and write ones at the entries that lie below A_i and to left of A_{i-1} .

Proof of Theorem 2.8. The proof is very similar to the one of Proposition 2.5. At first we write P_i as $\{x \in \mathbb{R}^{n_i} \mid A_i x \leq b_i\}$, such that A_i is a lower triangular matrix with ones on the diagonal. We can even prescribe the entries under the diagonal that belong to the first n_i rows. Then, in the construction of A' , we also write to the left of each of the lower blocks A_i row vectors g_1, \dots, g_j , but we are allowed to do some repetitions. We obtain the following new matrix A' :



The “further g_i^T ’s” block, say G , is allowed to contain some g_i ’s repeating: In the facet-ridge graph of P_i , color the nodes corresponding to P_i ’s facets with χ_i^Δ colors. This also colors the rows of the block G . To see that we may use the same g_i in every color class, recall that P is simple, so that each pair of facets containing a fixed k -face of P has a ridge in common. Thus, the g_i ’s corresponding to these facets will be pairwise distinct. The further proof is analogous to that of Proposition 2.5. \square

3. TOPOLOGICAL OBSTRUCTION

In this section, we give lower bounds on the minimal dimension $\delta_{pr}(k, \underline{n})$ of a (k, \underline{n}) -PPSN polytope. The argument is a topological obstruction based on Sarkaria's criterion for the embeddability of a simplicial complex in terms of colorings of Kneser graphs [Mat03]. We proceed in a similar way to [San09] and generalize his results. For the convenience of the reader, we quickly recall in the Appendix most of the usual tools used in this section: we present Gale transform in Appendix 0.3, we give in Appendix 0.4 a characterization of strictly preserved faces under projections, and we recall in Appendix 0.6 Sarkaria's Coloring and Embedding Theorem.

3.1. Preserved faces, Gale transform and topological obstruction. We are interested in projections of polytopes. For the sake of simplicity, and without loss of generality, we will only consider orthogonal projections. Let $d > e$ be two integers. We denote (u_1, \dots, u_d) the canonical orthogonal basis of \mathbb{R}^d . Let $\pi : \mathbb{R}^d \rightarrow \mathbb{R}^e$ denote the orthogonal projection with image $U = \langle u_1, \dots, u_e \rangle$, and kernel $V = \langle u_{e+1}, \dots, u_d \rangle$. Let $\tau : \mathbb{R}^d \rightarrow \mathbb{R}^{d-e}$ denote the dual projection (with image V and kernel U).

Let P be a full-dimensional simple polytope in \mathbb{R}^d , with 0 in its interior. Let F_1, \dots, F_f denote the facets of P and for all $i \in [f]$, let \vec{f}_i denote the normal vector of F_i , and $\vec{g}_i = \tau(\vec{f}_i)$. For any face F of P , let $I(F)$ denote the set of indices of the facets of P containing F , *i.e.*, such that $F = \bigcap_{i \in I(F)} F_i$. These notations are convenient for the Projection Lemma (see Appendix 0.4), that we will use extensively, and which affirms that a face F of P is strictly preserved under the projection π if and only if $\{\vec{g}_i \mid i \in I(F)\}$ is positively spanning.

Throughout this section, we assume that the vertices of P are strictly preserved under the projection π . Let $G = \{\vec{g}_i \mid i \in [f]\} \subset \mathbb{R}^{d-e}$ denote all the (dual) projections of the normal vectors of the facets of P .

Lemma 3.1. *The vector configuration G is the Gale transform of the vertex set $X = \{x_1, \dots, x_f\}$ of a (full-dimensional) polytope Q of $\mathbb{R}^{f-d+e-1}$.*

Proof. Observe first that the normal vectors of the facets of a full dimensional polytope in \mathbb{R}^d span \mathbb{R}^d . Thus, their projections span the image space \mathbb{R}^{d-e} .

In order to make the barycenter of G is 0, we choose the length of the normal vector \vec{f}_i of the facet F_i as the $(d-1)$ -dimensional volume of F_i . Given any vector \vec{z} , we say that a facet F_i of P is \vec{z} -upper when

$\vec{z} \cdot \vec{f}_i > 0$ (and \vec{z} -lower when this scalar product is negative). Then, the sum of the components on \vec{z} of the normal vectors of the \vec{z} -upper facets (resp. of the \vec{z} -lower facets) is given by the volume of the projection of P on \vec{z}^\perp . Thus, $\sum_{i \in [f]} \vec{f}_i \cdot \vec{z} = 0$, for all $\vec{z} \in \mathbb{R}^d$. Consequently, $\sum_{i \in [f]} \vec{g}_i = \tau(\sum_{i \in [f]} \vec{f}_i) = 0$, and G is the Gale transforms of a vector configuration X (see Appendix 0.3, Remark 0.6).

Finally, given any facet F_i of P , there exists a vertex v not in F_i . Since we assume that all vertices of P are strictly preserved under the projection π , the normal vectors of the facets of P containing this vertex v positively span \mathbb{R}^{d-e} (Projection Lemma 0.8). Consequently, $G \setminus \{\vec{g}_i\}$ is positively spanning, for any $i \in [f]$. Thus X is the vertex set of a polytope Q of $\mathbb{R}^{f-d+e-1}$ (see Appendix 0.3, Remark 0.6). \square

Let F be a face of P , strictly preserved under the projection π . Then, $\{\vec{g}_i \mid i \in I(F)\}$ is positively spanning, which implies that the set of vertices $\{x_i \mid i \notin I(F)\}$ forms a face of Q .

Let \mathcal{F} denote a subset of the set of strictly preserved faces of P under the projection π . Let \mathcal{K} denote the abstract polytopal complex obtained as the “down closure” of $\{[f] \setminus I(F) \mid F \in \mathcal{F}\}$. Since we assumed that P is simple, \mathcal{K} is even a simplicial complex. Furthermore, it is (isomorphic to) a subcomplex of the face complex of the polytope Q , and thus, it is embeddable into $\mathbb{R}^{f-d+e-2}$ (since the face complex of an h -dimensional polytope is embedded into the sphere \mathbb{S}^{h-1}).

Thus, given the simple polytope $P \subset \mathbb{R}^d$, and a set \mathcal{F} of faces of P that we want to preserve under projection, the study of the embeddability of the corresponding abstract simplicial complex \mathcal{K} provides lower bounds on the dimension e in which we can project P . To get such a lower bound, we use the method presented in Appendix 0.6:

- (1) we first choose our subset \mathcal{F} of strictly preserved faces sufficiently simple to be understandable, and sufficiently large to provide an obstruction;
- (2) we then understand the system \mathcal{Z} of minimal non-faces of the simplicial complex \mathcal{K} ;
- (3) finally, we find a suitable coloring of the Kneser graph on \mathcal{Z} and apply Sarkaria’s Coloring and Embedding Theorem to bound the dimension in which \mathcal{K} can be embedded: a t -coloring of $\text{KG}(\mathcal{Z})$ ensures that \mathcal{K} is not embeddable into $|V(\mathcal{K})| - t - 2 = f - t - 2$, and thus, that the dimension e is at least $d - t + 1$.

In the next subsections, we apply this method to our problem (preserving the k -skeleton of a polytope combinatorially equivalent to a product $\Delta_n = \Delta_{n_1} \times \dots \times \Delta_{n_r}$ of simplices under projection).

3.2. Preserving the k -skeleton of a product of simplices. In this subsection, we understand the abstract simplicial complex \mathcal{K} corresponding to our problem, and we describe its system of minimal non-faces.

The facets of $\Delta_{\underline{n}}$ are exactly the products:

$$\psi_{i,j} = \Delta_{n_1} \times \dots \times \Delta_{n_{i-1}} \times (\Delta_{n_i} \setminus \{j\}) \times \Delta_{n_{i+1}} \times \dots \times \Delta_{n_r},$$

for $i \in [r]$ and $j \in [n_i + 1]$ (by $\Delta_{n_i} \setminus \{j\}$, we mean the facet of Δ_{n_i} opposite to the j th vertex of Δ_{n_i}). We identify the facet $\psi_{i,j}$ with the element $j \in [n_i + 1]$ of the disjoint union $[n_1 + 1] \uplus [n_2 + 1] \uplus \dots \uplus [n_r + 1]$.

Let $F = F_1 \times \dots \times F_r$ be a k -face of $\Delta_{\underline{n}}$. Then F is contained in a facet $\psi_{i,j}$ of $\Delta_{\underline{n}}$ if and only if $j \notin F_i$. Thus, the set of facets of $\Delta_{\underline{n}}$ that do not contain F is exactly $F_1 \uplus \dots \uplus F_r$. Consequently, if we want to preserve the k -skeleton of $\Delta_{\underline{n}}$, then the abstract simplicial complex we are interested in is the down closure \mathcal{K} of

$$\{F_1 \uplus \dots \uplus F_r \mid \forall i \in [r], \emptyset \neq F_i \subset [n_i + 1] \text{ and } \sum_{i \in [r]} (|F_i| - 1) = k\}.$$

The following lemma gives a description of its minimal non-faces:

Lemma 3.2. *The system of minimal non-faces of \mathcal{K} is*

$$\mathcal{Z} = \{G_1 \uplus \dots \uplus G_r \mid \forall i \in [r], |G_i| \neq 1 \text{ and } \sum_{i \mid G_i \neq \emptyset} (|G_i| - 1) = k + 1\}.$$

Proof. A subset $G = G_1 \uplus \dots \uplus G_r$ of $[n_1 + 1] \uplus [n_2 + 1] \uplus \dots \uplus [n_r + 1]$ is a face of \mathcal{K} when it is extendable into a subset $F_1 \uplus \dots \uplus F_r$ with $\sum (|F_i| - 1) = k$ and $\emptyset \neq F_i \subset [n_i + 1]$ (for all $i \in [r]$), that is, when

$$k \geq |\{i \in [r] \mid G_i = \emptyset\}| + \sum_{i \in [r]} (|G_i| - 1) = \sum_{i \mid G_i \neq \emptyset} (|G_i| - 1).$$

Thus, G is a non-face if and only if $\sum_{i \mid G_i \neq \emptyset} (|G_i| - 1) \geq k + 1$.

If $\sum_{i \mid G_i \neq \emptyset} (|G_i| - 1) > k + 1$, then removing any element provides a smaller non-face. If there is an i such that $|G_i| = 1$, then removing the element of G_i provides a smaller non-face. Thus, if G is a minimal non-face, then $\sum_{i \mid G_i \neq \emptyset} (|G_i| - 1) = k + 1$, and $|G_i| \neq 1$ for all $i \in [r]$.

Reciprocally, if G is a non-minimal non-face, then it is possible to remove one element keeping a non-face. Let $i \in [r]$ be such that we can remove one element from G_i , keeping a non-face. Then, either $|G_i| = 1$, or $\sum_{j \mid G_j \neq \emptyset} (|G_j| - 1) \geq 1 + (|G_i| - 2) + \sum_{j \neq i \mid G_j \neq \emptyset} (|G_j| - 1) \geq k + 2$ (since we keep a non-face). \square

3.3. Colorings of $\text{KG}(\mathcal{Z})$. The next step consists in providing a suitable coloring for the Kneser graph on \mathcal{Z} . Let $S = \{i \in [r] \mid n_i = 1\}$ denote the set of indices corresponding to the segments, and $R = \{i \in [r] \mid n_i \geq 2\}$ the set indices corresponding to the non-segments in the product $\Delta_{\underline{n}}$. We provide a coloring for two different situations, which proves the following theorem:

Theorem 3.3. (1) If $k \leq \sum_{i \in R} \lfloor \frac{n_i - 2}{2} \rfloor$, then $\delta_{pr}(k, \underline{n}) \geq 2k + |R| + 1$.
 (2) If $k \geq \lfloor \frac{\sum_i n_i}{2} \rfloor$, then $\delta_{pr}(k, \underline{n}) \geq \sum n_i$.

Proof of (1). Let $k_1, \dots, k_r \in \mathbb{N}$ be such that

$$\sum_{i \in [r]} k_i = k \quad \text{and} \quad \begin{cases} 0 \leq k_i \leq \frac{n_i - 2}{2} & \text{for } i \in R; \\ k_i = 0 & \text{for } i \notin R. \end{cases}$$

Observe that

- (1) such a tuple exists since $k \leq \sum_{i \in R} \lfloor \frac{n_i - 2}{2} \rfloor$.
- (2) for any minimal non face $G = G_1 \uplus \dots \uplus G_r$ of \mathcal{Z} , there exists $i \in [r]$ such that $|G_i| \geq k_i + 2$. Indeed, if $|G_i| \leq k_i + 1$ for all $i \in [r]$, then

$$k + 1 = \sum_{i \mid G_i \neq \emptyset} (|G_i| - 1) \leq \sum_{i \mid G_i \neq \emptyset} k_i \leq \sum_{i \in [r]} k_i = k,$$

which is impossible.

For all $i \in [r]$, we fix a proper coloring $\gamma_i : \binom{[n_i + 1]}{[k_i + 2]} \rightarrow [\chi_i]$ of $\text{KG}_{n_i + 1}^{k_i + 2}$, where $\chi_i = 1$ if $i \in S$ and $\chi_i = n_i - 2k_i - 1$ if $i \in R$ (see Appendix 0.5). Then, we define a coloring $\gamma : \mathcal{Z} \rightarrow [\chi_1] \uplus \dots \uplus [\chi_r]$ of the Kneser graph on \mathcal{Z} as follows. Let $G = G_1 \uplus \dots \uplus G_r$ be a given minimal non-face of \mathcal{Z} . We choose arbitrarily an $i \in [r]$ such that $|G_i| \geq k_i + 2$, and again arbitrarily a subset g of G_i with $k_i + 2$ elements. We color G with the color of g in $\text{KG}_{n_i + 1}^{k_i + 2}$, that is, we define $\gamma(G) = \gamma_i(g)$.

It is easy to see that γ is a proper coloring of the Kneser graph $\text{KG}(\mathcal{Z})$. Indeed, let $G = G_1 \uplus \dots \uplus G_r$ and $H = H_1 \uplus \dots \uplus H_r$ be two minimal non-faces of \mathcal{Z} related by an edge in $\text{KG}(\mathcal{Z})$, *i.e.*, they do not intersect. They may get the same color only if we choose the same $i \in [r]$ with $|G_i| \geq k_i + 2$ and $|H_i| \geq k_i + 2$, and color G with $\gamma_i(g)$ for a certain $g \subset G_i$ of size $k_i + 1$, and H with $\gamma_i(h)$ for a certain $h \subset H_i$ of size $k_i + 1$. However, since $G_i \cap H_i = \emptyset$, we know that $g \cap h = \emptyset$. Thus, g and h are related by an edge in $\text{KG}_{n_i + 1}^{k_i + 2}$ and consequently, they are colored with different colors by γ_i . Thus, G and H get different colors.

This provides a proper coloring of $\text{KG}(\mathcal{Z})$ with $\sum \chi_i$ colors. Consequently, we know that the dimension e of the projection is at least

$$\sum_{i \in [r]} n_i - \sum_{i \in [r]} \chi_i + 1 = 2k + |R| + 1.$$

This proves Part (1) of Theorem 3.3. \square

Proof of (2). Let $G = G_1 \uplus \dots \uplus G_r$ and $H = H_1 \uplus \dots \uplus H_r$ be two minimal non-faces of \mathcal{Z} . Let $A = \{i \in [r] \mid G_i \neq \emptyset \text{ or } H_i \neq \emptyset\}$. Then

$$\begin{aligned} \sum_{i \in A} (|G_i| + |H_i|) &\geq \sum_{G_i \neq \emptyset} (|G_i| - 1) + \sum_{H_i \neq \emptyset} (|H_i| - 1) + |A| \\ &= 2k + 2 + |A| > \sum_{i \in [r]} n_i + |A| \geq \sum_{i \in A} (n_i + 1). \end{aligned}$$

Thus, there exists $i \in A$ such that $|G_i| + |H_i| > n_i + 1$, which implies that $G_i \cap H_i \neq \emptyset$, and proves that $G \cap H \neq \emptyset$.

