

Decomposition of Geometric Set Systems and Graphs

Dissertation

Dömötör Pálvölgyi

Jury:

Prof. J. Pach, thesis director
Prof. F. Eisenbrand, president
Prof. A. Shokrollahi, rapporteur
Prof. G. Rote, rapporteur
Prof. P. Valtr, rapporteur



Abstract

We study two decomposition problems in combinatorial geometry. The first part of the thesis deals with the decomposition of multiple coverings of the plane. We say that a planar set is *cover-decomposable* if there is a constant m such that any m -fold covering of the plane with its translates is decomposable into two disjoint coverings of the whole plane. Pach conjectured that every convex set is cover-decomposable. We verify his conjecture for polygons. Moreover, if m is large enough, depending on k and the polygon, we prove that any m -fold covering can even be decomposed into k coverings. Then we show that the situation is exactly the opposite in three dimensions, for any polyhedron and any m we construct an m -fold covering of the space that is not decomposable. We also give constructions that show that concave polygons are usually not cover-decomposable. We start the first part with a detailed survey of all results on the cover-decomposability of polygons.

The second part of the thesis investigates another geometric partition problem, related to planar representation of graphs. Wade and Chu defined the *slope number* of a graph G as the smallest number s with the property that G has a straight-line drawing with edges of at most s distinct slopes and with no bends. We examine the slope number of bounded degree graphs. Our main results are that if the maximum degree is at least 5, then the slope number tends to infinity as the number of vertices grows but every graph with maximum degree at most 3 can be embedded with only five slopes. We also prove that such an embedding exists for the related notion called *slope parameter*. Finally, we study the *planar slope number*, defined only for planar graphs as the smallest number s with the property that the graph has a straight-line drawing in the plane without any crossings such that the edges are segments of only s distinct slopes. We show that the planar slope number of planar graphs with bounded degree is bounded.

Keywords. Multiple coverings, Decomposability, Sensor networks, Hypergraph coloring, Graph drawing, Slope number, Planar graphs.

Résumé

Nous étudions deux problèmes de décomposition de la géométrie combinatoire. La première partie de cette thèse s'intéresse à la décomposition des recouvrements multiples du plan. On dit qu'un ensemble planaire est *recouvrement-décomposable** s'il existe une constante m de telle sorte que tous les m -fois recouvrements du plan avec ses translatées sont décomposables en deux recouvrements disjoints du plan tout entier. Pach a conjecturé que tout ensemble convexe est recouvrement-décomposable. Nous vérifions sa conjecture pour les polygones. De plus, si m est assez grand, en fonction de k et du polygone, nous montrons que tous les m -fois recouvrements peuvent être décomposés même en k recouvrements. Ensuite, nous montrons qu'en trois dimensions la situation est exactement l'inverse: pour n'importe quel polyèdre et pour tout m , nous construisons une m -fois recouvrement de l'espace qui n'est pas décomposable. Nous donnons également des constructions qui montrent que les polygones concaves ne sont généralement pas recouvrement-décomposables. Nous commençons la première partie avec une étude détaillée de tous les résultats sur la recouvrement-décomposabilité de polygones.

La deuxième partie de la thèse étudie un autre problème de partition géométrique, lié à la représentation planaire des graphes. Wade et Chu ont défini le *nombre de pente*[†] d'un graphe G comme le plus petit nombre s avec la propriété que G peut être dessiné avec des segments ayant au plus s pentes distinctes. Nous examinons le nombre de pente des graphes de degré borné. Nos principaux résultats sont que, si le degré maximum du graphe est d'au moins 5, alors le nombre de pente tend vers l'infini quand le nombre de sommets croît, mais tout graphe de degré au plus 3 peut être plongé dans le plan avec seulement cinq pentes. Nous montrons aussi qu'un tel plongement existe pour la notion appelé *paramètre de pente*[‡]. Enfin, nous étudions le *nombre de pente planaire*[§], défini seulement pour les graphes planaires, comme le plus petit nombre s avec la propriété que le graphe admet un dessin linéaire dans le plan sans intersections et tel que les segments sont de seulement s pentes distinctes. Nous montrons que le nombre de pente planaire des graphes planaires de degré borné est borné.

Mots-clés. Recouvrements multiples, Décomposabilité, Réseaux de sensors, Coloration des hypergraphes, Dessinage des graphes, Nombre de pente, Graphes planaires.

*cover-decomposable

†slope number

‡slope parameter

§planar slope number

“I’ll keep it short and sweet. Family, religion, friendship. These are the three demons you must slay if you wish to succeed in business.”

C. Montgomery Burns

I would like to thank all those who in this way or that helped in the creation of the thesis. First of all, of course, my co-authors, Balázs, Géza and Jani, without whom the thesis would consist of at most one chapter (my CV). It would be impossible to enumerate how much the three of them helped, so I have to refrain to one sentence each. Balázs, thank you for buying beer for those unforgettable research sessions at your home, and for writing the first sketch of our proofs that I could rewrite into a paper. Géza, thank you for spending so much of your time with working with me, and for rewriting my always incomprehensible arguments. János, thank you for beguiling me into the unabandonable fields of discrete geometry, for inviting my future-in-the-past wife and for rewriting Géza’s version.

I am also indebted to my supervisor at ELTE, Zoltán Király. Most papers in this thesis were started when I was your student, and although together we usually worked on other topics, you always encouraged me to also pursue this direction. I could not count how many of my conjectures (some of which were related to cover-decomposition) you have disproved. I would also thank all the others from Budapest who helped my combinatorial career, including but not restricted to András, Balázs, Cory, Dani, Gábor, Gyula, Máté, Nathan and Zoli.

I would also like to thank Thomas for letting me copy-paste the tex files of his thesis and for showing me the very important procedure of how to attach the pages together. I am sorry that I replaced your text between with mine. I would like to thank all other members and visitors of the DCG and DISOPT groups for creating such a wonderful extraworkular environment. Thank you Adrian, Andreas, Andrew, Dave, Filip, Fritz, Genna, Jarek, Jocelyne, Judit, Laura, Martin, Nicolai, Paul, Rado, Rom, Saurabh and Scruffy. I wish you all the best and many more happy days in Lausanne!

From the third floor, I would also like to thank Dani for the time we spent together, and for simultaneously translating my abstract to french and correcting the mistakes in the english version. I would also thank here Bernadett, Panos and the anonymous referees of my papers for their valuable comments.

I would like to thank all my jury members, especially Günter and Pavel, for their careful reading and useful comments, most of which I implemented in this corrected version.

Finally, I would like to thank my family for their support. Thank you anya, apa, Déni*, Eszter, Bence. But most importantly, I would like to thank my wife, Padmini, who fits to several above mentioned categories. Thank you for leaving me a few seconds each day to work on my research, for reading and spotting so many grammatical mistakes and for all the time we spent together, but maybe the beginning of my thesis is not the best place to discuss this in more detail.

And now for something completely different...

*Since I know that you will read this before you leave for the bike tour, I wanted to remind you to buy some chocolate for my birthday. You know, just because we did not meet it does not mean you can skip my present. Damn, on the other hand, it is kinda hot, so it would melt. Or will you have a cooling box or something like that with you? You know what, just forget about it and get me something after you are back.

Contents

1	Introduction and Organization	1
I	Decomposition of Multiple Coverings	3
2	Introduction and Survey	3
2.1	Basic Tricks	4
2.1.1	Dualization method	4
2.1.2	Divide et impera – Reduction to wedges	5
2.1.3	Totalitarianism	5
2.2	Boundary Methods	6
2.2.1	Decomposition to $\Omega(\sqrt{m})$ parts for symmetric polygons	9
2.2.2	Triangles	9
2.3	Path Decomposition and Level Curves	11
2.3.1	Classification of wedges	11
2.3.2	Level curves and decomposition to $\Omega(k)$ parts for symmetric polygons	14
2.3.3	Decomposition to $\Omega(k)$ parts for all polygons	16
2.4	Indecomposable Constructions	19
2.4.1	Concave quadrilaterals	19
2.4.2	General concave polygons and polyhedra	21
3	Decomposition of Coverings by Convex Polygons	22
3.1	Preparation	22
3.2	NC wedges - Proof of Lemma 2.15 and 2.16	22
3.3	Summary of Proof of Theorem A	29
3.4	Concluding Remarks	30
4	Indecomposable Coverings by Concave Polygons	32
4.1	The Construction - Proof of Theorems E and E'	33
4.2	Versions of Cover-decomposability	35
4.2.1	Number of sets: Finite, infinite or more	36
4.2.2	Covering the whole plane	37
4.3	Higher Dimensions - Proof of Theorem B'	39
4.4	Concluding Remarks	41
II	Slope Number of Graphs	43
5	Introduction and a Lower Bound	43
5.1	A Lower Bound for the Slope Number of Graphs with Bounded Degree	45
6	Drawing Cubic Graphs with at most Five Slopes	49
6.1	Embedding Cycles	49
6.2	Subcubic Graphs - Proof of Theorem 6.2	51

6.3	Cubic Graphs - Proof of Theorem 6.1	56
6.4	Algorithm	58
7	Slope Parameter of Cubic Graphs	59
7.1	Subcubic Graphs - Proof of Theorem 7.2	59
7.2	Cubic Graphs - Proof of Theorem 7.1	64
7.3	Concluding Remarks	68
8	Drawing Planar Graphs with Few Slopes	69
8.1	Straight-line Drawings - Proof of Theorem 8.1	70
8.2	One Bend per Edge - Proof of Theorem 8.2	72
8.3	Two Bends per Edge - Proof of Theorem 8.3	74
8.4	Lower Bounds	79
III	Conjectures, Bibliography and CV	81
9	Summary of Interesting Open Questions and Conjectures	81
9.1	Questions about Decomposition of Coverings	81
9.2	Questions about Slope Number of Graphs	83
10	Bibliography	85
11	CURRICULUM VITÆ	89

1 Introduction and Organization

Partitions are one of the best studied and most important notions of combinatorial mathematics. The number theoretic partition function $p(n)$, which represents the number of possible partitions of a natural number n , was already studied by Euler. Later many mysterious identities about it were proved by Ramanujan and many more are still studied today, including properties of Young-tableaux. In combinatorics, partitions are often called colorings, which is just another more visual way to imagine the decomposition of a set. Coloring the vertices or edges of a graph is probably the problem that fascinated mathematicians more than any other graph theoretical question, from the four-color conjecture to Ramsey theory. The investigation of Property B* was popularized by Erdős. A set system is said to have Property B if the elements of its ground set can be colored with two colors such that no set is monochromatic, i.e. each set contains both colors. This is strongly connected to the following geometric problem.

Suppose we have a finite number of sensors in a planar region R , each monitoring some part of R , called the range of the sensor. Each sensor has a duration for which it can be active and once it is turned on, it has to remain active until this duration is over, after which it will stay inactive. A schedule for the sensors is a starting time for each sensor that determines when it starts to be active. The goal is to find a schedule to monitor R for as long as we can. For any instance of this problem, define a set system \mathcal{F} as follows. The sensors will be the elements of the ground set of \mathcal{F} and the points of R will be the sets. An element is contained in a set if the respective sensor monitors the respective point. In the special case when the duration of each sensor is 1 unit of time, we can monitor R for 2 units of time if and only if \mathcal{F} has Property B.

Pach posed the following related problem. Suppose that every point of the plane is covered by many translates of the same planar set. Is it always possible to decompose this covering into two coverings? Thus our goal is to partition/color the covering sets such that every point will be contained in both parts/color classes. This question is again equivalent to asking whether certain families have Property B. Pach conjectured that for every convex set there is a constant m such that any m -fold covering is decomposable into two coverings. Such sets are called *cover-decomposable*. The first part of this thesis is centered around this conjecture in the case when the underlying set is a polygon. We show that convex polygons are cover-decomposable. Moreover, if m is large enough, depending on k and the polygon, we prove that any m -fold covering can even be decomposed into k coverings. Then we show that the situation is exactly the opposite in three dimensions. For any polyhedron and any m , we construct an m -fold covering of the space that is not decomposable. We also give constructions that show that concave polygons[†] are not usually cover-decomposable. We start the first part with a detailed survey of all results on the cover-decomposability of polygons.

*Named after Felix Bernstein who first studied this property.

†A polygon is *concave* if it is not convex.

In the second part we investigate another geometric decomposition problem related to planar representation of graphs. Partitioning the edges of a graph to obtain nice drawings or to show that the graph is complex in some sense was studied under various constraints. The *thickness* of a graph G is defined as the smallest number of planar subgraphs it can be decomposed into. It is one of the several widely known graph parameters that measures how far G is from being planar. The *geometric thickness* of G , defined as the smallest number of *crossing-free* subgraphs of a straight-line drawing of G whose union is G , is another similar notion. In this thesis we investigate a related parameter introduced by Wade and Chu. The *slope number* of a graph G is the smallest number s with the property that G has a straight-line drawing with edges of at most s distinct slopes and with no bends. It follows directly from the definitions that the thickness of any graph is at most as large as its geometric thickness, which, in turn, cannot exceed its slope number. Therefore the slope number is always an upper bound for the other two parameters. The slope number is also important for the visualization of graphs. Graphs with slope number two can be embedded in the plane using only vertical and horizontal segments. Generally, the smaller the slope number is, the simpler the visualization becomes.

The second part examines the slope number of bounded degree graphs. Our main results are that if the maximum degree is at least 5, then the slope number tends to infinity as the number of vertices grow, but every graph with maximum degree at most 3 can be embedded with only five slopes. The degree 4 case remains a challenging open problem. We also prove that such an embedding exists for the related notion called slope parameter, which is defined in the second part. Finally, we study the planar slope number of bounded degree graphs. This parameter is only defined for planar graphs. It is the smallest number s with the property that the graph has a straight-line drawing in the plane without any crossings, such that the edges are segments of only s distinct slopes. We show that the planar slope number of planar graphs with bounded degree is bounded.

In the third part we summarize the interesting open questions and conjectures about cover-decomposition and the slope number. These are followed by the bibliography and my curriculum vitæ.

Part I

Decomposition of Multiple Coverings

2 Introduction and Survey

This section mainly follows our manuscript with János Pach and Géza Tóth, Survey on the Decomposition of Multiple Coverings [PPT10].

Let $\mathcal{P} = \{ P_i \mid i \in I \}$ be a collection of planar sets. We say that \mathcal{P} is an *m-fold covering* if every point in the plane is contained in at least m members of \mathcal{P} . The biggest such k is called the thickness of the covering. A 1-fold covering is simply called a *covering*.

Definition. A planar set P is said to be *cover-decomposable* if there exists a (minimal) constant $m = m(P)$ such that every m -fold covering of the plane with translates of P can be decomposed into two coverings.

We will also refer to the problem of decomposing a covering as the *Cover Decomposition problem*. Pach [P80] proposed the problem of determining all cover-decomposable sets in 1980 and made the following conjecture.

Conjecture. (Pach) *All planar convex sets are cover-decomposable.*

This conjecture has been verified for open polygons through a series of papers.

Theorem A. (i) [P86] *Every centrally symmetric open convex polygon is cover-decomposable.*

(ii) [TT07] *Every open triangle is cover-decomposable.*

(iii) [PT10] *Every open convex polygon is cover-decomposable.*

In fact, in [PT10] a slightly stronger result is proved. In particular, it is shown that the union of finitely many copies of the same open convex polygon is also cover-decomposable. See also Section 3 for details. There are several recent negative results as well.

Theorem B. [PTT05] *Concave quadrilaterals are not cover-decomposable.*

In [PTT05] it was also shown that certain type of concave polygons are not cover-decomposable either. This has been generalized to a much larger class of concave polygons in [P10], see Section 4 for details. One can ask analogous questions in higher dimensions, and in [P10] it is shown that the situation is quite different.

Theorem B'. [P10] *Polytopes are not cover-decomposable in the space and in higher dimensions.*

For a cover-decomposable set P , one can ask for the exact value of $m(P)$. In most of the cases, the best known upper and lower bounds are very far from each other. For example, for any open triangle T we have $3 \leq m(T) \leq 19$, for the best upper bound see Ács [A10].

Definition. Let P be a planar set and $k \geq 2$ integer. If it exists, let $m_k(P)$ denote the smallest number m with the property that every m -fold covering of the plane with translates of P can be decomposed into k coverings.

We conjecture that $m_k(P)$ exists for all cover-decomposable P , but we cannot prove it in general, see also Question 9.2. In [P86] it is shown that for any centrally symmetric

convex open polygon P , $m_k(P)$ exists and $m_k(P) \leq f(k, P)$ where $f(k, P)$ is an exponential function of k for any fixed P . In [TT07] a similar result was shown for open triangles and in [PT10] for open convex polygons. However, all these results were improved to the optimal linear bound in a series of papers.

Theorem C. (i) [PT07] *For any centrally symmetric open convex polygon P , $m_k(P) = O(k^2)$.*

(ii) [A08] *For any centrally symmetric open convex polygon P , $m_k(P) = O(k)$.*

(iii) [GV10] *For any open convex polygon P , $m_k(P) = O(k)$.*

The problem of determining $m_k(P)$ can be reformulated in a slightly different way: we try to decompose an m -fold covering into as many coverings as possible. This problem is closely related to the Sensor Cover problem. Gibson and Varadarajan in [GV10] proved their result in this more general context. See Section 2.3.3 for details.

The goal of this section is to sketch the methods used in the above theorems, and to state the most important open problems. The first paper in this topic was [P86], and its methods were used by all later papers. Therefore, we first concentrate on this paper, and then we turn to the others.

2.1 Basic Tricks

Suppose that we have a k -fold covering of the plane with a family of translates of an open polygon P . By a standard compactness argument, we can select a subfamily which still forms a k -fold covering, and is locally finite. That is, each point is covered finitely many times. Therefore, we will assume without loss of generality that all coverings are locally finite.

2.1.1 Dualization method

In [P86] the results are proved in the *dual setting*. Suppose we have a collection $\mathcal{P} = \{ P_i \mid i \in I \}$ of translates of P . Let O_i be the center of gravity of P_i . The collection \mathcal{P} is a k -fold covering of the plane if and only if every translate of \bar{P} , the reflection of P through the origin, contains at least k points of the collection $\mathcal{O} = \{ O_i \mid i \in I \}$.

The collection $\mathcal{P} = \{ P_i \mid i \in I \}$ can be decomposed into two coverings if and only if the set $\mathcal{O} = \{ O_i \mid i \in I \}$ can be colored with two colors, such that every translate of \bar{P} contains a point of each of the colors. Note that the cardinality of I can be arbitrary. Using that P and \bar{P} are either both cover-decomposable, or none of them is, we have proved the following.

Lemma 2.1. *P is cover-decomposable if and only if there is an m , such that given any point set S , with the property that any translate of P contains at least m points of S , can be colored with two colors such that any translate of P contains points of both colors.*

Note that the same argument applies if we want to decompose the covering into $k > 2$ coverings. All mentioned papers use the same approach, that is, they all investigate the covering problem in the dual setting. Thus, from now on we will also investigate the problem in the dual setting.

2.1.2 Divide et impera – Reduction to wedges

This approach is also from [P86] and it is used in all papers on the topic. Two halflines, both of endpoint O , divide the plane into two parts, W_1 and W_2 , which we call *wedges*. A *closed wedge* contains its boundary, an *open wedge* does not. Point O is called the *apex* of the wedges. The *angle* of a wedge is the angle between its two boundary halflines, measured inside the wedge.

Let P be a polygon of n vertices and we have a multiple covering of the plane with translates of P . Then, the cover decomposition problem can be reduced to wedges as follows.

Divide the plane into small regions, say, squares, such that each square intersects at most two consecutive sides of any translate of P . If a translate of P contains sufficiently many points of S , then it contains many points of S in one of the squares, because every translate can only intersect a bounded number of squares. We color the points of S separately in each of the squares such that if a translate of P contains sufficiently many of them, then it contains points of both colors. If we focus on the subset S' of S in just one of the squares, then any translate of P “looks like” a wedge corresponding to one of the vertices of P . That is, if we consider W_1, \dots, W_n , the wedges corresponding to the vertices of P , then any subset of S' that can be cut off from S by a translate of P , can also be cut off by a translate of one of W_1, \dots, W_n . Note that S' is finite because of the locally finiteness of our original covering.

Lemma 2.2. *P is cover-decomposable if there is an m , such that any finite point set S can be colored with two colors such that any translate of any wedge of P that contains at least m points of S , contains points of both colors.*

Again, the same argument can be repeated in the case when we want to decompose a covering into $k > 2$ coverings. Thus, from now on, we will be interested in coloring point sets with respect to wedges when proving positive results. But in fact coloring point sets with respect to wedges can also be very useful to prove negative results as is shown by the next lemma.

2.1.3 Totalitarianism

So far our definition only concerned coverings of the *whole* plane, but we could investigate coverings of any fixed planar point set.

Definition 2.3. *A planar set P is said to be totally-cover-decomposable if there exists a (minimal) constant $m^T = m^T(P)$ such that every m^T -fold covering of ANY planar point set with translates of P can be decomposed into two coverings. Similarly, let $m_k^T(P)$ denote the smallest number m^T with the property that every m^T -fold covering of ANY planar point set with translates of P can be decomposed into k coverings.*

This notion was only defined in [P10], however, the proofs in earlier papers all work for this stronger version because of Lemma 2.2. Sometimes, when it can lead to confusion, we will call cover-decomposable sets *plane-cover-decomposable*. By definition, if a set is totally-cover-decomposable, then it is also plane-cover-decomposable. On the other hand,

there are sets (maybe even polygons) which are plane-cover-decomposable, but not totally-cover-decomposable. E.g. the disjoint union of a concave quadrilateral and a far enough halfplane is such a set. For these sets the following stronger version of Lemma 2.2 is true.

Lemma 2.4. *The open polygon P is totally-cover-decomposable if and only if there is an m^T such that any finite point set S can be colored with two colors such that any translate of any wedge of P that contains at least m^T points of S , contains points of both colors.*

Note that if we want to show that a set is not plane-cover-decomposable, then we can first show that it is not totally-cover-decomposable using this lemma for a suitable point set S and then adding more points to S and using Lemma 2.1. Of course, we have to be careful not to add any points to the translates that show that P is not totally-cover-decomposable. This is the path followed in [PTT05] and also in [P10], but there the point set S cannot always be extended. This will be discussed in detail in Section 4.

2.2 Boundary Methods

Let W be a wedge, and s be a point in the plane. A translate of W such that its apex is at s , is denoted by $W(s)$. More generally, if W is convex, then for points s_1, s_2, \dots, s_k , $W(s_1, s_2, \dots, s_k)$ denotes the *minimal* translate of W (for containment) which contains s_1, s_2, \dots, s_k .

Here we sketch the proof of Theorem A (i) from [P86], in the special case when P is an axis-parallel square. This square has an *upper-left*, *lower-left*, *upper-right*, and *lower-right* vertex. To each vertex there is a corresponding wedge, whose apex is at this vertex and whose sides contain the sides of the square incident to this vertex. Denote the corresponding wedges by W_{ul} , W_{ll} , W_{ur} , and W_{lr} , respectively. We refer to these four wedges as *P-wedges*. Let S be a finite point set. By Lemma 2.2 it is sufficient to prove the following.

Lemma 2.5. *S can be colored with two colors such that any translate of a P -wedge which contains at least five points of S , contains points of both colors.*

It will be very useful to define the *boundary* of S with respect to the wedges of P . It is a generalization of the convex hull; a point s of S is on the convex hull if there is a halfplane which contains s on its boundary, but none of the points of S in its interior.

Definition 2.6. *The boundary of S with respect to a wedge W , $Bd^W(S) = \{s \in S : W(s) \cap S = \emptyset\}$. Two W -boundary vertices, s and t are neighbors if $W(s, t) \cap S = \emptyset$.*

It is easy to see that W -boundary points have a natural order where two vertices are consecutive if and only if they are neighbors. Observe also that any translate of W intersects the W -boundary in an interval. Now the boundary of S with respect to the four P -wedges is the union of the four boundaries.

The W_{lr} - and a W_{ll} -boundary meets at the highest point of S (the point of maximum y -coordinate, which does not have to be unique, but for simplicity let us suppose it is), the W_{ll} - and a W_{ul} -boundary meets at the rightmost point, the W_{ul} - and a W_{ur} -boundary meets at the lowest point, and the W_{ur} - and a W_{lr} -boundary meets at the leftmost point. See Figure 1. For simplicity, translates of P -wedges, W_{ul} , W_{ll} , W_{ur} , W_{lr} , are denoted by W_{ul} , W_{ll} , W_{ur} , W_{lr} , respectively.




Figure 1: The boundary of a point set.

Points of S which are not boundary vertices, are called *interior* points. The main difference between the convex hull and the boundary, with respect to P , is that in the cyclic enumeration of all boundary vertices obtained by joining the natural orders on the four parts of the boundary together, a vertex could occur twice. These are called *singular* vertices, the others are called *regular* vertices. However, it can be shown that no vertex can appear three times in the cyclic enumeration, and all singular vertices have the same type: either all of them belong to W_{ul} and W_{lr} , or all of them belong to W_{ur} and W_{ll} . This also holds for any centrally symmetric convex polygon, singular boundary vertices all belong to the same two opposite boundary pieces.

The most important observation is the following.

Observation 2.7. *If a translate of a P -wedge, say, W_{ll} , contains some points of S , then it is the union of three subsets: (i) an interval of the boundary which contains at least one point from the W_{ll} -boundary, (ii) an interval of the boundary which contains at least one point from the W_{ur} -boundary, (iii) interior points. Note that (i) is non-empty, while (ii) and (iii) could be empty. Analogous statements hold for the other three wedges, and also for other symmetric polygons.*

A first naive attempt for a coloring could be to color all the boundary blue, and the interior red. Clearly, it is possible that there is a wedge that contains lots of boundary vertices and no interior vertices, so this coloring is not always good. Another naive attempt could be to color boundary vertices alternately red and blue. There is an obvious parity-problem here, and a problem with the singular vertices. But there is another, more serious problem, that a translate of a wedge could contain just one boundary vertex, and lots of interior vertices. So, we have to say something about the colors of the interior vertices but this leads to further complications. It turns out that a “mixture” of these approaches works.

Definition 2.8. *We call a boundary vertex s r -rich if there is a translate W of a P -wedge,*

such that s is the only W -boundary vertex in W but W contains at least r points of S .*

This definition is used in different proofs with a different constant r , but when it leads to no confusion, then we simply write rich instead of r -rich. In this proof rich means 5-rich, thus a boundary vertex s is rich if there is a wedge that intersects the W -boundary in s and contains at least four other points.†

Our general coloring rule will be the following.

- (1) Rich boundary vertices are blue.
- (2) There are no two red neighbors.
- (3) Color as many points red as possible, that is, let the set of red points $R \subset S$ be *maximal* under condition (1) and (2).

Note that from (3) we can deduce

- (4) Interior points are red.

A coloring that satisfies these conditions is called a *proper coloring*. There could be many such proper colorings of the same point set, and for centrally symmetric polygons, each of them is good for us. In [P86] an explicit proper coloring is given.

Now we are ready to finish the argument.

Proof of Lemma 2.5. Suppose that S is colored properly and W is a translate of a P -wedge, such that it contains at least five points of S . We can assume without loss of generality that W contains exactly five points of S . By Observation 2.7, W intersects the W -boundary of S in an interval.

First we find a blue point in W . If the above interval contains just one point then this point is rich as the wedge contains at least five points, and rich points are blue according to (1). If the interval contains at least two points, then one of them should be blue according to (2).

Now we show that there is also a red point in W . If W contains any interior point, then we are done according to (4). So we can assume by Observation 2.7 that $W \cap S$ is the union of two intervals, and all points in W are blue. Since we have five points, one of them, say, x , is not the endpoint of any of the intervals. If it is not rich, then, according to (3), it or one of its neighbors, all contained in W , is red. So, x should be rich. But then there is a translate W' of a P -wedge, which contains only x as a boundary vertex, and contains five points. Using that S is centrally symmetric, it can be shown that $S \cap W'$ is a proper subset of $S \cap W$, a contradiction, since both contain exactly five points. This concludes the proof of Lemma 2.5. \square

If we only consider wedges with more points, we can guarantee more red points in them.

Lemma 2.9. *In a proper coloring of S , any translate of a P -wedge which contains at least $5k$ points of S , contains at least one blue point and at least k red points.*

*In [P86] and [PT07] a slightly different definition is used, there s is required to be the only vertex from the whole (and not only from the W -) boundary in the translate of W . For symmetric polygons both definitions work, but, for example, for triangles only the above given definition can be used.

†Note that instead of $r = 5$ we could also pick $r = 4$ to define rich points in this proof and only the last line would require a little more attention.

The proof is very similar to the proof of Lemma 2.5, the difference is that now we color $5k$ -rich points red and we have to be a little more careful when counting red points, especially because of the possible singular points. Then, we can recolor red points recursively by Lemma 2.9, and we obtain an exponential upper bound on $m_k(P)$. Analogous statement holds for any centrally symmetric open convex polygon, therefore, we have

Theorem 2.10. *For any symmetric open convex polygon P , there is a c_P such that any c_P^k -fold covering of the plane with translates of P can be decomposed into k coverings.*

2.2.1 Decomposition to $\Omega(\sqrt{m})$ parts for symmetric polygons

Here we sketch the proof of Theorem C (i), following the proof of [PT07], which is a modification of the previous proof. We still assume for simplicity that P is an axis parallel square. The basic idea is the same as in the previous proof. Let $k \geq 2$. We will color S with k colors such that any P -wedge that contains at least $m = 18k^2$ points contains all k colors. We define k boundary layers and denote them by B_1, B_2, \dots, B_k , respectively. That is, denote the boundary of S by B_1 and let $S_2 = S \setminus B_1$. Similarly, for any $i < k$, once we have S_i , let B_i be the boundary of S_i and let $S_{i+1} = S \setminus B_i$. Boundary layer B_i will be “responsible” for color i . Color i takes the role of blue from the previous proof, while red points are distributed “uniformly” among the other $k - 1$ colors.

Slightly more precisely, a vertex $v \in B_i$ is rich if there is a translate of a P -wedge that intersects S_i in at least $18k^2 - 18ki$ points, and v is the only boundary vertex in it. We color rich vertices of B_i with color i , and color first the remaining singular, then the regular points periodically: $1, i, 2, i, \dots, k, i, 1, \dots$. The main observation is that if a P -wedge intersects B_i (for any i) in at least $18k$ points, then it contains a long interval which contains a point of each color.* Otherwise, it has to intersect each of the boundary layers, but then for each i , its intersection with B_i contains a rich point of color i .

2.2.2 Triangles

The main difficulty with non-symmetric polygons is that Observation 2.7 does not hold here; the intersection with a translate of a P -wedge is not the union of two boundary intervals and some interior points. In the case of triangles Tardos and Tóth [TT07] managed to overcome this difficulty, with a particular version of a proper coloring, thus proving Theorem A (ii), we sketch their proof in this section. For other polygons a different approach was necessary, we will see it later.

