

Cycle Packing and Cycle Transversal in Graphs on Oriented Surfaces

DISSERTATION

ZUR

ERLANGUNG DES DOKTORGRADES (DR. RER. NAT.)

DER

MATHEMATISCH-NATURWISSENSCHAFTLICHEN FAKULTÄT

DER

RHEINISCHEN FRIEDRICH-WILHELMS-UNIVERSITÄT BONN

VORGELEGT VON

NIKLAS SCHLOMBERG

AUS

LIPPSTADT

BONN, JANUAR 2026

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen
Fakultät der Rheinischen Friedrich-Wilhelms-Universität Bonn

Gutachter/Betreuer: Professor Dr. Jens Vygen
Gutachter: Professor Dr. László A. Végh
Tag der Promotion: 17. April 2026
Erscheinungsjahr: 2026

Acknowledgements

I want to thank Prof. Dr. Jens Vygen for many years of great supervision. Also, thanks to my second reviewer, Prof. Dr. László Végh. I want to thank Dan Král' for inviting me to Brno. I am very grateful to Luise Puhmann for helpful discussions and \LaTeX support. Thanks to Mohit Singh for suggesting the Fractional Local Ratio Method. Finally, thanks to all my co-authors: Igor Balla, Marek Filakovský, Bartłomiej Kielak, Dan Král', Luise Puhmann, Hanjo Thiele, and Jens Vygen.

Contents

1	Introduction	1
1.1	The Erdős–Pósa property	3
1.2	Overview over our results	4
1.3	Related work	6
1.3.1	Planar graphs	6
1.3.2	More general graphs	8
2	Preliminaries	11
2.1	Basic notation	11
2.2	Problem definition	12
2.3	Uncrossable cycle families	13
2.3.1	Examples for uncrossable cycle families	13
2.3.2	A stronger uncrossing property	14
2.4	Oracles	16
2.5	Embedded graphs	18
2.5.1	Laminar cycle families	18
2.5.2	Euler’s formula	19
2.6	Uncrossing	21
2.6.1	Uncrossing in planar graphs	21
2.6.2	Uncrossing in bounded-genus graphs	22
2.6.3	Uncrossing an LP solution	23
2.7	Relations between the edge and vertex versions	25
2.8	NP-hardness	28
2.9	Lower bounds	29
2.9.1	Edge-disjoint packing and edge transversal	29
2.9.2	Vertex-disjoint packing and vertex transversal	30
3	Combinatorial approximation algorithms for planar cycle packing	33
3.1	A PTAS for face-minimal cycles	33
3.2	A $(3 + \varepsilon)$ -approximation	35
3.2.1	The algorithm	35
3.2.2	The edge-disjoint case	37
3.3	A $(2 + \varepsilon)$ -approximation for \mathcal{C}_{all}	38
3.4	A PTAS for laminar families	39

4	A simple LP-based approximation	45
4.1	Edge-disjoint packing in planar graphs	45
4.2	Vertex-disjoint packing in planar graphs	46
4.2.1	Nice paths	47
4.2.2	Proof of the Efficient Cycle Lemma	49
4.2.3	Tightness of the Efficient Cycle Lemma	51
4.3	Packing cycles in bounded-genus graphs	52
4.4	Weighted Cycle Packing	53
4.4.1	Rounding the LP with the Local Ratio Method	54
4.4.2	Applications of Weighted Cycle Packing	55
5	Bounding the integrality gap for planar cycle packing	59
5.1	A new rounding algorithm	59
5.2	Proof of the Structure Lemma	63
5.3	Improving the bounds below 3.5	68
6	A constant Erdős–Pósa ratio in bounded-genus graphs	73
6.1	Topological facts we need	74
6.2	Eliminating facial cycles	75
6.3	Bounding the integrality gap	78
7	The Erdős–Pósa ratio of odd cycles in planar graphs	85
7.1	Properties of odd cycles	85
7.2	Proof of the Main Theorem	88
7.2.1	High-level outline	88
7.2.2	Detailed proof	89
7.3	Structural properties of clouds	91
8	Maximum k-systems on the torus	97
8.1	Overview of the proof	98
8.2	Nice sets and k -systems	100
8.3	Nice sets with large height	102
8.4	Bounding the height of a nice set	105
8.4.1	Analysis of the linear program	108
8.4.2	Constant height	111
8.5	Nice sets with height at most three	112
9	Open Problems	117
A	Computer-assisted improvement over Theorem 8.16	119
A.1	Decreasing the threshold for maximum height-3 sets	119
A.2	Our algorithm for small k	123
A.3	Proof of Theorem 8.1	129