Consequently, the Kneser graph $\text{KG}(\mathcal{Z})$ is independent (and we can color it with only one color). We obtain that the dimension e of the projection is at least $\sum n_i$. In other words, in this extremal case, there is no better (k, \underline{n}) -PSN polytope than the product $\Delta_{\underline{n}}$ itself. \square

Remark 3.4. *Theorem 3.3 (2) can be strengthened sometimes a little: If $k = \frac{1}{2} \sum n_i - 1$, and $k + 1$ is not representable as a sum of a subset of $\{n_1, \dots, n_r\}$, then $\delta_{pr}(k, \underline{n}) = \sum n_i$. This follows by the observation that in this case $\text{KG}(\mathcal{Z})$ is independent.*

Note that there is a set of k for which Theorem 3.3 does not say anything explicitly. One could simply apply the trivial formula $\delta_{pr}(k', \underline{n}) \leq \delta_{pr}(k, \underline{n})$ for $k' \leq k$, but the better approach is to find a better coloring. One can extend the colorings of both proofs, but it is even better to merge both above coloring ideas as follows:

We partition $[r] = A \uplus B$ and choose $k_i \geq 0$ ($i \in A$) and $k_B \geq 0$ such that

$$(4) \quad \left(\sum_{i \in A} k_i \right) + k_B \leq k.$$

We will see later what is the best choice. Observe that for all minimal non-faces $G = G_1 \uplus \dots \uplus G_r$,

- either there is an $i \in A$ such that $|G_i| \geq k_i + 2$,
- or $\sum_{i \in B \mid G_i \neq \emptyset} (|G_i| - 1) \geq k_B + 1$.

Indeed, otherwise

$$k + 1 = \sum_{i|G_i \neq \emptyset} (|G_i| - 1) \leq \left(\sum_{i \in A} k_i \right) + k_B \leq k.$$

Therefore we can color $\text{KG}(\mathcal{Z})$ as follows. Let $n_B = \sum_{i \in B} n_i$. Color the Kneser graphs $\text{KG}_{n_i+1}^{k_i+2}$ ($i \in A$) and $\text{KG}_{n_B}^{k_B+1}$ with pairwise disjoint color sets with

$$\chi_i = \begin{cases} n_i - 2k_i - 1 & \text{if } 2k_i \leq n_i - 2, \\ 1 & \text{if } 2k_i \geq n_i - 2 \end{cases}$$

and

$$\chi_B = \begin{cases} n_B - 2k_B & \text{if } 2k_B \leq n_B - 1, \\ 1 & \text{if } 2k_B \geq n_B - 1 \end{cases}$$

colors respectively. Then for a minimal non-face $G = G_1 \uplus \dots \uplus G_r$ we either find an $i \in A$ such that $|G_i| \geq k_i + 2$, so we take an arbitrary subset of G_i with $k_i + 2$ elements and take its color in $\text{KG}_{n_i+1}^{k_i+2}$, or we have $\sum_{i \in B | G_i \neq \emptyset} (|G_i| - 1) \geq k_B + 1$ such that we take an arbitrary subset of $\biguplus_{i \in B} (G_i \setminus \{n_i + 1\}) \subset \biguplus_{i \in B} [n_i]$ with $k_B + 1$ elements and take its color in $\text{KG}_{n_B}^{k_B+1}$. It is easy to see that this provides a valid coloring of the Kneser graph $\text{KG}(\mathcal{Z})$. Hence,

$$\chi_{\text{sum}} = \chi_{\text{sum}}(A, B, \underline{k_i}, k_B) = \sum_{i \in A} \chi_i + \chi_B \geq \chi(\text{KG}(\mathcal{Z})).$$

Therefore we get our desired bound

$$d_k = d_k(A, B, \underline{k_i}, k_B) = \sum_i n_i + 1 - \chi_{\text{sum}} \geq \delta_{pr}(k, \underline{n}).$$

All we have to do now is to find the best parameters A , B , k_i ($i \in A$) and k_B . We will proceed algorithmically, by keeping first A and B fixed and see how to choose the k_i 's and k_B to make d_k large.

By the formulas for χ_i and χ_B , it makes sense to assume $2k_i \leq n_i - 1$ for all $i \in A_i$ and $2k_B \leq \sum_{i \in B} n_i$.

Now, if all $k_i = 0 = k_B$, then

$$\begin{aligned} \chi_{\text{sum}}(A, B, \underline{0}, 0) &= \sum_{i \in A} (n_i - 1) + |S \cap A| + n_B \\ &= \sum_{i \in A} n_i - |A| + |S \cap A| + \sum_{i \in B} n_i \\ &= \sum_{i \in [r]} n_i - r + |B \cup S|, \end{aligned}$$

and

$$d_k(A, B, \underline{0}, 0) = 1 + r - |B \cup S|$$

(where $S = \{i \in [r] \mid n_i = 1\}$ denotes the set of segments).

For $i \in A$, increasing k_i by one decreases χ_i by (i.e. increases d_k by):

$$\begin{cases} 2, & \text{if } 2k_i \leq n_i - 4, \\ 1, & \text{if } 2k_i = n_i - 3, \\ 0, & \text{if } 2k_i \geq n_i - 2. \end{cases}$$

Similarly, increasing k_B by one decreases χ_B by (i.e. increases d_k by):

$$\begin{cases} 2, & \text{if } 2k_B \leq n_B - 3, \\ 1, & \text{if } 2k_B = n_B - 2, \\ 0, & \text{if } 2k_B \geq n_B - 1. \end{cases}$$

We will only consider the case $S \subset B$, since at the beginning when $k_i = 0 = k_B$, it does not matter for χ_{sum} whether the segments lie in A or B , but incrementing their k_{\dots} is only useful if they lie in B .

By (4) we are allowed to do at most k of these increases. So algorithmically we should first increase the k_{\dots} 's that decrease χ_{sum} by two, and then these that decrease χ_{sum} by one. Hence we get a case distinction on k , which is also depending on A and B . Let

$$K_1(A, B) = \sum_{i \in A} \left\lfloor \frac{n_i - 2}{2} \right\rfloor + \max \left(\left\lfloor \frac{n_B - 1}{2} \right\rfloor, 0 \right),$$

$$K_2(A, B) = |\{i \in A \mid n_i \text{ is odd}\}| + \delta_{[n_B \text{ is even and non-zero}]}$$

Proposition 3.5. *Let $A \uplus B = [r]$, $B \supset S = \{i \in [r] \mid n_i = 1\}$, $K_1 = K_1(A, B)$, $K_2 = K_2(A, B)$.*

(1) *If $0 \leq k \leq K_1$ then*

$$\delta_{pr}(k, \underline{n}) \geq r + 1 - |B| + 2k.$$

(2) *If $K_1 \leq k \leq K_1 + K_2$ then*

$$\delta_{pr}(k, \underline{n}) \geq r + 1 - |B| + K_1 + k.$$

(3) *If $K_1 + K_2 \leq k$ then*

$$\delta_{pr}(k, \underline{n}) \geq r + 1 - |B| + 2K_1 + K_2.$$

Taking $B = S$, that is $A = R = \{i \in [r] \mid n_i \geq 2\}$, yields:

Corollary 3.6. *If*

$$0 \leq k \leq \sum_{i \in R} \left\lfloor \frac{n_i - 2}{2} \right\rfloor + \max \left(\left\lfloor \frac{|S| - 1}{2} \right\rfloor, 0 \right),$$

then $\delta_{pr}(k, \underline{n}) \geq 2k + |R| + 1$.

Taking instead $B = [r]$ yields:

Corollary 3.7. *If $k \geq \left\lfloor \frac{\sum n_i}{2} \right\rfloor$ then $\delta_{pr}(k, \underline{n}) \geq \sum_i n_i$.*

Proposition 3.5 can be made more explicit by finding A and B that give the best results, which we will do next using the following algorithm, where \underline{n} is fixed and we iterate on k . The lower bound on $\delta_{pr}(k, \underline{n})$ will be denoted by d_k .

We define d_k iteratively:

- (1) For $k = 0$, take $A = R$ and $B = S$, so we are in case 1 of the proposition. Define d_0 to be the lower bound that the proposition tells us.
- (2) While we stay in case 1 of the proposition for $k' = k + 1$ with same A and B , we do the following:
 - Increase k by one, keep the same A and B , and calculate d_k according to case 1 of the proposition.
- (3) Then do the following as long as A contains i for which n_i is odd:
 - Increase k by one, move an i from A to B for which n_i is odd. Then we are again in case 1 and we calculate d_k .
 - If now n_B has become odd and if $|B|$ has become > 1 , then increase k by one and we still are in case 1, so we let A and B as they are and calculate d_k .
- (4) Then do the following as long as A is non-empty:
 - Increase k once more, move an i from A to B (for which n_i is even). Calculate d_k .
- (5) Then increase k once more. Since $A = \emptyset$, $B = [r]$, $k = K_1 + 1$ and $K_2 \in \{0, 1\}$, we arrived in case 3 of the proposition. We calculate its d_k , which is just $\sum n_i$.

If we draw the graph of d_k , then by increasing k , d_k increases by two except at the points where we move an element i from A to B . So the algorithm makes clear how the graph looks like:

- In step 2, d_k is always increasing by two.
- In step 3 part 1, d_k is always increasing by one.
- In step 3 part 2, d_k is always increasing by two (but we do this part of course only if the if-clause is fulfilled).
- In step 4, d_k is always increasing by one, except for the case when B is empty at the beginning of step 4 (which means that all n_i are even): Then d_k is staying the same during the first increment (we arrived in case 2 and case 3 of the proposition: $k = K_1 + K_2$), and during all further increments on k it is increasing by one.
- In step 5 we arrived at what Corollary 3.7 says.

For the rest of this section, we fix $A = R$, $B = S$, $K_1 = K_1(R, S)$ and $K_2 = K_2(R, S)$. Let $d_0 = r + 1 - |S|$ and define $n = \sum_i n_i$. We summarize our lower bound d_k explicitly with a case distinction (see Fig. 2):

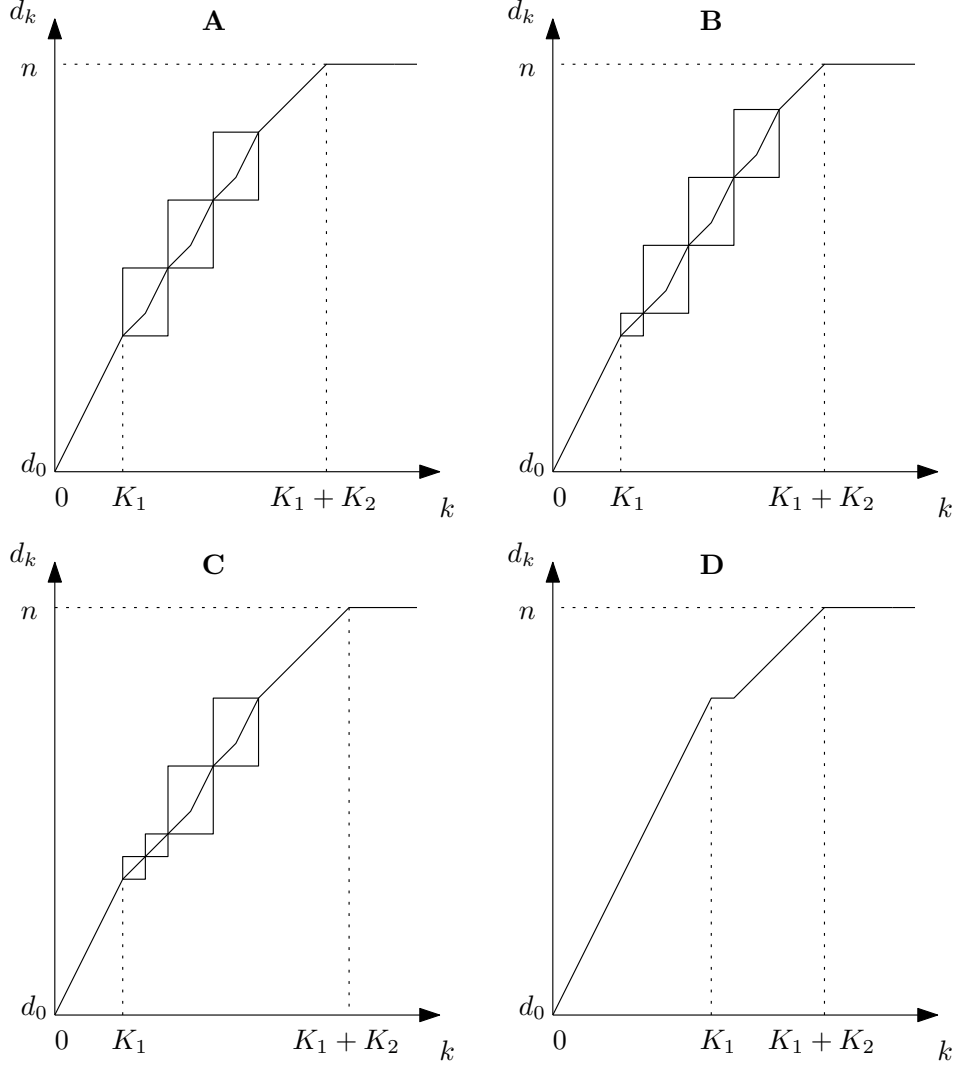


FIGURE 2. Four different situations for the lower bound.

Case **A** $|S|$ is even and non-zero: d_k is increasing by two for $0 \leq k \leq K_1$, then for each odd $n_i \geq 3$ we get first an increment by one and then another increment by two. Then all increments are one until we reach the trivial bound $d_k = n = \sum_i n_i$.

Case **B** $|S|$ is odd: As in case **A**, except that the first block corresponding to an odd $n_i \geq 3$ consists only of one increment by one.

- Case **C** $|S| = 0$ and there is an odd n_i : As in case **B**, except that the first two blocks corresponding to odd n_i 's consists only of one increment. The case that we have only one odd n_i is included.
- Case **D** All n_i are even: d_k is increasing by two for $0 \leq k \leq K_1$, the next increment is zero, and all further increments are one until we reach $d_k = n$.

Theorem 3.8. *The above defined d_k is a lower bound on $\delta_{pr}(k, \underline{n})$.*

Remark 3.9. *There are cases in which Remark 3.4 gives a better bound, e.g., for $\underline{n} = (4, 2)$ and $k = 2$.*

ACKNOWLEDGMENTS

This paper was written during our participation in the *DocCourse on Combinatorics and Geometry 2009* in Barcelona. We are extremely grateful to the *Centre de Recerca Matemàtica* for having given us the opportunity of working together during three months in a very stimulating environment.

REFERENCES

- [AZ99] N. Amenta and G. M. Ziegler. Deformed products and maximal shadows of polytopes. In *Advances in Discrete and Computational Geometry*, Amer. Math. Soc., Providence, Contemporary Mathematics 223, pages 57–90, 1999.
- [JS07] M. Joswig and T. Schröder. Neighborly cubical polytopes and spheres. *Israel J. Math.*, 159:221–242, 2007.
- [Mat02] J. Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
- [Mat03] J. Matoušek. *Using the Borsuk-Ulam theorem*. Universitext. Springer-Verlag, Berlin, 2003. Lectures on topological methods in combinatorics and geometry, Written in cooperation with A. Björner and G. M. Ziegler.
- [San09] R. Sanyal. Topological obstructions for vertex numbers of Minkowski sums. *J. Comb. Theory*, 116:168–179, 2009.
- [SZ09] R. Sanyal and G. M. Ziegler. Construction and analysis of projected deformed products. to appear in *Discrete Comput. Geom.*, 2009.
- [Zie95] G. M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.
- [Zie04] G. M. Ziegler. Projected products of polygons. *Electron. Res. Announc. Amer. Math. Soc.*, 10:122–134, 2004.

0. APPENDIX: BACKGROUND MATERIAL

In this Appendix, we present some well known facts about polytopes, Gale duality and embeddability of simplicial complexes. We refer to [Zie95], [Mat02] and [Mat03] for much more detailed discussions on these subjects.

0.1. Products of polytopes. A *polytope* is the convex hull of a finite point set of \mathbb{R}^n , or equivalently a bounded intersection of finitely many half-spaces in \mathbb{R}^n : we write $P = \text{conv}(X) = \{x \in \mathbb{R}^n \mid Ax \leq b\}$, where $X \subset \mathbb{R}^n$ is finite, $A \in \mathbb{R}^{n \times f}$ and $b \in \mathbb{R}^f$.