Suppose that P is a triangle with vertices A, B, C . There are three P -wedges, W_A, W_B , and W_C . We define the boundary just like before, it has three parts, the A -, B -, and C -boundary, each of them is an interval in the cyclic enumeration of the boundary vertices. Here comes the first difficulty, there could be a singular boundary vertex which appears three times in the cyclic enumeration of boundary vertices, once in each boundary. It is easy to see that there is at most one such vertex, and we can get rid of it by decomposing S into at most four subsets, such that in each of them singular boundary points all belong to the same two boundaries, just like in the case of centrally symmetric polygons. For simplicity of the description, assume that S has only regular boundary vertices.

*This 18 could be improved with a more careful analysis.




Figure 2: On the left, x is singular, on the right, there are only regular boundary vertices.

Again, we call a boundary vertex s rich if there is a translate W of a P -wedge, such that s is the only W -boundary vertex in W but W contains at least five points of S .

Our coloring will still be a proper coloring, that is

- (1) Rich boundary vertices are blue.
- (2) There are no two red neighbors.
- (3) Color as many points red as possible, that is, let the set of red points $R \subset S$ be *maximal* under condition (1) and (2).
- (4) Interior points are red.

But in this case, we will describe explicitly, how to obtain the set of red points. The coloring will be a kind of greedy algorithm. Consider the linear order on the lines of the plane that are parallel to the side BC , so that the line through A defined smaller than the line BC . We define the partial order $<_A$ on the points with $x <_A y$ if the line through x is smaller than the line through y . We have $A <_A B$ and $A <_A C$. Similarly define the partial order $<_B$ according to the lines parallel to AC with $B <_B C$ and $B <_B A$, and the partial order $<_C$ according to the lines parallel to AB with $C <_C A$ and $C <_C B$.

First, color all rich boundary vertices blue. Now take the A -boundary vertices of S and consider them in *increasing* order according to $<_A$. If we get to a point that is not colored, we color it red and we color every neighbor of it blue. These neighbors may have already been colored blue (because they are rich, or because of an earlier red neighbor) but they are not colored red since any neighbor of any red point is immediately colored blue. Continue, until all of the A -boundary is colored. Color the B - and C -boundaries similarly, using the other two partial orders.

Suppose that W is a translate of a P -wedge, such that it contains at least five points of S . We can assume without loss of generality that W contains exactly five points of S . Assume that W is a translate of W_A . The other two cases are exactly the same. To find a blue point, we proceed just like in the previous section, and it works for any proper coloring. We know that W intersects the A -boundary of S in an interval. If this interval contains just one point, then it is rich, so it is blue. If the interval contains at least two points, then one of them should be blue.

Now we show that there is also a red point in W . If W contains any interior point, then we are done. Therefore, we assume that all five points in W are boundary vertices. Since there are five points in W , one of them, say, x , is *not* (i) the first or last A -boundary vertex in W , and (ii) not the $<_A$ -minimal B -boundary point in W , and (iii) not the $<_A$ -minimal C -boundary point in W .

Suppose that x is rich. Then there is a translate W' of a P -wedge, which contains only x as a boundary vertex, and contains five points. It can be shown by some straightforward geometric observations that $S \cap W'$ is a proper subset of $S \cap W$, a contradiction, since $S \cap W$ both contain five points. So, x can not be rich. But then why would it be blue? The only reason could be that in the coloring process one of its neighbors on the boundary, y , was colored red earlier. But then again, some geometric observations show that $y \in W$, which shows that there is a red point in W . This concludes the proof.

The same idea works if we have singular boundary vertices which all belong to, say, to the A - and B -boundaries. The only difference is that we have to synchronize the coloring processes on the A - and B -boundaries, so that we get to the common vertices at the same time.

By a slightly more careful argument we obtain

Lemma 2.11. *The points of S can be colored with red and blue such that any translate of a P -wedge which contains at least $5k + 3$ of the points, contains a blue point and at least k red points.*

If we apply Lemma 2.11 recursively, we get an exponential bound on $m_k(P)$.

Lemma 2.12. *For any open triangle P , every $\frac{7 \cdot 5^k - 15}{20}$ -fold covering of the plane with translates of P can be decomposed into k coverings.*

2.3 Path Decomposition and Level Curves

In this chapter we present two generalizations of the boundary method that are used to prove the other positive theorems, Theorem A (iii), C (ii) and C (iii).

2.3.1 Classification of wedges

In order to prove Theorem A (iii), that says all open convex polygons are cover-decomposable, in [PT10] some new ideas were developed. In the previous results we colored a point set with respect to P -wedges, for some polygon P . In this paper, point sets are colored with respect to an *arbitrary* set of wedges.

Definition. Suppose that $\mathcal{W} = \{ W_i \mid i \in I \}$ is a collection of wedges. \mathcal{W} is said to be *non-conflicting* or simply NC, if there is a constant m with the following property. Any finite set of points S can be colored with two colors such that any translate of a wedge $W \in \mathcal{W}$ that contains at least m points of S , contains points of both colors.*

Note that if a collection of wedges, $\mathcal{W} = \{ W_i \mid i \in I \}$ is NC, then so is $\mathcal{W} = \{ W_i^ \mid i \in I \}$ where W_i^* is the closure or interior of the wedge W_i . This is true because if we perturbate any S such that the segment determined by any two points becomes non-parallel to any side of any of the wedges, then the collection of sets of points that can be cut off from S by a translate of a wedge from \mathcal{W} will not decrease.

It turns out that a single wedge is always NC. Then pairs of wedges which are NC are characterized. Finally, it is shown that a set of wedges is NC if and only if each pair is NC. From this characterization it follows directly that for any convex polygon P , the set of P -wedges is NC.

Lemma 2.13. *A single wedge is NC.*

A very important tool in the the proof of Lemma 2.13, and the following lemmas, is the *path decomposition* which is the generalization of the concept of the boundary. We give the proof of Lemma 2.13 to illustrate this method.

Proof. Let S be a finite point set and W a wedge. We prove the statement with $k = 3$, that is, S can be colored with two colors such that any translate of W that contains at least 3 points of S , contains a point of both colors. Suppose first that the angle of W is at least π . Then W is the union of two halfplanes, A and B . Take the translate of A (resp. B) that contains exactly two points of S , say, A_1 and A_2 (resp. B_1 and B_2). There might be coincidences between A_1, A_2 and B_1, B_2 , but still, we can color the set $\{A_1, A_2, B_1, B_2\}$ such that A_1 and A_2 (resp. B_1 and B_2) are of different colors. Now, if a translate of W contains three points, it contains either A_1 and A_2 , or B_1 and B_2 , and we are done.

Suppose now that the angle of W is less than π . We show that in this case the NC property holds with $k = 2$. We can assume that the positive x -axis is in W , this can be achieved by an appropriate rotation. For simplicity, also suppose that no direction determined by two points of S is parallel to the sides of W as with a suitable perturbation this can be achieved.

For any fixed y , let $W(2; y)$ be the translate of W which

- (1) contains at most two points of S ,
- (2) its apex has y -coordinate y , and
- (3) its apex has minimal x -coordinate.

It is easy to see that for any y , $W(2; y)$ is uniquely defined. Examine, how $W(2; y)$ changes as y runs over the real numbers. If y is very small (smaller than the y -coordinate of the points of S), then $W(2; y)$ contains two points, say X and Y , and one more, Z , on its boundary. As we increase y , the apex of $W(2; y)$ changes continuously. How can the set $\{X, Y\}$, of the two points in $W(2; y)$ change? For a certain value of y , one of them, say, X , moves to the boundary. At this point we have Y inside, and two points, X , and Z on the boundary. If we slightly further increase y , then Z replaces X , that is, Y and Z will be in $W(2; y)$ (see Figure 3). As y increases to infinity, the set $\{Z, Y\}$ could change several times, but each time it changes in the above described manner. Define a directed graph whose vertices are the points of S , and there is an edge from u to v if v replaced u during the procedure. We get two paths, P_1 and P_2 . The pair (P_1, P_2) is called the *path decomposition of S with respect to W , of order two* (see Figure 4).

Color the vertices of P_1 red, the vertices of P_2 blue. Observe that each translate of W that contains at least two points, contains at least one vertex of both P_1 and P_2 . This completes the proof. \square

We can define the *path decomposition of S with respect to W , of order k* very similarly. Let $W(k; y)$ be the translate of W which (1) contains at most k points of S , (2) its apex has y -coordinate y , and (3) its apex has minimal x -coordinate. Suppose that for y very small,




Figure 3: Z replaces X in $W(2; y)$.




Figure 4: Path decompositions of order two. $P_1 = X_1X_2\dots$, $P_2 = Y_1Y_2\dots$

$W(k; y)$ contains the points r_1, r_2, \dots, r_k , and at least one more on its boundary. Just like in the previous description, as we increase y , the set $\{r_1, r_2, \dots, r_k\}$ changes several times, such that one of its elements is *replaced* by some other vertex. Define a directed graph on the vertices of S such that there is an edge from r to s if s replaced r at some point. We get the union of k directed paths, $P_1^W, P_2^W, \dots, P_k^W$, which is called the *order k path decomposition of S with respect to W* . Note that the order 1 path decomposition is just the W -boundary of S , so this notion is a generalization of the boundary.*

Observation 2.14. (i) Any translate of W contains an interval of each of $P_1^W, P_2^W, \dots, P_k^W$, and (ii) if a translate of W contains k points of S , then it contains exactly one point of each of $P_1^W, P_2^W, \dots, P_k^W$.

Now we investigate the case when we have *two* wedges. We distinguish several cases according to the relative position of the two wedges, V and W .

Type 1 (Big): One of the wedges has angle at least π .

*But in general, $P_1(r)$ is not the boundary for higher order path decompositions! Although the union of the paths contains the boundary, the points of the boundary do not necessarily form a path.

For the other cases, we can assume without loss of generality that W contains the positive x -axis. Extend the boundary halflines of W to lines, they divide the plane into four parts, Upper, Lower, Left, and Right, which latter is W itself. See Figure 5.

Type 2 (Halfplane): One side of V is in Right and the other one is in Left. That is, the union of the wedges cover a halfplane. See Figure 6.

Type 3 (Contain): Either (i) one side of V is in Upper, the other one is in Lower, or (ii) both sides are in Right or (iii) both sides are in Left. See Figure 7.

Type 4. (Hard): One side of V is in Left and the other one is in Upper or Lower. This will be the hardest case. See Figure 8.

Type 5. (Special): Either (i) one side of V is in Right and the other one is in Upper or Lower, or (ii) both sides are in Upper, or (iii) both sides are in Lower. That is, the union of the wedges is in an open halfplane whose boundary contains the origin, but none of them contain the other. See Figure 9.

It is not hard to see that there are no other possibilities.

Lemma 2.15. *Let $\mathcal{W} = \{V, W\}$ be a set of two wedges, of Type 1, 2, 3, or 4. Then \mathcal{W} is NC.*

This lemma is proved in Section 3, for each case separately. It is also shown in [P10] that if $\mathcal{W} = \{V, W\}$ is a set of two wedges of Type 5 (Special), then \mathcal{W} is *not* NC. For the proof and its consequences see Section 4. In case of several wedges we have

Lemma 2.16. *A set of wedges $\mathcal{W} = \{W_1, W_2, \dots, W_t\}$ is NC if and only if any pair $\{W_i, W_j\}$ is NC.*

It is obvious that if two wedges are not NC then \mathcal{W} can not be NC. Therefore, a set of wedges is NC if and only if none of the pairs is of Type 5 (Special). The proof of Lemma 2.16 can again be found in Section 3. In fact, a somewhat stronger statement is true. At the end of Section 3 it is shown that if \mathcal{W} is NC, then for any k there is an m_k such that any finite point set can be colored with k colors such that if a translate of a wedge from \mathcal{W} contains at least m_k points, then it contains all k colors.

To finish the proof of Theorem A (iii), observe that two wedges corresponding to the vertices of a convex polygon cannot be of Type 1 (Big) or of Type 5 (Special). A summary of the whole proof of the theorem can be found at the end of Section 3.

2.3.2 Level curves and decomposition to $\Omega(k)$ parts for symmetric polygons

The *level curve* method was invented by Aloupis et. al. [A08] at the same time and independently from the path decomposition. Again suppose that the angle of W is less than π and W contains the positive x -axis. Now define the level curve of depth $r + 1$, $\mathcal{C}(r)$, as the collection of the apices of $W(r; y)$.^{*} Another equivalent way to define $\mathcal{C}(r)$ is as the boundary of the union of all the translates of W containing at most r points.

Note that this curve consists of straight line segments that are parallel to the sides of W . $\mathcal{C}(0)$ goes through all the boundary points, this shows that this notion is a generalization

^{*}In [A08] they denote this by $\mathcal{C}(r + 1)$ because they work with closed wedges.




Figure 5: Wedge W




Figure 6: Type 2 (Halfplane)




Figure 7: Type 3 (Contain)




Figure 8: Type 4 (Hard)




Figure 9: Type 5 (Special)

of the boundary. If $p \in \mathcal{C}(r)$, then $|W(p) \cap S|$ is either r or $r - 1$ and it is $r - 1$ only a finite number of times, when the respective translate has a point on both of its sides.

Now with the level curve method, we prove Theorem C (ii), as was done in [A08].

Suppose our symmetric polygon P has $2n$ vertices. We denote the P -wedges belonging to them in clockwise order by W_0, \dots, W_{2n-1} . All the indices should be considered in this section modulo $2n$. We call two wedges W_i and W_j *antipodal* if $i + n \equiv j$ modulo $2n$, that is, if they are the wedges belonging to two opposite vertices of the polygon. A crucial observation is (already used in [P86]) that any two P -wedges that are not antipodal, cover a half-space.

For every side of P , take two lines parallel to it that cut off $2r + 2$ points from each side of S . Denote the intersection of the n stripes formed by these lines by \mathcal{T} . Any large enough wedge has to intersect \mathcal{T} , thus it is enough to care about the wedges whose apex lies in \mathcal{T} . Now if we consider the level curves $\mathcal{C}^{W_i}(r)$, a simple geometric observation shows that only level curves belonging to antipodal wedges may cross inside \mathcal{T} and some further analysis shows that in fact there can be only one such pair (note the similarity to the singular points in case of symmetric polygons). This means that the regions cut off from \mathcal{T} by the curves $\mathcal{C}^{W_i}(r)$ are all disjoint with the possible exception of one pair. Without loss of generality, these are the curves of W_0 and W_n .

Another easy observation shows that any translate of W_i that contains at least $3r + 5$ points, must contain a point from $\mathcal{C}^{W_i}(r) \cap \mathcal{T}$, thus also a translate of W_i whose apex is on the level curve inside \mathcal{T} , containing r points from S . Therefore it is enough to care about these wedges, whose apex lies on the respective level curve. It is possible to parametrize these wedges with the circle parameterized by $[0, 2n)$ such that $W(t)$ is a translate of $W_{\lfloor t \rfloor}$. A crucial geometric observation is that if $p \in W(\lfloor t \rfloor + x) \cap W(\lfloor t \rfloor + z)$, where $0 \leq x \leq 1$ and $0 \leq z \leq n$, then $p \in W(\lfloor t \rfloor + y)$ for all $x \leq y \leq z$. If $p \in W(\lfloor t \rfloor + x) \cap W(\lfloor t \rfloor + z)$, where $0 \leq x \leq 1$ and $n \leq z \leq n + 1$, then p is contained in two antipodal wedges implying that it is contained in translates of W_0 and W_n but in no translates of any other other wedge from $W(t)$. Therefore, every p is contained either in an interval of the circle $[0, 2n)$, or in two intervals, one of which is a subinterval of $[0, 1]$, the other of $[n, n + 1]$. The simplest is if we take care of these two types separately, as any big wedge contains a lot of points from one of these groups. The first type forms a circular interval graph, if every point of the circle is covered m' -fold, then we can decompose this to $m'/3$ coverings with a simple greedy algorithm. In the second type, we want to color points with respect to a wedge and its rotation with 180 degrees. The greedy algorithm again gives a good decomposition from an m'' -fold covering into $m''/3$ coverings. Putting the numbers together this implies that $m_k \leq 18k + 5$ for any system of wedges derived from a symmetric polygon. This has to be multiplied by a constant depending on the shape of the polygon that comes from Lemma 2.2 to get a bound for the multiple-cover-decomposability function m_k of the polygon.

2.3.3 Decomposition to $\Omega(k)$ parts for all polygons

The decomposition to multiple coverings is also motivated by the following problem, called *Sensor Cover problem*.

Suppose we have a finite number of sensors in a region R , each monitoring some part of R , which is called the range of the sensor. Each sensor has a duration for which it can

be active and once it is turned on, it has to remain active until this duration is over, after which it will stay inactive. The load of a point is the sum of the durations of all ranges that contain it, and the load of the arrangement of sensors is the minimum load of the points of R . A schedule for the sensors is a starting time for each sensor that determines when it starts to be active.

The goal is to find a schedule to monitor the given area, R , for as long as we can. Clearly, the cover decomposability problem is a special case of the Sensor Cover problem, when the duration of each sensor is the same. Gibson and Varadarajan in [GV10] proved their result in this more general context.

Theorem D. [GV10] *For any open convex polygon P there is a $c(P)$ such that for any instance of the Sensor Cover problem with load $k \cdot c(P)$ where each range is a translate of P , there is a polynomial time computable schedule such that every point is monitored for k time units.*

In the special case where the duration of each sensor is 1 unit of time and R is the whole plane, this is equivalent to Theorem C (iii). As the proof is essentially the same, we will only sketch the proof for this special case to avoid changing terminology. In their proof they use the usual dualization and reduction to wedges, because of which it is enough to prove the following theorem (for the special case).

Theorem D'. [GV10] *If $\mathcal{W} = \{ W_i \mid i \in [n] \}$ is a system of P -wedges, then there is an α_n depending only on n , such that any point set S can be colored with k colors such that any translate of a wedge from \mathcal{W} that contains at least $\alpha_n k$ points, contains all k colors.*

Note that any two P -wedges are of Type 2 (Halfplane), 4 (Hard) or a special case of 3 (Contain). Their main lemma is the following easy observation.

Lemma 2.17. *For any point set $Q \subset P$, any wedge W , any k and any $L \geq 2k$, we can partially color the points of Q with k colors such that any translate of W that contains L points of P and at least $2k$ points of Q contains at most $2k$ colored points but contains all k colors. Moreover, if a point z is colored, then all points in $Q \cap W(z)$ are colored.*

Proof. In every step take a point from Q that covers a maximal, yet uncovered interval of $\mathcal{C}^W(L)^*$ until the whole curve is covered, then color these points with one color and repeat. \square

The trick is that we obtain a partial coloring using this lemma for a carefully chosen subset of P , any one of the wedges, $W_1 \in \mathcal{W}$, k and an $L = f(n)k$ (constant depending on n to be specified later) such that $\alpha_{n-1}k$ points remain uncolored in any translate of any wedge from $\mathcal{W}' = \mathcal{W} \setminus W_1$ that previously contained at least $\alpha_n k$ points. After applying this partial coloring n times, we are done.

Before we can specify Q , we need to define an order on the plane for every line that is parallel to the side of a wedge. (Thus together this gives at most $2n$ orders for general wedges, for P -wedges it would give n .) The order is very similar to the one used for triangles in [TT07]. For a wedge W and a line ℓ parallel to one of its sides, define $p <_\ell q$ if the line parallel to ℓ through q intersects $W(p)$. (So in the special case of P -wedges, p_i and p_{i+1}

*Here and later, the level curves are always with respect to P and not to Q .

are the minimal vertices of P according to $\langle_{p_i p_{i+1}}$.) For simplicity, we just refer to these orders as the $2n$ orders defined by the system \mathcal{W} .

Now we can define Q . A point $p \in P$ is in Q if there is a translate of W_1 containing exactly L points from P in which p is not among the first $\alpha_{n-1}k$ points in any of the $2n - 2$ orders defined by \mathcal{W} . Now let us apply Lemma 2.17 to this Q , W_1 , k and $L = ((2n - 1)\alpha_{n-1} + 6)k$. Note that each translate of W_1 whose apex lies on the curve $\mathcal{C}^W(2k)$ will contain at least $2k$ points of Q as $L \geq ((2n - 2)\alpha_{n-1} + 2)k$.

Claim 2.18. *If $W_i \in \mathcal{W}'$ contains $\alpha_n k$ points from P , where $\alpha_n \geq 3\alpha_{n-1} + 6$, then it contains $\alpha_{n-1}k$ uncolored points after applying the coloring of Lemma 2.17 to Q , W_1 , k and $L = ((2n - 1)\alpha_{n-1} + 6)k$.*

Proof. The proof depends on the type of W_1 and W_i . First suppose they are of Type 2 (Halfplane). Take a translate of W_i , $W_i(x)$ containing $\alpha_n k$ points of P . If it does not intersect the level curve $\mathcal{C}^{W_1}(L)$, then we did not color any of its points, we are done. If it intersects this level curve, then the intersection can be only one point, z . Moreover, $W_1(z)$ contains all the colored points contained in $W_i(x)$. Since $W_1(z)$ contains at most $2k$ colored points, we are done if $\alpha_n k \geq (\alpha_{n-1} + 2)k$.

The second case we consider is, if they are of Type 3 (Contain), such that a translate of W_i is (not necessarily properly) contained in $-W_1$ (the wedge obtained by reflecting W_1 to the origin). Take a translate of W_i , $W_i(x)$ containing $\alpha_n k$ points of P . If it does not intersect the level curve $\mathcal{C}^{W_1}(L)$, then we did not color any of its points, we are done. If there is a $z \in \mathcal{C}^{W_1}(L)$ for which $W_1(z) \cap W_i(x)$ contains at least $(\alpha_{n-1} + 2)k$ points, then we are done as only $2k$ of these can be colored. Otherwise, for any $z \in \mathcal{C}^{W_1}(L) \cap W_i(x)$ denote by a_z and b_z the points where the boundary of $W_1(z)$ and $W_i(x)$ meet. So if z is one of the two ends of the interval $\mathcal{C}^{W_1}(L) \cap W_i(x)$, we have $a_z = z$ or, respectively, $b_z = z$. We also know that for any $z \in \mathcal{C}^{W_1}(L)$, the wedge $W_1(z)$ contains at least $L - ((2n - 2)\alpha_{n-1})k$ points of Q . A continuity argument shows that there is a z' for which both $W_1(a_{z'})$ and $W_1(b_{z'})$ contain at least $(L - ((2n - 2)\alpha_{n-1})k - (\alpha_{n-1} + 2)k) / 2$ points. If this number is at least $2k$, then $W_i(x) \setminus W_1(z')$ cannot contain any colored points because of the moreover part of Lemma 2.17. This implies that $W_i(x)$ can contain only the at most $2k$ colored points of $W_1(z')$. So for this case we need the additional condition $L \geq ((2n - 1)\alpha_{n-1} + 6)k$.

Finally, notice that in the remaining cases, W_i can be cut into three parts, W_i^1 , W_i^2 and W_i^3 , such that W_i^1 and W_i^3 have a side parallel to one of the sides of W_1 and for each there is a halfplane that contains it with W_1 , while W_i^2 is contained in $-W_1$. If $W_i(x)$ contains at least $\alpha_n k$ points of P , then at least one of these three wedges must contain at least $\alpha_n k / 3$ points. If it is W_i^2 , then we are done as in the previous case if $\alpha_n k \geq 3(\alpha_{n-1} + 2)k$. If it is one of the other two wedges, then we arrive to our last case.

Suppose $W_i(x)$ contains at least $\alpha_n k / 3$ points of P , and there is a triangle such that W_i and W_1 are among its three wedges. Without loss of generality, suppose that their parallel side is the horizontal, they are contained in the ‘‘upward’’ halfplane and W_1 looks right, W_i left. Again, if $W_i(x)$ does not intersect the level curve $\mathcal{C}^{W_1}(L)$, then we did not color any of its points, we are done. If it does, then consider the colored points in $W_i(x)$ in increasing order with respect to the order defined by the non-horizontal side of W_1 . If there are at most $2k$ colored points in $W_i(x)$, then we are done. Otherwise, denote the $2k + 1^{\text{st}}$ colored point according to this order by y . Since y is colored, there is a $z \in W_i(x) \cap \mathcal{C}^{W_1}(L)$ for which

$y \in W_1(z)$. If $z \notin W_i(x)$, then $W_1(z)$ must contain all the points that are smaller than y in the order, a contradiction, as it can contain at most $2k$ colored points. If $z \in W_i(x)$, then we can use the property that $y \in Q$. A point was selected to Q from P only if there is a translate of W_1 containing exactly L points from P in which p is not among the first $\alpha_{n-1}k$ points in any of the $2n-2$ orders defined by $\mathcal{W}' = \mathcal{W} \setminus W_1$. The apex of this translate of W_1 must be on $W_i(x) \cap \mathcal{C}^{W_1}(L)$. But then $W_i(x)$ also contains the $\alpha_{n-1}k$ points smaller than y in the order defined by the non-horizontal side of W_i , which are necessarily uncolored, thus we are done. \square

Therefore we are also done with the proof of the theorem. Note that the bound that we get for α_n grows superexponentially with n because apart from $\alpha_n \geq 3\alpha_{n-1} + 6$, we must also guarantee $\alpha_n \geq L = ((2n-1)\alpha_{n-1} + 6)k$ to make sure that also the translates of W_1 contain all k colors. We would like to remark that this bound can be made exponential by introducing a more sophisticated notation and demanding a different “ α ” for each wedge in each step (so when there are j wedges left, then the “ α ” of W_i should be approximately $2^i 3^j$).

2.4 Indecomposable Constructions

In this section we survey results about coverings that cannot be decomposed into two coverings. The first such example was given in [MP86], where it was shown that the unit ball is not cover-decomposable. Thus for any k there is a covering of \mathbb{R}^3 with unit balls such that every point is covered by at least k balls, but the covering cannot be decomposed into two coverings. Later in [PTT05] several other constructions were given, all based on the geometric realization of the same hypergraph not having Property B*. It was shown by Erdős [E63] that the smallest number of sets of size k that do not have Property B is at least 2^{k-1} , so any indecomposable construction must be exponentially big. With a standard application of the Lovász Local Lemma [EL75] it can also be shown for “nice” geometric sets that if every point is covered by less than exponentially many translates, then the covering is decomposable.

We start by presenting the construction of [PTT05] using concave quadrilaterals proving Theorem B. Then we briefly preview the results of Section 4.

2.4.1 Concave quadrilaterals

We present the construction in the dual case. We suppose that the vertices of the quadrilateral, Q , are A, B, C and D in this order, the obtuse angle being at D . This implies that W_A and W_C are of Type 5 (Special), moreover, they belong to an even more special subclass: When we translate the wedges such that their apices are in the origin, then they are disjoint and there is an open halfplane that contains both of their closures (see the two right examples in Figure 9). For simplicity, let us suppose that W_A is a very thin wedge that contains a horizontal segment and W_C is a very thin wedge that contains

*We say that a hypergraph has Property B if the elements of the ground set can be colored with two colors such that any hyperedge contains both colors.

a vertical segment, the construction would work for any other two wedges that are derived from a concave quadrilateral.

First we give a finite set of points and a finite number of translates that show that Q is not totally-cover-decomposable. Then we show how this construction is extendable to give a covering of the whole plane. The construction is based on a construction using translates of the wedges W_A and W_C . We will use these wedges to realize the following k -uniform hypergraph, \mathcal{H} . The vertices of the hypergraph are sequences of length less than k consisting of the numbers from 1 through k : $V(\mathcal{H}) = [k]^{<k}$. There are two kinds of hyperedges. The first kind contains sequences of length l whose restriction to their first $l - 1$ members is the same. The second kind consists of a length $k - 1$ sequence and all its possible restrictions. So \mathcal{H} has roughly k^k vertices and edges.




Figure 10: Indecomposable covering with two special wedges of a concave quadrilateral.

The hyperedges of the first kind are realized by translates of W_A , the second kind by translates of W_C . The vertices of the hypergraph are all very close to a vertical line. Also, vertices that belong to a hyperedge of the first kind are all on a horizontal line, for each edge on a different one (see Figure 10). It is easy to see that this is indeed a geometric realization of \mathcal{H} , so the points cannot be colored with two colors such that every translate of W_A and W_C of size k contains both colors.

Now we need to extend the corresponding covering to the whole plane. Before we do, notice that it can be achieved that the centers of all translates of Q used in the construction lie on the same line. After going back from the dual to the primal, this means that we have a set of points, S , on a line, ℓ , and an indecomposable, k -fold covering of them with translates of Q . Add all translates of Q to our covering that are disjoint from S (see Figure 11). It is clear that the resulting covering remains indecomposable. Moreover, now every point not in S will be covered infinitely many times*, because Q cannot have two sides that are parallel to ℓ , so we can “go in” between any two points of S . Note that this statement is not necessarily true for arbitrary concave polygons, this is why the construction of [P10] is not always extendable this way.

*The construction can of course be modified to get a locally finite covering using a standard compactness argument.




Figure 11: Extending the original 2-fold covering of the four points by the solid quadrilaterals to a 2-fold covering of the whole plane by adding the dotted quadrilaterals.

2.4.2 General concave polygons and polyhedra

The construction for concave polygons differs from the quadrangular case because it is no longer true that any pair of Type 5 (Special) wedges have the property that they can be translated such that their apices are in the origin and they are disjoint (see the two left examples in Figure 9). Because of this a different hypergraph (also not having Property B) is realized. This construction has less points (about 4^k) and is more general as it can be realized by any pair of Type 5 (Special) wedges. The details can be found in Section 4.

However, this construction is not always extendable to give an indecomposable covering of the whole plane. Different notions of cover-decomposability and their connections are also studied in Section 4. Finally, as a corollary of the construction, it is also shown at the end of Section 4 that polyhedra (both convex and concave) are not cover-decomposable. This construction is extendable, thus we obtain that polyhedra are not space-cover-decomposable.

3 Decomposition of Coverings by Convex Polygons

This section is based on our paper with Géza Tóth, Convex polygons are cover-decomposable [PT10].

Our main result is the strongest statement, (iii), of Theorem A, which claims

Theorem A. *Every open convex polygon is cover-decomposable.*

We start by recalling some old definitions and making some new ones. Then we establish the earlier unproved Lemma 2.15 and 2.16, and finally, we summarize the proof of Theorem A.

3.1 Preparation

Now let W be a wedge, and X be a point in the plane. A translate of W such that its apex is at X , is denoted by $W(X)$. More generally, for points X_1, X_2, \dots, X_k , $W(X_1, X_2, \dots, X_k)$ denotes the *minimal* translate of W (for containment) whose closure contains X_1, X_2, \dots, X_k . The set of all translates of W is denoted by Tr^W and the set of those translates that contain exactly k points from a point set S is denoted by $Tr_k^W(S)$. The reflection of W about the origin is denoted by $-W$.