B Source code of our programs	131
B.1 Implementation of Algorithm 4	131
B.2 Implementation of Algorithm 5	133
B.3 Python script to compute the values $\rho_\ell, \gamma_\ell, \alpha_\ell$ and β_ℓ	137
Bibliography	141

Chapter 1

Introduction

This thesis is about two dual problems on cycles in (directed or undirected) graphs: The CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM. As input, both problems take a graph G and an implicitly given family \mathcal{C} of cycles in G . The CYCLE PACKING PROBLEM asks for a maximum-cardinality subset $\mathcal{S} \subseteq \mathcal{C}$ of pairwise vertex- or edge-disjoint cycles. On the other hand, the task of the (unweighted) CYCLE TRANSVERSAL PROBLEM is to find a minimum-cardinality subset T of vertices (or edges, respectively) that hits all cycles of \mathcal{C} , i.e., each cycle in \mathcal{C} contains some element of T .

Both problems are quite natural and well-studied. They contain several other important problems of Combinatorial Optimization as a special case. For example, the CYCLE PACKING PROBLEM contains the ODD CYCLE PACKING PROBLEM, the INDEPENDENT SET PROBLEM and the DISJOINT PATHS PROBLEM, while some special cases of the CYCLE TRANSVERSAL PROBLEM are the FEEDBACK VERTEX SET PROBLEM and the GRAPH BIPARTIZATION PROBLEM. All of these problems are NP-hard, even if the underlying graph G is planar.

Since our problems are quite hard in this generality, we will consider the case where G is planar or embedded in an orientable surface of constant genus. Furthermore, we will restrict the possibilities for our cycle family \mathcal{C} to those with a certain property, which is called *uncrossability* or *uncrossing property* due to Goemans and Williamson:

Definition 1.1 (Goemans, Williamson [41]). A family \mathcal{C} of cycles in a graph is called *uncrossable* if the following property holds.

Let $C_1, C_2 \in \mathcal{C}$ and let P_2 be a path in C_2 such that P_2 shares only its endpoints with C_1 . Then there is a path P_1 in C_1 between these endpoints such that $P_1 + P_2 \in \mathcal{C}$ and $(C_1 - P_1) + (C_2 - P_2)$ contains a cycle in \mathcal{C} (as an edge set). See Figure 1.1 for an example.

In fact, the class of uncrossable cycle families is quite rich: The most important examples for uncrossable cycle families in an undirected graph G are:

1. All cycles in G
2. All directed cycles in an orientation of G
3. All shortest cycles in G
4. All odd cycles in G




Figure 1.1: On the left, the path P_2 in the cycle $C_2 \in \mathcal{C}$ (green) shares only its endpoints with the cycle $C_1 \in \mathcal{C}$ (blue). As shown on the right, adding the path P_1 to P_2 yields a cycle in \mathcal{C} (orange), while the other edges contain another cycle in \mathcal{C} (violet).

5. All cycles in G that contain at least one edge from a given set $D \subseteq E(G)$ of demand edges
6. All cycles in G that contain exactly one edge from a given set $D \subseteq E(G)$ of demand edges (D -cycles)

An overview on the uncrossing property and more examples can be found in Section 2.3. Since both the CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM are still NP-hard for uncrossable cycle families in planar graphs, we study constant-factor approximation algorithms for both problems. Some of our results will be LP-based: The VERTEX-DISJOINT CYCLE PACKING PROBLEM and the VERTEX CYCLE TRANSVERSAL PROBLEM admit the natural LP relaxations

$$\max \left\{ \sum_{C \in \mathcal{C}} x_C : \sum_{C \in \mathcal{C}: v \in V(C)} x_C \leq 1 \ (v \in V(G)), \ x_C \geq 0 \ (C \in \mathcal{C}) \right\} \quad (1.1)$$

and

$$\min \left\{ \sum_{v \in V(G)} y_v : \sum_{v \in V(C)} y_v \geq 1 \ (C \in \mathcal{C}), \ y_v \geq 0 \ (v \in V(G)) \right\}. \quad (1.2)$$