A *face* of a polytope P is any subset of P of the form $F = P \cap H$, where H is an hyperplane such that P is contained in one of the closed half-spaces defined by H . The hyperplane H is *supporting* the face F . If $H = \{x \in \mathbb{R}^n \mid h \cdot x = c\}$ and $h \cdot x < c$ for all $x \in P$, then we say that h is a *normal vector* of the face F . The *dimension* of a face is the dimension of its affine span. The *f-vector* of an n -dimensional polytope P is the vector $f(P) = (f_0(P), \dots, f_n(P))$, where $f_k(P)$ denotes the number of k -dimensional faces of P .

Example 0.1. An n -dimensional simplex is the convex hull of $n + 1$ affinely independent points (for example, $\Delta_n = \text{conv}\{u_k \mid k \in [n + 1]\}$, where (u_1, \dots, u_{n+1}) denotes the canonical orthogonal basis of \mathbb{R}^{n+1}). A face of the simplex is the convex hull of any subset of its vertices, and in particular, the *f-vector* is given by $f_k(\Delta_n) = \binom{n+1}{k+1}$.

Let $k \in \mathbb{N}$. The *k-skeleton* of a polytope is the polygonal complex formed by all its faces of dimension $\leq k$. The 1-skeleton of a polytope is its *graph*. The graph of a polytope P – or more generally, the k -skeleton of P (for small k) – is not sufficient to determine P . For example, there exists polytopes whose graph is the complete graph, but which are not simplices: such polytopes are called *neighborly* (see Appendix 0.2).

The *product* of two polytopes $P \subset \mathbb{R}^n$ and $Q \subset \mathbb{R}^m$ is the polytope $P \times Q = \{(p, q) \in \mathbb{R}^{n+m} \mid p \in P, q \in Q\}$. The dimension of $P \times Q$ is the sum of the dimensions of P and Q . The non-empty faces of $P \times Q$ are precisely products of non-empty faces of P by non-empty faces of Q (and normal vectors are obtained as products of the corresponding normal vectors). In particular, the graph of $P \times Q$ is the product of the graphs of P and Q : its vertex set is $V(P \times Q) = V(P) \times V(Q)$ while its edge set is $E(P \times Q) = \{(p, q), (p', q) \mid (p, p') \in E(P)\} \cup \{(p, q), (p, q') \mid (q, q') \in E(Q)\}$. The *f-vector* of $P \times Q$ is given by $f_k(P \times Q) = \sum_{i+j=k} f_i(P)f_j(Q)$.




FIGURE 3. The graph of the product $\Delta_{(3,6)} = \Delta_3 \times \Delta_6$.

Example 0.2. Let $\underline{n} = (n_1, \dots, n_r)$ be a tuple of positive integers. The polytope $\Delta_{\underline{n}} = \Delta_{n_1} \times \dots \times \Delta_{n_r}$ has dimension $\sum n_i$ and its f -vector is given by

$$f_k(\Delta_{\underline{n}}) = \sum_{\substack{0 \leq k_i \leq n_i \\ k_1 + \dots + k_r = k}} \prod_{i \in [r]} \binom{n_i + 1}{k_i + 1}.$$

In particular, $f_0(\Delta_{\underline{n}}) = \prod n_i$ and $f_1(\Delta_{\underline{n}}) = \sum_{i \in [r]} \binom{n_i + 1}{2} \prod_{j \neq i} n_j$. The graph of $\Delta_{(3,6)}$ is represented in Fig. 3.

0.2. Cyclic polytopes. Let $t \mapsto \mu_d(t) = (t, t^2, \dots, t^d)^T$ denote the *moment curve* in \mathbb{R}^d , $t_1 < t_2 < \dots < t_n$ be n distinct real numbers and $C_d(n) = \text{conv}\{\mu_d(t_i) \mid i \in [n]\}$ denote the *cyclic polytope* (the convex hull of any n distinct points on the moment curve always leads to the same combinatorial polytope; and in fact, any d -dimensional algebraic curve of degree d will also give rise to the same combinatorial polytope). We naturally identify a face $\text{conv}\{\mu(t_i) \mid i \in F\}$ with the subset F of $[n]$.

Gale's evenness criterion (see Proposition 0.3) characterizes combinatorially which subsets of $[n]$ are facets of $C_d(n)$. For this, we call *blocs* of a subset F of $[n]$ the maximal subsets of consecutive elements of F . The *initial* (resp. *final*) bloc is the bloc containing 1 (resp.

n) – when it exists. Other blocs are called *inner* blocs. For example, $\{1, 2, 3, 6, 7, 9, 10, 11\} \subset [11]$ has 3 blocs: $\{1, 2, 3\}$ (initial), $\{6, 7\}$ (inner) and $\{9, 10, 11\}$ (final).

Proposition 0.3 (Gale evenness criterion). *A subset F of $[n]$ is a facet of $C_d(n)$ if and only if $|F| = d$ and all inner blocs of F have even size. Furthermore, such a facet F is supported by the hyperplane*

$$H_F = \{z \mid (\gamma_i(F))_{i \in [d]} \cdot z = -\gamma_0(F)\},$$

where $\gamma_0(F), \dots, \gamma_d(F)$ are defined as the coefficients:

$$\Pi_F(t) = \prod_{i \in F} (t - t_i) = \sum_{i=0}^d \gamma_i(F) t^i.$$

Finally, according to the parity of the size ℓ of the final bloc of F :

- (1) if ℓ is odd, then the whole cyclic polytope lies below H_F (according to the last coordinate); we say that F is an upper facet; its normal vector is $(\gamma_i(F))_{i \in [d]}$;
- (2) if ℓ is even, then the whole cyclic polytope lies above H_F ; we say that F is a lower facet; its normal vector is $(-\gamma_i(F))_{i \in [d]}$.

Proof. Observe first that

- (i) Vandermonde determinant

$$\det \begin{pmatrix} 1 & 1 & \dots & 1 \\ \mu(x_0) & \mu(x_1) & \dots & \mu(x_d) \end{pmatrix} = \prod_{0 \leq i < j \leq d} (x_j - x_i)$$

ensures that any $d+1$ points on the d -dimensional moment curve are affinely independent, and thus, that the cyclic polytope is simplicial.

- (ii) For any $t \in \mathbb{R}$,

$$(\gamma_i(F))_{i \in [d]} \cdot \mu_d(t) + \gamma_0(F) = \sum_{i=0}^d \gamma_i(F) t^i = \Pi_F(t).$$

- (iii) the coefficient $\gamma_d(F)$ is always 1, and thus, the vector $(\gamma_i(F))_{i \in [d]}$ points upwards (according to the last coordinate).

Let F be a subset of $[n]$ of size d . Then,

- (i) for all $j \in F$, $\Pi_F(t_j) = 0$; thus, H_F is the affine hyperplane spanned by F ;
- (ii) for all $j \notin F$, the sign of $\Pi_F(t_j)$ is $(-1)^{|F \cap \{j+1, \dots, n\}|}$.

In particular, if F has an odd inner bloc $\{a, a+1, \dots, b\}$, then $\Pi_F(t_{a-1})$ and $\Pi_F(t_{b+1})$ have different signs, and F is not a facet. Reciprocally, if all inner blocs have even size, then the sign of all $\Pi_F(t_j)$ is $(-1)^\ell$,




FIGURE 4. The cyclic polytopes $C_2(6)$ (with a unique upper facet $\{1, 6\}$) and $C_3(6)$ (with its upper facets $\{i, i + 1, 6\}$, $i \in [4]$).

where ℓ is the size of the final bloc. Thus, F is an upper facet when ℓ is odd, and a lower facet when ℓ is even. \square

Observe that cyclic polytopes in odd dimension have as many upper facets as lower facets, while they have very few in even dimension.

From this criterion, it is easy to understand the complete combinatorics of the faces of $C_d(n)$. In particular:

Corollary 0.4. *The cyclic polytope $C_d(n)$ is $\lfloor \frac{d}{2} \rfloor$ -neighborly: any subset $F \subset [n]$ with $|S| \leq \lfloor \frac{d}{2} \rfloor$ forms a face of $C_d(n)$.*

0.3. Gale duality. We start with two definitions in the linear world.

Let $p \geq q$ denote two integers. Let $A = \{\vec{a}_1, \dots, \vec{a}_p\}$ be a *vector configuration* in \mathbb{R}^q . Then

- (1) a *linear dependence* of A is a $\lambda \in \mathbb{R}^p$ such that $\sum \lambda_i \vec{a}_i = 0$;
- (2) a *linear evaluation* on A is a $\mu \in \mathbb{R}^p$ such that there exists a linear function $f : \mathbb{R}^q \rightarrow \mathbb{R}$ whose values on $\{\vec{a}_1, \dots, \vec{a}_p\}$ give precisely μ , i.e., such that $\mu_i = f(\vec{a}_i)$, for all $i \in [p]$.

In other words, if U denotes the matrix whose column vectors are the \vec{a}_i 's, a linear dependence is a vertex of the kernel of U . Similarly, if V denotes the matrix whose row vectors are the \vec{a}_i 's, a linear evaluation is a vector in the image of V .

Remark 0.5. Let $A = \{\vec{a}_1, \dots, \vec{a}_p\}$ be a (full rank) vector configuration of \mathbb{R}^q , U denote the matrix whose column vectors are the \vec{a}_i 's, and $\vec{u}_1, \dots, \vec{u}_q$ denote the row vectors of U . Let $(\vec{v}_1, \dots, \vec{v}_{p-q})$ be a basis of the kernel of U , V denote the matrix whose column vectors are the

\vec{v}_i 's, and $B = \{\vec{b}_1, \dots, \vec{b}_p\}$ be the row vectors of V . The following observation will be the key observation for Gale duality to be presented next:

- (1) the linear dependencies on A are precisely the linear evaluations on B ; and
- (2) the linear evaluations on A are precisely the linear dependencies on B .

For our purpose, we will also need the following definition: the vector configuration $A \subset \mathbb{R}^q$ is said to be *positively spanning* if one of the following equivalent properties holds:

- (i) any point of \mathbb{R}^q is a positive combination of A ;
- (ii) A spans \mathbb{R}^q and has a positive linear dependence;
- (iii) 0 is an interior point of the convex hull of A .

Observe in particular that if A has a positively spanning subset, then A itself is positively spanning.

We are now moving to the affine world. We define here similarly the notions corresponding to linear dependencies and evaluations.

Let $X = \{x_1, \dots, x_p\}$ be a *point configuration* in \mathbb{R}^q . Then

- (1) an *affine dependence* on X is a vector $\lambda \in \mathbb{R}^q$ such that both $\sum \lambda_i a_i = 0$ and $\sum \lambda_i = 0$;
- (2) an *affine evaluation* on X is a vector $\mu \in \mathbb{R}^p$ such that there exists an affine function $f : \mathbb{R}^q \rightarrow \mathbb{R}$ whose values on $\{x_1, \dots, x_p\}$ give precisely μ , *i.e.*, such that $\mu_i = f(x_i)$, for all $i \in [p]$.

Observe that affine evaluations give many information on the point configuration. For example, if X is the vertex set of a polytope P , then the faces of P are (convex hulls) of subsets of X for which there exists a non-negative affine evaluation that vanishes exactly on X . The same observation can be made for affine dependencies.

Let us also observe that affine and linear dependencies and evaluation are related by homogenization. For all $x \in \mathbb{R}^q$, we denote $\vec{x} = (x, 1)$ the corresponding homogenized vector in \mathbb{R}^{q+1} . This embedding $\mathbb{R}^q \rightarrow \mathbb{R}^{q+1}$ transforms affine notions in \mathbb{R}^q into linear notions in \mathbb{R}^{q+1} : the affine dependences on $X = \{x_1, \dots, x_p\}$ are precisely the linear dependences on $\vec{X} = \{\vec{x}_1, \dots, \vec{x}_p\}$ and the affine evaluations on X are precisely the linear evaluations on \vec{X} .

We are now ready to present *Gale duality*. We will see it as a transformation from an affine configuration of p points in \mathbb{R}^q into a linear configuration of p vectors in \mathbb{R}^{p-q-1} , that *dualizes* the notions of dependencies and evaluations.

We start from a point configuration $X = \{x_1, \dots, x_p\} \subset \mathbb{R}^q$. We denote U the $p \times (q+1)$ -matrix whose column vectors are the homogenized vectors \vec{x}_i 's. We choose a basis $(\vec{v}_1, \dots, \vec{v}_{p-q-1})$ of the kernel of U , and denote V the matrix whose column vectors are the \vec{v}_i 's, and $G = \{\vec{g}_1, \dots, \vec{g}_p\} \subset \mathbb{R}^{p-q-1}$ the row vectors of V . The sequence G is called the *Gale transform* of X .

Remark 0.6. (1) G is defined up to linear isomorphism.

- (2) The affine dependencies on X are precisely the linear evaluations on G ; the affine evaluations on X are precisely the linear dependencies on G (see Remark 0.5).
- (3) If X is the vertex set of a polytope P , then the faces of P corresponds to the zero sets of non-negative affine evaluations on X , and thus to zero sets of non-negative linear dependencies of G .
- (4) A vector configuration $G \subset \mathbb{R}^{p-q-1}$ is a Gale transform of a vector configuration X of \mathbb{R}^q if and only if it spans \mathbb{R}^{p-q-1} and has 0 as center of gravity: $\sum \vec{g}_i = 0$. Furthermore, this vector configuration X is the vertex set of a polytope if and only if $G \setminus \vec{g}_i$ spans \mathbb{R}^{p-q-1} and has a positive linear dependence (*i.e.*, is positively spanning), for all $i \in [p]$.

0.4. Projections of polytopes and strictly preserved faces. Let $d > e$ be two integers. We denote (u_1, \dots, u_d) the canonical orthogonal basis of \mathbb{R}^d . Let $\pi : \mathbb{R}^d \rightarrow \mathbb{R}^e$ denote the orthogonal projection with image $U = \langle u_1, \dots, u_e \rangle$, and kernel $V = \langle u_{e+1}, \dots, u_d \rangle$. Let $\tau : \mathbb{R}^d \rightarrow \mathbb{R}^{d-e}$ denote the dual projection (with image V and kernel U). Finally, let P be a full-dimensional simple polytope in \mathbb{R}^d , with 0 in its interior.

We want to characterize what faces of P are preserved under the projection π (see Fig. 5):

Definition 0.7. A proper face F of a polytope P is strictly preserved under the projection π if

- (i) $\pi(F)$ is a face of $\pi(P)$,
- (ii) F and $\pi(F)$ are combinatorially isomorphic, and
- (iii) $\pi^{-1}(\pi(F))$ equals F .

The characterization of strictly preserved faces of P uses the normal vectors of the facets of P . Let F_1, \dots, F_f denote the facets of P and for all $i \in [f]$, let \vec{f}_i denote the normal vector of F_i , and $\vec{g}_i = \tau(\vec{f}_i)$. For any face F of P , let $I(F)$ denote the set of indices of the facets of P containing F , *i.e.*, such that $F = \bigcap_{i \in I(F)} F_i$.




FIGURE 5. (a) Projection of a tetrahedron into \mathbb{R}^2 : the edge pq is strictly preserved, while neither the edge qr , nor the face qrs , nor the edge qs are (because of conditions (i), (ii) and (iii) respectively). (b) Projection of a tetrahedron into \mathbb{R} : only the vertex p is strictly preserved.

Lemma 0.8 (Projection Lemma [AZ99, Zie04]). *A face F of the polytope P is strictly preserved under the projection π if and only if $\{\vec{g}_i \mid i \in I(F)\}$ is positively spanning.*

Proof. Let F be a face of P , and let $\text{aff}(F) = x + L$, where x is any point of $\text{aff}(F)$, and L is the directing linear subspace of $\text{aff}(F)$. Then

- (1) F and $\pi(F)$ are combinatorially equivalent $\Leftrightarrow \pi|_L$ is injective $\Leftrightarrow \tau|_{L^\perp}$ is surjective $\Leftrightarrow \{\vec{g}_i \mid i \in I(F)\}$ is spanning.
- (2) the projection $\pi(F)$ is a face of $\pi(P)$ and $\pi^{-1}(\pi(F)) = F$ $\Leftrightarrow \exists z \in \text{im}(\pi)$ such that $z \cdot x = z \cdot \pi(x) = 1$ if $x \in F$ and $z \cdot x < 1$ if $x \in P \setminus F$ \Leftrightarrow there exists in $\text{im}(\pi) = \ker(\tau)$ a positive linear combination of $\{\vec{f}_i \mid i \in I(F)\} \Leftrightarrow 0$ is a positive linear combination of $\{\vec{g}_i \mid i \in I(F)\}$.