We can assume without loss of generality that the positive x -axis is in W , and that no two points from our point set, S , have the same y -coordinate. Both of these can be achieved by an appropriate rotation. We say that $X <_y Y$ if the y -coordinate of X is smaller than the y -coordinate of Y . This ordering is called the y -ordering. A subset I of S is an *interval* of S if $\forall X <_y Y <_y Z \in S : X, Z \in I \rightarrow Y \in I$.

The *boundary* of S with respect to W , $Bd^W(S) = \{X \in P : W(X) \cap S = \emptyset\}$. Note that a translate of W always intersects the boundary in an interval. For each $X \in Bd^W(S)$ the *shadow* of X is $Sh^W(X) = \{Y \in S : W(Y) \cap Bd^W(S) = X\}$. Observe that $\forall X, Y \in Bd^W(S) : Sh^W(X) \cap Sh^W(Y) = \emptyset$.

Now we give another proof using these notions for Lemma 2.13 that claims that any single wedge, W , is NC. In fact, we show that if the angle of W is less than π , then the points of S can be colored with two colors such that any wedge that has at least two points contains both colors (so the NC property holds with $k = 2$).

Proof. Color the points of the boundary alternating, according to the order $<_y$. For every boundary point X , color every point in the shadow of X to the other color than X . Color the rest of the points arbitrarily. Any translate of W that contains at least two points, contains one or two boundary points. If it contains one boundary point, then the other point is in its shadow, so they have different colors. If it contains two boundary points, then they are consecutive points according to the y -order, so they have different colors again. \square

3.2 NC wedges - Proof of Lemma 2.15 and 2.16

Now we can turn to the case when we have translates of two or more wedges at the same time. Remember that any pair of wedges belong to a Type determined by their relative position. The different types are classified in Section 2.3.1 and are depicted in Figures 6,

7, 8, 9. It is shown in [P10] that if $\mathcal{W} = \{V, W\}$ is a set of two wedges of Type 5 (Special), then \mathcal{W} is *not* NC. In a series of lemmas we show that all other pairs are NC, thus proving Lemma 2.15.

Lemma 3.1. *Let $\mathcal{W} = \{V, W\}$ be a set of two wedges, of Type 3 (Contain). Then \mathcal{W} is NC.*




Figure 12: Type 3 (Contain)

Note that this proof could be made slightly simpler with an argument similar to the one used in the proof of Lemma 2.17, here we reproduce the original proof.

Proof. We can assume that $W \supset V$ or $W \supset -V$ and W contains the positive x -axis, just like on the two right diagrams of Figure 12. Let $(P_1^W, P_2^W, \dots, P_k^W)$ be the path decomposition of S with respect to W , of order k .

Observe that any translate of V intersects any P_i^W in an interval of it. Indeed, if $X_1 <_y X_2 <_y X_3 \in P_i^W$, then $X_2 \in W(X_1, X_3) \cap -W(X_1, X_3)$, which is a subset of $V(X_1, X_3) \cap -V(X_1, X_3)$. See Figure 13.




Figure 13: $W(X_1, X_3) \cap -W(X_1, X_3) \subset V(X_1, X_3) \cap -V(X_1, X_3)$.

We show that we can color the points of S with red and blue such that any translate of W which contains at least 4 points, and any translate of V which contains at least 14 points, contains points of both colors. Consider $(P_1^W, P_2^W, P_3^W, P_4^W)$, the path decomposition of S with respect to W , of order 4. We color P_1^W and P_2^W such that every $W' \in Tr_4^W(S)$ contains a blue point of them, and every $V' \in Tr_7^V(P_1^W \cup P_2^W)$ contains points of both colors. Similarly, we color P_3^W and P_4^W such that every $W' \in Tr_4^W(S)$ contains a red point of them, and every $V' \in Tr_7^V(P_3^W \cup P_4^W)$ contains points of both colors. Finally, we color the rest of the points $R = S \setminus (P_1^W \cup P_2^W \cup P_3^W \cup P_4^W)$ such that every $V' \in Tr_2^V(R)$ contains points of both colors.

Recall that for any $W' \in Tr_4^W(S)$, $|W' \cap P_1^W| = |W' \cap P_2^W| = |W' \cap P_3^W| = |W' \cap P_4^W| = 1$. For any $X \in P_1^W$, $Y \in P_2^W$, if there is a $W' \in Tr_4^W(S)$ with $W' \cap P_1^W = \{X\}$ and $W' \cap P_2^W = \{Y\}$, then we say that X and Y are *friends*. If X (resp. Y) has only one friend Y (resp. X), then we call it a *fan* (of Y , resp. of X). If X or Y has at least one fan, then we say that it is a *star*. Those points that are neither fans, nor stars are called *regular*.

For an example, see Figure 4. On the left figure, Y_1 is a star, its fans are X_2 and X_3 , the other points are regular. On the right, Y_2 is a star, its fan is X_2 , the other points are regular.

Suppose first that all points of P_1^W and P_2^W are regular. Color every third point of P_1^W , red and the others blue. In P_2^W , color the friends of the red points blue, and color the rest of the points of P_2^W (every third) red. For any $W' \in Tr_4^W$, $W' \cap P_1^W$ and $W' \cap P_2^W$ are friends, therefore, at least one of them is blue. On the other hand, any $V' \in Tr_7^V(P_1^W \cup P_2^W)$ contains three consecutive points of P_1^W or P_2^W , and they have both colors.

Suppose now that not all the points of P_1^W and P_2^W are regular. Color all stars blue. The first and last friend of a star, in the y -ordering, is either a star or a regular vertex, the others are fans. Color the friends of each star alternately, according to the y -ordering, starting with blue, except the last two friends; color the last one blue, the previous one red. The so far uncolored regular points of P_1^W and P_2^W form pairs of intervals. We color each such pair of interval the same way as we did in the all-regular case, coloring the first point of each pair of intervals red. See Figure 14.

Clearly, if $W' \in Tr_4^W$ then it contains at least one blue point of $P_1^W \cup P_2^W$. If $V' \in Tr_7^V(P_1^W \cup P_2^W)$, then it contains four consecutive points of P_1^W or P_2^W , say, X_1, X_2, X_3, X_4 , in P_1^W . If $X <_y Y <_y Z \in P_1^W \cap V'$ and Y is a star, then V' must contain all fans of Y as well. Indeed, the fans of Y are in $W(X, Z) \setminus (W(X) \cup W(Z))$, and by our earlier observations, this is in $V(X, Z) \subset V'$. So, if either X_2 or X_3 is a star, then V' contains a red point, since every star has a red fan. Since the star itself is blue, we are done in this case. If X_1, X_2, X_3, X_4 contains three consecutive regular vertices then we are done again, by the coloring rule for the regular intervals. So we are left with the case when X_1 and X_4 are stars, X_2 and X_3 are regular. But in this case V' also contains the common friend Y of X_2 and X_3 in P_2^W , which is also a regular vertex. By the coloring rule for the regular intervals, one of Y, X_2 and X_3 is red, the other two are blue, so we are done.




Figure 14: Two examples of coloring of $P_1^W \cup P_2^W$. Friends are connected by edges.

For $P_3^W \cup P_4^W$ we use the same coloring rule as for $P_1^W \cup P_2^W$ but we switch the roles of the colors. So any $W' \in Tr_4^W$ contains at least one red point of $P_3^W \cup P_4^W$ and any $V' \in Tr_7^V(P_3^W \cup P_4^W)$ contains both colors.

Finally, we have to color the rest of the points $R = S \setminus (P_1^W \cup P_2^W \cup P_3^W \cup P_4^W)$ such that every $V' \in Tr_2^V(R)$ contains points of both colors. This can be achieved by the first proof of Lemma 2.13.

Now any $W' \in Tr_4^W$ contains at least one blue and at least one red point. If $V' \in Tr_{14}^V$, then either it contains at least two points of $R = P \setminus (P_1^W \cup P_2^W \cup P_3^W \cup P_4^W)$, or at least seven points of $P_1^W \cup P_2^W$, or at least seven points of $P_3^W \cup P_4^W$, and in all cases it contains points of both colors. This completes the proof of Lemma 3.1. \square

Definition 3.2. Suppose that $\mathcal{W} = \{V, W\}$ is a pair of wedges. \mathcal{W} is said to be asymmetric non-conflicting or simply ANC, if there is a constant k with the following property. Any finite set of points S can be colored with red and blue such that any translate of V that contains at least k points of S , contains a red point, and any translate of W that contains at least k points of S , contains a blue point.

The next technical result allows us to simplify all following proofs.

Lemma 3.3. If a pair of wedges is not of Type 5 (Special), and ANC, then it is also NC.




Figure 15: Type 5 (Special)

Proof. We can assume without loss of generality that V contains the positive x -axis, and W contains either the positive or the negative x -axis. Suppose that $\{V, W\}$ is ANC, let $k > 0$ arbitrary, and let S be a set of points. First we color $Bd^V(S)$. Let U be a wedge that also contains the positive x -axis, but has a very small angle. Then translates of V and translates of U both intersect $Bd^V(S)$ in its intervals. Clearly, the pair $\{U, W\}$ is of Type 3 (Contain), therefore, by Lemma 3.1, we can color $Bd^V(S)$ such that any translate of W , $W' \in Tr_4^W(Bd^V(S))$ and any translate of U , $U' \in Tr_{14}^U(Bd^V(S))$ contains both colors. But then any translate of V , $V' \in Tr_{14}^V(Bd^V(S))$ contains both colors as well.

Now we have to color $S \setminus Bd^V(S)$. We divide it into three parts as follows.

$$S_b = \{X \in S \setminus Bd^V(S) \mid \forall Y \in V(X) \cap Bd^V(S), Y \text{ is blue}\},$$

$$S_r = \{X \in S \setminus Bd^V(S) \mid \forall Y \in V(X) \cap Bd^V(S), Y \text{ is red}\},$$

$$S_0 = S \setminus (Bd^V(S) \cup S_b \cup S_r).$$

Any translate $V' \in Tr^V$ that intersects S_b in at least one point, must contain at least one blue point, from $Bd^V(S)$, so we only have to make sure that it contains a red point

too. Similarly, any $V' \in Tr^V$ that intersects S_r in at least one point, must contain a red point, and any $V' \in Tr^V$ that intersects S_0 must contain points of both colors.

Thus, we can simply color S_0 such that any $W' \in Tr_2^W(S_0)$ contains both colors, which can be done by Lemma 2.13.

With S_b , and with S_r , respectively, we proceed exactly the same way as we did with S itself, but now we change the roles of V and W . We get the (still uncolored) subsets $S_{b,b}$, $S_{b,r}$, $S_{b,0}$, $S_{r,b}$, $S_{r,r}$, $S_{r,0}$ with the following properties.

- Any translate $V' \in Tr^V$ or $W' \in Tr^W$, that intersects $S_{b,b}$ (resp. $S_{r,r}$) in at least one point, must contain at least one blue (resp. red) point.
- Any translate $V' \in Tr^V$ that intersects $S_{b,r}$ (resp. $S_{r,b}$) contains a blue (resp. red) point, and any translate $W' \in Tr^W$ that intersects $S_{b,r}$ (resp. $S_{r,b}$) contains a red (resp. blue) point.
- Any translate $V' \in Tr^V$ that intersects $S_{b,0}$ (resp. $S_{r,0}$) contains a blue (resp. red) point, and any translate $W' \in Tr^W$ that intersects $S_{b,0}$ (resp. $S_{r,0}$) contains points of both colors.

Color all points of $S_{b,b}$ and $S_{b,0}$ red, color all points of $S_{r,r}$ and $S_{r,0}$ blue. Finally, color $S_{b,r}$ using the ANC property of the pair (V, W) , and similarly, color $S_{r,b}$ also using the ANC property, but the roles of red and blue switched. Now it is easy to check that in this coloring any translate of V or W that contains sufficiently many points of S , contains a point of both colors. \square

Remark 3.4. *In [P10] it has been proved that if $\{V, W\}$ is a Special pair, then $\{V, W\}$ is not ANC. So, the following statement holds as well.*

Lemma 3.3'. *If a pair of wedges is ANC, then it is also NC.*

Lemma 3.5. *Let $\mathcal{W} = \{V, W\}$ be a set of two wedges, of Type 1 (Big). Then \mathcal{W} is NC.*

Proof. By Lemma 3.3, it is enough to show that $\{V, W\}$ is ANC. Let W be the wedge whose angle is at least π . Then W is the union of two halfplanes, say, H_1 and H_2 . Translate both halfplanes such that they contain exactly one point of S , denote them by X_1 and X_2 , respectively. Note that X_1 may coincide with X_2 . Color X_1 and X_2 red, and all the other points blue. Then any translate of W that contains at least one point, contains a red point, and any translate of V that contains at least three points, contains a blue point. \square

Lemma 3.6. *Let $\mathcal{W} = \{V, W\}$ be a set of two wedges, of Type 2 (Halfplane). Then \mathcal{W} is NC.*

Proof. Again, it is enough to show that they are ANC. Since $\{V, W\}$ is of Type 2 (Halfplane), $Bd^V(S)$ and $Bd^W(S)$ have at most one point in common. If $Bd^V(S)$ and $Bd^W(S)$ are disjoint, then color $Bd^V(S)$ blue, $Bd^W(S)$ red, and the other points arbitrarily. Then any nonempty translate of V (resp. W) contains a blue (resp. red) point.




Figure 16: Type 2 (Halfplane)

Otherwise, let X be their common point. Let $P = Bd^V(S) \cup Bd^W(S) \setminus X$, and consider its V -boundary, $Bd^V(P)$, and W -boundary, $Bd^W(P)$. Clearly, each point in $P = Bd^V(S) \setminus X$ belongs to $Bd^V(P)$, and each point in $P = Bd^W(S) \setminus X$ belongs to $Bd^W(P)$.

If $Bd^V(P)$ and $Bd^W(P)$ are disjoint, then color $Bd^V(S)$ blue, $Bd^W(P)$ and the other points red. Then any nonempty translate of V contains a blue point. Suppose that we have a translate of W with two points, both blue. Then it should contain X , and a point of $Bd^V(P)$. But this contradicts our assumption that $Bd^V(P)$ and $Bd^W(P)$ are disjoint. So, any translate of W which contains at least two points of S , contains a red point.

If $Bd^V(P)$ and $Bd^W(P)$ are not disjoint, then they have one point in common, let Y be their common point. If Y belongs to $Bd^W(P)$, then color $Bd^V(S)$ blue, $Bd^W(P)$ and the other points red. Then, by the same argument as before, any nonempty translate of V contains a blue point, and any translate of W which contains at least two points of S , contains a red point. Finally, if Y belongs to $Bd^V(P)$, then we proceed analogously, but the roles of V and W , and the colors, are switched. \square

Lemma 3.7. *Let $\mathcal{W} = \{V, W\}$ be a set of two wedges, of Type 4 (Hard). Then \mathcal{W} is NC.*




Figure 17: Type 4 (Hard)

Proof. As usual, we only prove that $\{V, W\}$ is ANC. Assume that W contains the positive x -axis. Just like in the definition of the different types, extend the boundary halflines of W to lines, they divide the plane into four parts, Upper, Lower, Left, and Right, latter of which is W itself. We can assume without loss of generality that V contains the negative x -axis, one side of V is in Upper, and one side is in Left, just like on the left of Figure 17.

Observe that if a translate of V and a translate of W intersect each other, then one of them contains the apex of the other one.

Claim 3.8. *For any point set P and $X \in P$, either $Bd^V(P \setminus X) \setminus Bd^V(P) = \emptyset$ or $Bd^W(P \setminus X) \setminus Bd^W(P) = \emptyset$.*

Proof. Suppose on the contrary that $Y \in Bd^V(P \setminus X) \setminus Bd^V(P)$ and $Z \in Bd^W(P \setminus X) \setminus Bd^W(P)$. Then $X \in V(Y)$ and $X \in W(Z)$, so $V(Y)$ and $W(Z)$ intersect each other, therefore, one of them contains the other one's apex, say, $Z \in V(Y)$. But this is a contradiction, since Y is a boundary point of $P \setminus X$. \square

Return to the proof of Lemma 3.7. Color $Bd^V(S) \setminus Bd^W(S)$ red, and $Bd^W(S) \setminus Bd^V(S)$ blue, the interior points arbitrarily. Now consider the points of $Bd^V(S) \cap Bd^W(S)$. For any $X \in Bd^V(S) \cap Bd^W(S)$, if $Bd^V(S \setminus X) \setminus Bd^V(S) \neq \emptyset$, then color it red, if $Bd^W(S \setminus X) \setminus Bd^W(S) \neq \emptyset$, then color it blue. For each of the remaining points Y we have $Bd^V(S \setminus Y) \setminus Bd^V(S) = Bd^W(S \setminus Y) \setminus Bd^W(S) = \emptyset$. Color each of these points such that they have the *opposite* color than the the previous point of $Bd^V(S) \cap Bd^W(S)$, in the y -ordering.

To prove that this coloring is good, let $V' \in Tr_2^V$, $V' \cap S = \{X, Y\}$. If it intersects $Bd^V(S) \setminus Bd^W(S)$, we are done. So assume that $V' \cap Bd^V(S) \subset Bd^V(S) \cap Bd^W(S)$. Let $X \in V' \cap Bd^V(S)$. If X is red, then by the coloring rule, $Bd^V(S \setminus X) \setminus Bd^V(S) = \emptyset$. But then Y is also a V -boundary point, so we have $Y \in Bd^V(S) \cap Bd^W(S)$. Again we can assume that Y is red, so $Bd^V(S \setminus Y) \setminus Bd^V(S) = \emptyset$. Suppose that $X <_y Y$. Since $V' \cap S = \{X, Y\}$, X and Y are consecutive points of $Bd^V(S) \cap Bd^W(S)$. Now it is not hard to see that $Bd^W(S \setminus Y) \setminus Bd^W(S) = \emptyset$. Therefore, by the coloring rule, X and Y have different colors. For the translates of W the argument is analogous, with the colors switched. \square

Now we turn to the case when we have more than two wedges.

Lemma 3.9. *For any $s, t > 0$ integers, there is a number $f(s, t)$ with the following property.*

Let $\mathcal{W} = \{W_1, W_2, \dots, W_t\}$ be a set of t wedges, such that any pair $\{W_i, W_j\}$ is NC, and let S be a set of points. Then S can be decomposed into t parts, S_1, S_2, \dots, S_t , such that for $i = 1, 2, \dots, t$, for any translate W'_i of W_i , if $|W'_i \cap S| \geq f(s, t)$ then $|W'_i \cap S_i| \geq s$.

Proof. The existence of $f(1, 2)$ is equivalent to the property that the corresponding two wedges are ANC. Now we show that $f(s, 2)$ exists for every s . Let V and W be two wedges that form a NC pair. Let $P_1^V, P_2^V, \dots, P_{s^2 f(1, 2)}^V$ be the path decomposition of S of order $s^2 f(1, 2)$, with respect to V . For $i = 1, 2, \dots, s$, let

$$H_i = \cup_{j=(i-1)s^2 f(1, 2)+1}^{is^2 f(1, 2)} P_j^V.$$

For each H_i , take the W -path decomposition, $P_1^W(H_i), \dots, P_{sf(1, 2)}^W(H_i)$, and for each $j = 1, 2, \dots, s$, let

$$H_i^j = \cup_{k=(j-1)f(1, 2)+1}^{jf(1, 2)} P_k^W(H_i).$$

For every $i, j = 1, 2, \dots, s$, color H_i^j , such that any translate of V (resp. W) that intersects it in at least $f(1, 2)$ points, contains at least one red (resp. blue) point of it. This is possible, since the pair $\{V, W\}$ is ANC.

Consider a translate V' of V that contains at least $s^2 f(1, 2)$ points of S . For every i , V' intersects H_i in $sf(1, 2)$ points, so there is a j such that it intersects H_i^j in at least $f(1, 2)$ points. Therefore, V' contains at least one red point of H_i^j , so at least s red points of S .

Consider now a translate W' of W that contains at least $s^2 f(1, 2)$ points of S . There is an i such that W' intersects H_i in at least $sf(1, 2)$ points. Therefore, it intersects each of $P_1^W(H_i), \dots, P_{sf(1, 2)}^W(H_i)$, in at least one point, so for $j = 1, 2, \dots, s$, W' intersects H_i^j

in at least $f(1, 2)$ points. Consequently, it contains at least one blue point of each H_i^j , so at least s blue points of S .

Now let $s, t > 2$ fixed and suppose that $f(s', t-1)$ exists for every s' . Let $\{W_1, \dots, W_t\}$ be our set of wedges, such that any pair of them is NC. Let $s' = f(s, 2)$. Partition our point set S into $S'_1, S'_2, \dots, S'_{t-1}$ such that for $i = 1, 2, \dots, t-1$, for any translate W'_i of W_i , if $|W'_i \cap S| \geq f(s', t-1)$ then $|W'_i \cap S'_i| \geq s' = f(s, 2)$. For each $i = 1, 2, \dots, t-1$, partition S'_i into two parts, S''_i and S^t_i , such that for any translate W'_i of W_i , if $|W'_i \cap S'_i| \geq f(s, 2)$ then $|W'_i \cap S''_i| \geq s$, and for any translate W'_t of W_t , if $|W'_t \cap S'_i| \geq f(s, 2)$ then $|W'_t \cap S^t_i| \geq s$. Finally, for $i = 1, 2, \dots, t-1$, let $S_i = S''_i$ and let $S_t = \cup_{j=1}^{t-1} S^t_j$. For $i = 1, 2, \dots, t-1$, any translate W'_i of W_i , if $|W'_i \cap S| \geq f(s', t-1)$ then $|W'_i \cap S'_i| \geq s' = f(s, 2)$, so $|W'_i \cap S_i| \geq s$,

And for any translate W'_t of W_t , if $|W'_t \cap S| \geq f(s', t-1)$, then for some $i = 1, 2, \dots, t-1$, $|W'_t \cap S'_i| \geq \frac{f(s', t-1)}{t-1} \geq f(s, 2)$, therefore, $|W'_t \cap S^t_i| \geq s$, so $|W'_t \cap S_t| \geq s$. This concludes the proof of Lemma 3.9. \square

Remark 3.10. *The proofs of Lemmas 3.1, 3.5, 3.6, and 3.7 imply that $f(1, 2) \leq 8$. Combining it with the proof of Lemma 3.9 we get the bound $f(s, t) \leq (8s)^{2^{t-1}}$.*

As a corollary, we have can now prove Lemma 2.16.

Lemma 2.16. *A set of wedges $\mathcal{W} = \{W_1, W_2, \dots, W_t\}$ is NC if and only if any pair $\{W_i, W_j\}$ is NC.*

Proof. Clearly, if some pair $\{W_i, W_j\}$ is *not* NC, then the whole set \mathcal{W} is not NC either. Suppose that every pair $\{W_i, W_j\}$ is NC. Decompose S into t parts S_1, S_2, \dots, S_t with the property that for $i = 1, 2, \dots, t$, for any translate W'_i of W_i , if $|W'_i \cap S| \geq f(3, t)$ then $|W'_i \cap S_i| \geq 3$. Then, by Lemma 2.13, each S_i can be colored with red and blue such that if $|W'_i \cap S_i| \geq 3$ then W'_i contains points of both colors. So this coloring of S has the property that for $i = 1, 2, \dots, t$, for any translate W'_i of W_i , if $|W'_i \cap S| \geq f(3, t)$ then it contains points of both colors. \square

3.3 Summary of Proof of Theorem A

Although we have already established Theorem A, we find it useful to give another summary of the complete proof.

Suppose that P is an open convex polygon of n vertices and $\mathcal{P} = \{P_i \mid i \in I\}$ is a collection of translates of P which forms an M -fold covering of the plane. We will set the value of M later. Let d be the minimum distance between any vertex and non-adjacent side of P . Take a square grid \mathcal{G} of basic distance $d/2$. Obviously, any translate of P intersects at most $K = 4\pi(\text{diam}(P) + d)^2/d^2$ basic squares. For each (closed) basic square B , using its compactness, we can find a *finite* subcollection of the translates such that they still form an M -fold covering of B . Take the union of all these subcollections. We have a *locally finite* M -fold covering of the plane. That is, every compact set is intersected by finitely many of the translates. It is sufficient to decompose this covering. For simplicity, use the same notation $\mathcal{P} = \{P_i \mid i \in I\}$ for this subcollection.

We formulate and solve the problem in its dual form. Let O_i be the center of gravity of P_i . Since \mathcal{P} is an M -fold covering of the plane, every translate of \bar{P} , the reflection of P through the origin, contains at least M points of the locally finite set $\mathcal{O} = \{O_i \mid i \in I\}$.

The collection $\mathcal{P} = \{ P_i \mid i \in I \}$ can be decomposed into two coverings if and only if the set $\mathcal{O} = \{ O_i \mid i \in I \}$ can be colored with two colors, such that every translate of \bar{P} contains a point of both colors.

Let $\mathcal{W} = \{ W_1, W_2, \dots, W_n \}$ be the set of wedges that correspond to the vertices of \bar{P} . By the convexity of \bar{P} , no pair $\{ W_i, W_j \}$ is of Type 5 (Special), therefore, by the previous Lemmas, each pair is NC. Consequently, by Lemma 2.16, \mathcal{W} is NC as well. So there is an m with the following property.

* Any set of points S can be colored with two colors such that any translate of W_1, \dots, W_n that contains at least m points of S , contains points of both colors.

Choose M such that $M \geq mK$, and color the points of \mathcal{O} in each basic square separately, with property *.

Since any translate P' of \bar{P} intersects at most K basic squares of the grid \mathcal{G} , P' contains at least $M/K \geq m$ points of \mathcal{O} in the same basic square B' . By the choice of the grid \mathcal{G} , B' contains at most one vertex of P' , hence $B' \cap P' = B' \cap W$, where W is a translate of some $W_i \in \mathcal{W}$. So, by property *, P' contains points of $\mathcal{O} \cap B'$ of both colors. This concludes the proof of Theorem A. \square

3.4 Concluding Remarks

Throughout this section we made no attempt to optimize the constants. However, it may be an interesting problem to determine (asymptotically) the smallest m in the proof of Theorem A.

Another interesting question is to decide whether this constant depends only on the number of vertices of the polygon, or on the shape as well. In particular, we cannot verify the following.

Conjecture 3.11. *There is a constant m such that any m -fold covering of the plane with translates of a convex quadrilateral can be decomposed into two coverings.*

With a slight modification of our proof of Theorem A, we get the following more general result about decomposition to k coverings.

Theorem A'. *For any open convex polygon (or concave polygon without Type 5 (special) wedges), P , and any k , there exists a (smallest) number $m_k(P)$, such that any $m_k(P)$ -fold covering of the plane with translates of P can be decomposed into k coverings.*

Our proof gives $m_k(P) < K_P(8k)^{2^{n-1}}$, where K_P is the constant K from the proof of Theorem A and n is the number of vertices of P . The best known lower bound on $m_k(P)$ is $\lfloor 4k/3 \rfloor - 1$ [PT07]. Recently Gibson and Varadarajan [GV10] proved Theorem C (iii), which is a linear upper bound for all convex polygons. However, their proof does not work for any cover-decomposable polygon, because they handle only one case of Type 3 (Contain). For a summary of their result see the end of Section 2. We conjecture that a linear upper bound also holds for cover-decomposable concave polygons.

Conjecture 3.12. *For any cover-decomposable polygon P , $m_k(P) = O(k)$.*

For more related conjectures see Section 9.

Our proofs use the assumption that the covering is locally finite, and for open polygons we could find a locally finite subcollection which is still a m -fold covering. Still, we strongly believe that Theorem A holds for closed convex polygons as well.

4 Indecomposable Coverings by Concave Polygons

This section is based on my paper, Indecomposable coverings with concave polygons [P10].

The main goal of this section is to prove results about *non-cover-decomposable* polygons. To understand some of the results, we need to recall the notions introduced in Section 2.1.3.

Definition 2.3. A planar set P is said to be *totally-cover-decomposable* if there exists a (minimal) constant $m^T = m^T(P)$ such that every m^T -fold covering of ANY planar point set with translates of P can be decomposed into two coverings. Similarly, let $m_k^T(P)$ denote the smallest number m^T with the property that every m^T -fold covering of ANY planar point set with translates of P can be decomposed into k coverings.

When we want to emphasize the difference from the original definition, we will call the cover-decomposable sets *plane-cover-decomposable*. By definition, if a set is totally-cover-decomposable, then it is also plane-cover-decomposable. On the other hand, we cannot rule out the possibility that there are sets, or even polygons, which are plane-cover-decomposable, but not totally-cover-decomposable.

The results of Theorem A all remain true if we write totally-cover-decomposable instead of cover-decomposable and similarly, in Theorem C we can replace $m_k(P)$ with $m_k^T(P)$. In fact in the proofs of these theorems, these more general claims are proved. The proof of Theorem B establishes first that concave quadrilaterals are not totally-cover-decomposable and then extends the covering, proving that they are also not plane-cover-decomposable. The main result of this section is a generalization of Theorem B. We show that *almost* all (open or closed) concave polygons are *not* totally-cover-decomposable and prove that most of them are also not plane-cover-decomposable. We need the “almost” because Theorem A’ implies that any concave polygon without Type 5 (Special) wedges is totally-cover-decomposable.

Our main result is the following

Theorem E. *If a polygon has a pair of Type 5 (Special) wedges, then it is not totally-cover-decomposable.*




Figure 18: Type 5: Special pair of wedges

Together with the previous theorem, this gives a complete characterization of totally-cover-decomposable open polygons; an open polygon is totally-cover-decomposable if and only if it does not have a Special pair of wedges.

We show that every concave polygon with no parallel sides has a pair of Special wedges, therefore we have

Theorem E’. *If a concave polygon has no parallel sides, then it is not totally-cover-decomposable.*

The proof of these theorems can be found in Section 4.1. The problem of deciding plane-cover-decomposability for concave polygons is still open. However, in Section 4.2, we prove that a large class of concave polygons are not plane-cover-decomposable. We also show that any “interesting” covering of the plane uses only countably many translates. (However, we do not consider here the problem when we want to decompose into infinitely many coverings; the interested reader is referred to the paper of Elekes, Mátrai and Soukup [EMS10].)

Finally, in Section 4.3, we investigate the problem in three or more dimensions. The notion of totally-cover-decomposability extends naturally and we can also introduce space-cover-decomposability. Previously, the following result was known.

Theorem F. Mani-Levitska, Pach [MP86] *The unit ball is not space-cover-decomposable.*

Using our construction, we establish the first theorem for polytopes which shows that the higher dimensional case is quite different from the two dimensional one.

Theorem B’. *Polytopes are not cover-decomposable in the space and in higher dimensions.*

4.1 The Construction - Proof of Theorems E and E’

In this section, for any k and any polygon C that has a Special pair of wedges, we present a (finite) point set and an indecomposable k -fold covering of it by (a finite number of) the translates of the polygon. We formulate (and solve) the problem in its dual form, like we did before. Here we recall how the dualization goes. Fix O , the center of gravity of C as our origin in the plane. For the planar set C and a point p in the plane we use $C(p)$ to denote the translate of C by the vector \vec{Op} . Let \bar{C} be the reflection through O of C . For any point x , $x \in C(p_i)$ if and only if $p_i \in \bar{C}(x)$. To see this, apply a reflection through the midpoint of the segment xp_i . This switches $C(p_i)$ and $\bar{C}(x)$, and also switches p_i and x .