We refer to (1.1) as the *(vertex-disjoint) cycle packing LP* and to (1.2) as the *(vertex) cycle transversal LP*. Similarly, the EDGE-DISJOINT CYCLE PACKING PROBLEM and the EDGE CYCLE TRANSVERSAL PROBLEM admit the LP relaxations

$$\max \left\{ \sum_{C \in \mathcal{C}} x_C : \sum_{C \in \mathcal{C}: e \in E(C)} x_C \leq 1 \ (e \in E(G)), \ x_C \geq 0 \ (C \in \mathcal{C}) \right\} \quad (1.3)$$

and

$$\min \left\{ \sum_{e \in E(G)} y_e : \sum_{e \in E(C)} y_e \geq 1 \ (C \in \mathcal{C}), \ y_e \geq 0 \ (e \in E(G)) \right\}, \quad (1.4)$$

which we refer to as the *edge-disjoint cycle packing LP* and the *edge cycle transversal LP*. Note that the cycle packing LP and cycle transversal LP are dual LPs.

The CYCLE TRANSVERSAL PROBLEM also admits a weighted version, the WEIGHTED CYCLE TRANSVERSAL PROBLEM, where additionally non-negative weights w on the vertices (or edges) are given. The task now is to find a minimum-weight set of vertices (or edges) that hits all cycles in \mathcal{C} . A straightforward modification of the cycle transversal LPs yields the *weighted cycle transversal LPs*:

$$\min \left\{ \sum_{v \in V(G)} w(v)y_v : \sum_{v \in V(C)} y_v \geq 1 \ (C \in \mathcal{C}), \ y_v \geq 0 \ (v \in V(G)) \right\} \quad (1.5)$$

and

$$\min \left\{ \sum_{e \in E(G)} w(e)y_e : \sum_{e \in E(C)} y_e \geq 1 \ (C \in \mathcal{C}), \ y_e \geq 0 \ (e \in E(G)) \right\}. \quad (1.6)$$

The integrality gaps of these six LPs and resulting min-max inequalities also constitute a major research topic of this thesis.

1.1 The Erdős–Pósa property

Let \mathcal{C} be a family of cycles in an (undirected) graph G . We denote the *packing number*, i.e., the maximum cardinality of a set of pairwise disjoint cycles from \mathcal{C} , by $\nu(\mathcal{C})$. The *transversal number*, i.e. the size of a minimum-cardinality transversal for \mathcal{C} is denoted by $\tau(\mathcal{C})$. Note that there are two versions of packing and transversal numbers: More precisely, we will use the notation $\nu_e(\mathcal{C})$ and $\tau_e(\mathcal{C})$ for the edge-disjoint packing and edge transversal number, while the vertex-disjoint packing and vertex transversal numbers are denoted by $\nu_v(\mathcal{C})$ and $\tau_v(\mathcal{C})$, respectively.

Since any feasible transversal needs to hit each cycle in a maximum-cardinality cycle packing, clearly $\tau(\mathcal{C}) \geq \nu(\mathcal{C})$ holds. A natural and well-studied question from Combinatorial Optimization asks about relations in the other direction, i.e. to bound the transversal number from above by some function of the packing number.

To this end, a famous result by Erdős and Pósa [32] states that there exists a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that $\tau_v(\mathcal{C}_{\text{all}}) \leq f(\nu_v(\mathcal{C}_{\text{all}}))$, where \mathcal{C}_{all} is the set of all cycles in a graph G . This property is known as *Erdős–Pósa property*. The same property also holds for the set of all directed cycles in a digraph [73]. However, the property does not hold for the sets of odd cycles [72, 70] or D -cycles [39] in undirected graphs, even if the underlying graph is embedded in the projective plane.

In planar graphs, however, we will see that the Erdős–Pósa property holds for any uncrossable family of cycles. In fact, in this case the function f can even be chosen linear, so the ratio $\frac{\tau(\mathcal{C})}{\nu(\mathcal{C})}$ is upper bounded by a constant. The supremum of this ratio is called the *Erdős–Pósa ratio*. A main topic of this thesis is to investigate bounds for the Erdős–Pósa ratio for uncrossable cycle families in planar graphs and graphs that are embedded in an orientable surface of constant genus.

One possibility to bound the Erdős–Pósa ratio is by using the LPs: For any LP the supremum of the ratio between optimum integral and optimum fractional solutions is called the *integrality gap* or *integrality ratio* for the LP. Since the LPs (1.1) and (1.2) as well as (1.3) and (1.4) are dual LPs their fractional optimum values coincide by LP duality. Thus, if we

bound the integrality gaps for the LPs (1.1) and (1.2) by α and β , respectively, then these bounds induce an upper bound of $\alpha \cdot \beta$ on the Erdős–Pósa ratio. Similarly we can bound the edge version of the Erdős–Pósa ratio by using the LPs (1.3) and (1.4).