□

0.5. Kneser graphs. Remind that a k -coloring of a graph $G = (V, E)$ is a map $c : V \rightarrow [k]$ such that $c(u) \neq c(v)$ for $(u, v) \in E$. As usual, we denote $\chi(G)$ the *chromatic number* of G (i.e., the minimal k such that G admits a k -coloring). We are interested in the chromatic number of so-called Kneser graphs.

Let \mathcal{Z} be a subset of the power set $2^{[n]}$ of $[n]$. The *Kneser graph* on \mathcal{Z} , denoted $\text{KG}(\mathcal{Z})$, is the graph with vertex set \mathcal{Z} and where two




FIGURE 6. The Kneser graph KG_5^2 is the Petersen graph.

vertices $X, Y \in \mathcal{Z}$ are related if and only if $X \cap Y = \emptyset$:

$$\text{KG}(\mathcal{Z}) = (\mathcal{Z}, \{(X, Y) \in \mathcal{Z}^2 \mid X \cap Y = \emptyset\}).$$

Let $\text{KG}_n^k = \text{KG}\left(\binom{[n]}{k}\right)$ denote the Kneser graph on the set of k -tuples of $[n]$. For example, the graph KG_n^1 is the complete graph K_n (of chromatic number n) and the graph KG_5^2 is the Petersen graph (of chromatic number 3) – see Fig. 6.

Remark 0.9. (1) If $n \leq 2k - 1$, then any two k -tuples of $[n]$ intersect and the Kneser graph KG_n^k is independent (*i.e.*, it has no edge). Thus its chromatic number is $\chi(\text{KG}_n^k) = 1$.
 (2) If $n \geq 2k + 1$, then $\chi(\text{KG}_n^k) \leq n - 2k + 2$. Indeed, the map $c : \binom{[n]}{k} \rightarrow [n - 2k + 2]$ defined by $c(F) = \min(F \cup \{n - 2k + 2\})$ is a $(n - 2k + 2)$ -coloring of KG_n^k .

In fact, it turns out that this upper bound is the exact chromatic number of the Kneser graph: $\chi(\text{KG}_n^k) = n - 2k + 2$. This result has been conjectured by Kneser in 1955, and proved by Lovász in 1978 applying the Borsuk-Ulam Theorem (see [Mat03] for more details). However, we will only need the upper bound in our topological obstruction.

0.6. Embeddability of simplicial complexes. In this last Appendix, we give a short summary of a standard technique to study the embeddability of simplicial complexes. More precisely, this technique provides a lower bound on the dimension in which a given simplicial complex can be embedded. Refer to [Mat03] for more details.

\mathbb{Z}_2 -spaces and \mathbb{Z}_2 -index. A \mathbb{Z}_2 -space is a topological space X together with a free action σ of \mathbb{Z}_2 on X (i.e., a fixed-point free involution on X). A \mathbb{Z}_2 -map between two \mathbb{Z}_2 -spaces (X, σ) and (Y, τ) is a continuous map $f : X \rightarrow Y$ that commutes with the \mathbb{Z}_2 -actions (i.e., such that $f \circ \sigma = \tau \circ f$). We write $X \rightarrow_{\mathbb{Z}_2} Y$ if there is a \mathbb{Z}_2 -map from X to Y .

First and foremost examples of \mathbb{Z}_2 -spaces are the spheres $\mathbb{S}^d = \{x \in \mathbb{R}^{d+1} \mid \|x\| = 1\}$ with the antipodal action $- : x \mapsto -x$. Borsuk-Ulam Theorem affirms that there is no \mathbb{Z}_2 -map from \mathbb{S}^d to \mathbb{S}^{d-1} . In general, given any \mathbb{Z}_2 -space (X, σ) , it is interesting to know what is the minimal sphere in which it can be \mathbb{Z}_2 -mapped. The dimension of this sphere is called the \mathbb{Z}_2 -index of X , and denoted $\text{Ind}_{\mathbb{Z}_2}(X)$:

$$\text{Ind}_{\mathbb{Z}_2}(X) = \min\{d \in \mathbb{N} \mid X \rightarrow_{\mathbb{Z}_2} \mathbb{S}^d\}.$$

Deleted joins. Let \mathcal{K} be a simplicial complex. The *deleted join* of \mathcal{K} is the complex

$$\mathcal{K}_{\Delta}^{*2} = \{F \uplus G \mid F, G \in \mathcal{K}, F \cap G = \emptyset\},$$

(where $F \uplus G$ denotes the disjoint union of F and G). The polyhedron of this complex can be written as the set of “formal convex combinations”

$$\|\mathcal{K}_{\Delta}^{*2}\| = \{tx \oplus (1-t)y \mid t \in [0, 1], x, y \in \|\mathcal{K}\|, \text{supp}(x) \cap \text{supp}(y) = \emptyset\}.$$

Together with the natural \mathbb{Z}_2 -action $\sigma : tx \oplus (1-t)y \rightarrow (1-t)y \oplus tx$ of exchange of coordinates, the deleted join forms a \mathbb{Z}_2 -space.

The embeddability question of \mathcal{K} can be understood via the \mathbb{Z}_2 -index of the deleted join. Indeed, if f is an embedding of \mathcal{K} into \mathbb{R}^d , then the join-function $f^{*2} : tx \oplus (1-t)y \rightarrow tf(x) \oplus (1-t)f(y)$ \mathbb{Z}_2 -maps $\|\mathcal{K}_{\Delta}^{*2}\|$ into $\mathcal{L} = (\mathbb{R}^d)^{*2} \setminus \{\frac{1}{2}x \oplus \frac{1}{2}x \mid x \in \mathbb{R}^d\}$. Since there exists a \mathbb{Z}_2 -map from \mathcal{L} to \mathbb{S}^d , this implies that if \mathcal{K} is embeddable in \mathbb{R}^d , then $\mathcal{K}_{\Delta}^{*2} \rightarrow_{\mathbb{Z}_2} \mathbb{S}^d$. In other words, if $\text{Ind}_{\mathbb{Z}_2}(\mathcal{K}_{\Delta}^{*2}) > d$, then \mathcal{K} is not embeddable into \mathbb{R}^d .

A lower bound via Kneser colorings. In order to give bounds on the possible dimension of an embedding of the simplicial complex \mathcal{K} , we thus “only” have to bound the \mathbb{Z}_2 -index of $\mathcal{K}_{\Delta}^{*2}$. A combinatorial method for this is provided by Sarkaria’s Coloring and Embedding Theorem (there are other methods, but we only use this one here; see [Mat03] for a detailed discussion).

We associate to the simplicial complex \mathcal{K} the set system \mathcal{Z} of *minimal non-faces* of \mathcal{K} , that is, the inclusion-minimal sets of $2^{V(\mathcal{K})} \setminus \mathcal{K}$. For example, the minimal non-faces of the k -skeleton of the n -dimensional simplex is $\binom{[n+1]}{k+2}$. Sarkaria’s Theorem bounds the \mathbb{Z}_2 -index of $\mathcal{K}_{\Delta}^{*2}$ in terms of the chromatic number of the Kneser graph of \mathcal{Z} :

Theorem 0.10. *Let \mathcal{K} be a simplicial complex, \mathcal{Z} be the system of minimal non-faces of \mathcal{K} , and $\text{KG}(\mathcal{Z})$ be the Kneser graph on \mathcal{Z} . Then*

$$\text{Ind}_{\mathbb{Z}_2}(\mathcal{K}_{\Delta}^{*2}) \geq |V(\mathcal{K})| - \chi(\text{KG}(\mathcal{Z})) - 1.$$

Therefore, if \mathcal{K} can be realized in \mathbb{R}^d then

$$d \geq |V(\mathcal{K})| - \chi(\text{KG}(\mathcal{Z})) - 1.$$

In other words, we get good lower bounds on the possible embedding dimension of \mathcal{K} when we obtain good colorings of the Kneser graph on the minimal non-faces of \mathcal{K} . In the topological obstruction argument of Section 3, we thus look for suitable colorings of the minimal non-faces of the simplicial complex that we want not to be embeddable in a too small dimensional space.

Realization of the Delaunay Triangulation

Victor Alvarez*

Marek Krčál†

April 20, 2009

Abstract

1 Preliminaries and previous work

Let $T = (V, E, C, \phi)$ be a *maximal* planar graph such that:

- V is the set of vertices,
- E is the set of edges,
- C is the set of cycles that bound the inner cells (counterclockwise), and
- ϕ is a mapping from V to \mathbb{R}^2 such that $\phi(v_i) \neq \phi(v_j)$ whenever $i \neq j$.

The function ϕ can actually be thought as the drawing of T in \mathbb{R}^2 . In addition, ϕ will be called *non-degenerate* if and only if the following four conditions hold:

1. The boundary of the outermost cell forms a convex polygon.
2. The circumcircle of each inner cell does not contain any vertex in the interior.
3. No three vertices of the outermost cell are collinear.
4. No four vertices lie on a common circle that circumscribes an inner cell.

If only the first two conditions hold, then ϕ is called degenerate.

*Universität des Saarlandes, Saarbrücken, Germany. alvarez@cs.uni-sb.de. Supported by CRM Barcelona.

†ITI Prague, Czech Republic. krcal@kam.mff.cuni.cz. Supported by CRM Barcelona.




Figure 1:

1.1 The problem

Given a *maximal* plane graph $G = (V, E, C)$, we would like to find a *drawing* function ϕ (if possible) such that G is a Delaunay triangulation on the embedded set of points $V \subset \mathbb{R}^2$.

The problem of finding ϕ or even deciding if ϕ exists was believed to be time consuming roughly twenty years ago. However, near five years after that, an answer for the decision question was given based on Linear Programming. Hence, the decision question can be answered in polynomial time. In this project we are searching for an algorithm for finding the actual ϕ that works in polynomial time.

1.2 Recognizing the Delaunay triangulation in polynomial time

Given $G = (V, E, C)$, see Figure 1 for example, we observe the following:

- We can assign a variable x_i per internal angle, and therefore we will have $3|C|$ variables.
- Each $x_i > 0$.
- For each face C_j , the sum of the three corresponding angle variables must add up to 180.
- For each inner vertex, the sum of the surrounding angle variables must add up to 360.
- For each outer vertex, the sum of the surrounding angle variables must be less than 180.
- For each inner edge, the sum of the two opposite angle variables must be at most 180 (*empty circle property*).

The previous set of conditions can be written as a set of inequalities, for the previous figure we would have the following:

$$\begin{aligned}
x_1 + x_2 + x_3 &= 180 \\
x_4 + x_5 + x_6 &= 180 \\
x_7 + x_8 + x_9 &= 180 \\
x_1 + x_4 + x_7 &= 360 \\
x_3 + x_5 &< 180 \\
x_6 + x_8 &< 180 \\
x_9 + x_2 &< 180 \\
x_2 + x_6 &< 180 \\
x_5 + x_9 &< 180 \\
x_3 + x_8 &< 180 \\
x_1 &> 0 \\
&\vdots \\
x_9 &> 0
\end{aligned}$$

And from there we can actually derive the following linear program:

$$\begin{aligned}
x_1 + x_2 + x_3 &= 180 \\
x_4 + x_5 + x_6 &= 180 \\
x_7 + x_8 + x_9 &= 180 \\
x_1 + x_4 + x_7 &= 360 \\
u + y_1 + x_3 + x_5 &= 180 \\
u + y_2 + x_6 + x_8 &= 180 \\
u + y_3 + x_9 + x_2 &= 180 \\
u + y_4 + x_2 + x_6 &= 180 \\
u + y_5 + x_5 + x_9 &= 180 \\
u + y_6 + x_3 + x_8 &= 180 \\
-u - y_7 + x_1 &= 0 \\
&\vdots \\
-u - y_{15} + x_9 &= 0
\end{aligned}$$

subject to the maximization of u (as the Delaunay triangulation of a set of points maximizes the minimum angle), which in case of Figure 1 would give the drawing on the left in the following Figure:




Figure 2:

It is clear that any Delaunay triangulation will satisfy the variable constraints. However, the converse is not necessarily true, we could have the following assignment and therefore the drawing on the right in Figure 1:

- $x_1 = x_4 = x_7 = 120$,
- $x_2 = x_3 = x_5 = x_6 = 30$,
- $x_8 = 33$,
- $x_9 = 27$.

However, not everything is lost, in [4] the following Theorem was shown:

Theorem 1 *Given a maximal plane graph $G = (V, E, C)$ and the corresponding linear program, then G can be drawn as the Delaunay triangulation of set of points if and only if the linear program is feasible.*

The proof of the previous Theorem is rather technical and messy and we will not show it here as it is not the purpose of this note. It is our purpose to show two rather different approaches that might lead to a solution, that is, to find the desired function ϕ , that works in polynomial time.

2 Two different approaches

As we mentioned before, in the following we will show two independent approaches that might give us at least some new insight on the problem at hand. The first is a simple but smart incremental construction, and the second is based on Semi-definite programming. Both, as we will see, have their own pros and cons.

2.1 An incremental construction

In this section we will see how the procedure that decides if a triangulation T can be realized as the Delaunay triangulation of some set of points can be used to produce a smart incremental construction that might work to find the desired drawing function.

Let us start with a trivial observation:

Observation 1 *If $T = (V, E, C)$ turns out to be outerplanar, then the solution to the linear program is a set of points in convex position, and that solution is correct.*

Let us make the non-trivial yet weak assumption that T has constant degree, as in general the Delaunay triangulation is unlikely to have vertices of large degree (except for the cases that one can concoct). Also, as we already have the deciding procedure, for the sake of simplicity we will directly assume that T can be realized as a Delaunay triangulation.

With the previous information we can start explaining how our simple incremental construction works.

Let us pick a vertex $v \in V$ of the outermost face of T and let us remove it. Now, if we wanted to realize $T' = T \setminus \{v\}$, most probably we would have to add some edges to T' as this triangulation itself might not be realizable as a Delaunay triangulation, since the removal of v also removes its adjacencies. However, as T is realizable, we know that we can find some edges that triangulate the set of neighbours of v , nevertheless, we do not know what set of adjacencies is the right one. Fortunately, here is where we are going to explode the fact that v has constant degree. We will try by brute force all possible sets of adjacencies between the neighbours of v , including the empty set, and we will feed the decision algorithm with each one. As such algorithm works in polynomial time, and the number of possible triangulations of the neighbourhood of v is of constant size, the total running time of this test is still polynomial. Moreover, as we know that T is realizable as Delaunay, we know that the algorithm will give positive for at least one of the input triangulations.

If we recurse with the previous procedure over and over again, eventually we will hit a set of points that is outerplanar, and by Observation 1 we know that what the deciding algorithm gives back is correct, so we can actually draw it, realize it, and then put back the vertices that we removed, in the opposite order that we remove them.

As we could see, the previous algorithm is quite simple and it runs in polynomial time, since per vertex we require polynomial time to decide the right triangulation for the next step. Unfortunately, as we have one positive side, we also have one negative side. We might find problems at the time we insert back the vertices we had previously removed. The problems we may

encounter are that we could potentially lose or gain adjacencies that were or were not part of the original triangulation.

The problems arise from the fact that not each feasible solution to the linear program is correct, hence, from iteration to iteration the values of angles might change to a point that, as we said, we lose or gain new adjacencies. However, it might be possible that we could circumvent such problems by adding new constraints to the linear program with each iteration or even by changing the objective function. We are still studying some of this methods and we believe that this algorithm is sound enough to lead to the first polynomial time drawing algorithm. Nonetheless, at the same time, we are exploring different paradigms, the most promising being the one using Semi-definite programming, as the one we will describe next.

2.2 Semidefinite Programming

Let us first recall the general form of a semidefinite programming problem. For more details we refer reader to [5].