Consider any collection $\mathcal{C} = \{C(p_i) \mid i \in I\}$ of translates of C and a point set X . The collection \mathcal{C} covers x at least k times if and only if $\bar{C}(x)$ contains at least k elements of the set $S = \{p_i \mid i \in I\}$. Therefore a k -fold covering of X transforms into a point set such that for every $x \in X$ the set $\bar{C}(x)$ contains at least k points of S . The required decomposition of \mathcal{C} exists if and only if the set S can be colored with two colors such that every translate $\bar{C}(x)$ that contains at least k elements of S contains at least one element of each color. Thus constructing a finite system of translates of \bar{C} and a point set where this latter property fails is equivalent to constructing an indecomposable covering using the translates of C .

If C has a Special pair of wedges, then so does \bar{C} . We will use the following theorem to prove Theorem E.

Theorem 4.1. *For any pair of Special wedges, V and W , and for every k, l , there is a point set of cardinality $\binom{k+l}{k} - 1$, such that for every coloring of S with red and blue, either*

there is a translate of V containing k red points and no blue points, or there is a translate of W containing l blue points and no red points.

Proof. Without loss of generality, suppose that the wedges are contained in the right halfplane.

For $k = 1$ the statement is trivial, just take l points such that any one is contained alone in a translate of W . Similarly k points will do for $l = 1$. Let us suppose that we already have a counterexample for all $k' + l' < k + l$ and let us denote these by $S(k', l')$. The construction for k and l is the following.

Place a point p in the plane and a suitable small scaled down copy of $S(k-1, l)$ left from p such that any translate of V with its apex in the neighborhood of $S(k-1, l)$ contains p , but none of the translates of W with its apex in the neighborhood of $S(k-1, l)$ does. Similarly place $S(k, l-1)$ such that any translate of W with its apex in the neighborhood of $S(k, l-1)$ contains p , but none of the translates of V with its apex in the neighborhood of $S(k, l-1)$ does. (See Figure 19.)




Figure 19: Sketch of one step of the induction and the first few steps.

If p is colored red, then

- either the $S(k-1, l)$ part already contains a translate of V that contains $k-1$ reds and no blues, and it contains p as well, which gives together k red points
- or the $S(k-1, l)$ part contains a translate of W that contains l blues and no reds and it does not contain p .

The same reasoning works for the case when p is colored blue.

Now we can calculate the number of points in $S(k, l)$. For $l = 1$ and for $k = 1$ we know that $|S(k, 1)| = k$ and $|S(1, l)| = l$, while the induction gives $|S(k, l)| = 1 + |S(k-1, l)| + |S(k, l-1)|$. From this we have $|S(k, l)| = \binom{k+l}{k} - 1$. \square

It is easy to see that if we use this theorem for a pair of Special wedges of \bar{C} and $k = l$, then for every coloring of (a possibly scaled down copy of) the above point set with two

colors there is a translate of \bar{C} that contains at least k points, but contains only one of the colors. This is because “ \bar{C} can locally behave like any of its wedges”. Therefore this construction completes the proof of Theorem E. \square

Remark 4.2. *Note that we can even give a finite collection of translates of V and W whose apices all lie on the same line such that one of them will satisfy the conclusion of the theorem. Moreover, this line can be any line that can touch a translate of each wedge in only its apex.*

Remark 4.3. *We note that for $k = l$ the cardinality of the point set is approximately $4^k/\sqrt{k}$, this significantly improves the previously known construction of Pach, Tardos and Tóth [PTT05] which used approximately k^k points and worked only for quadrilaterals, and in general, for “even more Special” pairs of wedges (the ones on the right side of Figure 18). It can be proved that this exponential bound is close to being optimal. Suppose that we have n points and $n < 2^{k-2}$. Since there are two kinds of wedges, there are at most $2n$ essentially different translates that contain k points. There are 2^n different colorings of the point set and each translate that contains k points is monochromatic for 2^{n-k+1} of the colorings. Therefore, there are at most $2n2^{n-k+1} < 2^n$ bad colorings, so there is a coloring with no monochromatic translates.*

Theorem E’ follows directly from the next result.

Lemma 4.4. *Every concave polygon that has no parallel sides, has a Special pair of wedges.*

Proof. Assume that the statement does not hold for a polygon C . There is a touching line ℓ to C such that the intersection of ℓ and C contains no segments and contains at least two vertices, v_1 and v_2 . (Here we use that C has no parallel sides.) Denote the wedges at v_i by W_i . Is the pair W_1, W_2 Special? They clearly fulfill the property (i), the only problem that can arise is that the translate of one of the wedges contains the other wedge. This means, without loss of generality, that the angle at v_1 contains the angle at v_2 . Now let us take the two touching lines to C that are parallel to the sides of W_2 . It is impossible that both of these lines touch v_2 , because then the touching line ℓ would touch only v_2 as well. Take a vertex v_3 from the touching line (or from one of these two lines) that does not touch v_2 . (See Figure 20.) This cannot be v_1 because then the polygon would have two parallel sides. Is the pair W_2, W_3 Special? They are contained in a halfplane (the one determined by the touching line). This means, again, that the angle at v_2 contains the angle at v_3 . Now we can continue the reasoning with the touching lines to C parallel to the sides of W_3 , if they would both touch v_3 , then the touching line ℓ would touch only v_3 . This way we obtain the new vertices v_4, v_5, \dots what contradicts the fact the C can have only a finite number of vertices. \square

4.2 Versions of Cover-decomposability

Here we consider different variants of cover-decomposability and prove relations between them.




Figure 20: How to find a Special pair of wedges.

4.2.1 Number of sets: Finite, infinite or more

We say that a set is *finite/countable-cover-decomposable*, if there exists a k such that every k -fold covering of any point set by a finite/countable number of its translates is decomposable. So by definition we have: totally-cover-decomposable \Rightarrow countable-cover-decomposable \Rightarrow finite-cover-decomposable. But which of these implications can be reversed? We will prove that the first can be for “nice” sets.

It is well-known that the plane is hereditary Lindelöf, i.e. if a point set is covered by open sets, then countably many of these sets also cover the point set. It is easy to see that the same holds for k -fold coverings as well. This observation implies the following lemma.

Lemma 4.5. *An open set is totally-cover-decomposable if and only if it is countable-cover-decomposable.* \square

The same holds for “nice” closed sets, such as polygons or discs. We say that a closed set C is *nice* if there is a t and a set \mathcal{D} of countably many closed halfdiscs such that if t different translates of C cover a point p , then their union covers a halfdisc from \mathcal{D} centered at p (meaning that p is halving the straight side of the halfdisc) and the union of their interiors covers the interior of the halfdisc. For a polygon, t can be the number of its vertices plus one, \mathcal{D} can be the set of halfdiscs whose side is parallel to a side of the polygon and has rational length. For a disc, t can be 2 and \mathcal{D} can be the set of halfdiscs whose side has a rational slope and a rational length. In fact every convex set is nice.

Claim 4.6. *Every closed convex set is nice.*

Proof. Some parts of the boundary of the convex set C might be segments, we call these sides. Trivially, every convex set can have only countably many sides. Choose $t = 5$ and let the set of halfdiscs \mathcal{D} be the ones whose side is either parallel to a side of C or its slope is rational and has rational length. Assume that 5 different translates of C cover a point p . Shifting these translates back to C , denote the points that covered p by p_1, \dots, p_5 . If any of these points is not on the boundary of C , we are done. The $p_1p_2p_3p_4p_5$ pentagon has two neighboring angles the sum of whose degrees is strictly bigger than 2π , without loss of generality, p_1 and p_2 . If p_1p_2 is also the side of C , then the 5 translates cover a halfdisc whose side is parallel to p_1p_2 , else they cover one whose side has a rational slope. \square

Taking a rectangle verifies that $t = 5$ is optimal in the previous proof.

Lemma 4.7. *A nice set is totally-cover-decomposable if and only if it is countable-cover-decomposable.*

Proof. We have to show that if we have an infinite covering of some point set S by the translates of our nice, countable-cover-decomposable set C , then we can suitably color the points of S . Denote by S^* the points that are covered by 2 copies of the same translate of the nice set C . Color one of these red, the other blue. Now we only have to deal with $S' = S \setminus S^*$ and we can suppose that there is only one copy of each translate. Now instead of coloring these translates, we rather show that we can choose countably many of them such that they still cover every point of S' many times. Using after this that the set is countable-cover-decomposable finishes the proof. So now we show that if there is a set of translates of C that cover every point of S' at least kt times, then we can choose countably many of these translates that cover every point of S' at least k times. It is easy to see that it is enough if we show this for $k = 1$ (since we can repeat this procedure k times).

Denote the points that are contained in the interior of a translate by S_0 . Because of the hereditary Lindelöf property, countably many translates cover S_0 . If a point $p \in S'$ is covered t times, then because of the nice property of C , a halfdisc from \mathcal{D} centered at p is covered by these translates. We say that this (one of these) halfdisc(s) belongs to p . Take a partition of $S' \setminus S_0$ into countably many sets $S_1 \cup S_2 \cup \dots$ such that the i^{th} halfdisc belongs to the points of S_i . Now it is enough to show that S_i can be covered by countably many translates. Denote the halfdisc belonging to the points of S_i by D_i . Using the hereditary Lindelöf property for S_i and open discs (not halfdiscs!) with the radius of D_i centered at the points of S_i , we obtain a countable covering of S_i . Now replacing the open discs with closed halfdiscs still gives a covering of S_i because otherwise we would have $p, q \in S_i$ such that p is in the interior of $q + D_i$, but interior of $q + D_i$ is covered by the interiors of translates of C , which would imply $p \in S_0$, contradiction. Finally we can replace each of the halfdiscs belonging to the points of S_i by t translates of C , we are done. \square

Unfortunately, we did not manage to establish any connection among the finite- and the countable-cover-decomposability. We conjecture that they are equivalent for nice sets (with a possible slight modification of the definition of nice). If one manages to find such a statement, then it would imply that considering cover-decomposability, it does not matter whether the investigated geometric set is open or closed, as long as it is nice. For example, it is unknown whether closed triangles are cover-decomposable or not. We strongly believe that they are.

4.2.2 Covering the whole plane

Remember that by definition if a set is totally-cover-decomposable, then it is also plane-cover-decomposable. However, the other direction is not always true. For example take the lower halfplane and “attach” to its top a pair of Special wedges (see Figure 21). Then the counterexample using the Special wedges works for a special point set, thus this set is not totally-cover-decomposable, but it is easy to see that a covering of the whole plane can always be decomposed.

For a given polygon C , our construction gives a set of points S and a non-decomposable k -fold covering of S by translates of C . It is not clear when we can extend this covering to a k -fold covering of the whole plane such that none of the new translates contain any point of S . This would be necessary to ensure that the covering remains non-decomposable.

We show that in certain cases it can be extended, but it remains an open problem to decide whether plane- and totally-cover-decomposability are equivalent or not for open polygons/bounded sets.




Figure 21: The lower halfplane with a Special pair of wedges at its top.

Theorem 4.8. *If a concave polygon C has two Special wedges that have a common locally touching line, and one of the two touching lines parallel to this line is touching C in only a finite number of points (i.e. does not contain a side), then it is not plane-cover-decomposable.*

Proof. Assume, without loss of generality, that this locally touching line is vertical. We just have to extend our construction with the Special wedges into a covering of the whole plane. Or, in the dual, we have to add more points to our construction such that every translate of C will contain at least k points. Of course, to preserve that the construction works, we cannot add more points into those translates that we used in the construction. Otherwise, our argument that the construction is correct, does not work. Because of Remark 4.2 we can suppose that the apices of the wedges all lie on the same vertical line. Therefore, the translates can all be obtained from each other via a vertical shift, because we had a vertical locally touching line to both wedges. Now we can simply add all points that are not contained in any of these original translates. Proving that every translate of C contains at least k points is equivalent to showing that the original translates do not cover any other translate of C . It is clear that they could only cover a translate that can be obtained from them via a vertical shift. On the touching vertical line each of the translates has only finitely many points. In the construction we have the freedom to perturbate the wedges a bit vertically, this way we can ensure that the intersection of each other translate (obtainable via a vertical shift) with this vertical line is not contained in the union of the original translates. \square

Corollary 4.9. *A pentagon is totally-cover-decomposable if and only if it is plane-cover-decomposable.*

Proof. All totally-cover-decomposable sets are also plane-cover-decomposable. (See an example on Figure 22.) If our pentagon is not totally-cover-decomposable, then it has a Special pair of wedges and it must also have a touching line that touches it in these Special wedges, thus we can use the previous theorem. \square




Figure 22: A pentagon that is cover-decomposable but is not the union of a finite number of translates of the same convex polygon.

The same argument does not work for hexagons, for example we do not know whether the hexagon depicted in Figure 23/c is plane-cover-decomposable or not.




Figure 23: Three different polygons. (a): totally-cover-decomposable (hence, also plane-cover-decomposable), (b): not plane-cover-decomposable, (hence neither totally-cover-decomposable), (c): not totally-cover-decomposable, but not known if plane-cover-decomposable.

4.3 Higher Dimensions - Proof of Theorem B'

The situation is different for the space. For any polytope and any k , one can construct a k -fold covering of the space that is not decomposable. First note that it is enough to prove this result for the three dimensional space, since for higher dimensions we can simply intersect our polytope with a three dimensional space, use our construction for this three dimensional polytope and then extend it naturally. To prove the theorem for three dimensional polytopes, first we need some observations about polygons. Given two polygons and one side of each of them that are parallel to each other, we say that these sides are *directedly parallel* if the polygons are on the “same side” of the sides (i.e. the halfplane which contains the first polygon and whose boundary contains this side of the first polygon can be shifted to contain the second polygon such that its boundary contains that side of the second polygon). We will slightly abuse this definition and say that a side is directedly parallel if it is directedly parallel to a side of the other polygon. We can similarly define *directedly parallel faces* for a single polytope. We say that a face is directedly parallel to another, if they are parallel to each other and the polytope is on the “same side” of the faces (e.g. every face is directedly parallel to itself and if the polytope is convex, then to no other face).

Lemma 4.10. *Given two convex polygons, both of which have at most two sides that are directedly parallel, there is always a Special pair among their wedges.*

Proof. Take the smallest wedge of the two polygons, excluding the ones that contain a wedge whose sides are both directedly parallel, if it exists. Without loss of generality, we can suppose that the right side of this minimal wedge is not directedly parallel and it is going to the right (i.e. its direction is $(1,0)$), while the left side goes upwards. Take a wedge of the other polygon both of whose sides go upwards (there always must be one since the right side of the first wedge was not directedly parallel). If this second wedge is not contained in the first, we found a Special pair. If the second wedge is contained in the first wedge, then because of the minimality of the first wedge, we get a contradiction. \square

Theorem B'. *Polytopes are not cover-decomposable in the space and in higher dimensions.*

Proof. We will, as usual, work in the dual case. This means that to prove that our polytope C is not totally-cover-decomposable, we will exhibit a point set for any k such that we cannot color it with two colors such that any translate of C that contains at least k points contains both colors. These points will be all in one plane, and the important translates of C will intersect this plane either in a concave polygon or in one of two convex polygons. It is enough to show that this concave polygon is not cover-decomposable or that among the wedges of these convex polygons there is a Special pair.

Take a plane π that is not parallel to any of the segments determined by the vertices of C . The touching planes of C parallel to π are touching C in one vertex each, A and B . Denote the planes parallel to π that are very close to A and B and intersect C , by π_A and π_B . Denote $C \cap \pi_A$ by C_A and $C \cap \pi_B$ by C_B . Now we will have two cases.

Case 1. C_A or C_B is concave.

Without loss of generality, assume C_A is concave. Then since no two faces of C incident to A can be parallel to each other, with a perturbation of π_A we can achieve that the sides of C_A are not parallel. After this, using Theorem E', we are done.

Case 2. Both C_A and C_B are convex.

Now by perturbing π , we cannot necessarily achieve that C_A and C_B have no parallel sides, but we can achieve that they have at most two directedly parallel sides. This is true because there can be at most two pairs of faces that are directedly parallel to each other and one of them is incident to A , the other to B , since A is touched from above, B from below by the plane parallel to π . Therefore C_A and C_B satisfy the conditions of Lemma 4.10, this finishes the proof of totally-cover-decomposability.

To prove non-space-cover-decomposability, just as in the proof of Theorem 4.8, we have to add more points to the constructions, such that every translate will contain at least k points, but we do not add any points to the original translates of our construction. This is the same as showing that these original translates do not cover any other translate. Note that there are two types of original translates (depending on which wedge of it we use) and translates of the same type can be obtained from each other via a shift that is parallel to the side of the halfplane in π that contains our Special wedges. This means that the centers of all the original translates lie in one plane. With a little perturbation of the construction, we can achieve that this plane is in general position with respect to the polytope. But in

this case it is clear that the translates used in our construction cannot cover any other translate, this proves space-cover-decomposability. \square

4.4 Concluding Remarks

A lot of questions remain open. In three dimensions, neither polytopes, nor unit balls are cover-decomposable. Is there any nice (e.g. open and bounded) set in three dimensions that is cover-decomposable? Maybe such nice sets exist only in the plane.

Conjecture 4.11. *Three-dimensional convex sets are not cover-decomposable.*

In Section 4.2.1 we have seen that interesting covers with translates of nice sets only use countably many translates. We could not prove, but conjecture, that every cover can be somehow reduced to a locally finite cover. Is it true that if a nice set is finite-cover-decomposable, then it is also countable-cover-decomposable? This would have implications about the cover-decomposability of closed sets.

Conjecture 4.12. *Closed, convex polygons are cover-decomposable.*

In Section 4.2.2 we have seen that our construction is not naturally extendable to give an indecomposable covering of the whole plane. Maybe the reason for this is that it is impossible to find such a covering.

Question 4.13. *Are there polygons that are not totally-cover-decomposable but plane-cover-decomposable?*

Part II

Slope Number of Graphs

5 Introduction and a Lower Bound

A planar layout of a graph G is called a *drawing* if the vertices of G are represented by distinct points in the plane and every edge is represented by a continuous arc connecting the corresponding pair of points and not passing through any other point representing a vertex [DETT99]. If it leads to no confusion, in notation and terminology we make no distinction between a vertex and the corresponding point, and between an edge and the corresponding arc. If the edges are represented by line segments, the drawing is called a *straight-line drawing*. The *slope* of an edge in a straight-line drawing is the slope of the corresponding segment.

Wade and Chu [WC94] introduced the following graph parameter: The *slope number* of a graph G is the smallest number s with the property that G has a straight-line drawing with edges of at most s distinct slopes. Let us compare this with two other well studied graph parameters. The *thickness* of a graph G is defined as the smallest number of planar subgraphs it can be decomposed into [MOS98]. It is one of the several widely known graph parameters that measures how far G is from being planar. The *geometric thickness* of G , defined as the smallest number of *crossing-free* subgraphs of a straight-line drawing of G whose union is G , is another similar notion. [K73]. It follows directly from the definitions that the thickness of any graph is at most as large as its geometric thickness, which, in turn, cannot exceed its slope number. For many interesting results about these parameters, consult [DEH00, DEK04, DSW04, DW06, E04, HSV99].

Obviously, if G has a vertex of degree d , then its slope number is at least $\lceil d/2 \rceil$, because, according to the above definitions, in a proper drawing two edges are not allowed to partially overlap. The question arises whether the slope number can be bounded from above by any function of the maximum degree d (see [DSW07]). Barát, Matoušek, and Wood [BMW06] and, independently, in our paper with Pach [PP06], we proved using a counting argument that the answer is no for $d \geq 5$. We present this latter proof, which gives a better bound, in Section 5.1. We show that for any $d \geq 5$ and n , there exist a graph with n vertices of maximum degree d , whose slope number is at least $n^{\frac{1}{2} - \frac{1}{d-2} - o(1)}$. Since then this bound was improved for $d \geq 9$ by Dujmović, Suderman and Wood [DSW07], they showed that the slope number is at least $n^{1 - \frac{8+\epsilon}{d+4}}$. Note that for smaller d 's our bound is still the best.

Trivially, every graph of maximum degree two has slope number at most three. The case $d = 3$ was solved in our paper [KPPT08], in which we prove that every cubic graph* can be drawn with 5 slopes. This proof is presented in Section 6. Later, Muckamala and Szegedy [MSz07] showed that 4 slopes suffice if the graph is connected. However, for disconnected graphs, still five slopes is the best bound. This is because it cannot be guaranteed that different components are drawn with the same four slopes. It would be interesting to decide whether four fixed directions always suffice.

*A graph is *cubic* if its maximum degree is at most 3.

The case $d = 4$ remains an interesting open problem.

Conjecture 5.1. *The slope number of graphs with maximum degree 4 is unbounded.*

In Section 7 we investigate a similar notion, called *slope parameter*, first defined as follows by Ambrus, Barát, and P. Hajnal [ABH06]. Given a set P of points in the plane and a set Σ of slopes, define $G(P, \Sigma)$ as the graph on the vertex set P , in which two vertices $p, q \in P$ are connected by an edge if and only if the slope of the line pq belongs to Σ . The *slope parameter* $s(G)$ of G is the size of the smallest set Σ of slopes such that G is isomorphic to $G(P, \Sigma)$ for a suitable set of points P in the plane. This definition was motivated by the fact that all connections (edges) in an electrical circuit (graph) G can be easily realized by the overlay of $s(G)$ finely striped electrically conductive layers.

The slope parameter, $s(G)$, is closely related to the slope number. For instance, for triangle-free graphs, $s(G)$ is at least as large as the slope number of G , thus also bigger than the thickness and the geometric thickness. Indeed, in the drawing realizing the slope parameter, there are no three points on a line, so this drawing proves that the slope number is smaller or equal to the slope parameter, the only difference being that in case of the slope number it is not obligatory to connect two vertices if their slope is in Σ . The slope parameter of a triangle-free graph is also at least its *edge chromatic number*, $\chi'(G)$, as there can be at most one edge with the same slope from any vertex.

On the other hand, the slope parameter sharply differs from other parameters in the sense that the slope parameter of a complete graph on n vertices is *one*, while the thickness, the geometric thickness, and the slope number of K_n tend to infinity as $n \rightarrow \infty$. Jamison [J86] proved that the slope number of K_n is n .

Our main result in Section 7 is that the slope parameter of every cubic graph is also bounded. In our drawing no three vertices will be collinear, therefore as a corollary we obtain another proof for the fact that the slope number of cubic graphs is bounded, with a worse constant.

Finally, in Section 8 we investigate the *planar slope number* of bounded degree planar graphs. The planar slope number of a planar graph G is the smallest number s with the property that G has a straight-line drawing with non-crossing edges of at most s distinct slopes. We prove that any bounded degree planar graph has a bounded planar slope number. Then we investigate planar drawings where we allow one or two bends on each edge, in which cases we prove better bounds. For the exact statement of our results, see the beginning of Section 8.

5.1 A Lower Bound for the Slope Number of Graphs with Bounded Degree

This section is based on our paper with János Pach, Bounded-degree graphs can have arbitrarily large slope numbers [PP06].

If it creates no confusion, the vertex (edge) of G and the point (segment) representing it will be denoted by the same symbol. Dujmović et al. [DSW04] asked whether the slope parameter of bounded-degree graphs can be arbitrarily large. The following short argument shows that the answer is yes for graphs of degree at most *five*.

Define a “frame” graph F on the vertex set $\{1, \dots, n\}$ by connecting vertex 1 to 2 by an edge and connecting every $i > 2$ to $i - 1$ and $i - 2$. Adding a perfect matching M between these n points, we obtain a graph $G_M := F \cup M$. The number of different matchings is at least $(n/3)^{n/2}$. Let G denote the huge graph obtained by taking the union of disjoint copies of all G_M . Clearly, the maximum degree of the vertices of G is *five*. Suppose that G can be drawn using at most S slopes, and fix such a drawing.

For every edge $ij \in M$, label the points in G_M corresponding to i and j by the slope of ij in the drawing. Furthermore, label each frame edge ij ($|i - j| \leq 2$) by its slope. Notice that no two components of G receive the same labeling. Indeed, up to translation and scaling, the labeling of the edges uniquely determines the positions of the points representing the vertices of G_M . Then the labeling of the vertices uniquely determines the edges belonging to M . Therefore, the number of different possible labelings, which is $S^{|F|+n} < S^{3n}$, is an upper bound for the number of components of G . On the other hand, we have seen that the number of components (matchings) is at least $(n/3)^{n/2}$. Thus, for any S we obtain a contradiction, provided that n is sufficiently large. \square

With some extra care one can refine this argument to obtain

Theorem 5.2. *For any $d \geq 5$ and n , there exist a graph with n vertices of maximum degree d , whose slope number is at least $n^{\frac{1}{2} - \frac{1}{d-2} - o(1)}$.*

Proof. Now instead of a matching, we add to the frame F in every possible way a $(d - 4)$ -regular graph R on the vertex set $\{1, \dots, n\}$. Thus, we obtain at least $(cn/d)^{(d-4)n/2}$ different graphs $G_R := F \cup R$, each having maximum degree at most d (here $c > 0$ is a constant; see e.g. [BC78]). Suppose that each G_R can be drawn using S slopes $\sigma_1 < \dots < \sigma_S$. Now we cannot insist that these slopes are the same for all G_R , therefore, these numbers will be regarded as variables.

Fix a graph $G_R = F \cup R$ and one of its drawings with the above properties, in which vertex 1 is mapped into the origin and vertex 2 is mapped into a point whose x -coordinate is 1. Label every edge belonging to F by the symbol σ_k representing its slope. Furthermore, label each vertex j with a $(d - 4)$ -tuple of the σ_k s: with the symbols corresponding to the slopes of the $d - 4$ edges incident to j in R (with possible repetition). Clearly, the total number of possible labelings of the frame edges and vertices is at most $S^{|F|+(d-4)n} < S^{(d-2)n}$. Now the labeling itself does not necessarily identify the graph G_R , because we do not know the *actual* values of the slopes σ_k .

However, we can show that the number of different G_R s that receive the same labeling cannot be too large. To prove this, first notice that for a fixed labeling of the edges of the

frame, the coordinates of every vertex i can be expressed as the ratio of two polynomials of degree at most n in the variables $\sigma_1, \dots, \sigma_S$. Indeed, let $\sigma(ij)$ denote the label of $ij \in F$, and let $x(i)$ and $y(i)$ denote the coordinates of vertex i . Since, by assumption, we have $x(1) = y(1) = 0$ and $x(2) = 1$, we can conclude that $y(2) = \sigma(12)$. We have the following equations for the coordinates of 3:

$$y(3) - y(1) = \sigma(13)(x(3) - x(1)), \quad y(3) - y(2) = \sigma(23)(x(3) - x(2)).$$

Solving them, we obtain

$$x(3) = \frac{\sigma(12) - \sigma(23)}{\sigma(13) - \sigma(23)}, \quad y(3) = \frac{\sigma(13)(\sigma(12) - \sigma(23))}{\sigma(13) - \sigma(23)},$$

and so on. In particular, $x(i) = \frac{Q_i(\sigma_1, \dots, \sigma_S)}{Q'_i(\sigma_1, \dots, \sigma_S)}$, for suitable polynomials Q_i and Q'_i of degree at most $i - 1$. Moreover, Q'_j is a multiple of Q'_i for all $j > i$.

Since

$$x(i) - x(j) = \frac{Q_i \frac{Q'_j}{Q'_i} - Q_j}{Q'_j},$$

we can decide whether the image of i is to the left of the image of $j > i$, to the right of it, or they have the same x -coordinate, provided that we know the “sign pattern” of the polynomials $P'_{ij} := Q_i \frac{Q'_j}{Q'_i} - Q_j$ and Q'_j , i.e., we know which of them are positive, negative, or zero.

Now if we also know that σ_k is one of the labels associated with vertex i , the condition that the line connecting i and j has slope σ_k can be rewritten as

$$\frac{y(i) - y(j)}{x(i) - x(j)} - \sigma_k = \frac{\sigma(1i)Q_i Q'_j - \sigma(1j)Q'_i Q_j}{Q_i Q'_j - Q'_i Q_j} - \sigma_k = 0,$$

that is, as a polynomial equation $P_{ijk}(\sigma_1, \dots, \sigma_S) = 0$ of degree at most $2n$. For a fixed labeling of the frame edges and vertices, there are $d - 4$ labels k associated with a vertex i , so that the number of these polynomials P_{ijk} is at most $(d - 4)n(n - 1)$. Thus, together with the $\binom{n}{2} + n$ polynomials P'_{ij} and Q'_j , we have fewer than dn^2 polynomials, each of degree at most $2n$.

It is easy to verify that, for any fixed labeling, the sign pattern of these polynomials uniquely determines the graph G_R . (Observe that if the label of a vertex i is a $(d - 4)$ -tuple containing the symbol σ_k , then from the sign pattern of the above polynomials we can reconstruct the sequence of all vertices that belong to the line of slope σ_k passing through i , from left to right. From this sequence, we can select all elements whose label contains σ_k , and determine all edges of R along this line.)

To conclude the proof, we need the Thom-Milnor theorem [BPR03]: Given N polynomials in $S \leq N$ variables, each of degree at most $2n$, the number of sign patterns determined by them is at most $(CNn/S)^S$, for a suitable constant $C > 0$.

In our case, the number of graphs G_R is at most the number of labelings ($< S^{(d-2)n}$) multiplied by the maximum number of sign patterns of the above $< dn^2$ polynomials of degree at most $2n$. By the Thom-Milnor theorem, this latter quantity is smaller than

$(Cdn^3)^S$. Thus, the number of G_R s is at most $S^{(d-2)n}(Cdn^3)^S$. Comparing this to the lower bound $(cn/d)^{(d-4)n/2}$ stated in the first paragraph of the proof, we obtain that $S \geq n^{\frac{1}{2} - \frac{1}{d-2} - o(1)}$, as required. Now instead of a matching, we add to the frame F in every possible way a $(d-4)$ -regular graph R on the vertex set $\{1, \dots, n\}$. Thus, we obtain at least $(cn/d)^{(d-4)n/2}$ different graphs $G_R := F \cup R$, each having maximum degree at most d (here $c > 0$ is a constant; see e.g. [BC78]). Suppose that each G_R can be drawn using S slopes $\sigma_1 < \dots < \sigma_S$. Now we cannot insist that these slopes are the same for all G_R , therefore, these numbers will be regarded as variables.

Fix a graph $G_R = F \cup R$ and one of its drawings with the above properties, in which vertex 1 is mapped into the origin and vertex 2 is mapped into a point whose x -coordinate is 1. Label every edge belonging to F by the symbol σ_k representing its slope. Furthermore, label each vertex j with a $(d-4)$ -tuple of the σ_k s: with the symbols corresponding to the slopes of the $d-4$ edges incident to j in R (with possible repetition). Clearly, the total number of possible labelings of the frame edges and vertices is at most $S^{|F|+(d-4)n} < S^{(d-2)n}$. Now the labeling itself does not necessarily identify the graph G_R , because we do not know the *actual* values of the slopes σ_k .