1.2 Overview over our results

We start in Chapter 2 with an introduction to notation and fundamental facts that we need later. Apart from introducing examples of uncrossable cycle families (Section 2.3.1), we also give a slightly stronger, but equivalent definition of uncrossability (Section 2.3.2). This will become useful in Section 2.6 where we show how to *uncross* integral and fractional solutions to the packing LPs (1.1) and (1.3) for uncrossable \mathcal{C} . In fact, in the planar case, we show how to obtain an optimum LP solution where the support is *laminar*, i.e., the embeddings of no two cycles in the support “cross”. If G is embedded in an orientable surface of constant genus $g \geq 1$, we can still compute a near-optimum solution to (1.1) or (1.3) where the embeddings of any two cycles in the support cross at most once. We finish Chapter 2 with some reductions between edge and vertex versions of our problems (Section 2.7), a quite general NP-hardness proof for the planar VERTEX-DISJOINT CYCLE PACKING PROBLEM (Section 2.8) and some lower bound examples for the integrality gaps and Erdős–Pósa ratios (Section 2.9).

In Chapter 3 we give our first constant-factor approximation algorithms for the (vertex- or edge-disjoint) CYCLE PACKING PROBLEM. Our first main result is a relatively simple polynomial-time $(3 + \varepsilon)$ -approximation for any constant $\varepsilon > 0$ for uncrossable cycle families in planar graphs, which is presented in Section 3.2. As a main tool for this result we establish a PTAS for the case where all cycles have disjoint interior (Section 3.1). The results from Section 3.1 and 3.2 (and part of Chapter 2) are joint work with Hanjo Thiele and Jens Vygen and part of a paper that is published in SIAM Journal on Computing [79]; a preliminary abstract appeared in proceedings of SODA 2023.

For the rather simple case where \mathcal{C} is the family of all cycles in a planar graph G we show that a similar approach even yields a $(2 + \varepsilon)$ -approximation for any constant $\varepsilon > 0$ for the VERTEX-DISJOINT CYCLE PACKING PROBLEM. This result can be found in Section 3.3.

Next, in Section 3.4 we show how to obtain a PTAS (i.e., a polynomial-time $(1 + \varepsilon)$ -approximation for any constant $\varepsilon > 0$) for the CYCLE PACKING PROBLEM when \mathcal{C} is laminar. This is particularly interesting as it implies an $(\alpha^* + \varepsilon)$ -approximation algorithm for any $\varepsilon > 0$ for uncrossable \mathcal{C} , where α^* is the integrality gap of the cycle packing LP (1.1) on laminar cycle families in planar graphs. We call this value, which we will see to constitute an upper bound on both integrality gaps of (1.1) and (1.3) for uncrossable cycle families (cf. Section 2.7), the *laminar cycle packing integrality gap*. There is no known example showing that α^* exceeds 2.

In Chapter 4 we will give some first constant upper bounds on the integrality gaps of the cycle packing LPs (1.1) and (1.3) for uncrossable cycle families \mathcal{C} in planar graphs. To this end, we observe that an upper bound of 4 for the integrality gap of the fully planar edge-disjoint paths LP (LP (1.3) for the family \mathcal{C} of all D -cycles) by Garg, Kumar and Sebő [37] easily extends to arbitrary uncrossable cycle families. For the vertex-disjoint version (LP (1.1)) we use new ideas to upper bound the integrality gap by 5.

Our methods from Chapter 4 allow for two generalizations: First, if G is not planar, but embedded in an orientable surface of constant genus g , we can combine our methods with observations from Huang et al. [48] to give an upper bound of $O(g^2)$ on the integrality gaps

of (1.1) and (1.3) (Section 4.3). Furthermore, we get a polynomial-time $O(g^2)$ -approximation algorithm for the CYCLE PACKING PROBLEM as the proof is constructive. The second generalization is the WEIGHTED CYCLE PACKING PROBLEM: Here, each cycle in \mathcal{C} comes with a weight and we want find a cycle packing of maximum total weight. In Section 4.4 we show that this problem admits a polynomial-time 4-approximation for edge-disjoint packing and a 5-approximation for vertex-disjoint packing if the weighted version of the cycle packing LP allows for an optimum (fractional) solution with laminar support. This is the case e.g. for the weighted version of the DISJOINT PATHS PROBLEM. All results presented in Chapter 4 are based on joint work with Hanjo Thiele and Jens Vygen [79].