Primal version of semidefinite program (Primal SDP):

$$\begin{array}{ll} \text{maximize} & C \cdot X \quad \text{over all symmetric real } n \times n \text{ matrices } X \\ \text{subject to} & X \succeq 0 \\ & D_1 \cdot X = d_1 \\ & \vdots \\ & D_k \cdot X = d_k \end{array}$$

where C and D_i s are given symmetric real $n \times n$ matrices and, for example, $C \cdot X$ is a dot product (when matrices are considered as vectors), formally $C \cdot X = \text{tr}(CX)$. It is a basic fact, that (without changing the expressive power nor solvability) we can allow linear inequalities in the definition of the Primal SDP.

The idea of the second approach is based on the following two steps:

- Considering Primal SDP as an optimization problem over set of vectors in \mathbb{R}^n that will represent the vertices of the triangulation and centers of circles circumscribing the faces of the triangulation.
- Expressing the Delaunay constraints as a linear (in)equalities in terms of scalar products of those vectors

In more details, the well known characterization for a symmetric matrix $X \in \mathbb{R}^{n \times n}$ to be positive semidefinite is that there are vectors $v_1, v_2, \dots, v_n \in \mathbb{R}^n$ such that X is their Gramm matrix, i.e. $X_{i,j} = v_i \cdot v_j$.

From now on, let n denote the number of vertices ($n = |V|$) and m the number of faces ($m = |C|$) in the triangulation. As we indicated, we will represent the set of vertices V and the set S of centers of circles circumscribing

the inner (triangular) faces as vectors in Euclidean space, i.e. $V, S \subset \mathbb{R}^{n+m}$. Now, for each face $(a, b, c) \in C$ there is one $s \in S$ such that the following holds:

$$\|a - s\| = \|b - s\| = \|c - s\|$$

$$\text{For each } d \in V \text{ holds } \|d - s\| \geq \|a - s\|$$

After raising all the terms to the power of two, the previous can be rewritten in the form of linear (in)equalities in terms of scalar products of the vectors in V and S :

$$a \cdot a - 2a \cdot s + s \cdot s = b \cdot b - 2b \cdot s + s \cdot s = c \cdot c - 2c \cdot s + s \cdot s$$


$$\text{For each } d \in V \text{ holds } d \cdot d - 2d \cdot s + s \cdot s \geq a \cdot a - 2a \cdot s + s \cdot s$$

Taking the previous constraints for each face $(a, b, c) \in C$ we would already get a Primal SDP. For the sake of simplicity, we allow degenerate realizations as well as nondegenerate. But still, we did not encode the conditions stating that $a \neq b$ for each $a, b \in V$ and that the outer face is convex. And at last but not least, we need to get the vertices lying in the plane, not in a higher dimensional Euclidean space. Let us address this issues separately:


Convexity of the outer face

Here we are able to come up with a relaxed constraints that, in the case that we will be able to get the solution in the plane, would force the outer face to be convex.

Whenever we have a combinatorial triangulation, that can be realized in the plane, we can add three vertices v_0, v_1 and v_2 in such a way that the realization lies in the interior of their convex hull. The new Delaunay triangulation is fully described by the choice of three vertices a_0, a_1 and a_2 of the outer face such that each vertex of the segment of consecutive vertices of the outer face between a_i and a_{i+1} is connected with v_i .



For each such a choice (there is only cubically many of them) we can ask for a realization of such an extended triangulation. Now each edge (a, b) of the former outer face has two adjacent triangles. Let s be the center of the circle circumscribing the triangle containing (a, b) and one of the vertices v_0, v_1 or v_2 and let t be the circle circumscribing the other triangle containing (a, b) .



The vector $t - s$ is perpendicular to the vector $a - b$. Hence we are able to enforce the convexity by adding the following constraint for each such vertices a and b :

$$\text{For each } v \in V \quad v \cdot (t - s) \geq a \cdot (t - s)$$

The only thing, that is left to be considered, is to enforce the centers s and t to be nonequal. Since we can ask the vertices v_i to be arbitrarily far away from the original triangulation we can also ask the center t to be arbitrarily far away from the original triangulation, so in particular the constraint

$$||s - t||^2 \geq 1$$

can be added.

Making the vertices different

If we could prove a polynomial lower bound on the minimal edge length such that any combinatorial triangulation has a realization with all edge lengths greater (polynomial in terms of the average length of edges, say) we could then add such a lower bound as a constraint for length of every edge without changing the polynomial solvability of the Primal SDP. In the current experiments we use some small fixed constant.

Making the solution planar

This seems to be a crucial part of the whole approach. The solution with all vectors in a plane corresponds to a positive semidefinite matrix of rank 2. Even asking for a solution of a rank by one smaller than the size of the matrix for a general Primal SDP is NP-hard ([5].) On the other hand there are special classes of programs where the solution optimal according to some appropriate objective function has a lower rank, see [3, 4]. In the spirit of these results we try to come up with an objective function such that the optimal solution would be of rank two.

We have made an experimental Python script that (by use of CSDP [2]) tries to solve the Delaunay-realization problem for a given combinatorial triangulation. We plan to encode a bunch of promising objective functions and run experiments for a set of triangulations. The promising objective functions include:

$$\sum_{a \in V} \left(\frac{1}{|D(a)|} \sum_{b \in D(a)} ||a - b|| - \frac{1}{|N(a)|} \sum_{b \in N(a)} ||a - b|| \right)$$


where the $N(a)$ is the set of neighbors of a and $D(a)$ is some appropriate set of vertices (or their linear combinations) that are in some sense distant from a . Additionally, we add a constraint $\sum_{\text{face } f} r_f^2 = m$ where r_f is a radius of the circle circumscribing the face f . We can do this since arbitrary solution can be arbitrarily scaled so for example we can ask for a solution with the average area of all the circles fixed.

The particular choices for $D(a)$ might be:

- The set of all vertices except a ($V \setminus \{a\}$).

- The set of all vertices except a and its neighbors ($V \setminus (\{a\} \cup N(a))$).
- The set of all vertices of the outer face, when a is an inner vertex (not on the outer face) and the set of all vertices otherwise.

The experiments made so far include only one rather simple and naive objective function (that we will not present here), yet most the solutions obtained had limited rank up to 5. Visualizations of some of the successful solutions are included:



References

- [1] A. I. Barvinok, *Problems of distance geometry and convex properties of quadratic maps*. Discrete Comput. Geom. **13**, no. 2, (1995), 189–202.

- [2] B. Borchers. *CSDP, A C Library for Semidefinite Programming*. Optimization Methods and Software, **11(1)**, (1999), 613–623.
- [3] Frank Gring, Christoph Helmberg, Markus Wappler, *Embedded in the shadow of the separator*. SIAM J. Optim. **19**, no. 1, (2008), 472–501.
- [4] T. Hiroshima, Y. Miyamoto and K. Sugihara *Another Proof of Polynomial-Time Recognizability of Delaunay Graphs*. IEICE Trans. Fundamentals, **E83-A**, No. 4, (2000), 627–638.
- [5] László Lovász, *Semidefinite programs and combinatorial optimization*, minicourse & lecture notes, Budapest, (1995)

Recoloring Directed Graphs

Stefan Felsner*

Clemens Huemer†

Maria Saumell‡

Abstract

Let G be a directed graph and k a positive integer. We define the k -color graph of G ($D_k(G)$ for short) as the directed graph having all k -colorings of G as node set, and where two k -colorings β and φ are joined by a directed edge $\beta \rightarrow \varphi$ if φ is obtained from β by choosing a vertex v and recoloring v so that its color is different from the colors of all its out-neighbors. We investigate reachability questions in $D_k(G)$. In particular we want to know whether a fixed legal k' -coloring ψ of G with $k' \leq k$ is reachable in $D_k(G)$ from every possible initial k -coloring β . Interesting instances of this problem arise when G is planar and the orientation is an arbitrary α -orientation for fixed α . Our main result is that reachability can be guaranteed if the orientation has maximal out-degree $\leq k - 1$ and an accessible pseudo-sink.

1 Introduction

Let $G = (V, A)$ be a directed graph and k a positive integer. A k -coloring of G is an assignment of colors from the set $\{1, 2, \dots, k\}$ to the vertices of G , that is, a mapping $\gamma : V \rightarrow \{1, 2, \dots, k\}$. A *legal k -coloring* is a k -coloring with the property that the end-vertices of every edge receive different colors.

Given G and k , we define the *k -color graph of G* , denoted by $D_k(G)$, as follows: its node set consists of all k -colorings of G , and two k -colorings β and φ are joined by a directed edge $\beta \rightarrow \varphi$ if they differ in color on precisely one vertex v of G and $\varphi(v) \neq \varphi(w)$ for all out-neighbors w of v , in other words, if φ is obtained from β by choosing a vertex v and recoloring v so that its color is different from the colors of all its out-neighbors. A recoloring of a vertex in this way will be called a *valid recoloring step*.

If β and φ are two k -colorings of G , we may ask whether there is a path from β to φ in $D_k(G)$. In this case, we say that φ is *reachable* from β . We have studied reachability questions in $D_k(G)$.

Traditionally the colorings that have received more attention in the literature are the legal ones, and one may consider non-legal colorings to have bad properties. Therefore an interesting question is to determine whether a non-legal coloring can be easily modified so as to become legal. Here the reachability issue we answer is precisely whether any coloring can be transformed into a legal coloring by a sequence of valid recoloring steps, that is, whether, given some coloring, it is always possible to reach a legal one in $D_k(G)$. We have settled this problem for a special class of directed graphs, namely, the α -orientations.

An α -orientation of a graph $G = (V, A)$ is an orientation of its edges such that the out-degree of every vertex $v \in V$ is given by a function $\alpha : V \rightarrow \mathbb{N}$ (see, for example, [9, 6, 10]). We deal with α -orientations in which the out-degree of each vertex is bounded from above. More precisely, in Sections 2 and 3 we look at two particular cases: 2-orientations of planar quadrangulations (of which we describe the 3-color graph) and 3-orientations of planar triangulations (of which we consider the 4-color graph). In Section 4 we generalize our results.

*Institut für Mathematik, Technische Universität Berlin, felsner@math.tu-berlin.de.

†Departament de Matemàtica Aplicada IV, Universitat Politècnica de Catalunya, Barcelona, Spain, clemens.huemer@upc.edu.

‡Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Spain, maria.saumell@upc.edu.

Before we proceed, let us mention that there exists some related work in the context of undirected graphs. If H is such a graph, the k -color graph of H has been defined as the graph with node set the legal k -colorings of H , and where two nodes are adjacent whenever the corresponding colorings differ in color on just one vertex v of H [5]. The connectivity of the k -color graph of undirected graphs has been studied in [3, 4, 5].

2 The 3-color graph of a 2-orientation of a planar quadrangulation

2.1 Definitions and results

A *planar quadrangulation* is a plane graph $Q = (V \cup \{s, t\}, E)$ such that Q is a maximal bipartite graph with $|V| = n$, and s, t are two non-consecutive vertices on the outer face. A *2-orientation* of Q is an orientation of its edges such that every vertex different from s and t has out-degree two. Given a 2-orientation of Q , it is easy to see that s and t are sinks, i.e., vertices with out-degree zero. It has been proved that every quadrangulation admits a 2-orientation [9, 11].

Since a quadrangulation is a bipartite graph, its vertices can be divided into two groups, the white ones and the black ones, in such a way that every edge is incident to two vertices from different groups. Let us assume that s and t are black vertices. The following definition is useful for our purposes:

Definition 2.1. A *separating decomposition* is an orientation and coloring of the edges of Q with colors red and blue satisfying that:

- (1) The edges incident to s are colored red and the edges incident to t are colored blue.
- (2) Every vertex $v \neq s, t$ is incident to an interval of red edges and an interval of blue edges. Among these edges, only one red edge and one blue edge are outgoing. If v is white, the two outgoing edges are the first ones in clockwise order in the interval of edges of their color. If v is black, the outgoing edges are the clockwise last in their color.

It can be shown that separating decompositions fulfill the additional property that the red edges form a directed tree rooted in s (T_r for short) and the blue edges form a tree rooted in t (T_b for short) [7]. See Figure 1 for an example. This ensures that there exist paths in Q from every vertex $v \neq s, t$ to s and t , a property that will be used later.




Figure 1: A quadrangulation with a separating decomposition.

For each non-special vertex $v \neq s, t$ we denote the directed paths in T_r and T_b from v to the corresponding sink by $P_r(v)$ and $P_b(v)$, respectively. Then one can prove that the paths $P_r(v)$ and

$P_b(v)$ only share vertex v , so they split the quadrangulation into two regions [8]. We define $R_r(v)$ (respectively, $R_b(v)$) to be the region to the right of $P_r(v)$ (respectively, $P_b(v)$) and including both paths. A nice observation here is that, for each vertex of the path, all incoming edges belonging to a fixed region have the same color.

Separating decompositions and 2-orientations are closely related. Indeed let $Q = (V \cup \{s, t\}, E)$ be a plane quadrangulation; then it has been proved that there is bijection between separating decompositions and 2-orientations of Q [7, 9]. To be precise, a separating decomposition yields a 2-orientation, while the edges of a 2-orientation can be colored so as to obtain a separating decomposition.

Now that we have all the necessary background, we can state our first result:

Theorem 2.2. *Let G be a 2-orientation of a planar quadrangulation. Then, given any 3-coloring β of G , it is always possible to reach a legal 2-coloring in $D_3(G)$.*

We prove Theorem 2.2 in three different ways. Our proofs are constructive, that is, we provide algorithms to obtain the path from β to the legal 2-coloring in $D_3(G)$, and we show that they are correct.

All proofs make use of Lemma 2.3. The idea behind it is as follows. Suppose that v is a non-special vertex such that v and all its out-neighbors have different colors in β (v is a *rainbow* in β). Although the two outgoing edges from v are legal, we might need to recolor v to reach a legal coloring in $D_3(G)$, which is not possible in one single valid recoloring step. Lemma 2.3 shows that v can nevertheless be recolored.

Suppose that ς, τ are two different colors in the set $\{1, 2, 3\}$. The expression $\overline{\{\varsigma, \tau\}}$ will denote the color in $\{1, 2, 3\}$ different from ς and τ .

Lemma 2.3. *Let $P = \{v_l, v_{l-1}, \dots, v_2, s\}$ be a directed path in G without any directed chord. Let $f(v_l)$ denote the color of the vertex pointed by v_l different from v_{l-1} . If v_l is colored ς and $f(v_l) = \tau$, there exists a sequence of valid recoloring steps involving only $v_l, v_{l-1}, \dots, v_2, s$ such that, at the end, v_l has color $\overline{\{\varsigma, \tau\}}$.*

Proof. We prove the statement by induction on the length l of P . First observe that, in any situation, we can assume that s has the most convenient color, as all edges incident to this vertex are incoming.

Let us look at the case $l = 2$, i. e., $P = \{v_2, s\}$. Suppose s is colored ς . Then, in a valid recoloring step, v_2 can be assigned color $\overline{\{\varsigma, \tau\}}$. Now assume that the lemma is valid for all paths of length at most $l - 1$. If v_{l-1} has color different from $\overline{\{\varsigma, \tau\}}$, then v_l can be recolored with this color. Otherwise, by hypothesis of induction, v_{l-1} can be assigned a color different from $\overline{\{\varsigma, \tau\}}$ by applying a sequence of valid recoloring steps that only involves vertices $v_{l-1}, v_{l-2}, \dots, v_2, s$. Since v_l does not point to any of the vertices $v_{l-2}, v_{l-3}, \dots, v_2, s$, these recoloring steps do not modify the value of $f(v_l)$. Therefore, in the new coloring v_l points to a vertex with color τ and a vertex with color different from $\overline{\{\varsigma, \tau\}}$, so it can be assigned color $\overline{\{\varsigma, \tau\}}$. \square

Observe that the previous proof already provides an algorithm to recolor vertex v_l in $O(l)$ valid recoloring steps. It consists of visiting the vertices of the path (starting from v_l) until a non-rainbow vertex is found, changing the color of this vertex and then recoloring backwards. Figure 2 shows an example of the steps executed by the algorithm.