However, we can show that the number of different G_R s that receive the same labeling cannot be too large. To prove this, first notice that for a fixed labeling of the edges of the frame, the coordinates of every vertex i can be expressed as the ratio of two polynomials of degree at most n in the variables $\sigma_1, \dots, \sigma_S$. Indeed, let $\sigma(ij)$ denote the label of $ij \in F$, and let $x(i)$ and $y(i)$ denote the coordinates of vertex i . Since, by assumption, we have $x(1) = y(1) = 0$ and $x(2) = 1$, we can conclude that $y(2) = \sigma(12)$. We have the following equations for the coordinates of 3:

$$y(3) - y(1) = \sigma(13)(x(3) - x(1)), \quad y(3) - y(2) = \sigma(23)(x(3) - x(2)).$$

Solving them, we obtain

$$x(3) = \frac{\sigma(12) - \sigma(23)}{\sigma(13) - \sigma(23)}, \quad y(3) = \frac{\sigma(13)(\sigma(12) - \sigma(23))}{\sigma(13) - \sigma(23)},$$

and so on. In particular, $x(i) = \frac{Q_i(\sigma_1, \dots, \sigma_S)}{Q'_i(\sigma_1, \dots, \sigma_S)}$, for suitable polynomials Q_i and Q'_i of degree at most $i-1$. Moreover, Q'_j is a multiple of Q'_i for all $j > i$.

Since

$$x(i) - x(j) = \frac{Q_i \frac{Q'_j}{Q'_i} - Q_j}{Q'_j},$$

we can decide whether the image of i is to the left of the image of $j > i$, to the right of it, or they have the same x -coordinate, provided that we know the “sign pattern” of the polynomials $P'_{ij} := Q_i \frac{Q'_j}{Q'_i} - Q_j$ and Q'_j , i.e., we know which of them are positive, negative, or zero.

Now if we also know that σ_k is one of the labels associated with vertex i , the condition that the line connecting i and j has slope σ_k can be rewritten as

$$\frac{y(i) - y(j)}{x(i) - x(j)} - \sigma_k = \frac{\sigma(1i)Q_i Q'_j - \sigma(1j)Q'_i Q_j}{Q_i Q'_j - Q'_i Q_j} - \sigma_k = 0,$$

that is, as a polynomial equation $P_{ijk}(\sigma_1, \dots, \sigma_S) = 0$ of degree at most $2n$. For a fixed labeling of the frame edges and vertices, there are $d - 4$ labels k associated with a vertex i , so that the number of these polynomials P_{ijk} is at most $(d - 4)n(n - 1)$. Thus, together with the $\binom{n}{2} + n$ polynomials P'_{ij} and Q'_j , we have fewer than dn^2 polynomials, each of degree at most $2n$.

It is easy to verify that, for any fixed labeling, the sign pattern of these polynomials uniquely determines the graph G_R . (Observe that if the label of a vertex i is a $(d - 4)$ -tuple containing the symbol σ_k , then from the sign pattern of the above polynomials we can reconstruct the sequence of all vertices that belong to the line of slope σ_k passing through i , from left to right. From this sequence, we can select all elements whose label contains σ_k , and determine all edges of R along this line.)

To conclude the proof, we need the Thom-Milnor theorem [BPR03]: Given N polynomials in $S \leq N$ variables, each of degree at most $2n$, the number of sign patterns determined by them is at most $(CNn/S)^S$, for a suitable constant $C > 0$.

In our case, the number of graphs G_R is at most the number of labelings ($< S^{(d-2)n}$) multiplied by the maximum number of sign patterns of the above $< dn^2$ polynomials of degree at most $2n$. By the Thom-Milnor theorem, this latter quantity is smaller than $(Cdn^3)^S$. Thus, the number of G_R s is at most $S^{(d-2)n}(Cdn^3)^S$. Comparing this to the lower bound $(cn/d)^{(d-4)n/2}$ stated in the first paragraph of the proof, we obtain that $S \geq n^{\frac{1}{2} - \frac{1}{d-2} - o(1)}$, as required. \square

6 Drawing Cubic Graphs with at most Five Slopes

This section is based on our paper with Balázs Keszegh, János Pach and Géza Tóth, Drawing cubic graphs with at most five slopes [KPPT08].

Our main result is

Theorem 6.1. *Every graph of maximum degree at most three has slope number at most five.*

Our terminology is somewhat unorthodox: by the *slope* of a line ℓ , we mean the angle α modulo π such that a counterclockwise rotation through α takes the x -axis to a position parallel to ℓ . The slope of an edge (segment) is the slope of the line containing it. In particular, the slopes of the lines $y = x$ and $y = -x$ are $\pi/4$ and $-\pi/4$, and they are called *Northeast* (or *Southwest*) and *Northwest* (or *Southeast*) lines, respectively.

For any two points $p_1 = (x_1, y_1), p_2 = (x_2, y_2) \in \mathbf{R}^2$, we say that p_2 is *to the North* (or *to the South* of p_1 if $x_2 = x_1$ and $y_2 > y_1$ (or $y_2 < y_1$)). Analogously, we say that p_2 is *to the Northeast* (to the *Northwest*) of p_1 if $y_2 > y_1$ and p_1p_2 is a *Northeast* (*Northwest*) line. Directions are often abbreviated by their first letters: N, NE, E, SE, etc. These four directions are referred to as *basic*. That is, a line ℓ is said to be of one of the four basic directions if ℓ is parallel to one of the axes or to one of the NE and NW lines $y = x$ and $y = -x$.

The main tool of our proof is the following result of independent interest.

Theorem 6.2. *Let G be a connected graph that is not a cycle and whose every vertex has degree at most three. Suppose that G has at least one vertex of degree less than three, and denote by v_1, \dots, v_m the vertices of degree at most two ($m \geq 1$).*

Then, for any sequence x_1, x_2, \dots, x_m of real numbers, linearly independent over the rationals, G has a straight-line drawing with the following properties:

- (1) Vertex v_i is mapped into a point with x -coordinate $x(v_i) = x_i$ ($1 \leq i \leq m$);
- (2) The slope of every edge is $0, \pi/2, \pi/4$, or $-\pi/4$.
- (3) No vertex is to the North of any vertex of degree two.
- (4) No vertex is to the North or to the Northwest of any vertex of degree one.

Before this theorem only the following special cases were known.

It was shown by Dujmović et al. [DESW07] that every planar graph with maximum degree three has a drawing with non-crossing straight-line edges of at most three different slopes, except that three edges of the outer-face may have a bend.

Max Engelstein [E05], a student from Stuyvesant High School, New York has shown that every graph of maximum degree *three* that has a Hamiltonian cycle can be drawn with edges of at most *five* different slopes.

6.1 Embedding Cycles

Let C be a straight-line drawing of a cycle in the plane. A vertex v of C is said to be a *turning point* if the slopes of the two edges meeting at v are not the same.

We start with two simple auxiliary statements.

Lemma 6.3. *Let C be a straight-line drawing of a cycle such that the slope of every edge is 0 , $\pi/4$, or $-\pi/4$. Then the x -coordinates of the vertices of C are not independent over the rational numbers.*

Moreover, there is a vanishing linear combination of the x -coordinates of the vertices, with as many nonzero (rational) coefficients as many turning points C has.

Proof. Let v_1, v_2, \dots, v_n denote the vertices of C in cyclic order ($v_{n+1} = v_1$). Let $x(v_i)$ and $y(v_i)$ be the coordinates of v_i . For any i ($1 \leq i \leq n$), we have $y(v_{i+1}) - y(v_i) = \lambda_i(x(v_{i+1}) - x(v_i))$, where $\lambda_i = 0, 1$, or -1 , depending on the slope of the edge $v_i v_{i+1}$. Adding up these equations for all i , the left-hand sides add up to zero, while the sum of the right-hand sides is a linear combination of the numbers $x(v_1), x(v_2), \dots, x(v_n)$ with integer coefficients of absolute value at most *two*.

Thus, we are done with the first statement of the lemma, unless all of these coefficients are zero. Obviously, this could happen if and only if $\lambda_1 = \lambda_2 = \dots = \lambda_n$, which is impossible, because then all points of C would be collinear, contradicting our assumption that in a proper *straight-line drawing* no edge is allowed to pass through any vertex other than its endpoints.

To prove the second statement, it is sufficient to notice that the coefficient of $x(v_i)$ vanishes if and only if v_i is not a turning point. \square

Lemma 6.3 shows that Theorem 6.2 does not hold if G is a cycle. Nevertheless, according to the next claim, cycles satisfy a very similar condition. Observe, that the main difference is that here we have an exceptional vertex, denoted by v_0 .

Lemma 6.4. *Let C be a cycle with vertices v_0, v_1, \dots, v_m , in this cyclic order.*

Then, for any real numbers x_1, x_2, \dots, x_m , linearly independent over the rationals, C has a straight-line drawing with the following properties:

- (1) Vertex v_i is mapped into a point with x -coordinate $x(v_i) = x_i$ ($1 \leq i \leq m$);
- (2) The slope of every edge is $0, \pi/4$, or $-\pi/4$.
- (3) No vertex is to the North of any other vertex.
- (4) No vertex has a larger y -coordinate than $y(v_0)$.

Proof. We can assume without loss of generality that $x_2 > x_1$. Place v_1 at any point $(x_1, 0)$ of the x -axis. Assume that for some $i < m$, we have already determined the positions of v_1, v_2, \dots, v_i , satisfying conditions (1)–(3). If $x_{i+1} > x_i$, then place v_{i+1} at the (unique) point Southeast of v_i , whose x -coordinate is x_{i+1} . If $x_{i+1} < x_i$, then put v_{i+1} at the point West of x_i , whose x -coordinate is x_{i+1} . Clearly, this placement of v_{i+1} satisfies (1)–(3), and the segment $v_i v_{i+1}$ does not pass through any point v_j with $j < i$.

After m steps, we obtain a noncrossing straight-line drawing of the path $v_1 v_2 \dots v_m$, satisfying conditions (1)–(3). We still have to find a right location for v_0 . Let R_W and R_{SE} denote the rays (half-lines) starting at v_1 and pointing to the West and to the Southeast. Further, let R be the ray starting at v_m and pointing to the Northeast. It follows from the construction that all points v_2, \dots, v_m lie in the convex cone below the x -axis, enclosed by the rays R_W and R_{SE} .

Place v_0 at the intersection point of R and the x -axis. Obviously, the segment $v_m v_0$ does not pass through any other vertex v_j ($0 < j < m$). Otherwise, we could find a drawing of the cycle $v_j v_{j+1} \dots v_m$ with slopes $0, \pi/4$, and $-\pi/4$. By Lemma 6.3, this would imply

that the numbers x_j, x_{j+1}, \dots, x_m are *not* independent over the rationals, contradicting our assumption. It is also clear that the horizontal segment v_0v_1 does not pass through any vertex different from its endpoints because all other vertices are below the horizontal line determined by v_0v_1 . Hence, we obtain a proper straight-line drawing of C satisfying conditions (1),(2), and (4).

It remains to verify (3). The only thing we have to check is that $x(v_0)$ does not coincide with any other $x(v_i)$. Suppose it does, that is, $x(v_0) = x(v_i) = x_i$ for some $i > 0$. By the second statement of Lemma 6.3, there is a vanishing linear combination

$$\lambda_0x(v_0) + \lambda_1x_1 + \lambda_2x_2 + \dots + \lambda_mx_m = 0$$

with rational coefficients λ_i , where the number of nonzero coefficients is at least the number of turning points, which cannot be smaller than *three*. Therefore, if in this linear combination we replace $x(v_0)$ by x_i , we still obtain a nontrivial rational combination of the numbers x_1, x_2, \dots, x_m . This contradicts our assumption that these numbers are independent over the rationals. \square

6.2 Subcubic Graphs - Proof of Theorem 6.2

First we settle Theorem 6.2 in a special case.

Lemma 6.5. *Let $m, k \geq 2$ and let G be a graph consisting of two disjoint cycles, $C = \{v_0, v_1, \dots, v_m\}$ and $C' = \{v'_0, v'_1, \dots, v'_m\}$, connected by a single edge $v_0v'_0$.*

Then, for any sequence $x_1, x_2, \dots, x_m, x'_1, x'_2, \dots, x'_k$ of real numbers, linearly independent over the rationals, G has a straight-line drawing satisfying the following conditions:

- (1) *The vertices v_i and v'_j are mapped into points with x -coordinates $x(v_i) = x_i$ ($1 \leq i \leq m$) and $x(v'_j) = x'_j$ ($1 \leq j \leq k$).*
- (2) *The slope of every edge is $0, \pi/2, \pi/4$, or $-\pi/4$.*
- (3) *No vertex is to the North of any vertex of degree two.*

Proof. Apply Lemma 6.4 to cycle C with vertices v_0, v_1, \dots, v_m and with assigned x -coordinates x_1, x_2, \dots, x_m , and analogously, to the cycle C' , with vertices v'_0, v'_1, \dots, v'_k and with assigned x -coordinates x'_1, x'_2, \dots, x'_k . For simplicity, the resulting drawings are also denoted by C and C' .

Let x_0 and x'_0 denote the x -coordinates of $v_0 \in C$ and $v'_0 \in C'$. It follows from Lemma 6.3 that x_0 is a linear combination of x_1, x_2, \dots, x_m , and x'_0 is a linear combination of x'_1, x'_2, \dots, x'_k with rational coefficients. Therefore, if $x_0 = x'_0$, then there is a nontrivial linear combination of $x_1, x_2, \dots, x_m, x'_1, x'_2, \dots, x'_k$ that gives 0, contradicting the assumption that these numbers are independent over the rationals. Thus, we can conclude that $x_0 \neq x'_0$. Assume without loss of generality that $x_0 < x'_0$. Reflect C' about the x -axis, and shift it in the vertical direction so that v'_0 ends up to the Northeast from v_0 . Clearly, we can add the missing edge $v_0v'_0$. Let D denote the resulting drawing of G . We claim that D meets all the requirements of the Theorem. Conditions (1), (2), and (3) are obviously satisfied, we only have to check that no vertex lies in the interior of an edge. It follows from Lemma 6.4 that the y -coordinates of v_1, \dots, v_m are all smaller than or equal to the y -coordinate of v_0 and the y -coordinates of v'_1, \dots, v'_k are all greater than or equal to the

y -coordinate of v'_0 . We also have $y(v_0) < y(v'_0)$. Therefore, there is no vertex in the interior of $v_0v'_0$. Moreover, no edge of C (resp. C') can contain any vertex of v'_0, v'_1, \dots, v'_k (resp. v_0, v_1, \dots, v_m) in its interior. \square

The rest of the proof is by induction on the number of vertices of G . The statement is trivial if the number of vertices is at most *two*. Suppose that we have already established Theorem 6.2 for all graphs with fewer than n vertices.

Suppose that G has n vertices, it is not a cycle and not the union of two cycles connected by one edge. Let v_1, v_2, \dots, v_m be the vertices of G with degree less than *three*, and let the x -coordinates assigned to them be x_1, x_2, \dots, x_m .

We distinguish several cases.

Case 1: G has a vertex of degree one.

Assume, without loss of generality, that v_1 is such a vertex. If G has no vertex of degree *three*, then it consists of a simple path $P = v_1v_2 \dots v_m$, say. Place v_m at the point $(x_m, 0)$. In general, assuming that v_{i+1} has already been embedded for some $i < m$, and $x_i < x_{i+1}$, place v_i at the point West of v_{i+1} , whose x -coordinate is x_i . If $x_i > x_{i+1}$, then put v_i at the point Northeast of v_{i+1} , whose x -coordinate is x_i . The resulting drawing of $G = P$ meets all the requirements of the theorem. To see this, it is sufficient to notice that if v_j would be Northwest of v_m for some $j < m$, then we could apply Lemma 6.3 to the cycle $v_jv_{j+1} \dots v_m$, and conclude that the numbers x_j, x_{j+1}, \dots, x_m are dependent over the rationals. This contradicts our assumption.

Assume next that v_1 is of degree *one*, and that G has at least one vertex of degree *three*. Suppose without loss of generality that $v_1v_2 \dots v_kv$ is a path in G , whose internal vertices are of degree *two*, but the degree of w is *three*. Let G' denote the graph obtained from G by removing the vertices v_1, v_2, \dots, v_k . Obviously, G' is a connected graph, in which the degree of w is *two*.

If G' is a *cycle*, then apply Lemma 6.4 to $C = G'$ with w playing the role of the vertex v_0 which has no preassigned x -coordinate. We obtain an embedding of G' with edges of slopes $0, \pi/4$, and $-\pi/4$ such that $x(v_i) = x_i$ for all $i > k$ and there is no vertex to the North, to the Northeast, or to the Northwest of w . By Lemma 6.3, the numbers $x(w), x_{k+1}, \dots, x_m$ are not independent over the rationals. Therefore, $x(w) \neq x_k$, so we can place v_k at the point to the Northwest or to the Northeast of w , whose x -coordinate is x_k , depending on whether $x(w) > x_k$ or $x(w) < x_k$. After this, embed v_{k-1}, \dots, v_1 , in this order, so that v_i is either to the Northeast or to the West of v_{i+1} and $x(v_i) = x_i$. According to property (4) in Lemma 6.3, the path $v_1v_2 \dots v_k$ lies entirely above G' , so that no point of G can lie to the North or to the Northwest of v_1 .

If G' is *not a cycle*, then use the induction hypothesis to find an embedding of G' that satisfies all conditions of Theorem 6.2, with $x(w) = x_k$ and $x(v_i) = x_i$ for every $i > k$. Now place v_k very far from w , to the North of it, and draw v_{k-1}, \dots, v_1 , in this order, in precisely the same way as in the previous case. Now if v_k is far enough, then none of the points v_k, v_{k-1}, \dots, v_1 is to the Northwest or to the Northeast of any vertex of G' . It remains to check that condition (4) is true for v_1 , but this follows from the fact that there is no point of G whose y -coordinate is larger than that of v_1 .

From now on, we can and will assume that G has *no vertex of degree one*.

A graph with *four* vertices and *five* edges between them is said to be a Θ -graph.

Case 2: G contains a Θ -subgraph.

Suppose that G has a Θ -subgraph with vertices a, b, c, d , and edges ab, bc, ac, ad, bd . If neither c nor d has a third neighbor, then G is identical to this graph, which can easily be drawn in the plane with all conditions of the theorem satisfied.

If c and d are connected by an edge, then all four points of the Θ -subgraph have degree *three*, so that G has no other vertices. So G is a complete graph of four vertices, and it has a drawing that meets the requirements.

Suppose that c and d have a common neighbor $e \neq a, b$. If e has no further neighbor, then a, b, c, d, e are the only vertices of G , and again we can easily find a proper drawing. Thus, we can assume that e has a third neighbor f . By the induction hypothesis, $G' = G \setminus \{a, b, c, d, e\}$ has a drawing satisfying the conditions of Theorem 6.2. In particular, no vertex of G' is to the North of f (and to the Northwest of f , provided that the degree of f in G' is *one*). Further, consider a drawing H of the subgraph of G induced by the vertices a, b, c, d, e , which satisfies the requirements. We distinguish two subcases.

If the degree of f in G' is *one*, then take a very small *homothetic* copy of H (i.e., similar copy in parallel position), and rotate it about e in the clockwise direction through $3\pi/4$. There is no point of this drawing, denoted by H' , to the Southeast of e , so that we can translate it into a position in which e is to the Northwest of $f \in V(G')$ and very close to it. Connecting now e to f , we obtain a drawing of G satisfying the conditions. Note that it was important to make H' very small and to place it very close to f , to make sure that none of its vertices is to the North of any vertex of G' whose degree is at most *two*, or to the Northwest of any vertex of degree *one* (other than f).

If the degree of f in G' is *two*, then we follow the same procedure, except that now H' is a small copy of H , rotated by π . We translate H' into a position in which e is to the North of f , and connect e to f by a vertical segment. It is again clear that the resulting drawing of G meets the requirements in Theorem 6.2. Thus, we are done if c and d have a common neighbor e .

Suppose now that only one of c and d has a third neighbor, different from a and b . Suppose, without loss of generality, that this vertex is c , so that the degree of d is *two*. Then in $G' = G \setminus \{a, b, d\}$, the degree of c is *one*. Apply the induction hypothesis to G' so that the x -coordinate originally assigned to d is now assigned to c (which had no preassigned x -coordinate in G). In the resulting drawing, we can easily reinsert the remaining vertices, a, b, d , by adding a very small square whose lowest vertex is at c and whose diagonals are parallel to the coordinate axes. The highest vertex of this square will represent d , and the other two vertices will represent a and b .

We are left with the case when both c and d have a third neighbor, other than a and b , but these neighbors are different. Denote them by c' and d' , respectively. Create a new graph G' from G , by removing a, b, c, d and adding a new vertex v , which is connected to c' and d' . Draw G' using the induction hypothesis, and reinsert a, b, c, d in a small neighborhood of v so that they form the vertex set of a very small square with diagonal ab . (See Figure 24.) As before, we have to choose this square sufficiently small to make sure that a, b, c, d are not to the North of any vertex $w \neq c', d', v$ of G' , whose degree is at most *two*, or to the Northwest of any vertex of degree *one*. Thus, we are done if G has a Θ -subgraph.

So, from now on we assume that G has *no* Θ -subgraph.




Figure 24: Replacing v by Θ .

Case 3: G has no cycle that passes through a vertex of degree two.

Since G is not three-regular, it contains at least one vertex of degree *two*. Consider a decomposition of G into two-connected blocks and edges. If a block contains a vertex of degree *two*, then it consists of a single edge. The block decomposition has a treelike structure, so that there is a vertex w of degree *two*, such that G can be obtained as the union of two graphs, G_1 and G_2 , having only the vertex w in common, and there is no vertex of degree *two* in G_1 .

By the induction hypothesis, for any assignment of rationally independent x -coordinates to all vertices of degree less than *three*, G_1 and G_2 have proper straight-line embeddings (drawings) satisfying conditions (1)–(4) of the theorem. The only vertex of G_1 with a preassigned x -coordinate is w . Applying a vertical translation, if necessary, we can achieve that in both drawings w is mapped into the same point. Using the induction hypothesis, we obtain that in the union of these two drawings, there is no vertex in G_1 or G_2 to the North or to the Northwest of w , because the degree of w in G_1 and G_2 is *one* (property (4)). This is stronger than what we need: indeed, in G the degree of w is *two*, so that we require only that there is no point of G to the North of w (property (3)).

The superposition of the drawings of G_1 and G_2 satisfies all conditions of the theorem. Only two problems may occur:

1. A vertex of G_1 may end up at a point to the North of a vertex of G_2 with degree *two*.
2. The (unique) edges in G_1 and G_2 , incident to w , may partially overlap.

Notice that both of these events can be avoided by enlarging the drawing of G_1 , if necessary, from the point w , and rotating it about w by $\pi/4$ in the clockwise direction. The latter operation is needed only if problem 2 occurs. This completes the induction step in the case when G has no cycle passing through a vertex of degree *two*.

It remains to analyze the last case.

Case 4: G has a cycle passing through a vertex of degree two.

By assumption, G itself is not a cycle. Therefore, we can also find a *shortest* cycle C whose vertices are denoted by v, u_1, \dots, u_k , in this order, where the degree of v is *two* and the degree of u_1 is *three*. The length of C is $k + 1$.

It follows from the minimality of C that u_i and u_j are not connected by an edge of G , for any $|i - j| > 1$. Moreover, if $|i - j| > 2$, then u_i and u_j do not even have a common neighbor ($1 \leq i \neq j \leq k$). This implies that any vertex $v \in V(G \setminus C)$ has at most *three* neighbors on C , and these neighbors must be consecutive on C . However, *three* consecutive vertices of C , together with their common neighbor, would form a Θ -subgraph in G (see Case 2). Hence, we can assume that every vertex belonging to $G \setminus C$ is joined to at most *two* vertices on C .

Let B_i denote the set of all vertices of $G \setminus C$ that have precisely i neighbors on C ($i = 0, 1, 2$). Thus, we have $V(G \setminus C) = B_0 \cup B_1 \cup B_2$. Further, $B_1 = B_1^2 \cup B_1^3$, where an element of B_1 belongs to B_1^2 or B_1^3 , according to whether its degree in G is *two* or *three*.

Consider the list v_1, v_2, \dots, v_m of all vertices of G with degree *two*. (Recall that we have already settled the case when G has a vertex of degree *one*.) Assume without loss of generality that $v_1 = v$ and that v_i belongs to C if and only if $1 \leq i \leq j$ for some $j \leq m$.

Let \mathbf{x} denote the *assignment* of x -coordinates to the vertices of G with degree *two*, that is, $\mathbf{x} = (x(v_1), x(v_2), \dots, x(v_m)) = (x_1, x_2, \dots, x_m)$. Given G , C , \mathbf{x} , and a real parameter L , we define the following so-called EMBEDDING PROCEDURE(G, C, \mathbf{x}, L) to construct a drawing of G that meets all requirements of the theorem, and satisfies the additional condition that the y -coordinate of every vertex of C is at least L higher than the y -coordinates of all other vertices of G .

STEP 1: If $G' := G \setminus C$ is *not* a cycle, then construct recursively a drawing of $G' := G \setminus C$ satisfying the conditions of Theorem 6.2 with the assignment \mathbf{x}' of x -coordinates $x(v_i) = x_i$ for $j < i \leq m$, and $x(u'_1) = x_1$, where u'_1 is the unique vertex in $G \setminus C$, connected by an edge to $u_1 \in V(C)$.

If $G' = G \setminus C$ is a cycle, then, by assumption, there are at least two edges between C and G' . One of them connects u_1 to u'_1 . Let $u_\alpha u'_\alpha$ be another such edge, where $u_\alpha \in C$ and $u'_\alpha \in G'$. Since the maximum degree is three, $u'_1 \neq u'_\alpha$. Now construct recursively a drawing of $G' := G \setminus C$ satisfying the conditions of Lemma 6.4, with the assignment \mathbf{x}' of x -coordinates $x(v_i) = x_i$ for $j < i \leq m$, $x(u'_1) = x_1$, and with exceptional vertex u'_α .

STEP 2: For each element of $B_1^2 \cup B_2$, take two rays starting at this vertex, pointing to the Northwest and to the North. Further, take a vertical ray pointing to the North from each element of B_1^3 and each element of the set $B_{\mathbf{x}} := \{(x_2, 0), (x_3, 0), \dots, (x_j, 0)\}$. Let \mathcal{R} denote the set of all of these rays. Choose the x -axis above all points of G' and all intersection points between the rays in \mathcal{R} .

For any u_h ($1 \leq h \leq k$) whose degree in G is *three*, define $N(u_h)$ as the unique neighbor of u_h in $G \setminus C$. If u_h has degree *two* in G , then $u_h = v_i$ for some $1 \leq i \leq j$, and let $N(u_h)$ be the point $(x_i, 0)$.

STEP 3: Recursively place u_1, u_2, \dots, u_k on the rays belonging to \mathcal{R} , as follows. Place u_1 on the vertical ray starting at $N(u_1) = u'_1$ such that $y(u_1) = L$. Suppose that for some $i < k$ we have already placed u_1, u_2, \dots, u_i , so that $L \leq y(u_1) \leq y(u_2) \leq \dots \leq y(u_i)$ and there is no vertex to the West of u_i . Next we determine the place of u_{i+1} .

If $N(u_{i+1}) \in B_1^2$, then let $r \in \mathcal{R}$ be the ray starting at $N(u_{i+1})$ and pointing to the Northwest. If $N(u_{i+1}) \in B_1^3 \cup B_{\mathbf{x}}$, let $r \in \mathcal{R}$ be the ray starting at $N(u_{i+1})$ and pointing to the North. In both cases, place u_{i+1} on r : if u_i lies on the left-hand side of r , then put u_{i+1} to the Northeast of u_i ; otherwise, put u_{i+1} to the West of u_i .

If $N(u_{i+1}) \in B_2$, then let $r \in \mathcal{R}$ be the ray starting at $N(u_{i+1})$ and pointing to the North, or, if we have already placed a point on this ray, let r be the other ray from $N(u_{i+1})$, pointing to the Northwest, and proceed as before.




Figure 25: Recursively place u_1, u_2, \dots, u_k on the rays belonging to \mathcal{R} .

STEP 4: Suppose we have already placed u_k . It remains to find the right position for $u_0 := v$, which has only two neighbors, u_1 and u_k . Let r be the ray at u_1 , pointing to the North. If u_k lies on the left-hand side of r , then put u_0 on r to the Northeast of u_k ; otherwise, put u_0 on r , to the West of u_k .

During the whole procedure, we have never placed a vertex on any edge, and all other conditions of Theorem 6.2 are satisfied \square .

Remark that the y -coordinates of the vertices $u_0 = v, u_1, \dots, u_k$ are at least L higher than the y -coordinates of all vertices in $G \setminus C$. If we fix G, C , and \mathbf{x} , and let L tend to infinity, the coordinates of the vertices given by the above EMBEDDING PROCEDURE(G, C, \mathbf{x}, L) change continuously.




Figure 26: Find the right position for u_0 .

6.3 Cubic Graphs - Proof of Theorem 6.1

We are going to show that any graph G with maximum degree three permits a straight-line drawing using only the four basic directions (of slopes $0, \pi/2, \pi/4$, and $-\pi/4$), and perhaps one further direction, which is almost vertical and is used for at most one edge in each connected component of G .

Denote the connected components of G by G_1, G_2, \dots, G_t . If a component G_s is not three-regular, or if it is a complete graph with *four* vertices, then, by Theorem 6.2, it can be drawn using only the four basic directions. If G_s has a Θ -subgraph, one can argue in the same way as in Case 2 of the proof of Theorem 6.2: Embed recursively the rest of the graph, and attach to it a small copy of this subgraph such that all edges of the Θ -subgraph, as well as the edges used for the attachment, are parallel to one of the four basic directions. Actually, in this case, G_s itself can be drawn using the four basic directions, so the fifth direction is not needed.

Thus, in the rest of the proof we can assume that G_s is three-regular, it has more than *four* vertices, and it contains no Θ -subgraph. For simplicity, we drop the subscript and we write G instead of G_s . Choose a shortest cycle $C = u_0 u_1 \dots u_k$ in G . Each vertex of C has precisely one neighbor in $G \setminus C$. On the other hand, as in the proof of the last case of Theorem 6.2, all vertices in $G \setminus C$ have at most two neighbors in C .

We distinguish two cases.