In Chapter 5 we give an improved upper bound of $\frac{20+\sqrt{130}}{9} < 3.5$ on the integrality gaps of both LP (1.1) and (1.3) for uncrossable cycle families in planar graphs. This result appeared in the proceedings of ICALP 2024 [78]. Our bounds on the integrality gaps of the cycle packing LPs are complemented by previous results on the integrality gaps of the cycle transversal LPs (1.2) and (1.4): Goemans and Williamson [41] gave an upper bound of 3 on the integrality gaps of these LPs for uncrossable cycle families \mathcal{C} in planar graphs. Later, Berman and Yaroslavtsev [16] improved the upper bound to 2.4. Both bounds even hold with respect to the weighted versions (1.5) and (1.6). As described in Section 1.1, combining this with our bound from Chapter 5 yields an upper bound of $2.4 \cdot \frac{20+\sqrt{130}}{9} < 8.38$ on the Erdős–Pósa ratio for uncrossable cycle families in planar graphs.

For graphs G that are embedded in an orientable surface of genus g there is less known about the cycle transversal LP. Sun [83] recently showed that the integrality gap is bounded by $O(g)$ for the family of directed cycles in an orientation of G . We extend ideas of [83] to show a similar bound of $O(g)$ for a more general class of (uncrossable) cycle families, which includes odd cycles or D -cycles, but not shortest cycles, for example (see Section 6.3). Together with our results from Section 4.3 this yields not only the Erdős–Pósa property, but even an upper bound of $O(g^3)$ on the Erdős–Pósa ratio for many of the most interesting cycle families in graphs of genus g .

One possibility to improve on the bounds for the Erdős–Pósa ratio from Chapter 5 is to consider a special family \mathcal{C} of cycles and use the structure of this particular family. In Chapter 7 we show that the Erdős–Pósa ratio for odd cycles in planar graphs does not exceed 4. This result is joint work with Luise Puhmann [69].

Finally, in Chapter 8 we consider a related, but slightly different problem: A theorem by Greene [44] states that any set of closed and simple curves on an orientable surface of genus g such that any two of the curves intersect in at most k points contains at most $O(g^{k+1} \log g)$ free homotopy classes. A recent improvement of this result for $k = 1$ by Aougab and Gaster [7] makes a key step for our $O(g^2)$ -approximation algorithm from Section 4.3. In Chapter 8 we review the simple case of $g = 1$, i.e., the surface is the torus. We show tight bounds for the maximum number of free homotopy classes of closed and simple curves on the torus that pairwise intersect in at most k points. This result is joint work with Igor Balla, Marek Filakovský, Bartłomiej Kielak, and Daniel Král’ [14]. Part of the result has been published at the European Conference on Combinatorics, Graph Theory and Applications (EuroComb) 2025.

\mathcal{C}	Packing		Transversal		Erdős–Pósa	
	approx	gap	approx	gap	ratio	
$\mathcal{C}_{\text{all}}^{\rightarrow}$	1 [59]	1 [59]	1 [59]	1 [59]	1 [59]	
\mathcal{C}_{all}	$2 + \varepsilon$ [20]	$[2, \mathbf{3.5}]_{\subset} [2, 4]$ [60]	1 [easy]	2 [easy]	4	Král’ [60]
\mathcal{C}_{odd}	2 [55]	2 [55]	1 [46, 31]	1 [31]	2	[55]
$\mathcal{C}_D^{\leftarrow}$	$\mathbf{3 + \varepsilon} < 4$ [37]	$[2, \mathbf{3.5}]_{\subset} [2, 4]$ [37]	$1 + \varepsilon$ [52]	$[1.5, 2]$ [47, 37]	$[2, 4]$	[36]
any	$\mathbf{3 + \varepsilon}$	$[2, \mathbf{3.5})$	2.4 [16]	$[2, 2.4]$ [16]	$[4, \mathbf{8.38})$	

Table 1.1: State of the art for *edge-disjoint* cycle packing and transversal of certain uncrossable families in *planar* graphs. The names of the cycle families are explained in Section 1.3.1; the last row refers to a general uncrossable family of cycles. The table shows the best known approximation ratios for cycle packing and cycle transversal and the known bounds on the integrality gaps of the LPs (1.3) and (1.4). The last column shows the known bounds on the worst ratio of transversal and packing number. Results marked [easy] are easy because minimal feedback edge sets are spanning trees in the planar dual. New results presented in this thesis are shown in bold blue; here we also show the previous best. See Section 1.3.1 for the theorems corresponding to the results marked in blue.