2.2 First proof of Theorem 2.2

Let β be a 3-coloring of G . We prove that, from β , it is always possible to reach a coloring where all white vertices have color 2 or 3. Our proof is based on the next lemma:

Lemma 2.4. *There exists a sequence of valid recoloring steps that increases the number of white vertices that have color 2 or 3 by one, unless all white vertices already have color 2 or 3.*

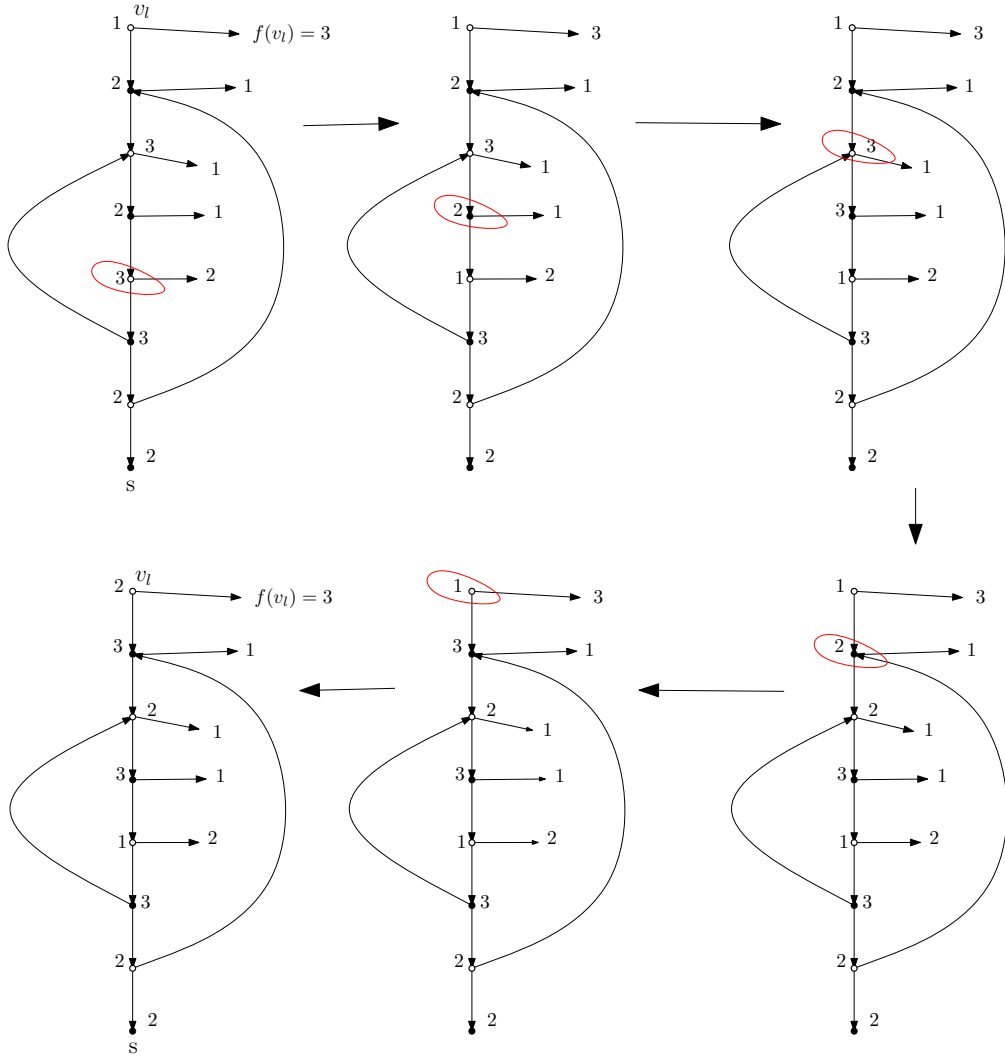


Figure 2: An example for Lemma 2.3. It is shown how v_l can be recolored.

Note that Lemma 2.4 then proves Theorem 2.2. We can repeatedly apply Lemma 2.4 until no white vertex has color 1. Then, since the graph is bipartite, each black vertex can be colored with 1 and we obtain a legal coloring. In a last step we can even reach a legal coloring with two colors by recoloring every white vertex with, say, color 2. This legal 2-coloring will be denoted by ψ .

Note also that in Lemma 2.4 we are only interested in the number of white vertices having colors 2 or 3; we might create many monochromatic edges (i.e., edges connecting points of the same color) during the recoloring process.

We need an auxiliary result in order to prove Lemma 2.4.

Lemma 2.5. *Let $P = \{v_l, v_{l-1}, \dots, v_2, s\}$ be a directed path in G without any directed chord. Then there exists a sequence of valid recoloring steps involving only the vertices of P such that, at the end, each white vertex in P has color 2 or 3.*

Proof. Again, we prove the lemma by induction on the length l of P . The base case is $l = 2$. Let $P = \{v_2, s\}$; we can assume that s has color 1. Vertex v_2 has two outgoing edges, where at least one points to a vertex of color 1. Thus v_2 does not point to both colors 2 and 3, and we can assign the desired color to v_2 .

Now suppose that the statement holds for each path P' of length $l-2$ (we go from $l-2$ to l because, if a path starts at a black vertex, we simply ignore this vertex). Consider a path P of length l . First we show that v_l can be recolored to color 2 or 3 (by probably messing up the coloring of the whole path). If v_l already has color 2 or 3, we only have to consider the path P' from v_{l-2} to s , which, by induction, can be well recolored without changing the color of vertex v_l . The same holds if v_l points to a vertex of color 1 - then we simply recolor v_l and apply induction to P' . Let us thus assume that v_l has color 1 and points to vertices of color 2 and 3. Without loss of generality, suppose that v_{l-1} has color 2. Then we apply Lemma 2.3 with $f(v_l) = 3$, that is, we apply the sequence of valid recoloring steps involving only $v_l, v_{l-1}, \dots, v_2, s$ such that, at the end, v_l has color 2. We are left with a path $P' = \{v_{l-2}, \dots, v_1\}$ which can be well recolored by induction. \square

Following the proof of Lemma 2.5, we find that a path of length l satisfying the hypothesis can be recolored so that each white vertex has color 2 or 3 in $O(l^2)$ steps.

Now we have all the ingredients to prove Lemma 2.4:

Proof of Lemma 2.4. Among all white vertices of color 1, pick some vertex v of minimum distance to the sink s (here we use the property that, in a separating decomposition, there exist paths from every vertex $v \neq s, t$ to the sinks). Consider a shortest directed path from v to s . This path fulfills the hypothesis of Lemma 2.5. Thus, by applying Lemma 2.5, we can recolor all vertices along the path so as that all white vertices have color 2 or 3. Since only vertices along the path have been recolored, the number of white vertices that are assigned 2 or 3 increases by one. \square

Algorithm 1 summarizes the first way we have presented to reach a legal 2-coloring in $D_3(G)$, given a 3-coloring β of G . Let us count the number of valid recoloring steps performed by the algorithm. Steps 3 and 4 entail a linear number of recolorings. As for Step 2, we have already pointed out that the algorithm in Lemma 2.5 for a given vertex v takes $O(d_G^2(v, s))$ recoloring steps, where the distance $d_G(v, w)$ between a vertex v and a vertex w is the length (i.e., the number of edges) of the shortest directed path from v to w in G ($d_G(v, w) = +\infty$ if there are no directed paths from v to w). Hence, if the set of white vertices in G is denoted by V_w , Algorithm 1 executes at most $O(n) + \sum_{v \in V_w} d_G^2(v, s) = O(n^3)$ valid recoloring steps.

Algorithm 1 Reaching a legal 2-coloring, first algorithm

- 1: Sort all white vertices of color 1 by increasing distance to the sink s . Let L be the ordered list containing these vertices.
 - 2: Let v be the first element in L . Remove v from L and apply the algorithm in Lemma 2.5 to assign color 2 or 3 to all white vertices in the shortest path in G between v and s . Repeat until L is empty.
 - 3: Assign color 1 to all black vertices.
 - 4: Recolor all white vertices with color 2.
-

Now let H be a directed graph and r a vertex of H . We say that r is *accessible* if H contains a directed path from every other vertex to r .

Corollary 2.6. *Let H be a bipartite graph with an accessible sink. Then ψ is reachable in $D_3(H)$ from all 3-colorings of H .*

2.3 Second proof of Theorem 2.2

In this subsection we propose a method that gives a shorter path in $D_3(G)$ from the coloring β to the coloring ψ . This method is explained in Algorithm 2.

Step 2 of Algorithm 2 involves postorder traversal in the red tree. In a binary tree, postorder traversal performs the following operations recursively at each node: (1) traverse the left subtree; (2) traverse the right subtree; (3) visit the root. This process can be generalized in the natural way to plane m -ary trees, as the notions of left and right are well-defined. The vertices in Figure 1 are labelled according to postorder in the red tree.

Let us introduce some notation. If $v \neq s, t$ is a vertex in G , then the red (respectively, blue) edge outgoing from v points to a vertex that will be denoted by v_r (respectively, v_b). If the path $P_r(v)$ is of the form $P_r(v) = \{v, v_{l-1}, \dots, v_2, s\}$, we define $P'_r(v)$ as one of the paths from v to s on a subset of $\{v, v_{l-1}, \dots, v_2, s\}$ which cannot be shortened by directed chords.

Algorithm 2 Reaching a legal 2-coloring, second algorithm

- 1: If necessary, recolor vertex t with 1.
 - 2: Sort all vertices different from s and t according to the postorder traversal sequence in the red tree. Let L be the ordered list containing the vertices.
 - 3: Let v be the first element in L . Remove v from L .
 If v is white and v_r has color 2, apply the algorithm in Lemma 2.3 to recolor v_r using the path $P'_r(v_r)$. Afterwards, assign color 2 to v .
 If v is black, apply (if necessary) the algorithm in Lemma 2.3 to make vertices v_r and v_b have color different from 1 by using the paths $P'_r(v_r)$ and $P'_r(v_b)$, respectively. Afterwards, assign color 1 to v .
 Repeat until L is empty.
 - 4: If necessary, recolor vertices s and t with 1.
-

Lemma 2.7. *Given β , Algorithm 2 gives a sequence of $O(n^2)$ valid recoloring steps to reach ψ .*

Proof. Let us look at Step 3. Let v be the first element in L at some point of the algorithm. If v is a white vertex, the red edge outgoing from v_b belongs to the region $R_b(v)$ (see Figure 3). This implies that v_b has already been processed and, as will be shown in a few lines, has its final color, which, in this case, is 1. Our aim is to assign to v its final color, namely, 2. If v_r is not colored 2, we can do it in one single valid recoloring step. Otherwise we first change the color of v_r by means of the algorithm in Lemma 2.3; we use the path $P_r(v_r) = \{v_r, v_{l-1}, \dots, v_2, s\}$, possibly shortened by directed chords. Since this algorithm only recolors vertices in the path and all of them are after v in L , we do not change the color of any vertex that has already been processed.

If v is black, the red edge outgoing from v_b belongs to $R_r(v)$ (see Figure 3). Since all red edges pointing to v are in $R_b(v)$, v_b is after v if T_r is traversed in postorder. So, in this case, neither v_r nor v_b have already been processed and we might change their color to be able to assign color 1 to v . Again, if we use the appropriate paths, we do not recolor any vertex that has already been treated.

To conclude, after an iteration of Step 3 the number of vertices that are assigned its final color and will not be recolored in future steps increases by one. Consequently, the algorithm is correct.

As for the number of valid recoloring steps performed by the algorithm, the assignment of the final color of each vertex v costs $O(l)$ recolorings, where l is the length of the path $P_r(v)$. Therefore, the total number of recoloring steps is $O(n^2)$. \square




Figure 3: Postorder in the red tree of the vertices v, v_r, v_b .

2.4 Third proof of Theorem 2.2

Finally, we provide one more algorithm to reach the coloring ψ in $O(n^2)$ steps. This algorithm has the property that can be generalized to the k -color graph of some α -orientations (see Sections 3 and 4).

Lemma 2.8. *Given β , Algorithm 3 gives a sequence of $O(n^2)$ valid recoloring steps to reach ψ .*

Proof. This is a particular case of Lemma 4.3, which is proved in Section 4. \square

Algorithm 3 Reaching a legal 2-coloring, third algorithm

- 1: Sort the non-special vertices by decreasing distance to s (choose any order if there are ties). Let L be the ordered list containing the vertices.
 - 2: Let v be the first element in L . Remove v from L .
 If v is white, apply (if necessary) the algorithm in Lemma 2.3 to make vertices v_r and v_b have color different from 2 by using their shortest paths in G to s . Afterwards, assign color 2 to v .
 If v is black, apply (if necessary) the algorithm in Lemma 2.3 to make vertices v_r and v_b have color different from 1 by using their shortest paths in G to s . Afterwards, assign color 1 to v .
 Repeat until L is empty.
 - 3: If necessary, recolor vertices s and t with 1.
-

3 The 4-color graph of a 3-orientation of a planar triangulation

A *plane triangulation* T is a maximal plane graph with n vertices and three special vertices a_1, a_2, a_3 in the outer face. A *3-orientation* of T is an orientation of its edges such that every vertex different from a_1, a_2, a_3 has out-degree three. It can be proved that every triangulation admits a 3-orientation [9].

Definition 3.1. A *Schnyder wood* is an orientation and coloring of the inner edges of T with colors red, green and blue satisfying that:

- (1) The edges incident to a_1 are colored red, the edges incident to a_2 are colored green and the edges incident to a_3 are colored blue.
- (2) Every vertex $v \neq a_1, a_2, a_3$ has three outgoing edges that, in clockwise order, have colors red, green and blue. The incoming edges in an interval between two outgoing edges have the third color.

An example of a triangulation with a Schnyder wood is shown in Figure 4.

Analogously to in the case of separating decompositions, Schnyder woods have the property that the edges of each color form a directed tree rooted at one of the vertices in the outer face [12]. Hence, there exist paths from every inner vertex to each of the special ones.

Another analogy with separating decompositions is the following theorem (see [7, 9]). If T is a plane triangulation with outer vertices a_1, a_2, a_3 , then Schnyder woods and 3-orientations of T are in bijection. Furthermore, the edges of a 3-orientation of T can be colored to obtain a Schnyder wood.

In the case of 3-orientations of a planar triangulations, we can prove:

Proposition 3.2. *Let G be a 3-orientation of a planar triangulation. Then, given any 4-coloring β of G , it is always possible to reach a legal 4-coloring in $D_4(G)$.*

Proof. The path in $D_4(G)$ from β to a given legal 4-coloring of G (which exists by the 4-colour theorem [1, 2]) is computed by Algorithm 4, which is shown to be correct in Lemma 4.3. \square




Figure 4: A triangulation with a Schnyder wood.

4 General result

In this section we generalize the previous results to a special class of α -orientations that includes 2-orientations of planar quadrangulations and 3-orientations of planar triangulations.

Let G be a directed graph in which the maximum out-degree of the vertices is $k - 1$, and let β be a k -coloring of G . As in the case $k = 3$, we will say that a vertex v is a *rainbow* in β if v has out-degree $k - 1$, and v and all its out-neighbors have different colors in β . We will say that v is a *pseudo-sink* if it is guaranteed to be a non-rainbow in every k -coloring of G .

Theorem 4.1. *Let G be a directed graph where the maximum out-degree of the vertices is $k - 1$. Assume that G has an accessible pseudo-sink s . Then, given any k -coloring β of G and any legal k' -coloring ψ of G with $k' \leq k$, it is possible to reach ψ from β in $D_k(G)$.*

The proof of Theorem 4.1 uses several ingredients we have already seen. First we need a generalization of Lemma 2.3. We omit its proof because its very similar to the one in Lemma 2.3.

Lemma 4.2. *Let $P = \{v_l, v_{l-1}, \dots, v_2, s\}$ be a directed path in G without any directed chord. If v_l is colored ς , there exists a sequence of valid recoloring steps involving only $v_l, v_{l-1}, \dots, v_2, s$ such that, at the end, v_l has a color different from ς .*

The method we propose to obtain a path from β to ψ in $D_k(G)$ is detailed in Algorithm 4.