Case 1. $G \setminus C$ is a cycle. Since G is three-regular, C and $G \setminus C$ are of the same size and the remaining edges of G form a matching between the vertices of C and the vertices of $G \setminus C$. For any i , $0 \leq i \leq k$, let u'_i denote the vertex of $G \setminus C$ which is connected to u_i . Denote the vertices of $G \setminus C$ by v_0, v_1, \dots, v_k , in cyclic order, so that $v_1 = u'_1$. Then we have $v_i = u'_0$, for some $i > 1$. Apply Lemma 6.4 to $G \setminus C$ with a rationally independent assignment \mathbf{x} of x -coordinates to the vertices v_1, \dots, v_k , such that $x(v_1) = 1$, $x(v_i) = \sqrt{2}$, and the x -coordinates of the other vertices are all greater than $\sqrt{2}$. (Recall that v_0 is an exceptional vertex with no assigned x -coordinate.) It is not hard to see that if we follow the construction described in the proof of Lemma 6.4, we also have $x(v_0) > \sqrt{2}$.

Case 2. $G \setminus C$ is *not* a cycle. Let u'_0 denote the neighbor of u_0 in $G \setminus C$. Since G has no Θ -subgraph, u'_0 cannot be joined to both u_1 and u_k . Assume without loss of generality that u'_0 is not connected to u_1 . Let u'_1 denote the neighbor of u_1 in $G \setminus C$.

Fix a rationally independent assignment \mathbf{x} of x -coordinates to the vertices of degree at most *two* in $G \setminus C$, such that $x(u'_0) = \sqrt{2}$, $x(u'_1) = 1$, and the x -coordinates of the other vertices are all greater than $\sqrt{2}$. Consider a drawing of $G \setminus C$, meeting the requirements of Theorem 6.2.

Now in both cases, let G' denote the graph obtained from G after the removal of the edge $u_0 u'_0$. Clearly $G \setminus C = G' \setminus C$, and for any L , EMBEDDING PROCEDURE(G', C, \mathbf{x}, L) gives a drawing of G' . It follows from the construction, that $x(u_0) = x(u_1) = x(u'_1) = 1$, $x(u'_0) = \sqrt{2}$. Therefore, for any sufficiently small $\varepsilon > 0$ there is an $L > 0$ such that EMBEDDING PROCEDURE(G', C, \mathbf{x}, L) gives a drawing of G' , in which the slope of the line connecting u_0 and u'_0 is $\frac{\pi}{2} + \varepsilon$.

We want to add the segment $u_0 u'_0$ to this drawing. Since there is no vertex with x -coordinate between 1 and $\sqrt{2}$, the segment $u_0 u'_0$ cannot pass through any vertex of G .

Summarizing: if ε is sufficiently small (that is, if L is sufficiently large), then each component of the graph has a proper drawing in which all edges are of one of the four basic directions, with the exception of at most one edge whose slope is $\frac{\pi}{2} + \varepsilon$. If we choose an $\varepsilon > 0$ that works for all components, then the whole graph can be drawn using only at most *five* directions. This concludes the proof of Theorem 6.1. \square

6.4 Algorithm

Based on the proof, it is not hard to design an algorithm to find a proper drawing, in quadratic time.

First, if our graph is a circle, we have no problem drawing it in $O(n)$ steps. If our graph has a vertex of degree *one* then the procedure of Case 1 of the proof of Theorem 6.2 requires at most $O(m)$ time when we reinsert v_1, \dots, v_m .

We can check if our graph has any Θ -subgraph in $O(n)$ time. If we find one, we can proceed by induction as in Case 2 of the proof of Theorem 6.2. We can reinsert the Θ -subgraph as described in Case 2 in $O(1)$ time.

Now assume that we have a vertex v of degree *two*. Execute a breadth first search from any vertex, and take a *minimal* vertex of degree two, that is, a vertex v of degree two, all of whose descendants are of degree three. If there is an edge in the graph connecting a descendant of v with a non-descendant, then there is a cycle through v ; we can find a minimal one with a breadth first search from it and proceed as in Case 4. Otherwise, v can play the role of w in Case 3, and we can proceed recursively.

Finally, if the graph is 3-regular, then we draw each component separately, except the last step, when we have to pick an ϵ small enough simultaneously for all components, this takes $O(n)$ steps. We only have to find the greatest slope and pick an ϵ such that $\frac{\pi}{2} + \epsilon$ is even steeper.

We believe that this algorithm is far from being optimal. It may perform a breadth first search for each induction step, which is probably not necessary. One may be able to replace this step by repeatedly updating the results of the first search. We cannot even rule out that the problem can be solved in linear time.

7 Slope Parameter of Cubic Graphs

This section is based on our paper with Balázs Keszegh, János Pach and Géza Tóth, Cubic graphs have bounded slope parameter [KPPT10].

Let us recall the definition of the *slope parameter*. Given a set P of points in the plane and a set Σ of slopes, define $G(P, \Sigma)$ as the graph on the vertex set P , in which two vertices $p, q \in P$ are connected by an edge if and only if the slope of the line pq belongs to Σ . The *slope parameter* $s(G)$ of G is the size of the smallest set Σ of slopes such that G is isomorphic to $G(P, \Sigma)$ for a suitable set of points P in the plane.

Any graph G of maximum degree *two* splits into vertex-disjoint cycles, paths, and possibly isolated vertices. Hence, for such graphs we have $s(G) \leq 3$. In contrast, as was shown by Barát *et al.* [BMW06], for any $d \geq 5$, there exist graphs of maximum degree d , whose slope parameters are arbitrarily large.

Remember that a graph is said to be *cubic* if the degree of each of its vertices is at most *three*. A cubic graph is *subcubic* if each of its connected components has a vertex of degree smaller than *three*.

The main result of this section is

Theorem 7.1. *Every cubic graph has slope parameter at most seven.*

This theorem is not likely to be tight. The best lower bound we are aware of is *four*. This bound is attained, for example, for the 8-vertex subcubic graph that can be obtained from the graph formed by the edges of a 3-dimensional cube by deleting one of its edges.

We will refer to the angles $i\pi/5$, $0 \leq i \leq 4$, as the *five basic slopes*. We start by proving the following statement, which constitutes the first step of the proof of Theorem 7.1.

Theorem 7.2. *Every subcubic graph has slope parameter at most five. Moreover, this can be realized by a straight-line drawing such that no three vertices are on a line and each edge has one of the five basic slopes.*

Using the fact that in the drawing guaranteed by Theorem 7.2 no *three* vertices are collinear, we can also conclude that the slope *number* of every subcubic graph is at most *five*. In the Section 6, however, it was shown that this number is at most *four* and for cubic graphs it is at most *five*.

7.1 Subcubic Graphs - Proof of Theorem 7.2

The proof is by induction on the number of vertices of the graph. Clearly, the statement holds for graphs with fewer than *three* vertices. Let n be fixed and suppose that we have already established the statement for graphs with fewer than n vertices. Let G be a subcubic graph of n vertices. We can assume that G is connected, otherwise we can draw each of its connected components separately and translate the resulting drawings through suitable vectors so that no two points in distinct components determine a line of basic slope.

To obtain a straight-line drawing of G , we have to find proper locations for its vertices. At each inductive step, we start with a drawing of a subgraph of G satisfying the conditions of Theorem 7.2 and extend it by adding a vertex. At a given stage of the procedure, for any vertex v that has already been added, consider the (basic) slopes of all edges adjacent

to v that have already been drawn, and let $\mathbf{sl}(v)$ denote the set of integers $0 \leq i < 5$ for which $i\pi/5$ is such a slope. That is, at the beginning $\mathbf{sl}(v)$ is undefined, then it gets defined, and later it may change (expand). Analogously, for any edge uv of G , denote by $\mathbf{sl}(uv)$ the integer $0 \leq i < 5$ for which the slope of uv is $i\pi/5$.

Case 1: G has a vertex of degree one.

Assume without loss of generality, that v is a vertex of degree *one*, and let w denote its only neighbor. Deleting v from G , the degree of w in the resulting graph G' is at most *two*. Therefore, by the induction hypothesis, G' has a drawing meeting the requirements. As w has degree at most *two*, there is a basic slope σ such that no other vertex of G' lies on the line ℓ of slope σ that passes through w . Draw all *five* lines of basic slopes through each vertex of G' . These lines intersect ℓ in finitely many points. We can place v at any other point of ℓ , to obtain a proper drawing of G .

From now on, assume that G has no vertex of degree one.

Case 2: G has no cycle that passes through a vertex of degree two.

Since G is subcubic, it contains a vertex w of degree *two* such that G is the union of two graphs, G_1 and G_2 , having only vertex w in common. Both G_1 and G_2 are subcubic and have fewer than n vertices, so by the induction hypothesis both of them have a drawing satisfying the conditions. Translate the drawing of G_2 so that the points representing w in the two drawings coincide. Since w has degree *one* in both G_1 and G_2 , by a possible rotation of G_2 about w through an angle that is a multiple of $\pi/5$, we can achieve that the two edges adjacent to w are not parallel. By scaling G_2 from w , if necessary, we can also achieve that the slope of no segment between a vertex of $G_1 \setminus w$ and a vertex of $G_2 \setminus w$ is a basic slope. Thus, the resulting drawing of G meets the requirements.

Case 3: G has a cycle passing through a vertex of degree two.

If G itself is a cycle, we can easily draw it. If it is not the case, let C be a *shortest* cycle which contains a vertex of degree two. Let u_0, u_1, \dots, u_k denote the vertices of C , in this cyclic order, such that u_0 has degree *two* and u_1 has degree *three*. The indices are understood mod $k+1$, that is, for instance, $u_{k+1} = u_0$. It follows from the minimality of C that u_i and u_j are not connected by an edge of G whenever $|i-j| > 1$.

Since $G \setminus C$ is subcubic, by assumption, it admits a straight-line drawing satisfying the conditions. Each u_i has at most *one* neighbor in $G \setminus C$. Denote this neighbor by t_i , if it exists. For every i for which t_i exists, we place u_i on a line passing through t_i . We place the u_i 's one by one, "very far" from $G \setminus C$, starting with u_1 . Finally, we arrive at u_0 , which has no neighbor in $G \setminus C$, so that it can be placed at the intersection of two lines of basic slope, through u_1 and u_k , respectively. We have to argue that our method does not create "unnecessary" edges, that is, we never place two independent vertices in such a way that the slope of the segment connecting them is a basic slope. In what follows, we make this argument precise.

The locations of the vertices u_0, u_1, \dots, u_k are determined by using the following algorithm, $\text{PROCEDURE}(G, C, u_0, u_1, x)$, where G is the input subcubic graph, C is a shortest cycle passing through a vertex of degree *two*, u_0 , that has a degree *three* neighbor, u_1 , and x is a real parameter. Note that $\text{PROCEDURE}(G, C, u_0, u_1, x)$ is a *nondeterministic* algorithm, as we have more than one choice at certain steps. (However, it is very easy to make it deterministic.)




Figure 27: The four possible locations of u_i .

PROCEDURE(G, C, u_0, u_1, x)

- STEP 0. Since $G \setminus C$ is subcubic, it has a representation with the *five* basic slopes. Take such a representation, scaled and translated in such a way that t_1 (which exists since the degree of u_1 is three) is at the origin, and all other vertices are within unit distance from it.

For any i , $2 \leq i \leq k$, for which u_i does not have a neighbor in $G \setminus C$, let t_i be any unoccupied point closer to the origin than 1, such that the slope of none of the lines connecting t_i to t_1, t_2, \dots, t_{i-1} or to any other already embedded point of $G \setminus C$ is a basic slope.

For any point p and for any $0 \leq i \leq 4$, let $\ell_i(p)$ denote the line with i th basic slope, $i\pi/5$, passing through p . Let ℓ_i stand for $\ell_i(O)$, where O denotes the origin.

We will place u_1, \dots, u_k recursively, so that u_j is placed on $\ell_i(t_j)$, for a suitable i . Once u_j has been placed on some $\ell_i(t_j)$, define $\text{ind}(u_j)$, the *index* of u_j , to be i . (The indices are taken mod 5. Thus, for example, $|i - i'| \geq 2$ is equivalent to saying that $i \neq i'$ and $i \neq i' \pm 1 \pmod{5}$.) Start with u_1 . The degree of t_1 in $G \setminus C$ is at most *two*, so that at the beginning the set $\mathbf{sl}(t_1)$ (defined in the first paragraph of this section) has at most *two* elements. Let $l \notin \mathbf{sl}(t_1)$. Direct the line $\ell_l(t_1)$ arbitrarily, and place u_1 on it at distance x from t_1 in the positive direction. (According to this rule, if $x < 0$, then u_1 is placed on $\ell_l(t_1)$ at distance $|x|$ from t_1 in the *negative* direction.)

Suppose that u_1, u_2, \dots, u_{i-1} have been already placed and that u_{i-1} lies on the line $\ell_l(t_{i-1})$, that is, we have $\text{ind}(u_{i-1}) = l$.

- STEP i . We place u_i at one of the following four locations (see Figure 27):
 - (1) the intersection of $\ell_{l+1}(t_i)$ and $\ell_{l+2}(u_{i-1})$;
 - (2) the intersection of $\ell_{l+2}(t_i)$ and $\ell_{l+3}(u_{i-1})$;
 - (3) the intersection of $\ell_{l-1}(t_i)$ and $\ell_{l-2}(u_{i-1})$;
 - (4) the intersection of $\ell_{l-2}(t_i)$ and $\ell_{l-3}(u_{i-1})$.

Choose from the above four possibilities so that the edge $u_i t_i$ is not parallel to any other edge already drawn and adjacent to t_i , i.e., before adding the edge $u_i t_i$ to the drawing, $\mathbf{sl}(t_i)$ did not include $\mathbf{sl}(u_i t_i)$.

It follows directly from (1)–(4) that $\mathbf{sl}(u_{i-1}) \subset \{l, l-1, l+1 \pmod{5}\}$, while $\mathbf{sl}(u_i u_{i-1}) \subset \{l-2, l+2 \pmod{5}\}$, that is, before adding the edge $u_i u_{i-1}$ to the drawing, we had $\mathbf{sl}(u_i u_{i-1}) \not\subset \mathbf{sl}(u_{i-1})$. Avoiding for $u_i t_i$ the slopes of the edges already incident to t_i , leaves available two of the choices (1), (2), (3), (4).

Let u'_{i-1} be the translation of u_{i-1} by the vector $\overrightarrow{t_{i-1}O}$, and similarly, let u'_i be the translation of u_i by the vector $\overrightarrow{t_i O}$. That is, $Ou'_{i-1}u_{i-1}t_{i-1}$ and $Ou'_i u_i t_i$ are parallelograms. We have

$$\begin{aligned} \overline{Ou_{i-1}} - 1 &< \overline{Ou'_{i-1}} < \overline{Ou_{i-1}} + 1, \\ \overline{Ou_i} - 1 &< \overline{Ou'_i} < \overline{Ou_i} + 1, \end{aligned}$$

and

$$2 \cos\left(\frac{\pi}{5}\right) \overline{Ou'_{i-1}} = \overline{Ou'_i}.$$

Therefore, for any possible location of u_i , we have

$$1.6\overline{Ou_{i-1}} - 4 < 2 \cos\left(\frac{\pi}{5}\right) \overline{Ou_{i-1}} - 4 < \overline{Ou_i} < 2 \cos\left(\frac{\pi}{5}\right) \overline{Ou_{i-1}} + 4 < 1.7\overline{Ou_{i-1}} + 4.$$

Suppose that $|x| \geq 50$. Clearly, $|x| - 1 < \overline{Ou_1}$, and by the previous calculations it is easy to show by induction that $|x| - 1 < \overline{Ou_i}$ for all $i \leq k$. Therefore, $1.5\overline{Ou_{i-1}} < 1.6\overline{Ou_{i-1}} - 4$ so we obtain

$$1.5\overline{Ou_{i-1}} < \overline{Ou_i}. \quad (1)$$

We have to verify that the above procedure does not produce “unnecessary” edges, that is, the following statement is true.

Claim 7.3. *Suppose that $|x| \geq 50$.*

- (i) The slope of $u_i u_j$ is not a basic slope, for any $j < i - 1$.
- (ii) The slope of $u_i v$ is not a basic slope, for any $v \in V(G \setminus C)$, $v \neq t_i$.

Proof. (i) Suppose that the slope of $u_i u_j$ is a basic slope for some $j < i - 1$. By repeated application of inequality (1), we obtain that $\overline{Ou_i} > 1.5^{i-j} \overline{Ou_j} > 2\overline{Ou_j}$. On the other hand, if $u_i u_j$ has a basic slope, then easy geometric calculations show that $\overline{Ou_i} < 2 \cos\left(\frac{\pi}{5}\right) \overline{Ou_j} + 4 < 2\overline{Ou_j}$, a contradiction.

(ii) Suppose for simplicity that $t_i u_i$ has slope 0, i.e., it is horizontal. By the construction, no vertex v of $G \setminus C$ determines a horizontal segment with t_i , but all of them are within distance 2 from t_i . As $\overline{Ou_i} > x - 1$, segment vu_i is almost, but not exactly horizontal. That is, we have $0 < |\angle t_i u_i v| < \pi/5$, contradiction. \square

Suppose that STEP 0, STEP 1, \dots , STEP k have already been completed. It remains to determine the position of u_0 . We need some preparation. The notation $|x| \geq 2 \pmod{5}$ means that $x = 2$ or $x = 3 \pmod{5}$.

Claim 7.4. *There exist two integers $0 \leq \alpha, \beta < 5$ with $|\alpha - \beta| \geq 2 \pmod{5}$ such that starting the PROCEDURE with $\text{ind}(u_1) = \alpha$ and with $\text{ind}(u_1) = \beta$, we can continue so that $\text{ind}(u_2)$ is the same in both cases.*

Proof. Suppose that the degrees of t_1 and t_2 in $G \setminus C$ are *two*, that is, there are two forbidden lines for both u_1 and u_2 . In the other cases, when the degree of t_1 or the degree of t_2 is less than *two*, or when $t_1 = t_2$, the proof is similar, but simpler. We can place u_1 on $\ell_l(t_1)$ for any $l \notin \mathbf{sl}(t_1)$. Therefore, we have three choices, two of which, $\ell_\alpha(t_1)$ and $\ell_\beta(t_1)$, are not consecutive, so that $|\alpha - \beta| \geq 2 \pmod{5}$.

The vertex u_2 cannot be placed on $\ell_m(t_2)$ for any $m \in \mathbf{sl}(t_2)$, so there are three possible lines for u_2 : $\ell_x(t_2)$, $\ell_y(t_2)$, $\ell_z(t_2)$, say. For any fixed location of u_1 , we can place u_2 on four possible lines, so on at least two of the lines $\ell_x(t_2)$, $\ell_y(t_2)$, and $\ell_z(t_2)$. Therefore, at least one of them, say $\ell_x(t_2)$, can be used for both locations of u_1 . \square

Claim 7.5. *We can place the vertices u_1, u_2, \dots, u_k using the PROCEDURE so that for all k we have $|\text{ind}(u_1) - \text{ind}(u_k)| \geq 2 \pmod{5}$.*

Proof. By Claim 7.4, there are two placements of the vertices of $C \setminus \{u_0, u_k\}$, denoted by u_1, u_2, \dots, u_{k-1} and by $u'_1, u'_2, \dots, u'_{k-1}$ such that $|\text{ind}(u_1) - \text{ind}(u'_1)| \geq 2 \pmod{5}$, and $\text{ind}(u_i) = \text{ind}(u'_i)$ for all $i \geq 2$. That is, we can start placing the vertices on two non-neighboring lines so that from the second step of the PROCEDURE we use the same lines. We show that we can place u_k so that u_1 and u_k , or u'_1 and u_k are on non-neighboring lines. Having placed u_{k-1} (or u'_{k-1}), we have four choices for $\text{ind}(u_k)$. Two of them can be ruled out by the condition $\text{ind}(u_k) \notin \mathbf{sl}(t_k)$. We still have two choices. Since u_1 and u'_1 are on non-neighboring lines, there is only one line which is a neighbor of both of them. Therefore, we still have at least one choice for $\text{ind}(u_k)$ such that $|\text{ind}(u_1) - \text{ind}(u_k)| \geq 2$ or $|\text{ind}(u'_1) - \text{ind}(u_k)| \geq 2$. \square




Figure 28: $\ell_{i+1}(u_1)$, does not separate the vertices of $G \setminus C$ from u_k , $\ell_{i-1}(u_1)$ does.

- STEP $k + 1$. Let $i = \text{ind}(u_1)$, $j = \text{ind}(u_k)$, and assume, by Claim 7.5, that $|i - j| \geq 2 \pmod{5}$. Consider the lines $\ell_{i-1}(u_1)$ and $\ell_{i+1}(u_1)$. One of them, $\ell_{i+1}(u_1)$, say, does not separate the vertices of $G \setminus C$ from u_k , the other one does. See Fig. 28.

Place u_0 at the intersection of $\ell_{i+1}(u_1)$ and $\ell_i(u_k)$.

Claim 7.6. *Suppose that $|x| \geq 50$.*

- (i) The slope of u_0u_j is not a basic slope, for any $1 < j < k$.
- (ii) The slope of u_0v is not a basic slope, for any $v \in V(G \setminus C)$.

Proof. (i) Denote by u_{k+1} the intersection of $\ell_{i+1}(O)$ and $\ell_i(u_k)$. Suppose that the slope of u_0u_j is a basic slope for some $1 < j < k$. As in the proof of Claim 7.3, by repeated application of inequality (1), we obtain that $\overline{Ou_{k+1}} > 1.5^{k+1-j}\overline{Ou_j} > 2\overline{Ou_j}$. On the other




Figure 29: The four possible locations of u_0 .

hand, by an easy geometric argument, if the slope of u_0u_j is a basic slope, then $\overline{Ou_{k+1}} < 2 \cos\left(\frac{\pi}{5}\right) \overline{Ou_j} + 4 < 2\overline{Ou_j}$, a contradiction, provided that $|x| \geq 50$.

(ii) For any vertex $v \in G \setminus C$, the slope of the segment u_0v is strictly between $i\pi/5$ and $(i+1)\pi/5$, therefore, it is not a basic slope. See Figure 29. This concludes the proof of the claim and hence Theorem 7.2. \square

7.2 Cubic Graphs - Proof of Theorem 7.1

First we note that if G is connected, then Theorem 7.1 is an easy corollary to Theorem 7.2. Indeed, delete any vertex, and then put it back using two extra directions. If G is not connected, the only problem that may arise is that these extra directions can differ for different components. We will define a family of drawings for each component G^i of G , depending on parameters ε_i , and then choose the values of these parameters in such a way that the extra directions will coincide.

Suppose that G is a cubic graph. If a connected component is not 3-regular then, by Theorem 7.2, it can be drawn using the *five* basic slopes. If a connected component is a complete graph K_4 on *four* vertices, then it can also be drawn using the basic slopes. For the sake of simplicity, suppose that we do not have such components, i. e. each connected component G^1, \dots, G^m of G is 3-regular and none of them is isomorphic to K_4 .

First we concentrate on G^1 . Let C be a shortest cycle in G^1 . We distinguish two cases.

Case 1: C is not a triangle. Denote by u_0, \dots, u_k the vertices of C , and let t_0 be the neighbor of u_0 not belonging to C . Delete the edge u_0t_0 , and let \bar{G} be the resulting graph.

Case 2: C is a triangle. Every vertex of C has precisely *one* neighbor that does not belong to C . If all these neighbors coincide, then G^1 is a complete graph on *four* vertices, contradicting our assumption. So one vertex of C , u_0 , say, has a neighbor t_0 which does not belong to C and which is not adjacent to the other two vertices, u_1 and u_2 , of C . Delete the edge u_0t_0 , and let \bar{G} be the resulting graph.

Observe that in both cases, u_k and t_0 are not connected in G^1 . Indeed, suppose for a contradiction that they are connected. In the first case, G^1 would contain the triangle $u_0u_k t_0$, contradicting the minimality of C . In the second case, the choice of u_0 would be violated.

There will be exactly two edges with extra directions, u_0u_k and u_0t_0 . The slope of u_0u_k will be very close to a basic slope and the slope of u_0t_0 will be decided at the end, but we will show that almost any choice will do.

For any nonnegative ε and real x , $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ is defined as follows. STEPS 0, 1, \dots , k are identical to STEPS 0, 1, \dots , k of $\text{PROCEDURE}(\bar{G}, C, u_0, u_1, x)$.

- STEP $k + 1$. If there is a segment, determined by the vertices of $G \setminus C$, of slope $i\pi/5 + \varepsilon$ or $i\pi/5 - \varepsilon$, for any $0 \leq i < 5$, then STOP. In this case, we say that ε is *1-bad* for \bar{G} .

Otherwise, when ε is *1-good*, let $i = \text{ind}(u_1)$ and $j = \text{ind}(u_k)$. We can assume by Claim 7.5 that $|i - j| \geq 2 \pmod{5}$. Consider the lines $\ell_{i-1}(u_1)$ and $\ell_{i+1}(u_1)$. One of them does not separate the vertices of $G \setminus C$ from u_k , the other one does.

If $\ell_{i-1}(u_1)$ separates $G \setminus C$ from u_k , then place u_0 at the intersection of $\ell_{i+1}(u_1)$ and the line through u_k with slope $i\pi/5 + \varepsilon$. If $\ell_{i+1}(u_1)$ separates $G \setminus C$ from u_k , then place u_0 at the intersection of $\ell_{i-1}(u_1)$ and the line through u_k with slope $i\pi/5 - \varepsilon$.

Since STEPS 0, 1, \dots , k are identical in $\text{PROCEDURE}(\bar{G}, C, u_0, u_1, x)$ and in $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$, Claims 7.3, 7.4, and 7.5 are also true for the MODIFIEDPROCEDURE .

Moreover, it is easy to see that an analogue of Claim 7.6 also holds with an identical proof, provided that ε is sufficiently small: $0 < \varepsilon < 1/100$.

Claim 7.6'. *Suppose that $|x| \geq 50$ and $0 < \varepsilon < 1/100$.*

- (i) *The slope of u_0u_j is not a basic slope, for any $1 < j < k$.*
- (ii) *The slope of u_0v is not a basic slope, for any $v \in V(\bar{G} \setminus C)$.* □

Perform $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ for a fixed ε , and observe how the drawing changes as x varies. For any vertex u_i of C , let $u_i(x)$ denote the position of u_i , as a function of x . For every i , the function $u_i(x)$ is linear, that is, u_i moves along a line as x varies.

Claim 7.7. *With finitely many exceptions, for every value of x , $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ produces a proper drawing of \bar{G} , provided that ε is 1-good.*

Proof. Claims 7.3, 7.4, 7.5, and 7.6' imply Claim 7.7 for $|x| \geq 50$. Let u and v be two vertices of \bar{G} . Since $u(x)$ and $v(x)$ are linear functions, their difference, $\vec{uv}(x)$, is also linear.

If uv is an edge of \bar{G} , then the direction of $\vec{uv}(x)$ is the same for all $|x| \geq 50$. Therefore, it is the same for all values of x , with the possible exception of one value, for which $\vec{uv}(x) = 0$ holds.

If uv is not an edge of \bar{G} , then the slope of $\vec{uv}(x)$ is not a basic slope for any $|x| \geq 50$. Therefore, with the exception of at most *five* values of x , the slope of $\vec{uv}(x)$ is never a basic slope, nor does $\vec{uv}(x) = 0$ hold. □

Take a closer look at the relative position of the endpoints of the missing edge, $u_0(x)$ and $t_0(x)$. Since $t_0 \in \bar{G} \setminus C$, $t_0 = t_0(x)$ is the same for all values of x . The position of $u_0 = u_0(x)$ is a linear function of x . Let ℓ be the line determined by the function $u_0(x)$. If ℓ passes through t_0 , then we say that ε is *2-bad* for \bar{G} . If ε is 1-good and it is not 2-bad for \bar{G} , then we say that it is *2-good* for \bar{G} . If ε is 2-good, then by varying x we can achieve almost any slope for the edge t_0u_0 . This will turn out to be crucially important, because we want to attain that these slopes coincide in all components.

Claim 7.8. *Suppose that the values $\varepsilon \neq \delta$, $0 < \varepsilon, \delta < 1/100$, are 1-good for \bar{G} . Then at least one of them is 2-good for \bar{G} .*

Proof. Suppose, for simplicity, that $\text{ind}(u_1) = 0$, $\text{ind}(u_k) = 2$, and that u_1 and u_k are in the right half-plane (of the vertical line through the origin). The other cases can be settled analogously. To distinguish between $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ and $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \delta)$, let $u_0^\varepsilon(x)$ denote the position of u_0 obtained by the first procedure and $u_0^\delta(x)$ its position obtained by the second. Let ℓ^ε and ℓ^δ denote the lines determined by the functions $u_0^\varepsilon(x)$ and $u_0^\delta(x)$. Suppose that x is very large. Since, by inequality (1), we have $\overline{u_k(x)O} > 1.5\overline{u_1(x)O}$, both $u_0^\varepsilon(x)$ and $u_0^\delta(x)$ are on the line $\ell_1(u_1(x))$, very far in the positive direction. Therefore, both of them are above the line $\ell_{\pi/10}$. On the other hand, if $x < 0$ is very small (i.e., if $|x|$ is very big), both $u_0^\varepsilon(x)$ and $u_0^\delta(x)$ lie below the line $\ell_{\pi/10}$. It follows that the slopes of ℓ^ε and ℓ^δ are larger than $\pi/10$, but smaller than $\pi/5$.

Suppose that neither ε nor δ is 2-good. Then both ℓ^ε and ℓ^δ pass through t_0 . That is, for a suitable value of x , we have $u_0^\varepsilon(x) = t_0$. We distinguish two cases.

Case 1: $u_0^\varepsilon(x) = t_0 = u_k(x)$. Then, as x varies, the line determined by $u_k(x)$ coincides with $\ell_2(t_0)$. Consequently, t_0 and u_k are connected in G^1 , a contradiction.

Case 2: $u_0^\varepsilon(x) = t_0 \neq u_k(x)$. In order to get a contradiction, we try to determine the position of $u_0^\delta(x)$. If we consider $\text{STEP } k+1$ in $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ and in $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \delta)$, we can conclude that $u_1(x)$ lies on $\ell_1(u_0^\varepsilon) = \ell_1(t_0)$, $u_0^\delta(x)$ lies on $\ell_1(u_1(x))$, therefore, $u_0^\delta(x)$ lies on $\ell_1(t_0)$. On the other hand, $u_0^\delta(x)$ lies on ℓ^δ , and, by assumption, ℓ^δ passes through t_0 . However, we have shown that ℓ^δ and $\ell_1(t_0)$ have different slopes, therefore, $u_0^\delta(x)$ must be at their intersection point, so we have $u_0^\delta(x) = u_0^\varepsilon(x) = t_0$.

Considering again $\text{STEP } k+1$ in $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ and in $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \delta)$, we can conclude that the point $u_0^\delta(x) = t_0 = u_0^\varepsilon(x)$ belongs to both $\ell_\varepsilon(u_k(x))$ and $\ell_\delta(u_k(x))$. This contradicts our assumption that $u_k(x)$ is different from $u_0^\delta(x) = t_0 = u_0^\varepsilon(x)$. \square

By Claim 7.7, for every $\varepsilon < 1/100$ and with finitely many exceptions for every value of x , $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ produces a proper drawing of \bar{G} . When we want to add the edge u_0t_0 , the slope of $u_0(x)t_0$ may coincide with the slope of $u(x)u'(x)$, for some $u, u' \in \bar{G}$. The following statement guarantees that this does not happen “too often”. We use $\alpha(\vec{u})$ to denote the *slope* of a vector \vec{u} .