1.3 Related work

Except for the aforementioned results on the integrality gap of the cycle transversal LPs by Goemans and Williamson [41] and Berman and Yaroslavtsev [16], previous works on the CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM mostly only consider particular examples for the underlying cycle family \mathcal{C} . In this section we review some previous work on both problems (and the Erdős–Pósa ratio) for the most interesting and well-researched examples for the cycle family \mathcal{C} .

1.3.1 Planar graphs

We first review some previous work on the CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM in planar graphs for certain cycle families \mathcal{C} . An overview on the state of the art and our improvements is given in Tables 1.1 and 1.2. Note that even in planar graphs, due to a result by Yannakakis [87] the VERTEX CYCLE TRANSVERSAL PROBLEM is NP-hard for a large class of cycle families, including all examples presented in the tables.

Our first cycle family of interest is the set of all (undirected) cycles in G , which we denote by \mathcal{C}_{all} . In this case, the edge versions of our problems are relatively easy: Minimal edge transversals correspond to spanning trees in the planar dual of G . In particular, the CYCLE TRANSVERSAL PROBLEM can be solved optimally in polynomial time, and the transversal integrality gap is exactly 2. The Erdős–Pósa ratio here is known to be 4 (the upper bound follows from the 4-color theorem due to Ma, Yu and Zang [60], tightness was shown by an example by Král’, see [60]). Clearly, this bound also provides an upper bound of 4 on the integrality gap of the cycle packing LP. Theorem 5.1 improves this upper bound below 3.5. Caprara, Panconesi and Rizzi [20] showed that the corresponding CYCLE PACKING PROBLEM is NP-hard and complemented this result with a $(2 + \varepsilon)$ -approximation algorithm.

\mathcal{C}	Packing		Transversal				Erdős–Pósa ratio
	approx	gap	approx		gap		
$\mathcal{C}_{\text{all}}^{\rightarrow}$	$3 + \varepsilon < 15.95$ [19]	$[2, \mathbf{3.5}] \subset [2, 15.95]$ [19]	2.4	[16]	$[1.5, 2.4]$	[16]	$[2, \mathbf{8.38}] \subset [2, 38.28]$
\mathcal{C}_{all}	$2 + \varepsilon < 3$ [24, 60]	$[1.5, 3]$ [24, 60]	$1 + \varepsilon$	[53]	$[1.5, 2.4]$	[16]	$[2, 3]$ [24, 60]
\mathcal{C}_{odd}	$3 + \varepsilon < 6$ [54]	$[2, \mathbf{3.5}] \subset [2, 6]$ [54]	2.4	[16]	$[1.5, 2.4]$	[16]	$[2, \mathbf{4}] \subset [2, 6]$ [54]
$\mathcal{C}_D^{\leftarrow}$	$3 + \varepsilon$	$[2, \mathbf{3.5}]$	2.4	[16]	$[1.5, 2.4]$	[16]	$[2, \mathbf{8.38}]$
any	$3 + \varepsilon$	$[2, \mathbf{3.5}]$	2.4	[16]	$[2, 2.4]$	[16]	$[\mathbf{4}, \mathbf{8.38}]$

Table 1.2: State of the art for *vertex-disjoint* cycle packing and transversal of certain uncrossable families in *planar* graphs. The names of the cycle families are explained in Section 1.3.1; the last row refers to a general uncrossable family of cycles. The table shows the best known approximation ratios for cycle packing and cycle transversal and the known bounds on the integrality gaps of the LPs (1.1) and (1.2). The last column shows the known bounds on the worst ratio of transversal and packing number. New results presented in this thesis are shown in bold blue; here we also show the previous best. See Section 1.3.1 for the theorems corresponding to the results marked in blue.

For vertex-disjoint packing of arbitrary cycles in a planar graph [24] and [60] derived an upper bound of 3 on the Erdős–Pósa ratio from Euler’s formula (see Remark 2.25). In terms of approximating a maximum cycle packing, Theorem 3.10 improves on this by giving a $(2 + \varepsilon)$ -approximation algorithm for any $\varepsilon > 0$. The core argument is very similar to the corresponding edge-disjoint result in [20]. For the CYCLE TRANSVERSAL PROBLEM, which in this case is also called FEEDBACK VERTEX SET PROBLEM, Kleinberg und Kumar [53] gave a PTAS (in the unweighted case).