Algorithm 4 Reaching a legal k' -coloring

- 1: Sort all vertices by decreasing distance to s (choose any order if there are ties). Let L be the ordered list containing the vertices. For each $v \in L$, let $\delta_{\leq}(v)$ be the set of out-neighbors of v that are after v in L .
 - 2: Let v be the first element in L . Remove v from L .
Apply (if necessary) the algorithm in Lemma 4.2 to make the vertices in $\delta_{\leq}(v)$ have color different from $\psi(v)$ by using their shortest paths in G to s . Afterwards, assign color $\psi(v)$ to v .
Repeat until L is empty.
-

Lemma 4.3. *Given β and ψ , Algorithm 4 gives a sequence of $O(n^2)$ valid recoloring steps to reach ψ from β .*

Proof. The vertices are processed in such a way that the recolorings necessary to assign the definite color to a vertex do not involve the vertices that have already been treated. Thus, after each iteration of Step 2, the number of vertices that are assigned its final color and will not be recolored increases.

The number of valid recoloring steps performed by the algorithm can be counted as in the proof of Lemma 2.7. \square

Acknowledgements

This work has been carried out in the context of the I-Math Winter School DocCourse Combinatorics and Geometry 2009: Discrete and Computational Geometry, organized by Centre de Recerca Matemàtica. We thank the organizers of the course and the other participants for providing a stimulating working environment.

References

- [1] Kenneth Appel and Wolfgang Haken. Every planar map is four colorable. I. Discharging. *Illinois Journal of Mathematics*, 21(3):429-490, 1977.
- [2] Kenneth Appel, Wolfgang Haken and John Koch. Every planar map is four colorable. II. Reducibility. *Illinois Journal of Mathematics*, 21(3):491-567, 1977.
- [3] Paul Bonsma and Luis Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. In Mathematical Foundations of Computer Science 2007 (MFCS 07). Lecture Notes in Computer Science, 4708:738-749, 2007.
- [4] Luis Cereceda, Jan van den Heuvel and Matthew Johnson. Mixing 3-colourings in bipartite graphs. In Graph-Theoretic Concepts in Computer Science, Proceedings of the 33rd International Workshop on Graph Drawing (WG 2007). Lecture Notes in Computer Science, 4769:166-177, 2007.
- [5] Luis Cereceda, Jan van den Heuvel and Matthew Johnson. Connectedness of the graph of vertex-colourings. *Discrete Mathematics*, 308(5-6):913-919, 2008.
- [6] Stefan Felsner. Lattice structures from planar graphs. *Electronic Journal of Combinatorics*, 11(1), Research paper R15, 24 pp., 2004.
- [7] Stefan Felsner, Éric Fusy, Marc Noy and David Orden. Bijections for Baxter families and related objects. *Submitted*, arXiv:0803.1546, 2008.
- [8] Stefan Felsner, Clemens Huemer, Sarah Kappes and David Orden. Binary labelings for plane quadrangulations and their relatives. *Submitted*, arXiv:math.CO/0612021, 2008.
- [9] Hubert de Fraysseix and Patrice Ossona de Mendez. On topological aspects of orientations. *Discrete Mathematics*, 229:57-72, 2001.
- [10] Patrice Ossona de Mendez. Orientations bipolaires. PhD thesis, École Des Hautes Études en Sciences Sociales, Paris, 1994.
- [11] Atsuhiko Nakamoto and Mamoru Watanabe. Cycle reversals in oriented plane quadrangulations and orthogonal plane partitions. *Journal of Geometry*, 68:200-208, 2000.
- [12] Walter Schnyder. Planar graphs and poset dimension. *Order*, 5(4):323-343, 1989.

Sets with Small Neighborhood in the Integer Lattice

Aaron Dall, Frederik von Heymann, Birgit Vogtenhuber

April 20, 2009

Abstract

For a given cardinality we want to find lattice point configurations such that the number of lattice points with distance 1 to the set is small. Sets with the smallest number possible are called optimal. It is known that sets of points with coordinate sum less or equal to some integer k are optimal. We show that they are unique for their cardinalities. Also we will discuss the question of how to characterize optimal sets in general, and if by adding the points of distance 1 to an optimal set we will always get an optimal set. In both cases the answer is positive for the plane.

Acknowledgements

We thank the organizers of the DocCourse and the CRM for supporting us and our research from Jan-Mar 2009, and Oriol Serra for proposing the problems considered in this paper.

1 Introduction

Consider the d -dimensional lattice \mathbb{Z}^d with the distance $d(x, y) = \sum_i |x_i - y_i|$, where $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$. Given a finite set $X \subset \mathbb{Z}^d$, the *neighborhood*, $N(X)$, of X is

$$N(X) = \left\{ y \in \mathbb{Z}^d \setminus X : d(y, X) = 1 \right\},$$

and the size of this neighborhood is $n(X) = |N(X)|$. Further, the *boundary*, ∂X , of X are the points of X that are next to some point of $N(X)$, and the *interior*, $\text{int } X$, of X are the points of X for which all neighbors are in X , i.e.,

$$\partial X = \{x \in X : d(x, N(X)) = 1\},$$

$$\text{int } X = \{x \in X : d(x, N(X)) > 1\} = X \setminus \partial X.$$

For any $r \in \mathbb{N}$, the set $B_r^d = \{x \in \mathbb{Z}^d : \sum_i x_i < r\}$ is called the *ball* (of radius r). Note that the neighborhood $N(B_r^d)$ of a ball B_r^d contains exactly the points of \mathbb{Z}^d with coordinate sum r .




Figure 1: The ball B_4^2 (blue) and its neighborhood (yellow).

A finite set $X \subset \mathbb{Z}^d$ is called *optimal* if the size of its neighborhood $n(X)$ is minimal among all sets $Y \subset \mathbb{Z}^d$ with $|Y| = |X|$.

Background. In this paper we consider the problem of characterizing point sets with minimum neighborhood size among all sets of fixed cardinality. Isoperimetric problems of this kind have arisen in a number of different contexts, with several definitions of neighborhood, and several different underlying finite and infinite lattices. One approach to a solution is providing an ordering of the lattice points such that the first j of them form an optimal set of their cardinality for every j . For the Boolean lattice (chains of length two) this is the celebrated theorem of Harper [Har66]. Kruskal-Katona, and Clements and Lindström [CL69] solve this for chains of arbitrary length l . Beruzkov and Serra [BS02] consider the problem for cartesian powers of graphs.

Macaulay [Mac27] presents an ordering of the nonnegative d -dimensional integer points \mathbb{Z}_+^d having coordinate sum $\leq k$ such that the first j of them have a minimum number of neighbors with coordinate sum $k+1$ among all sets of k -sum points. Wang and Wang [WW77a] present, as a similar (and equivalent) result, sets in \mathbb{Z}_+^d that minimize the number of neighbors in \mathbb{Z}_+^d , and extended this to an ordering of the points of \mathbb{Z}^d such that the first j of them minimize the number of neighbors in \mathbb{Z}^d . They call these optimal

sets *standard spheres*, and we adopt this terminology. Basically, a standard sphere is a ball plus possibly some points of the neighborhood of the ball.

As it happens, the sequence of standard spheres enjoys the property that it is closed under the operation of adding the neighborhood. In particular, balls B_r^d are optimal sets in \mathbb{Z}^d , which is the analogous result to the classical isoperimetric problem in Euclidean space. In contrast with its Euclidean counterpart, standard spheres are not, in general, the only optimal sets in lattices.

In the case of the Boolean lattice, all optimal sets have been characterized by Beruzkov [Bez89], while for general lattices the complete characterization is still an open problem.

For a survey about different types of isoperimetric problems, as well as a thorough list of references, see [Bez94].

In this paper we address the following questions.

Problem 1.1. *Let $X \subset \mathbb{Z}^d$ be an optimal set. Is it true that $N(X)$ is also an optimal set?*

The answer is yes in the Boolean lattice (see [WW77b]). In Section 3 we give a proof for \mathbb{Z}^2 .

Problem 1.2. *Is it true that balls are the only optimal sets of their cardinality in the d -dimensional lattice?*

The answer is again yes for the Boolean lattice (see [Bez94] and references therein). We prove that it is also true for the d -dimensional lattice (see Section 2).

Problem 1.3. *What are necessary and sufficient conditions for sets $X \subset \mathbb{Z}^d$ to be optimal?*

In Section 3 we explain some necessary conditions for optimality in \mathbb{Z}^2 .

1.1 Some Basic Observations

We call a set $X \subset \mathbb{Z}^d$ *connected* if, for any $x, y \in X$, there exists a path $(v_j)_{1 \leq j \leq k}$ from x to y in X such that any two consecutive points in the path differ by at most one in each coordinate, i.e. $v_j - v_{j+1} = \sum_{i \leq d} c_i e_i$ with $c_i \in \{-1, 0, 1\}$.

Proposition 1.4. *A necessary condition for a finite set $X \subset \mathbb{Z}^d$ to be optimal is that X is connected.*

Proof. Consider a finite set X consisting of two connected components U and V . Further let u_{\max} be some point of U with maximal first coordinate, and v_{\min} a point of V with minimal first coordinate. Then $u_{\max} + (1, 0, \dots, 0) \in N(X)$. Thus by translating V such that $u_{\max} + (1, 0, \dots, 0) = v_{\min}$ we get a set X' with $|X'| = |X|$ and $n(X') \leq n(X) - 1$. \square

In the following, all considered sets $X \subset \mathbb{Z}^d$ will be finite and connected.

Proposition 1.5. *The optimal neighborhood size is increasing with the cardinality of the point set: if X and Y are optimal sets with $|X| < |Y|$, then $n(X) \leq n(Y)$.*

Wang and Wang show this for standard spheres [WW77a], and thus it has to be true for all optimal sets.

Proposition 1.6. *Consider the ball B_r^d , with $|B_r^d| = s$. Then the neighborhood of any set $X \subset \mathbb{Z}^d$ with $|X| = s + 1$ has size $n(X) \geq n(B_r^d) + d - 1$.*

Proof. This follows again directly from the results in [WW77a]. The standard sphere with $s + 1$ points is B_r^d plus a point $x \in N(B_r^d)$ in the positive orthant. The grid point x has coordinate sum r , and thus it has in each coordinate direction one neighbor with coordinate sum $r + 1$. All these neighbors must lie in $N(B_r^d \cup \{x\}) \setminus N(B_r^d)$. Thus the neighborhood of $B_r^d \cup \{x\}$ has at least size $n(B_r^d) + d - 1$.

As standard spheres are optimal, any set of cardinality $s + 1$ needs to have at least this neighborhood size. \square

Given a set $X \in \mathbb{Z}^d$, we want to consider slices of X with respect to some coordinate direction. To this end we denote by

$$L_{k,l}(X) = \{x \in X : x_k = l\}$$

the l^{th} k -level of X . Every k -level lies in a $(d - 1)$ -dimensional hyperplane, and we denote by $L_{k,l}^{d-1}(X) \in \mathbb{Z}^{d-1}$ the $(d - 1)$ -dimensional set that results from $L_{k,l}(X)$ by omitting the k -th coordinate:

$$L_{k,l}^{d-1}(X) = \left\{ (x_1, \dots, x_{k-1}, x_{k+1}, \dots, x_d) \in \mathbb{Z}^{d-1} : (x_1, \dots, x_d) \in L_{k,l}(X) \right\}.$$

In the other direction, given some set $Y \in \mathbb{Z}^{d-1}$ then we define the l^{th} k -extension of Y as the set $Y_{k,l}^d \in \mathbb{Z}^d$ that is obtained by adding a (new) k -th coordinate with value l :

$$Y_{k,l}^d = \left\{ (x_1, \dots, x_d) \in \mathbb{Z}^d : x_k = l, (x_1, \dots, x_{k-1}, x_{k+1}, \dots, x_d) \in Y \subset \mathbb{Z}^{d-1} \right\}$$

Observation 1.7. *The size of the neighborhood of some k -level, $n(L_{k,l}(X))$, equals the size of the neighborhood of $L_{k,l}^{d-1}(X)$ (in $d - 1$ dimensions) plus two times the size of X .*

Observation 1.8. *Let $X \subset \mathbb{Z}^d$ and let $L_{k,l}(X)$ be some level (in any coordinate-direction) with nonempty interior $I = \text{int } L_{k,l}^{d-1}(X)$. Then adding (any subset of) the k -extensions $I_{k,l-1}^d$ and $I_{k,l+1}^d$ to X does not increase the size of the neighborhood $n(X)$.*

By adding we here mean taking the union of the point sets.

Observation 1.9. *If $X \subset \mathbb{Z}^d$ contains some level $L_{k,l}(X)$ such that the following two statements hold:*

1. $L_{k,j+1}^{d-1}(X) \subseteq \text{int } L_{k,j}^{d-1}(X)$ for all $j \geq l$;
2. $L_{k,j-1}^{d-1}(X) \subseteq \text{int } L_{k,j}^{d-1}(X)$ for all $j \leq l$,

then the size of the neighborhood of X is $n(X) = n(L_{k,l}(X))$.




Figure 2: The highlighted level, $L_{2,1}(X)$, satisfies Observation 1.9.

Observation 1.10. *The cardinality of any ball $|B_r^d|$ is odd.*

Proof. The result is trivial for $d = 1$. For $d > 1$ a ball consists of an odd number of balls of dimension $d - 1$. \square

2 Uniqueness for Balls

2.1 Dimension 2

Given a set $X \subset \mathbb{Z}^2$, assume w.l.o.g. that the origin $(0,0)$ is part of X , and consider the following four *diagonal tangents* of X :

$$\begin{aligned} t_{++} : \quad & +x_1 + x_2 = c_{++} \geq 0 \\ t_{+-} : \quad & +x_1 - x_2 = c_{+-} \geq 0 \\ t_{-+} : \quad & -x_1 + x_2 = c_{-+} \geq 0 \\ t_{--} : \quad & -x_1 - x_2 = c_{--} \geq 0 \end{aligned}$$

where $c_{\pm,\pm}$ are chosen so that for each equality there is some point in X that satisfies it and every point in X satisfies the inequalities obtained by replacing $=$ with \leq . The *diagonal hull* $DH(X)$ of X is the set of all integer points in the region bounded by these tangents. Further, X is called *diagonal convex* if $X = DH(X)$.




Figure 3: A set with its diagonals (left) and the diagonal hull of the set (right).

If a set $X \subset \mathbb{Z}^2$ is diagonal-convex, then each point on the boundary lies on at most two of the diagonal tangents, and X resembles a parallelogram. However, if two diagonal tangents intersect in a non-integer point, then there is a (connected) portion of ∂X of cardinality 2 that is parallel to some coordinate direction. We will refer to such parts as *axis-aligned* components (of ∂X). Accordingly, for each of the diagonal tangents t we call $\partial X \cap t$ a *diagonal* (of ∂X).

Note that ∂X might have none, two, or four axis-aligned components. For example, the diagonal hull of the set X in Figure 3 has two axis-aligned components, both of which are parallel to the horizontal axis.

Proposition 2.1. *Let $X \subset \mathbb{Z}^2$, then $n(DH(X)) \leq n(X)$.*

Proof. $DH(X)$ is obtained from X by repeatedly applying the operation from Observation 1.8 until there are no points in any direction that can be added. \square

Theorem 2.2. *Balls B_r^2 are unique optimal sets of their cardinality.*

Proof. Assume that $X \subset \mathbb{Z}^2$ is a set with $|X| = |B_r^2|$ and $n(X) = n(B_r^2)$ for some $r \in \mathbb{Z}_+$. Then X has to be diagonal convex. Otherwise $|DH(X)| > |B_r^2|$ and $n(DH(X)) \leq n(X) = n(B_r^2)$, which contradicts Proposition 1.6.

There are four cases for the possible number and relative positions of the axis-aligned components: ∂X can have four, two parallel (opposite), none, or two orthogonal axis-aligned components of size 2 (see Figure 4).

Our main strategy will be to transfer parts of ∂X to some part of $N(X)$ without increasing the size of the neighborhood. For these new sets it is then easy to see that they cannot simultaneously be optimal and have the cardinality of a ball.




Figure 4: The four basic shapes of diagonal convex sets in \mathbb{Z}^2 .

Case 1. ∂X has four axis-aligned components.