Claim 7.9. *Let $\vec{u}(x)$ and $\vec{v}(x): R \rightarrow R^2$ be two linear functions, and let $\ell(u)$ and $\ell(v)$ denote the lines determined by $\vec{u}(x)$ and $\vec{v}(x)$. Suppose that for some $x_1 < x_2 < x_3$, the vectors \vec{u}, \vec{v} do not vanish and that their slopes coincide, that is, $\alpha(\vec{u}(x_1)) = \alpha(\vec{v}(x_1))$, $\alpha(\vec{u}(x_2)) = \alpha(\vec{v}(x_2))$, and $\alpha(\vec{u}(x_3)) = \alpha(\vec{v}(x_3))$. Then $\ell(u)$ and $\ell(v)$ must be parallel.*

Proof. If $\ell(u)$ passes through the origin, then for every value of x , $\vec{u}(x)$ has the same slope. In particular, $\alpha(\vec{v}(x_1)) = \alpha(\vec{v}(x_2)) = \alpha(\vec{v}(x_3))$. Therefore, $\ell(v)$ also passes through the origin and is parallel to $\ell(u)$. (In fact, we have $\ell(u) = \ell(v)$.) We can argue analogously if $\ell(u)$ passes through the origin. Thus, in what follows, we can assume that neither $\ell(u)$ nor $\ell(v)$ passes through the origin.




Figure 30: $\ell(u)$ and $\ell(v)$ must be parallel.

Suppose that $\alpha(\vec{u}(x_1)) = \alpha(\vec{v}(x_1))$, $\alpha(\vec{u}(x_2)) = \alpha(\vec{v}(x_2))$, and $\alpha(\vec{u}(x_3)) = \alpha(\vec{v}(x_3))$. For any x , define $\vec{w}(x)$ as the intersection point of $\ell(v)$ and the line connecting the origin to $\vec{u}(x)$, provided that they intersect. Clearly, $\vec{v}(x) = \vec{w}(x)$ for $x = x_1, x_2, x_3$, and $\vec{u}(x)$ and $\vec{w}(x)$ have the same slope for every x . The transformation $\vec{u}(x) \rightarrow \vec{w}(x)$ is a projective transformation from $\ell(u)$ to $\ell(v)$, therefore, it preserves the cross ratio of any four points. That is, for any x , we have

$$(\vec{u}(x_1), \vec{u}(x_2); \vec{u}(x_3), \vec{u}(x)) = (\vec{w}(x_1), \vec{w}(x_2); \vec{w}(x_3), \vec{w}(x)).$$

Since both $\vec{u}(x)$ and $\vec{v}(x)$ are linear functions, we also have

$$(\vec{u}(x_1), \vec{u}(x_2); \vec{u}(x_3), \vec{u}(x)) = (\vec{v}(x_1), \vec{v}(x_2); \vec{v}(x_3), \vec{v}(x)).$$

Hence, we can conclude that $\vec{v}(x) = \vec{w}(x)$ for all x . However, this is impossible, unless $\ell(u)$ and $\ell(v)$ are parallel. Indeed, suppose that $\ell(u)$ and $\ell(v)$ are not parallel, and set x in such a way that $\vec{u}(x)$ is parallel to $\ell(v)$. Then $\vec{w}(x)$ cannot have the same slope as $\vec{u}(x)$, a contradiction. \square

Suppose that ε is 2-good and let us fix it. As above, let $u_0^\varepsilon(x)$ be the position of u_0 obtained by $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$, and let ℓ^ε be the line determined by $u_0^\varepsilon(x)$.

Suppose also that there exist two independent vertices of \bar{G} , $u, u' \neq u_0$, such that the line determined by $uu'(x)$ is parallel to ℓ^ε . Then we say that ε is 3-bad for \bar{G} . If ε is 2-good and it is not 3-bad for \bar{G} , then we say that it is 3-good for \bar{G} .

It is easy to see that, for any $0 < \varepsilon, \delta < 1/100$, ℓ^ε and ℓ^δ are not parallel, therefore, for any fixed u, u' , there is at most one value of ε for which the line determined by $uu'(x)$ is parallel to ℓ^ε . Thus, with finitely many exceptions, all values $0 < \varepsilon < 1/100$ are 3-good.

Summarizing, we have obtained the following.

Claim 7.10. *Suppose that ε is 3-good for \bar{G} . With finitely many exceptions, for every value of x , $\text{MODIFIEDPROCEDURE}(\bar{G}, C, u_0, u_1, x, \varepsilon)$ gives a proper drawing of G^1 . \square*

Now we are in a position to complete the proof of Theorem 7.1. Proceed with each of the components as described above for G^1 . For any fixed i , let $u_0^i v_0^i$ be the edge deleted from G^i , and denote the resulting graphs by $\bar{G}^1, \dots, \bar{G}^m$. Let $0 < \varepsilon < 1/100$ be fixed in such a way that ε is 3-good for all graphs $\bar{G}^1, \dots, \bar{G}^m$. This can be achieved, in view of the fact that there are only finitely many values of ε which are not 3-good. Perform $\text{MODIFIEDPROCEDURE}(\bar{G}^i, C^i, u_0^i, u_1^i, x^i, \varepsilon)$. Now the line ℓ^i determined by all possible locations of u_0^i does not pass through t_0^i .

Notice that when $\text{MODIFIEDPROCEDURE}(\bar{G}^i, C^i, u_0^i, u_1^i, x^i, \varepsilon)$ is executed, apart from edges with basic slopes, we use an edge with slope $r\pi/5 \pm \varepsilon$, for some integer $r \pmod 5$. By using rotations through $\pi/5$ and a reflection, if necessary, we can achieve that each component G^i is drawn using the basic slopes and one edge of slope ε .

It remains to set the values of x_i and draw the missing edges $u_0^i v_0^i$. Since the line ℓ^i determined by the possible locations of u_0^i does not pass through t_0^i , by varying the value of x^i , we can attain any slope for the missing edge $t_0^i u_0^i$, except for the slope of ℓ^i . By Claim 7.10, with finitely many exceptions, all values of x^i produce a proper drawing of G^i . Therefore, we can choose x^1, \dots, x^m so that all segments $t_0^i u_0^i$ have the same slope and every component G^i is properly drawn using the same seven slopes. Translating the resulting drawings through suitable vectors gives a proper drawing of G , this completes the proof of Theorem 7.1.

7.3 Concluding Remarks

In the proof of Theorem 7.1, the slopes we use depend on the graph G . However, the proof shows that one can simultaneously embed all cubic graphs using only *seven* fixed slopes.

It is unnecessary to use $|x| \geq 50$, in every step, we could pick any x , with finitely many exceptions.

It seems to be only a technical problem that we needed *two* extra directions in the proof of Theorem 7.1. We believe that *one* extra direction would suffice.

The most interesting problem that remains open is to decide whether the number of slopes needed for graphs of maximum degree *four* is bounded.

Another not much investigated question is to estimate the complexity of computing the slope parameter of a graph. A related problem is to decide under what conditions a graph can be drawn on a polynomial sized grid using a fixed number of slopes.

Question 7.11. *Is it possible to draw all cubic graphs with a bounded number of slopes on a polynomial sized grid?*

8 Drawing Planar Graphs with Few Slopes

This section is based on our paper with Balázs Keszegh and János Pach, Drawing planar graphs of bounded degree with few slopes [KPP10].

In this section, we will be concerned with drawings of planar graphs. Unless it is stated otherwise, all *drawings* will be *non-crossing*, that is, no two arcs that represent different edges have an interior point in common.

Every planar graph admits a straight-line drawing [F48]. From the practical and aesthetic point of view, it makes sense to minimize the number of slopes we use [WC94]. Remember that the *planar slope number* of a planar graph G is the smallest number s with the property that G has a straight-line drawing with edges of at most s distinct slopes. If G has a vertex of degree d , then its planar slope number is at least $\lceil d/2 \rceil$, because in a straight-line drawing no two edges are allowed to overlap.

Dujmović, Eppstein, Suderman, and Wood [DESW07] raised the question whether there exists a function f with the property that the planar slope number of every planar graph with maximum degree d can be bounded from above by $f(d)$. Jelinek et al. [JJ10] have shown that the answer is yes for *outerplanar* graphs, that is, for planar graphs that can be drawn so that all of their vertices lie on the outer face. In Section 8.2, we answer this question in full generality. We prove the following.

Theorem 8.1. *Every planar graph with maximum degree d admits a straight-line drawing, using segments of $O(d^2(3 + 2\sqrt{3})^{12d}) \leq K^d$ distinct slopes.*

The proof is based on a paper of Malitz and Papakostas [MP94], who used Koebe’s theorem [K36] on disk representations of planar graphs to prove the existence of drawings with relatively large angular resolution. As the proof of these theorems, our argument is nonconstructive; it only yields a nondeterministic algorithm with running time $O(dn)$. However, if one combines our result with a polynomial time algorithm that computes the ϵ -approximation of the disk representation (see e.g. Mohar [1]), then one can obtain a deterministic algorithm running in time exponential in d but polynomial in n .

For $d = 3$, much stronger results are known than the one given by our theorem. Dujmović et al. [DESW07] showed that every planar graph with maximum degree 3 admits a straight-line drawing using at most 3 different slopes, except for at most 3 edges of the outer face, which may require 3 additional slopes. This complements Ungar’s old theorem [U53], according to which 3-regular, 4-edge-connected planar graphs require only 2 slopes and 4 extra edges.

The exponential upper bound in Theorem 8.1 is probably far from being optimal. However, we were unable to give any superlinear lower bound for the largest planar slope number of a planar graph with maximum degree d . The best constructions we are aware of are presented in Section 8.5.

We also show that significantly fewer slopes are sufficient if we are allowed to represent the edges by short noncrossing polygonal paths. If such a path consists of $k + 1$ segments, we say that the edge is drawn by k *bends*. In Section 8.3, we show if we allow one bend per edge, then every planar graph can be drawn using segments with $O(d)$ slopes.

Theorem 8.2. *Every planar graph G with maximum degree d can be drawn with at most 1 bend per edge, using at most $2d$ slopes.*

Allowing *two* bends per edge yields an optimal result: almost all planar graphs with maximum degree d can be drawn with $\lceil d/2 \rceil$ slopes. In Section 8.4, we establish

Theorem 8.3. *Every planar graph G with maximum degree $d \geq 3$ can be drawn with at most 2 bends per edge, using segments of at most $\lceil d/2 \rceil$ distinct slopes. The only exception is the graph formed by the edges of an octahedron, which is 4-regular, but requires 3 slopes. These bounds are best possible.*

It follows from the proof of Theorem 8.3 that in the cyclic order of directions, the slopes of the edges incident to any given vertex form a contiguous interval. Moreover, the $\lceil d/2 \rceil$ directions we use can be chosen to be equally spaced in $[0, 2\pi)$. We were unable to guarantee such a nice property in Theorem 8.2: even for a fixed d , as the number of vertices increases, the smallest difference between the $2d - 2$ slopes we used tends to zero. We suspect that this property is only an unpleasant artifact of our proof technique.

8.1 Straight-line Drawings - Proof of Theorem 8.1

Note that it is sufficient to prove the theorem for triangulated planar graphs, because any planar graph can be triangulated by adding vertices and edges so that the degree of each vertex increases only by a factor of at most three [PT06], so at the end we will lose this factor.

We need the following result from [MP94], which is not displayed as a theorem there, but is stated right above Theorem 2.2.

Lemma 8.4. (Malitz-Papakostas) *The vertices of any triangulated planar graph G with maximum degree d can be represented by nonoverlapping disks in the plane so that two disks are tangent to each other if and only if the corresponding vertices are adjacent, and the ratio of the radii of any two disks that are tangent to each other is at least α^{d-2} , where $\alpha = \frac{1}{3+2\sqrt{3}} \approx 0.15$.*

Lemma 8.4 can be established by taking any representation of the vertices of G by tangent disks, as guaranteed by Koebe's theorem, and applying a conformal mapping to the plane that takes the disks corresponding to the three vertices of the outer face to disks of the same radii. The lemma now follows by the observation that any internal disk is surrounded by a ring of at most d mutually touching disks, and the radius of none of them can be much smaller than that of the central disk.

The idea of the proof of Theorem 8.1 is as follows. Let G be a triangulated planar graph with maximum degree d , and denote its vertices by v_1, v_2, \dots . Consider a disk representation of G meeting the requirements of Lemma 8.4. Let D_i denote the disk that represents v_i , and let O_i be the center of D_i . By properly scaling the picture if necessary, we can assume without loss of generality that the radius of the smallest disk D_i is sufficiently large. Place an integer grid on the plane, and replace each center O_i by the nearest grid point. Connecting the corresponding pairs of grid points by segments, we obtain a straight-line drawing of G . The advantage of using a grid is that in this way we have control of the slopes of the edges. The trouble is that the size of the grid, and thus the number of slopes used, is very large. Therefore, in the neighborhood of each disk D_i , we use a portion of a grid whose side length is proportional to the radius of the disk. These grids will nicely fit

together, and each edge will connect two nearby points belonging to grids of comparable sizes. Hence, the number of slopes used will be bounded. See Figure 31.




Figure 31: Straight-line graph from disk representation

Now we work out the details. Let r_i denote the radius of D_i ($i = 1, 2, \dots$), and suppose without loss of generality that r^* , the radius of the smallest disk is

$$r^* = \min_i r_i = \sqrt{2}/\alpha^{d-2} > 1,$$

where α denotes the same constant as in Lemma 8.4.

Let $s_i = \lfloor \log_d(r_i/r^*) \rfloor \geq 0$, and represent each vertex v_i by the integer point nearest to O_i such that both of its coordinates are divisible by d^{s_i} . (Taking a coordinate system in general position, we can make sure that this point is unique.) For simplicity, the point representing v_i will also be denoted by v_i . Obviously, we have that the distance between O_i and v_i satisfies

$$\overline{O_i v_i} < \frac{d^{s_i}}{\sqrt{2}}.$$

Since the centers O_i of the disks induce a (crossing-free) straight-line drawing of G , in order to prove that moving the vertices to v_i does not create a crossing, it is sufficient to verify the following statement.

Lemma 8.5. *For any three mutually adjacent vertices, v_i, v_j, v_k in G , the orientation of the triangles $O_iO_jO_k$ and $v_iv_jv_k$ are the same.*

Proof. By Lemma 8.4, the ratio between the radii of any two adjacent disks is at least α^{d-2} . Suppose without loss of generality that $r_i \geq r_j \geq r_k \geq \alpha^{d-2}r_i$. For the orientation to change, at least one of $\overline{O_iv_i}$, $\overline{O_jv_j}$, or $\overline{O_kv_k}$ must be at least half of the smallest altitude of the triangle $O_iO_jO_k$, which is at least $\frac{r_k}{2}$.

On the other hand, as we have seen before, each of these numbers is smaller than

$$\frac{d^{s_i}}{\sqrt{2}} \leq \frac{r_i/r^*}{\sqrt{2}} = \frac{\alpha^{d-2}r_i}{2} \leq \frac{r_k}{2}$$

which completes the proof. \square

Now we are ready to complete the proof of Theorem 8.1. Take an edge v_iv_j of G , with $r_i \geq r_j \geq \alpha^{d-2}r_i$. The length of this edge can be bounded from above by

$$\begin{aligned} \overline{v_iv_j} &\leq \overline{O_iO_j} + \overline{O_iv_i} + \overline{O_jv_j} \leq r_i + r_j + \frac{d^{s_i}}{\sqrt{2}} + \frac{d^{s_j}}{\sqrt{2}} \leq 2r_i + \sqrt{2}d^{s_i} \leq 2r_i + \sqrt{2}r_i/r^* \\ &\leq r_i/r^*(2r^* + \sqrt{2}) \leq \frac{r_j/r^*}{\alpha^{d-2}}(2r^* + \sqrt{2}) < \frac{d^{s_j+1}}{\alpha^{d-2}}\left(\frac{2\sqrt{2}}{\alpha^{d-2}} + \sqrt{2}\right). \end{aligned}$$

According to our construction, the coordinates of v_j are integers divisible by d^{s_j} , and the coordinates of v_i are integers divisible by $d^{s_i} \geq d^{s_j}$, thus also by d^{s_j} .

Thus, shrinking the edge v_iv_j by a factor of d^{s_j} , we obtain a segment whose endpoints are integer points at a distance at most $\frac{d}{\alpha^{d-2}}\left(\frac{2\sqrt{2}}{\alpha^{d-2}} + \sqrt{2}\right)$. Denoting this number by $R(d)$, we obtain that the number of possible slopes for v_iv_j , and hence for any other edge in the embedding, cannot exceed the number of integer points in a disk of radius $R(d)$ around the origin. Thus, the planar slope number of any triangulated planar graph of maximum degree d is at most roughly $R^2(d)\pi = O(d^2/\alpha^{4d})$, which completes the proof. \square

Our proof is based on the result of Malitz and Papakostas that does not have an algorithmic version. However, with some reverse engineering, we can obtain a nondeterministic algorithm for drawing a triangulated planar graph of bounded degree with a bounded number of slopes. Because of the enormous constants in our expressions, this algorithm is only of theoretical interest. Here is a brief sketch.

Nondeterministic algorithm. First, we guess the three vertices of the outer face and their coordinates in the grid scaled according to their radii. Then embed the remaining vertices one by one. For each vertex, we guess the radius of the corresponding disk as well as its coordinates in the proportionally scaled grid. This algorithm runs in nondeterministic $O(dn)$ time.

8.2 One Bend per Edge - Proof of Theorem 8.2

In this section, we represent edges by noncrossing polygonal paths, each consisting of at most two segments. Our goal is to establish Theorem 8.2, which states that the total number of directions assumed by these segments grows at most linearly in d .

The proof of Theorem 8.2 is based on a result of Fraysseix et al. [FOR94], according to which every planar graph can be represented as a contact graph of T -shapes. A T -shape consists of a vertical and a horizontal segment such that the upper endpoint of the vertical segment lies in the interior of the horizontal segment. The vertical and horizontal segments of T are called its *leg* and *hat*, while their point of intersection is the *center* of the T -shape. The two endpoints of the hat and the bottom endpoint of the leg are called *ends* of the T -shape.

Two T -shapes are *noncrossing* if the interiors of their segments are disjoint. Two T -shapes are *tangent* to each other if they are noncrossing but they have a point in common.

Lemma 8.6. (Fraysseix et al.) *The vertices of any planar graph with n vertices can be represented by noncrossing T -shapes such that*

1. *two T -shapes are tangent to each other if and only if the corresponding vertices are adjacent;*
2. *the centers and the ends of the T -shapes belong to an $n \times n$ grid.*

Moreover, such a representation can be computed in linear time.

The proof of the lemma is based on the canonical ordering of the vertices of a planar graph, introduced in [FPP89].

Proof of Theorem 8.2. Consider a representation of G by T -shapes satisfying the conditions in the lemma. See Figure 32(a). For any $v \in V(G)$, let T_v denote the corresponding T -shape. We define a drawing of G , in which the vertex v is mapped to the center of T_v . To simplify the presentation, the center of T_v is also denoted by v . For any $uv \in E(G)$, let p_{uv} denote the point of tangency of T_u and T_v . The polygonal path $up_{uv}v$ consists of a horizontal and a vertical segment, and these paths together almost form a drawing of G with one bend per edge, using segments of two different slopes. The only problem is that these paths partially overlap in the neighborhoods of their endpoints. Therefore, we modify them by replacing their horizontal and vertical pieces by almost horizontal and almost vertical ones, as follows.

For any $1 \leq i \leq d$, let α_i denote the slope of the (almost horizontal) line connecting the origin $(0, 0)$ to the point $(2in, -1)$. Analogously, let β_i denote the slope of the (almost vertical) line passing through $(0, 0)$ and $(1, 2in)$.

Fix a T -shape T_v in the representation of G . It is tangent to at most d other T -shapes. Starting at its center v , let us pass around T_v in the counterclockwise direction, so that we first visit the upper left side of its hat, then its lower left side, then the left side and right side of its leg, etc. We number the points of tangencies along T_v in this order. (Note that there are no points of tangencies on the lower side of the hat.)

Suppose now that the hat of a T -shape T_u is tangent to the leg of T_v , and let p_{uv} be their point of tangency. Assume that p_{uv} was the number i point of tangency along T_u and the number j point of tangency along T_v . Let p'_{uv} denote the unique point of intersection of the (almost horizontal) line through u with slope α_i and the (almost vertical) line through v with slope β_j . In our drawing of G , the edge uv will be represented by the polygonal path $up'_{uv}v$. See Figure 32(c) for the resulting drawing and Figure 32(b) for a version distorted for the human eye to show the underlying structure.




Figure 32: Representation with T -shapes and the drawing with one bend per edge

Since the segments we used are almost horizontal or vertical, the modified edges $up'_{uv}v$ are very close (within distance $1/2$) of the original polygonal paths $up_{uv}v$. Thus, no two nonadjacent edges can cross each other. On the other hand, the order in which we picked the slopes around each v guarantees that no two edges incident to v will cross or overlap. This completes the proof. \square

8.3 Two Bends per Edge - Proof of Theorem 8.3

In this section, we draw the edges of a planar graph by polygonal paths with at most *two* bends. Our aim is to establish Theorem 8.3.

Note that the statement is trivially true for $d = 1$ and is false for $d = 2$. It is sufficient to prove Theorem 8.3 for even values of d . For $d = 4$, the assertion was first proved by Liu et al. [LMS91] and later, independently, by Biedl and Kant [BK98] (also that the only exception is the octahedral graph). The latter approach is based on the notion of *st-ordering* of biconnected (2-connected) graphs from Lempel et al. [LEC67]. We will show that this method generalizes to higher values of $d \geq 5$. As it is sufficient to prove the statement for even values of d , from now on we suppose that $d \geq 6$ even. We will argue that it is enough to consider *biconnected* graphs. Then we review some crucial claims from [BK98] that will enable us to complete the proof. We start with some notation.

Take $d \geq 5$ lines that can be obtained from a vertical line by clockwise rotation by $0, \pi/d, 2\pi/d, \dots, (d-1)\pi/d$ degrees. Their slopes are called the d regular slopes. We will use these slopes to draw G . Since these slopes depend only on d and not on G , it is enough to prove the theorem for connected graphs. If a graph is not connected, its components can be drawn separately.

In this section we always use the term “slope” to mean a regular slope. The *directed slope* of a directed line or segment is defined as the angle (mod 2π) of a clockwise rotation that takes it to a position parallel to the upward directed y -axis. Thus, if the directed slopes of two segments differ by π , then they have the same slope. We say that the slopes of the segments incident to a point p form a *contiguous interval* if the set $\vec{S} \subset \{0, \pi/d, 2\pi/d, \dots, (2d-1)\pi/d\}$ of directed slopes of the segments directed away from p , has the property that for all but at most one $\alpha \in \vec{S}$, we have that $\alpha + \pi/d \pmod{2\pi} \in \vec{S}$

(see Figure 35).

Finally, we say that G admits a *good drawing* if G has a planar drawing such that every edge has at most 2 bends, every segment of every edge has one of the $\lceil d/2 \rceil$ regular slopes, and the slopes of the segments incident to any vertex form a contiguous interval. If t is a vertex whose degree is at least two but less than d , then we can define the two *extremal segments* at t as the segments corresponding to the slopes at the two ends of the contiguous interval formed by the slopes of all the segments incident to t . Also define the *t -wedge* as the infinite cone bounded by the extension of the two extremal segments, which contains all segments incident to t and none of the “missing” segments. See Figure 33. For a degree one vertex t we define the *t -wedge* as the infinite cone bounded by the extension of the rotations of the segment incident to t around t by $\pm\pi/2d$.




Figure 33: The t -wedge

To prove Theorem 8.3, we show by induction that every connected planar graph with maximum degree $d \geq 6$ with an arbitrary t vertex whose degree is strictly less than d admits a good drawing that is contained in the t -wedge. Note that such a vertex always exist because of Euler’s polyhedral formula, thus Theorem 8.3 is indeed a direct consequence of this statement. First we show how the induction step goes for graphs that have a cut vertex, then (after a lot of definitions) we prove the statement also for biconnected graphs (without the induction hypothesis).

Lemma 8.7. *Let G be a connected planar graph of maximum degree d , let $t \in V(G)$ be a vertex whose degree is strictly smaller than d , and let $v \in V(G)$ be a cut vertex. Suppose that for any connected planar graph G' of maximum degree d , which has fewer than $|V(G)|$ vertices, and for any vertex $t' \in V(G')$ whose degree is strictly smaller than d , there is a good drawing of G' that is contained in the t' -wedge. Then G also admits a good drawing that is contained in the t -wedge.*

Proof. Let G_1, G_2, \dots denote the connected components of the graph obtained from G after the removal of the cut vertex v , and let G_i^* be the subgraph of G induced by $V(G_i) \cup \{v\}$.

If $t = v$ is a cut vertex, then by the induction hypothesis each G_i^* has a good drawing in the v -wedge*. After performing a suitable rotation for each of these drawings, and

*Of course the v -wedges for the different components are different.

identifying their vertices corresponding to v , the lemma follows because the slopes of the segments incident to v form a contiguous interval in each component.

If $t \neq v$, then let G_j be the component containing t . Using the induction hypothesis, G_j^* has a good drawing. Also, each G_i^* for $i \geq 2$ has a good drawing in the v -wedge. As in the previous case, the lemma follows by rotating and possibly scaling down the components for $i \neq j$ and again identifying the vertices corresponding to v . \square

In view of Lemma 8.7, in the sequel we consider only biconnected graphs. We need the following definition.

Definition 8.8. *An ordering of the vertices of a graph, v_1, v_2, \dots, v_n , is said to be an st -ordering if $v_1 = s$, $v_n = t$, and if for every $1 < i < n$ the vertex v_i has at least one neighbor that precedes it and a neighbor that follows it.*

In [LEC67], it was shown that any biconnected graph has an st -ordering, for any choice of the vertices s and t . In [BK98], this result was slightly strengthened for planar graphs, as follows.

Lemma 8.9. (Biedl-Kant) *Let D_G be a drawing of a biconnected planar graph, G , with vertices s and t on the outer face. Then G has an st -ordering for which $s = v_1$, $t = v_n$ and v_2 is also a vertex of the outer face and v_1v_2 is an edge of the outer face.*

We define G_i to be the subgraph of G induced by the vertices v_1, v_2, \dots, v_i . Note that G_i is connected. If i is fixed, we call the edges between $V(G_i)$ and $V(G) \setminus V(G_i)$ the *pending edges*. For a drawing of G , D_G , we denote by D_{G_i} the drawing restricted to G_i and to an initial part of each pending edge connected to G_i .

Proposition 8.10. *In the drawing D_G guaranteed by Lemma 8.9, v_{i+1}, \dots, v_n and the pending edges are in the outer face of D_{G_i} .*

Proof. Suppose for contradiction that for some i and $j > i$, v_j is not in the outer face of D_{G_i} . We know that v_n is in the outer face of D_{G_i} as it is on the outer face of D_G , thus v_n and v_j are in different faces of D_{G_i} . On the other hand, by the definition of st -ordering, there is a path in G between v_j and v_n using only vertices from $V(G) \setminus V(G_i)$. The drawing of this path in D_G must lie completely in one face of D_{G_i} . Thus, v_j and v_n must also lie in the same face, a contradiction. Since the pending edges connect $V(G_i)$ and $V(G) \setminus V(G_i)$, they must also lie in the outer face. \square

By Lemma 8.9, the edge v_1v_2 lies on the boundary of the outer face of D_{G_i} , for any $i \geq 2$. Thus, we can order the pending edges connecting $V(G_i)$ and $V(G) \setminus V(G_i)$ by walking in D_G from v_1 to v_2 around D_{G_i} on the side that does not consist of only the v_1v_2 edge, see Figure 34(a). We call this the *pending-order* of the pending edges between $V(G_i)$ and $V(G) \setminus V(G_i)$ (this order may depend on D_G). Proposition 8.10 implies

Proposition 8.11. *The edges connecting v_{i+1} to vertices preceding it form an interval of consecutive elements in the pending-order of the edges between $V(G_i)$ and $V(G) \setminus V(G_i)$.*

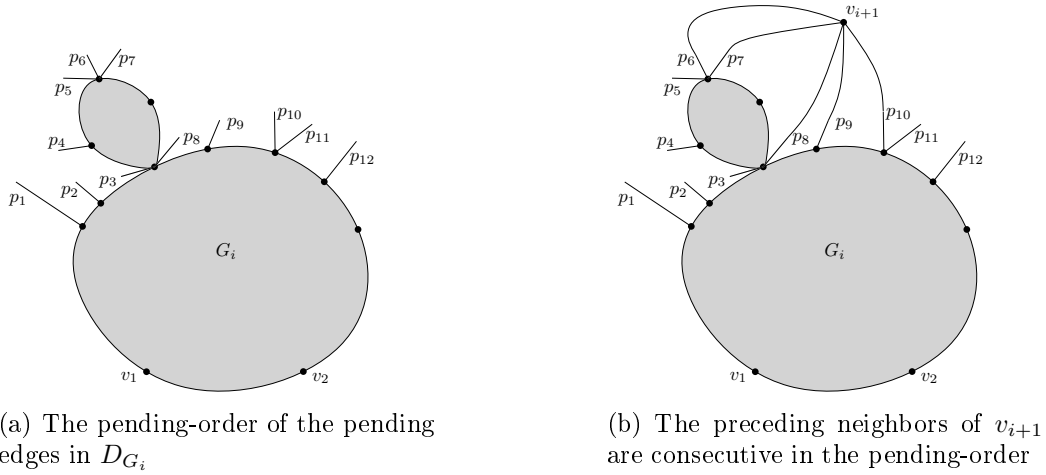


Figure 34: Properties of the st -ordering

For an illustration see Figure 34(a).

Two drawings of the same graph are said to be *equivalent* if the circular order of the edges incident to each vertex is the same in both drawings. Note that in this order we also include the pending edges (which are differentiated with respect to their yet not drawn end).

Now we are ready to finish the proof of Theorem 8.3, as the following lemma is the only missing step.