Next, we consider the set of all directed cycles in a (planar) digraph G . We call this cycle family $\vec{\mathcal{C}}_{\text{all}}$. The famous Lucchesi–Younger Theorem [59] shows that the edge version has Erdős–Pósa ratio 1, i.e. equality is attained in the inequality $\tau(\mathcal{C}) \geq \nu(\mathcal{C})$. For the vertex version, Reed and Shepherd [74] gave the first constant upper bound of 28 on the Erdős–Pósa ratio. After Fox and Pach (see [19]) improved this bound to 16.31, Cames van Batenburg, Esperet and Müller [19] gave a bound of 15.95. This work now decreases the bound below 8.38. We also improve the matching best known approximation guarantee of 15.95 [19] to $(3 + \varepsilon)$ and show that the corresponding CYCLE PACKING PROBLEM is NP-hard (Corollary 2.41).

Let us now consider the set \mathcal{C}_{odd} of all odd cycles in a planar graph G . As in the case for arbitrary cycles in G , there is a simpler characterization of edge transversals: They are given by T -joins in the planar dual of G , where T is the set of odd faces of G . In particular, the transversal LP corresponds to the T -join polyhedron and has integrality gap 1 [31]. For the Erdős–Pósa ratio, Král’ and Voss [55] used the same characterization to show a tight bound of 2. Note that a simple reduction by Reed [72] shows that both the corresponding edge-disjoint and vertex-disjoint planar CYCLE PACKING PROBLEM is NP-hard. For the vertex version, Fiorini et al. [33] gave a bound of 10 on the Erdős–Pósa ratio. After an improvement to 6 by Král’, Sereni and Stacho [54], Theorem 7.1 yields a bound of 4. Our methods also give the best known approximation guarantee and integrality gap bound for the CYCLE PACKING PROBLEM.

Finally, one of the most interesting and well-studied cycle families for both of our problems

is the family of D -cycles: Given a set D of demand edges in G , then a D -cycle is a cycle in G that contains exactly one demand edge. We denote the family of all D -cycles in G by $\mathcal{C}_D^=1$. Since removing the demand edge from a D -cycle results in a path in $G - D$ between the endpoints of the demand edge, the D -CYCLE PACKING PROBLEM is equivalent to the DISJOINT PATHS PROBLEM, and D -Cycle Packing in planar graphs corresponds to the DISJOINT PATHS PROBLEM in fully planar instances.

Both the edge-disjoint and the vertex-disjoint version of the Fully Planar DISJOINT PATHS PROBLEM are NP-hard [62]. For the edge-disjoint version, the first constant-factor approximations and bounds on the integrality gap were given by Huang et al. [47] and Garg, Kumar and Sebő [37]; the best upper bound on the integrality gap is 4 [37]. Garg and Kumar [36] showed that also the Erdős–Pósa ratio for this problem is at most 4. We improve on both the approximation guarantee (Theorem 3.5) and the upper bound for the integrality gap (Theorem 5.1). For the EDGE CYCLE TRANSVERSAL PROBLEM the best upper bound on the integrality gap is 2 [47, 37]. Klein, Mathieu and Zhou [52] gave a PTAS for the problem. Due to a result by Middendorf and Pfeiffer [62], the Fully Planar VERTEX-DISJOINT PATHS PROBLEM contains the edge-disjoint version as a special case. Here, our results yield the first constant-factor approximation algorithms.

For most of the uncrossable families discussed above, a lower bound of 2 on the integrality gaps of (1.1) and (1.3) is known, which is also the best-known lower bound for general uncrossable families. Most of the corresponding examples can be constructed by modifying a complete graph on 4 vertices. Regarding the Erdős–Pósa ratio the best known lower bound for most of our uncrossable cycle families is 2. For the family of all cycles in G , Král’ (see [60]) gave a lower bound of 4 in the edge-disjoint case. We show that this implies a lower bound of 4 also in the vertex-disjoint case, but for a different uncrossable cycle family (Proposition 2.43). See Section 2.9 for examples that prove the lower bounds presented in Tables 1.1 and 1.2.

There exist other examples of uncrossable cycle families that have been studied. For example, Rautenbach and Regen [71] considered the CYCLE PACKING PROBLEM with the family of shortest cycles in G , which is also uncrossable. Furthermore, the (uncrossable) family of all cycles that contain at least one vertex from a specified set $S \subseteq V(G)$ has been considered, for example by Goemans and Williamson [41]. For both problems the results presented in this thesis yield the best-known upper bounds for the integrality gaps of the corresponding cycle packing LPs.