Consider two opposite diagonals. They both have the same length k , and are adjacent to $k + 1$ points of the neighborhood. Remove all points from one of these diagonals, and add them along the other diagonal (from bottom to top, see Figure 5). This results in a set X' with $n(X') = n(X)$ that is not diagonal convex (as it contains

an axis-aligned component of size 3), which is a contradiction to the assumption that X is optimal and has the cardinality of a ball.




Figure 5: Case 1: Arranging all points from a diagonal along the opposite diagonal gives the same neighborhood-size and a non-diagonal of size 3.

Case 2. ∂X has two parallel non-diagonals.

Consider the levels in the coordinate direction k in which the axis-aligned components both constitute a level. Then every k -level has even cardinality, and thus $|X|$ is even. By Observation 1.10, this is a contradiction to X having the cardinality of a ball.




Figure 6: Case 2: Even parity.

Case 3. ∂X has no axis-aligned components.

Let k and l the lengths of the diagonals (each pair of opposite diagonals has the same length).

Case 3.1 If $k = l$, then X is a ball.




Figure 7: Case 3: $k = l$ (left), $k = l + 1$ (middle), and $k > l + 1$ (right).

Case 3.2 If $k > l$, then remove all points from one of the shorter diagonals and add them along one of the (at least before) longer diagonals. The resulting set X' has $n(X') = n(X)$. Further, for $k = l + 1$ it is a diagonal convex set with two parallel axis-aligned components of size 2, while for $k > l + 1$, it is not diagonal convex. In either case, we get a contradiction to the assumption that X is optimal and has the cardinality of a ball.




Figure 8: Case 3: changed point sets. Turquoise points are additional ones.

Case 4. ∂X has two orthogonal axis-aligned components.

Consider the diagonal that connects the two axis-aligned components, and say it has length k . Then the opposite diagonal has length $k + 1$ and the two other diagonals both have length l .

Case 4.1. If $k = l$, then this is exactly a ball minus one diagonal. These are optimal sets, but obviously cannot have the cardinality of a ball.

Case 4.2. If $k = l - 1$, then this is exactly a ball plus one diagonal. Again,




Figure 9: Case 4: left to right: $k = l$, $k = l - 1$, $k < l - 1$, and $k > l$.

these are optimal sets, but cannot have the cardinality of a ball.

Case 4.3. If $k < l - 1$, then we remove the diagonal of length k and add it again along one of the diagonals of length l .

In the resulting point set there is at least one point missing in this new diagonal, and adding it does not increase the size of the neighborhood. By Proposition 1.6 this contradicts our assumptions.

Case 4.4. If $k > l$, then we remove a diagonal of length l and add it along a diagonal of length k . Again, there is at least one point missing in the new diagonal and adding it does not increase the size of the neighborhood.

□

2.2 The General Case: d Dimensions

In [WW77a] the authors define a transformation of one set to another of the same cardinality: The *k -normalization* $N_k(X)$ of a set $X \subset \mathbb{Z}^d$ is obtained by:

1. replacing all (nonempty) k -levels of X by $(d - 1)$ -dimensional standard spheres of the same cardinality, and
2. changing the order of the levels such that the largest one is the 0^{th} k -level and the remaining ones are arranged half of them above and half of them below the 0^{th} k -level in a way that $|L_{k,i}| \geq |L_{k,i+1}|$ for $i \geq 0$, and $|L_{k,j}| \geq |L_{k,j-1}|$ for $j \leq 0$.




Figure 10: Case 4: changed point sets. Turquoise points are additional ones.

To show that standard spheres are optimal sets, they take an arbitrary set and repeatedly apply 1-normalization and d -normalization to it. They prove that this series of transformations terminates with a standard sphere after a finite number of steps, and that it does not increase the size of the neighborhood.

We are going to use their proof to show the uniqueness of balls as optimal sets.

Theorem 2.3. *Balls B_r^d are unique optimal sets of their cardinality.*

Proof. The proof is by induction on the dimension d , the induction base $d = 2$ being Theorem 2.2 from the last section.

For the induction step consider some optimal set $X \subset \mathbb{Z}^d$ with $|X| = |B_r^d|$ for some r . If we repeatedly apply 1-normalization and d -normalization to X as in [WW77a], then X is transformed to B_r^d . Consider the set Y that occurs in this transformation process exactly before the last normalization step. We assume the last step is in direction 1 to reduce the number of variables in the following.

Then there is a one-to-one correspondence between the 1-levels of Y and the 1-levels of B_r^d , such that all cardinalities match. Note that the levels of B_r^d are $(d - 1)$ -dimensional balls. Define

$$\begin{aligned} l_{\min} &:= \min\{l \in \mathbb{Z} : L_{1,l}(Y) \neq \emptyset\}, \text{ and} \\ l_{\max} &:= \max\{l \in \mathbb{Z} : L_{1,l}(Y) \neq \emptyset\}. \end{aligned}$$

Then as a lower bound for the size of the neighborhood of Y we have

$$n(Y) \geq |L_{1,l_{\max}}(Y)| + |L_{1,l_{\min}}(Y)| + \sum_{l=l_{\min}}^{l_{\max}} n(L_{1,l}^{d-1}(Y))$$

Now as every level $L_{1,l}(Y)$ has the cardinality of a ball $B_{r_l}^{d-1}$ for some r_l , it follows by the induction hypothesis that they all really have to be these balls (as otherwise $n(L_{1,l}^{d-1}(Y)) > n(B_{r_l}^{d-1})$ and thus $n(X) \geq n(Y) > n(B_r^d)$, which would be a contradiction to the assumption that X is optimal).

Finally consider the order of the 1-levels in Y . If they are not ordered in the same way as for the according ball B_r^d , then there exist two adjacent levels $L_{1,i}(Y) = B_{r_i}^{d-1}$ and $L_{1,j}(Y) = B_{r_j}^{d-1}$ such that $r_j < r_i - 1$. But then

$$I = \text{int} \left(L_{1,i}^{d-1}(Y) \right) \setminus L_{1,j}^{d-1}(Y) \neq \emptyset$$

and thus by Observation 1.8 we can add points to Y without increasing the size of the neighborhood $n(Y)$. This is, once again, a contradiction to Proposition 1.6 since Y was assumed to be optimal and of the same cardinality as some B_r^d .

Thus $X = Y = B_r^d$ which completes the proof. \square

3 Necessary Conditions for Optimal Sets

3.1 Back to Dimension 2

In Proposition 1.4 we showed that connectedness is a necessary condition for a set $X \subset \mathbb{Z}^2$ to be optimal. As a first step towards further conditions we consider the shapes of X and $Y = X \cup N(X)$ for diagonal convex sets, and the relation between $n(X)$ and $n(Y)$.

Proposition 3.1. *If a set $X \subset \mathbb{Z}^2$ is diagonal convex, then $Y = X \cup N(X)$ is again diagonal convex and $n(Y) = n(X) + 4$.*

Proof. The lengths of the axis-aligned parts of $N(Y)$ and $N(X)$ are identical, while the lengths of the diagonal parts of $N(Y)$ are the lengths of the diagonal parts of $N(X)$ plus one. \square

This proposition tells us that diagonal convex sets have a shape that is “stable” under the operation of adding the neighborhood, and that the size of the neighborhood behaves in a nice way. But there is a far larger class of sets that behaves in essentially the same way:




Figure 11: A diagonal convex set X , together with $N(X)$ and $N(X \cup N(X))$.

Consider some set $X \subset \mathbb{Z}^2$, such that the neighborhood $N(X)$ is a simple cycle. By this we mean that for any two points in $N(X)$ there are exactly two disjoint paths between them (where paths are defined in the same way as previously for the definition of connectedness).

Note that if we regard the cycle $N(X)$ as a (not necessarily convex) polygon, the grid points in the interior of the polygon are all points of X . Observe that the inside and outside angles at any $v \in N(X)$ are at least 90 degrees, and at least 135 degrees if one of the polygon edges adjacent to v is in a coordinate direction. An interior angle smaller than this would give $d(v, X) > 1$, and an exterior angle would lead to a subcycle of length 3 in $N(X)$.

We will proceed through the cycle $N(X)$ in the counter-clockwise direction and consider the occurring direction changes with respect to the oriented coordinate directions (e_i, σ) with $i \in \{1, 2\}$ and $\sigma \in \{+, -\}$:

- Choose as starting (and ending) point the topmost point of the tangent $t_{++} : x_1 + x_2 = c > 0$.
- Choose as starting (and ending) direction $(e_1, -)$.
- When proceeding through the cycle, remember the current coordinate direction (e_i, σ) , and the number of occurred direction changes. In every step x_k to x_{k+1}
 1. keep the direction (e_i, σ) if $\sigma = \text{sign}(x_{k+1} - x_k)_i$,
 2. otherwise change the direction to (e_j, σ') with $j \neq i$ and $\sigma' = \text{sign}(x_{k+1} - x_k)_j$.

- as the starting point is reached with a direction different from $(e_1, -)$, the turn to this direction (according to the turning rule above) counts as an occurring turn.

Every considered (and counted) direction change is a (left or right) turn by 90 degrees. As we require the starting direction to be identical to the ending direction, we turn 360 degrees in total. Thus we count an even number $k \geq 4$ of direction changes (each of the oriented directions at least once).

Also, by the above observations about the occurring angles, we notice that for every turn there is a diagonal part of $N(X)$ of size at least 2, that could be seen as being in either of the two directions before and after the turn. We will call such a diagonal part of $N(X)$ a *connecting diagonal*.

We call a set X *close-to-convex* if $N(X)$ is a cycle and if in the process described above there are four changes in direction (i.e. every oriented coordinate direction appears exactly once).




Figure 12: Proceeding through the neighborhood of a close-to-convex set. Direction changes occur at the black arrows.

Observation 3.2. *If a set $X \subset \mathbb{Z}^2$ is close-to-convex, then $Y = X \cup N(X)$ is again close-to-convex, and $n(Y) = n(X) + 4$. Moreover, Y has the same shape as X despite for the four connecting diagonals which each get longer by one. See Figure 13 for an example.*

Note that each connecting diagonal is identical to the intersection of $X \cup N(X)$ with one of its diagonal tangents.

Observation 3.3. *The standard spheres that are presented in [WW77a] are optimal sets that are close-to-convex, and for any standard sphere S , $S \cup N(S)$*




Figure 13: A close-to-convex set and two layers of neighborhood.

is again a standard sphere (and thus optimal). This implies that for any optimal set $X \subset \mathbb{Z}^2$ the inequality $n(Y) \geq n(X) + 4$ holds for $Y = X \cup N(X)$.

Let us go back to general connected sets $X \subset \mathbb{Z}^2$. We denote the cycle $C(X) \subseteq N(X)$ for which X lies in the interior of the polygon defined by this cycle as the *cycle that surrounds* X . Further we call the (finite) set $\text{cl}(X)$ of grid points enclosed by $C(X)$ the *closure* of X .

An ordered subset $\{x_1, x_2, \dots, x_n\} \subset \mathbb{Z}^2$ is a *lattice path* if the elements are distinct and $d(x_i, x_{i+1}) = 1$ for all i . For emphasis, we note that every lattice path is a path, as defined in section 1, but not conversely.

A *hole* in X is a subset $H \subset \text{cl}(X) \setminus X$ such that

1. H is connected, and
2. for every $h \in H$, every lattice path from h to $C(X)$ contains some element of X .

Proposition 3.4. *If $X \subset \mathbb{Z}^2$ is connected, then for $Y = X \cup N(X)$ we have $n(Y) \leq n(X) + 4$.*

Proof. For any close-to-convex set this is obviously true.

Now if X is not close-to-convex, then either its neighborhood $N(X)$ does not form a cycle or we will get more than four turns when proceeding through the cycle like described above.

In the latter case assume we have counted $2k + 4$ turns. Then we have counted exactly $k + 4$ left turns and k right turns. For every left turn the

connecting diagonal gets longer by (at most) one, while for each right turn the connecting diagonal gets shorter by (at least) one (see Figure 14). All parts in between are just translated by 1 and are thus (at most) as long as they were.




Figure 14: All possible right turns (up to symmetry).

Note that $N(Y)$ need not be a cycle, and that the *at most* and *at least* statements from above stem from the fact that points of $N(Y)$ might be created duplicately from more than one part of $N(X)$, see Figures 15 and 16 for examples.




Figure 15: $N(Y)$ does not form a cycle: creating a hole.




Figure 16: $N(Y)$ does not form a cycle: Right turns with short connecting diagonal.

Now what is left is the situation when $N(X)$ is not a cycle. Here, consider the cycle $C(X) \subset N(X)$ that surrounds X .

For the part of the neighborhood of $N(Y)$ that lies outside of $C(X)$ the same arguments as above can be applied.

For the inside part observe that every component of $N(Y)$ corresponds to a hole of $X \cup N(X)$, see Figure 17, and thus the size of the neighborhood inside $C(X)$ is decreasing.

□




Figure 17: A connected set X visualizing some possible shapes. The lines indicate all paths in $N(X)$.

Proposition 3.5. *If $X \subset \mathbb{Z}^2$ is not close-to-convex, then X cannot be optimal.*

Proof. Assume that X is connected but not close-to-convex, and consider again the cycle $C(X) \subset N(X)$ that surrounds X .

If $C(X) \subsetneq N(X)$, then the set Y containing all points inside this cycle $C(X)$ has $|Y| > |X|$ and $n(Y) < n(X)$. Thus X cannot be optimal.

If $C(X) = N(X)$ then consider $Z = X \cup N(X)$. From the proof of Proposition 3.4 we know that if $N(Z)$ is a cycle again, then the length of every right turn connecting diagonal is shorter than the according one in $N(X)$. Thus, after finitely many steps of adding the neighborhood we obtain a set that is not optimal (as its neighborhood is not a cycle). But as Proposition 3.4 holds for every step of adding the neighborhood, X cannot be optimal either. □

Next we show that adding the neighborhood to a set does not carry us away from optimality. More precisely, we have:

Proposition 3.6. *Consider any connected set $X \subset \mathbb{Z}^2$ and its union with its neighborhood $X' = X \cup N(X)$. Let Y and Y' be the standard spheres of*

cardinality $|Y| = |X|$ and $|Y'| = |X'|$, respectively. If $n(X) = n(Y) + k$ for some $k \geq 0$, then $n(X') \leq n(Y') + k$.

Proof. $|Y'| = |X'| = |X| + n(X) = |Y| + n(Y) + k \geq |Y \cup n(Y)|$. Thus from Observation 1.5, it follows that

$$n(Y') \geq n(Y \cup n(Y)) = n(Y) + 4 = n(X) - k + 4 = n(X') - k.$$

□

Now we can state the answer to Problem 1.1, i.e., we have shown that, at least in dimension 2, the set of optimal sets is closed under the operation of adding the neighborhood. We record this solution in the following corollary.

Corollary 3.7. *If a set $X \in \mathbb{Z}^2$ is optimal, then $n(X \cup N(X)) = n(X) + 4$ and $X \cup N(X)$ is also optimal.*

References

- [Bez89] S L Bezrukov, *On the construction of solutions of a discrete isoperimetric problem in hamming space*, Mathematics of the USSR-Sbornik **63** (1989), no. 1, 81–96.
- [Bez94] S.L. Bezrukov, *Isoperimetric problems in discrete spaces*, Extremal Problems for Finite Sets **3** (1994), 59–91.
- [BS02] Sergei L. Bezrukov and Oriol Serra, *A local-global principle for vertex-isoperimetric problems*, Discrete Math. **257** (2002), no. 2-3, 285–309.
- [CL69] G.F. Clements and B. Lindström, *A generalization of a combinatorial theorem of Macaulay*, J. Combinatorial Theory Ser. A **7** (1969), 230–238.
- [Har66] L.H. Harper, *Optimal numberings and isoperimetric problems on graphs*, J. Combin. Theory **1** (1966), 385–393.
- [Mac27] F. S. Macaulay, *Some Properties of Enumeration in the Theory of Modular Systems*, Proc. London Math. Soc. **s2-26** (1927), no. 1, 531–555.
- [WW77a] Da-Lun Wang and Ping Wang, *Discrete isoperimetric problems*, SIAM Journal on Applied Mathematics **32** (1977), no. 4, 860–870.

- [WW77b] ———, *Extremal configurations on a discrete torus and a generalization of the generalized macaulay theorem*, SIAM Journal on Applied Mathematics **33** (1977), no. 1, 55–59.