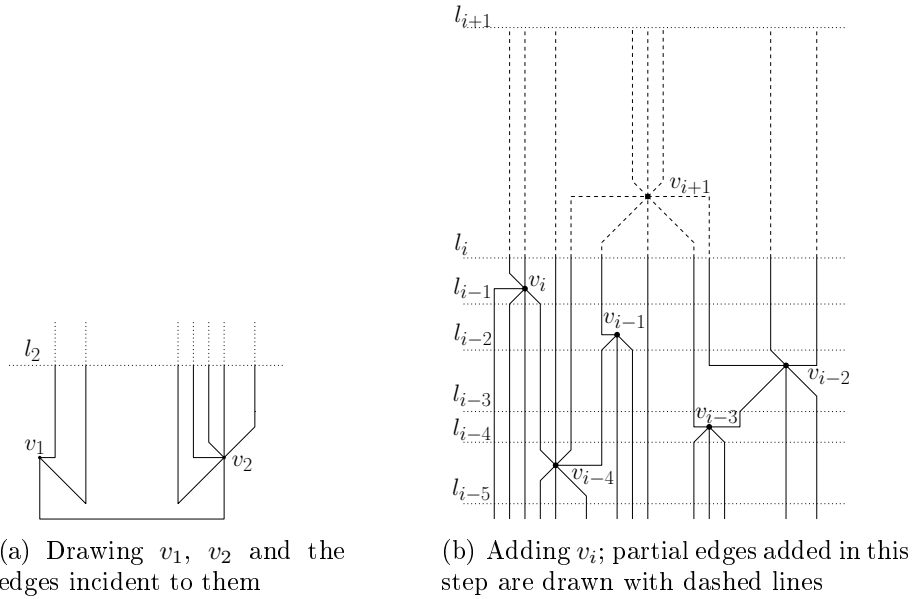


Figure 35: Drawing with at most two bends

Lemma 8.12. *For any biconnected planar graph G with maximum degree $d \geq 6$ and for any vertex $t \in V(G)$ with degree strictly less than d , G admits a good drawing that is contained in the t -wedge.*

Proof. Take a planar drawing D_G of G such that t is on the outer face and pick another vertex, s , from the outer face. Apply Lemma 8.9 to obtain an st -ordering with $v_1 = s, v_2$, and $v_n = t$ on the outer face of D_G such that v_1v_2 is an edge of the outer face. We will build up a good drawing of G by starting with v_1 and then adding v_2, v_3, \dots, v_n one by one to the outer face of the current drawing. As soon as we add a new vertex v_i , we also draw the initial pieces of the pending edges, and we make sure that the resulting drawing is equivalent to the drawing D_{G_i} .

Another property of the good drawing that we maintain is that every edge consists of precisely three pieces. (Actually, an edge may consist of fewer than 3 segments, because two consecutive pieces are allowed to have the same slope and form a longer segment) The middle piece will always be vertical, except for the middle piece of v_1v_2 .

Suppose without loss of generality that v_1 follows directly after v_2 in the clockwise order of the vertices around the outer face of D_G . Place v_1 and v_2 arbitrarily in the plane so that the x -coordinate of v_1 is smaller than the x -coordinate of v_2 . Connect v_1 and v_2 by an edge consisting of three segments: the segments incident to v_1 and v_2 are vertical and lie below them, while the middle segment has an arbitrary non-vertical regular slope. Draw a horizontal auxiliary line l_2 above v_1 and v_2 . Next, draw the initial pieces of the other (pending) edges incident to v_1 and v_2 , as follows. For $i = 1, 2$, draw a short segment from v_i for each of the edges incident to it (except for the edge v_1v_2 , which has already been drawn) so that the directed slopes of the edges (including v_1v_2) form a contiguous interval and their circular order is the same as in D_G . Each of these short segments will be followed by a vertical segment that reaches above l_2 . These vertical segments will belong to the middle pieces of the corresponding pending edges. Clearly, for a proper choice of the lengths of the short segments, no crossings will be created during this procedure. So far this drawing, including the partially drawn pending edges between $V(G_2)$ and $V(G) \setminus V(G_2)$, will be equivalent to the drawing D_{G_2} . As the algorithm progresses, the vertical segments will be further extended above l_2 , to form the middle segments of the corresponding edges. For an illustration, see Figure 35(a).

The remaining vertices $v_i, i > 2$, will be added to the drawing one by one, while maintaining the property that the drawing is equivalent to D_{G_i} and that the pending-order of the actual pending edges coincides with the order in which their vertical pieces reach the auxiliary line l_i . At the beginning of step $i + 1$, these conditions are obviously satisfied. Now we show how to place v_{i+1} .

Consider the set X of intersection points of the vertical (middle) pieces of all pending edges between $V(G_i)$ and $V(G) \setminus V(G_i)$ with the auxiliary line l_i . By Proposition 8.11, the intersection points corresponding to the pending edges incident to v_{i+1} must be consecutive elements of X . Let m be (one of) the median element(s) of X . Place v_{i+1} at a point above m , so that the x -coordinates of v_{i+1} and m coincide, and connect it to m . (In this way, the corresponding edge has only one bend, because its second and third piece are both vertical.) We also connect v_{i+1} to the upper endpoints of the appropriately extended vertical segments passing through the remaining elements of X , so that the directed slopes of the segments leaving v_{i+1} form a contiguous interval of regular slopes. For an illustration see Figure 35(b). Observe that this step can always be performed, because, by the definition of st -orderings, the number of edges leaving v_{i+1} is strictly smaller than d . This is not necessarily true in the last step, but then we have $v_n = t$, and we assumed that the degree of t was

smaller than d . To complete this step, draw a horizontal auxiliary line l_{i+1} above v_{i+1} and extend the vertical portions of those pending edges between $V(G_i)$ and $V(G) \setminus V(G_i)$ that were not incident to v_{i+1} until they hit the line l_{i+1} . (These edges remain pending in the next step.) Finally, in a small vicinity of v_{i+1} , draw as many short segments from v_{i+1} using the remaining directed slopes as many pending edges connect v_{i+1} to $V(G) \setminus V(G_{i+1})$. Make sure that the directed slopes used at v_{i+1} form a contiguous interval and the circular order is the same as in D_G . Continue each of these short segments by adding a vertical piece that hits the line l_{i+1} . The resulting drawing, including the partially drawn pending edges, is equivalent to $D_{G_{i+1}}$.

In the final step, if we place the auxiliary line l_{n-1} high enough, then the whole drawing will be contained in the v_n -wedge and we obtain a drawing that meets the requirements. \square

8.4 Lower Bounds

In this section, we construct a sequence of planar graphs, providing a nontrivial lower bound for the planar slope number of bounded degree planar graphs. They also require more than the trivial number $(\lceil d/2 \rceil)$ slopes, even if we allow one bend per edge. Remember that if we allow *two* bends per edge, then, by Theorem 8.3, for all graphs with maximum degree $d \geq 3$, except for the octahedral graph, $\lceil d/2 \rceil$ slopes are sufficient, which bound is optimal.

Theorem 8.13. *For any $d \geq 3$, there exists a planar graph G_d with maximum degree d , whose planar slope number is at least $3d - 6$. In addition, any drawing of G_d with at most one bend per edge requires at least $\frac{3}{4}(d - 1)$ slopes.*




Figure 36: Lower bounds

Proof. The construction of the graph G_d is as follows. Start with a graph of 6 vertices, consisting of two triangles, abc and $a'b'c'$, connected by the edges aa' , bb' , and cc' (see Figure

36(a)). Add to this graph a cycle C of length $3(d-3)$, and connect $d-3$ consecutive vertices of C to a , the next $d-3$ of them to b , and the remaining $d-3$ to c . Analogously, add a cycle C' of length $3(d-3)$, and connect one third of its vertices to a' , one third to b' , one third to c' . In the resulting graph, G_d , the maximum degree of the vertices is d .

In any crossing-free drawing of G_d , either C lies inside the triangle abc or C' lies inside the triangle $a'b'c'$. Assume by symmetry that C lies inside abc , as in Figure 36(a).

If the edges are represented by straight-line segments, the slopes of the edges incident to a, b , and c are all different, except that aa', bb' , and cc' may have the same slope as some other edge. Thus, the number of different slopes used by any straight-line drawing of G_d is at least $3d-6$.

Suppose now that the edges of G_d are represented by polygonal paths with at most one bend per edge. Assume, for simplicity, that every edge of the triangle abc is represented by a path with exactly one bend (otherwise, an analogous argument gives an even better result). Consider the $3(d-3)$ polygonal paths connecting a, b , and c to the vertices of the cycle C . Each of these paths has a segment incident to a, b , or c . Let S denote the set of these segments, together with the 6 segments of the paths representing the edges of the triangle abc .

Claim 8.14. *The number of segments in S with any given slope is at most 4.*

Proof. The sum of the degrees of any polygon on k vertices is $(k-2)\pi$. Every direction is covered by exactly $k-2$ angles of a k -gon (counting each side $1/2$ times at its endpoints). Thus, if we take every other angle of a hexagon, then, even including its sides, every direction is covered at most 4 times. (See Figure 36(b).) \square

The claim now implies that for any drawing of G with at most one bend per edge, we need at least $(3(d-3) + 6)/4 = \frac{3}{4}(d-1)$ different slopes. \square

Part III

Conjectures, Bibliography and CV

9 Summary of Interesting Open Questions and Conjectures

9.1 Questions about Decomposition of Coverings

Conjecture. (Pach) *All planar convex sets are cover-decomposable.*

Conjecture 3.11. *There is a constant m such that any m -fold covering of the plane with translates of a convex quadrilateral can be decomposed into two coverings.*

Conjecture 3.12. *For any cover-decomposable polygon P , $m_k(P) = O(k)$.*

In fact, the following, more general version seems to be open.

Question 9.1. *Is it true that for any set P , $m_k(P) = O(k)$?*

We cannot even say anything about decomposition into three coverings.

Question 9.2. *Is it true that for any set P , if $m_2(P)$ exists then $m_3(P)$ also exists?*

This latter can be asked about abstract sets instead of geometric sets in the following way.

Suppose we have a finite system of sets, \mathcal{F} . We say that a multiset \mathcal{M} is a multiset of \mathcal{F} if its elements are from \mathcal{F} . We say that \mathcal{M} is t -fold if for every element of the ground set there are at least t sets from \mathcal{M} that contain it. We say that \mathcal{M} is k -wise decomposable if we can color the sets from \mathcal{M} with k colors such that every element is contained in a set of each color. Define m_k as the smallest number such that if a multiset of \mathcal{F} is m_k -fold, then it is also k -wise decomposable. This number always exists as it is easy to see that $m_k \leq (k-1)|\mathcal{F}| + 1$.

Question 9.3. *Can we bound $m_3(\mathcal{F})$ with some function of $m_2(\mathcal{F})$? (Independently from \mathcal{F} .)*

Question 9.4. *Is it true that $m_k(\mathcal{F}) = O(k \cdot m_2(\mathcal{F}))$?*

Tardos [T09] proved a strongly related result. For any m he constructed a set system, \mathcal{F} , that covers the ground set m -fold and any 2-fold covering of the ground set with a subsystem of \mathcal{F} is decomposable into two coverings but \mathcal{F} cannot be decomposed into three coverings.

Conjecture 4.11. *Three-dimensional convex sets are not cover-decomposable.*

Conjecture 4.12. *Closed, convex polygons are cover-decomposable.*

Question 4.13. *Are there polygons that are not totally-cover-decomposable but plane-cover-decomposable?*

Finally, another finite problem that has a strong connection to cover-decomposition.

Question 9.5. For $A \subset [n]$ denote by a_i the i^{th} smallest element of A .

For two k -element sets, $A, B \subset [n]$, we say that $A \leq B$ if $a_i \leq b_i$ for every i .

A k -uniform hypergraph $\mathcal{H} \subset [n]$ is called a shift-chain if for any hyperedges, $A, B \in \mathcal{H}$, we have $A \leq B$ or $B \leq A$. (So a shift-chain has at most $k(n - k) + 1$ hyperedges.)

Is it true that shift-chains have Property B^* if k is large enough?

An affirmative answer would be a huge step towards Pach's conjecture, that all planar convex sets are cover-decomposable. To see this, for any fixed convex set C and natural k , and any y real number, define $C(k; y)$ as the translate of C which

- (1) contains exactly k points of a given point set, S ,
- (2) the center of C has y -coordinate y ,
- (3) the center of C has minimal x -coordinate,

if such a translate exists. If we associate $i \in [n]$ to the element of S with the i^{th} smallest y -coordinate, then an easy geometric argument shows that $\mathcal{H} = \{C(k; y) \cap S \mid y \in \mathbb{R}\}$ is a shift-chain.

For $k = 2$ there is a trivial counterexample to the question: (12),(13),(23).

A magical counterexample was found for $k = 3$ by a computer program by Fulek [F10]:

(123),(124),(125),(135),(145),(245),(345),(346),(347),(357),
 (367),(467),(567),(568),(569),(579),(589),(689),(789).

If we allow the hypergraph to be the union of two shift-chains (with the same order), then the construction in Section 4 gives a counterexample for any k , so arguments using that the average degree is small (like the Lovász Local Lemma) probably fail.

*A hypergraph \mathcal{H} has Property B if we can color its vertices with two colors such that no hyperedge is monochromatic.

9.2 Questions about Slope Number of Graphs

Conjecture 5.1. *The slope number of graphs with maximum degree 4 is unbounded.*

Question 9.6. *Can every cubic graph be drawn with the four basic directions*?
What if no three vertices can be collinear?*

Question 7.11. *Is it possible to draw all cubic graphs with a bounded number of slopes on a polynomial sized grid?*

Question 9.7. *Does the planar slope number of planar graphs with maximum degree d grow exponentially or polynomially with d ?*

Conjecture 9.8. *We can fix $O(d)$ slopes such that any planar graph with maximum degree d can be drawn with these slopes if each edge can have one bend.*

*Vertical, horizontal and the two diagonal (45°) directions.

10 Bibliography

References

- [A10] B. Ács: Síkfedések szétbonthatósága, Master Thesis (in Hungarian).
- [A08] G. Aloupis, J. Cardinal, S. Collette, S. Langerman, D. Orden, P. Ramos: Decomposition of Multiple Coverings into More Parts, to appear in *Discrete and Computational Geometry*, arXiv:0807.0552v1.
- [ABH06] G. Ambrus, J. Barát, and P. Hajnal: The slope parameter of graphs, *Acta Sci. Math. (Szeged)* **72** (2006), no. 3-4, 875–889.
- [BMW06] J. Barát, J. Matoušek, and D. Wood: Bounded-degree graphs have arbitrarily large geometric thickness, *Electronic J. Combinatorics* **13**/1 (2006), R3.
- [BPR03] S. Basu, R. Pollack, and M.-F. Roy: *Algorithms in Real Algebraic Geometry*, Springer-Verlag, Berlin, 2003.
- [BC78] E. A. Bender and E. R. Canfield: The asymptotic number of labeled graphs with given degree sequences, *J. Combin. Theory Ser. A* **24** (1978), 296–307.
- [BK98] T. Biedl and G. Kant: A better heuristic for orthogonal graph drawings, *Comput. Geom.* **9** (1998), pp. 159-180.
- [B07] A. L. Buchsbaum, A. Efrat, S. Jain, S. Venkatasubramanian and K. Yi: Restricted strip covering and the sensor cover problem, SODA 2007: 1056-1063.
- [BMP05] P. Brass, J. Pach, and W. Moser: *Research Problems in Discrete Geometry*, Springer, Berlin, 2005.
- [DETT99] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis: *Graph Drawing*, Prentice Hall, Upper Saddle River, N.J., 1999.
- [DEH00] M. B. Dillencourt, D. Eppstein, and D. S. Hirschberg: Geometric thickness of complete graphs, *J. Graph Algorithms Appl.* **4(3)** (2000), 5–17.
- [DESW07] V. Dujmović, D. Eppstein, M. Suderman, and D. R. Wood: Drawings of planar graphs with few slopes and segments, *Comput. Geom.* **38** (2007), no. 3, 194–212.
- [DSW04] V. Dujmović, M. Suderman, and D. R. Wood: Really straight graph drawings, in: *Graph Drawing, Lecture Notes in Computer Science 3383*, Springer, Berlin, 2004, 122–132.
- [DSW07] V. Dujmović, M. Suderman, and D. R. Wood: Graph drawings with few slopes, *Comput. Geom.* **38** (2007), 181–193.
- [DW06] V. Dujmović and D. R. Wood: Graph treewidth and geometric thickness parameters, *Proc. 13th Internat. Symp. on Graph Drawing (GD '05)* (P. Healy and N.S. Nikolov, eds.), *Lecture Notes in Comput. Sci.* **3843**, Springer, Berlin, 2006, 129–140.

- [DEK04] C. A. Duncan, D. Eppstein, and S. G. Kobourov: The geometric thickness of low degree graphs, *Proc. 20th ACM Symp. on Computational Geometry (SoCG'04)*, ACM Press, 2004, 340–346.
- [EMS10] M. Elekes, T. Mátrai and L. Soukup: On splitting infinite-fold covers, to appear in *Fund. Math.*
- [E05] M. Engelstein: Drawing graphs with few slopes, Research paper submitted to the Intel Competition for high school students, New York, October 2005.
- [E04] D. Eppstein: Separating thickness from geometric thickness, *Towards a Theory of Geometric Graphs (J. Pach, ed.)*, *Contemporary Mathematics* **342**, Amer. Math. Soc, Providence, 2004, 75–86.
- [E63] P. Erdős: On a combinatorial problem, *Nordisk Mat. Tidskr.* **11** 5-10.
- [EL75] P. Erdős and L. Lovász: Problems and results on 3-chromatic hypergraphs and some related questions, *Infinite and Finite Sets* (to Paul Erdős on his 60th birthday), II. North-Holland. pp. 609-627.
- [F48] I. Fáry: On straight line representation of planar graphs, *Acta Univ. Szeged. Sect. Sci. Math.* **11** (1948), 229–233.
- [FOR94] H. de Fraysseix, P. Ossona de Mendez, P. Rosenstiehl: On triangle contact graphs, *Combinatorics, Probability and Computing* **3** (1994), 233–246.
- [FPP89] H. de Fraysseix, J. Pach, and R. Pollack: How to draw a planar graph on a grid, *Combinatorica* **10**(1) (1990), 41–51.
- [F10] R. Fulek, personal communication.
- [GV10] M. Gibson and K. Varadarajan: Decomposing Coverings and the Planar Sensor Cover Problem, arXiv:0905.1093v1.
- [GT94] A. Garg and R. Tamassia: Planar drawings and angular resolution: Algorithms and bounds, *Second Annual European Symposium, Lecture Notes in Computer Science* **855**, Springer, 1994, 12–23.
- [HSV99] J. P. Hutchinson, T. C. Shermer, and A. Vince: On representations of some thickness-two graphs, *Comput. Geom.* **13** (1999), 161–171.
- [J86] R. E. Jamison: Few slopes without collinearity, *Discrete Math.* **60** (1986), 199–206.
- [JJ10] V. Jelinek, E. Jelinková, J. Kratochvíl, B. Lidický, M. Tesař, and T. Vyskočil: The planar slope number of planar partial 3-trees of bounded degree, *Graph Drawing, Lecture Notes in Computer Science* **5849**, Springer, Berlin, 2010, 304–315.
- [K73] P. C. Kainen: Thickness and coarseness of graphs, *Abh. Math. Sem. Univ. Hamburg* **39** (1973), 88–95.

- [K36] P. Koebe: Kontaktprobleme der konformen Abbildung, *Berichte Verhand. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Klasse* **88** (1936) 141–164.
- [KPP10] B. Keszegh, J. Pach and D. Pálvölgyi: Drawing planar graphs of bounded degree with few slopes, in preparation.
- [KPPT10] B. Keszegh, J. Pach, D. Pálvölgyi, and G. Tóth: Cubic graphs have bounded slope parameter, *J. Graph Algorithms Appl.* **14**(1): 5-17 (2010).
- [KPPT08] B. Keszegh, J. Pach, D. Pálvölgyi, and G. Tóth: Drawing cubic graphs with at most five slopes. *Comput. Geom.* **40** (2008), no. 2, 138–147.
- [LEC67] A. Lempel, S. Even, I. Cederbaum: An algorithm for planarity testing of graphs, *Theory of Graphs* (P. Rosenstiehl, ed.), Gordon and Breach, New York, 1967, 215–232 .
- [LMS91] Y. Liu, A. Morgana, B. Simeone: General theoretical results on rectilinear embeddability of graphs, *Acta Math. Appl. Sinica* **7** (1991), 187–192.
- [MP94] S. Malitz and A. Papakostas: On the angular resolution of planar graphs, *SIAM J. Discrete Math.* **7** (1994), 172–183.
- [MP86] P. Mani-Levitska, J. Pach: Decomposition problems for multiple coverings with unit balls, manuscript, 1986.
- [1] B. Mohar: A polynomial time circle packing algorithm, *Discrete Math.* **117** (1993), 257–263.
- [MSz07] P. Mukkamala and M. Szegedy: Geometric representation of cubic graphs with four directions, *Comput. Geom.* **42** (2009), 842–851.
- [MOS98] P. Mutzel, T. Odenthal, and M. Scharbrodt: The thickness of graphs: a survey, *Graphs Combin.* **14** (1998), 59–73.
- [P80] J. Pach: Decomposition of multiple packing and covering, *Diskrete Geometrie*, 2. Kolloq. Math. Inst. Univ. Salzburg, 169-178, 1980.
- [P86] J. Pach: Covering the plane with convex polygons, *Discrete and Computational Geometry* **1** (1986), 73-81.
- [PP06] J. Pach and D. Pálvölgyi: Bounded-degree graphs can have arbitrarily large slope numbers, *Electronic J. Combinatorics* **13**/1 (2006), N1.
- [PPT10] J. Pach, D. Pálvölgyi, and G. Tóth: Survey on the Decomposition of Multiple Coverings, in preparation.
- [PTT05] J. Pach, G. Tardos, G. Tóth: Indecomposable coverings, *The China–Japan Joint Conference on Discrete Geometry, Combinatorics and Graph Theory (CJCDGCGT 2005), Lecture Notes in Computer Science* **4381**, Springer, Berlin, 2007.
- [PT06] J. Pach and G. Tóth: Crossing number of toroidal graphs, *Graph drawing, Lecture Notes in Comput. Sci.* **3843**, Springer, Berlin, 2006, 334–342.

- [PT07] J. Pach, G. Tóth: Decomposition of multiple coverings into many parts, *Proc. 23rd ACM Symposium on Computational Geometry, (SoCG07) 2007*, 133-137. Also in: *Comput. Geom. Theory and Applications*, accepted.
- [P10] D. Pálvölgyi: Indecomposable coverings with concave polygons, to appear in *Discrete and Computational Geometry*.
- [PT10] D. Pálvölgyi and G. Tóth: Convex polygons are cover-decomposable, *Discrete and Computational Geometry* **43**(3): 483–496 (2010).
- [T09] G. Tardos, personal communication.
- [TT07] G. Tardos, G. Tóth: Multiple coverings of the plane with triangles, *Discrete and Computational Geometry* **38** (2007), 443-450.
- [U53] P. Ungar: On diagrams representing maps, *J. London Math. Soc.* **28** (1953), 336–342.
- [WC94] G. A. Wade and J.-H. Chu: Drawability of complete graphs using a minimal slope set, *The Computer J.* **37**/2 (1994), 139–142.

11 CURRICULUM VITÆ

PERSONAL DATA

First name: **Dömötör**

Surname: **Pálvölgyi**

Date of birth: 4th July, 1981.

Nationality: Hungarian

Other languages: English (fluent), French (good), Italian, German, Telugu (basic)

e-mail: *doemoetoer.palvoelgyi@epfl.ch*

STUDIES

2008– **Ecole Polytechnique Fédérale de Lausanne (EPFL)**

Assistant-doctorant of János Pach

2000–2008 Eötvös Loránd University (ELTE), Budapest

Graduate Studies in Pure Mathematics, Theoretical Computer Science

-Supervisor: Zoltán Király.

-State scholarship for research in Communication Complexity, Complexity Theory, Combinatorial Geometry, Extremal Combinatorics, Network Information Flows since 2005.

-Member of the Communication Networks Laboratory.

Undergraduate Studies at Faculty of Science, Department of Mathematics

-Master thesis in Communication Complexity 2005.

-Holder of the State Scholarship Award (2004–2005).

-Holder of the Outstanding Student of the Faculty Award (2004).

1994–2000 Fazekas Mihály High School, Budapest

Class specialized in mathematics

-Hungarian baccalaureate exams (qualification: outstanding), 2000.

-International Mathematical Olympiad, South Korea, silver medal, 2000.

-National Mathematics Competition (OKTV), Hungary, 2nd place, 2000

-International Hungarian Mathematics Competition, Slovakia, 1st prize, 2000.

TEACHING EXPERIENCE

2008–

-Teaching Assistant for Linear Algebra, Calculus, Geometric Graphs at EPFL.

2005–2008

-Lecturer and Teaching Assistant for Complexity Theory, Theory of Algorithms at ELTE.

2002–2005

-Teaching Assistant for Calculus, Real Analysis, Measure Theory at ELTE.

PUBLICATIONS

Submitted/Accepted:

- Bin Packing via Discrepancy of Permutations (with Friedrich Eisenbrand and Thomas Rothvoß), submitted.
- Almost optimal pairing strategy for Tic-Tac-Toe with numerous directions (with Padmini Mukkamala), submitted.
- Unique-maximum and conflict-free colorings for hypergraphs and tree graphs. (with Panagiotis Cheilaris and Balázs Keszegh), submitted.
- Indecomposable coverings with concave polygons, to appear in: Discrete and Computational Geometry.
- Permutations, hyperplanes and polynomials over finite fields (with András Gács, Tamás Héger and Zoltán Lóránt Nagy), to appear in: Finite Fields and Their Applications. (Earlier version: 22nd British Combinatorial Conference.)

2010

- On weakly intersecting pairs of sets (with Zoltán Király, Zoltán Lóránt Nagy and Mirkó Visontai), presented at: 7th International Conference on Lattice Path Combinatorics and Applications.
- Vectors in a Box (with Kevin Buchin, Jiri Matousek and Robin A. Moser), presented at: Coimbra Meeting on 0-1 Matrix Theory and Related Topics.
- Consistent digital line segments (with Tobias Christ and Milos Stojakovic), in: SoCG 2010.
- Convex polygons are cover-decomposable (with G. Tóth), in: Discrete and Computational Geometry.
- Testing additive integrality gaps (with Friedrich Eisenbrand, Nicolai Hähnle and Gennady Shmonin), in: SODA 2010.
- Cubic Graphs Have Bounded Slope Parameter (with B. Keszegh, J. Pach, and G. Tóth), in: J. Graph Algorithms Appl. 14(1): 5-17 (2010). (Earlier version: Proceedings of Graph Drawing 2008, 50–60.)
- Finding the biggest and smallest element with one lie (with D. Gerbner, B. Patkós and G. Wiener), in: Discrete Applied Mathematics 158(9): 988-995 (2010). (Earlier version: International Conference on Interdisciplinary Mathematical and Statistical Techniques - IMST 2008 / FIM.)
- Polychromatic Colorings of Arbitrary Rectangular Partitions (with D. Gerbner, B. Keszegh, N. Lemons, C. Palmer and B. Patkós), in: Discrete Math. 310, No. 1, 21-30 (2010). (Earlier version: 6th Japanese-Hungarian Symposium on Discrete Mathematics and Its Applications.)

2009

- 2D-TUCKER is PPAD-complete, in: WINE 2009.
- Combinatorial necklace splitting, in: Electronic J. Combinatorics 16 (1) (2009), R79.

-Deciding Soccer Scores and Partial Orientations of Graphs, in: Acta Universitatis Sapientiae, Mathematica, 1, 1 (2009) 35–42. (Earlier version in: EGRES Technical Reports 2008.)

2008

-Drawing cubic graphs with at most five slopes (with B. Keszegh, J. Pach, and G. Tóth), in: Comput. Geom. 40 (2008), no. 2, 138–147. (Earlier version: Graph Drawing 2006, Lecture Notes in Computer Science 4372, Springer, 2007, 114–125.)

2007

-Revisiting sequential search using question-sets with bounded intersections, in: Journal of Statistical Theory and Practice Vol 1, Num 2 (2007).

2006

-P2T is NP-complete, in: EGRES Quick-Proofs 2006.
-Bounded-degree graphs can have arbitrarily large slope numbers (with J. Pach), Electronic J. Combinatorics 13 (1) (2006), N1.

2005

-Balj6 S rnyak, in: Matematikai Lapok (in Hungarian). English title: A short proof of the Kruskal-Katona theorem.
-Communication Complexity (Master’s Thesis). Supervisor: Z. Kiraly.

ATTENDED CONFERENCES AND WORKSHOPS

2010

-1st Emlektabla Workshop in Gyongyostarjan, July 26–29.
-8th Gremo’s Workshop on Open Problems in Morschach (SZ), June 30–July 2.
-Coimbra Meeting on 0–1 Matrix Theory and Related Topics in Coimbra, June 17–19.
-3eme cycle romand de Recherche Operationnelle in Zinal, January 17–21.

2009

-WINE 2009 in Rome, December 16–18.
-Workshop on Combinatorics: Methods and Applications in Mathematics and Computer Science at IPAM, UCLA, September 7–November 7.
-The 14th International Conference on Random Structures and Algorithms in Poznan, August 3–7.
-7th Gremo’s Workshop on Open Problems in Stels (GR), July 6–10.
-Algorithmic and Combinatorial Geometry in Budapest, June 15–19.
-6th Japanese-Hungarian Symposium on Discrete Mathematics and Its Applications in Budapest, May 16–19.
-3eme cycle romand de Recherche Operationnelle in Zinal, January 18–22.

2008

-16th International Symposium on Graph Drawing (GD 2008) in Heraklion, September

21–24.

-Building Bridges (L. Lovász 60), Fete of Combinatorics and Computer Science in Budapest and in Keszthely, Lake Balaton, August 5 – 15.

-Bristol Summer School on Probabilistic Techniques in Computer Science, July 6–11.

-International Conference on Interdisciplinary Mathematical and Statistical Techniques (IMST 2008 / FIM XVI) at University of Memphis, May 15–18.

-Expanders in Pure and Applied Mathematics at IPAM, UCLA, February 11–15.

2007

-HSN Spring Workshop in Balatonkenese, May 31–June 1.

-Extremal Combinatorics Workshop at Rényi Institute of Mathematics in Budapest, June 4–8.

-Advanced Course on Analytic and Probabilistic Techniques in Combinatorics at Centre de Recerca Matemàtica in Barcelona, January 15–26.

2006

-Final Combstru Workshop at Universitat Politècnica de Catalunya in Barcelona, September 26–28.

-14th International Symposium on Graph Drawing (GD 2006) in Karlsruhe, September 18–20.

-Horizon of Combinatorics, EMS Summer School and Conference at Rényi Institute of Mathematics in Budapest and in Balatonalmádi, July 10–22.

-HSN Spring Workshop in Balatonkenese, May 23–24.

2005

-European Conference on Combinatorics, Graph Theory, and Applications (EuroComb 2005) at Technische Universität in Berlin, September 5–9.

Summer 2004 RIPS, UCLA

I participated in an REU (Research & Education for Undergraduates) and I was assigned to a project sponsored by BioDiscovery. Our goal was to create an algorithm that would reveal which genes are responsible for cancer from an input database.

Spring 2002 Hungarian-Dutch Exchange Program, Eindhoven

The program was sponsored by Philips, and we were asked to work on various projects.