For cycle families \mathcal{C} that are not uncrossable surprisingly little is known. For example, the set of all even cycles is not uncrossable. Here Göke et al. [42] generalized Goemans and Williamson’s [41] technique to get a constant upper bound on the vertex transversal LP; for the CYCLE PACKING PROBLEM no constant-factor approximation algorithm is known.

1.3.2 More general graphs

Now let us review known results on the CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM in more general graphs.

For several cycle families the CYCLE PACKING PROBLEM in general graphs is hard to approximate: If n is the number of vertices of the underlying graph then the problem of packing directed cycles in a digraph is quasi-NP-hard to approximate within a factor of $O(\log^{1-\varepsilon} n)$ for any $\varepsilon > 0$ [57]. Packing (arbitrary) cycles in undirected graphs is still quasi-NP-hard to approximate within a factor of $O(\log^{\frac{1}{2}-\varepsilon})$ [34]. Similarly, the CYCLE TRANSVERSAL PROBLEM

is hard to approximate: The VERTEX COVER PROBLEM, which is APX-hard [30], can be reduced to the CYCLE TRANSVERSAL PROBLEM where $\mathcal{C} = \mathcal{C}_{\text{all}}$ is the set of all cycles in G . Under the Unique Games Conjecture also the case where $\mathcal{C} = \overrightarrow{\mathcal{C}}_{\text{all}}$ is the set of directed cycles in a digraph is hard to approximate within a constant factor [45].

There are only few positive results for our problems in more general graphs. For the FEEDBACK VERTEX SET PROBLEM, i.e. the CYCLE TRANSVERSAL PROBLEM, where \mathcal{C} is the set of all cycles in a graph G , there exists a 2-approximation [15] which works even in the weighted case.

For the case where \mathcal{C} is the set of all odd cycles in G more is known: If G is highly connected, then the Erdős–Pósa ratio for this cycle family is at most 2 [70][86]. Also, the Erdős–Pósa property holds if G can be embedded in an orientable surface of bounded genus, which was first proved by Kawarabayashi and Nakamoto [51]. Conforti et al. [28] gave a bound on the Erdős–Pósa ratio of 19^{g+1} , which we improve to $O(g^3)$ in Chapter 6.

Finally, there exist some previous works considering the D -CYCLE PACKING PROBLEM which is, as noted above, equivalent to the DISJOINT PATHS PROBLEM: Chekuri, Khanna and Shepherd [22] showed how to obtain an $O(\sqrt{n})$ -approximation, where $n = |V(G)|$. On the other hand, Chuzhoy, Kim and Nimavat [27] showed that both the edge-disjoint and the vertex-disjoint versions cannot be approximated within a factor of $2^{O(\log^{1-\varepsilon} n)}$ for any $\varepsilon > 0$ unless $\text{NP} \subseteq \text{DTIME}(n^{\text{poly} \log n})$.

Recently, Huang et al. [48] gave an $O(g^2)$ -approximation for the edge-disjoint case if G can be embedded in an orientable surface of genus g . We extend this to the vertex-disjoint case in Section 4.3.

The integrality gap of the EDGE D -CYCLE TRANSVERSAL PROBLEM is also called the *multicut gap*. In general graphs, this multicut gap is in $\Theta(\log |D|)$ [38]. If $G - D$ is a tree then the multicut gap is 2 [39]. For planar G , Garg, Kumar and Sebő [37] proved an upper bound of 2. In the case that $G - D$ contains no $K_{r,r}$ -minor, Tardos and Vazirani [84] showed that the multicut gap is at most $O(r^3)$. Note that this implies an $O(g^{\frac{3}{2}})$ -upperbound in the case that $G - D$ has genus g and thus also an $O(1)$ -bound if $G - D$ is planar. Our results from Chapter 6 improve on this bound if G has constant genus.

In Chapter 8 we consider maximum k -systems on the torus. These are not directly related to the CYCLE PACKING PROBLEM or CYCLE TRANSVERSAL PROBLEM, but only to tools we use in Section 4.3 to extend ideas from the planar to the bounded-genus case. Therefore, we refer for an introduction to k -systems and related work to Chapter 8.

Chapter 2

Preliminaries

In this chapter we establish some fundamental facts and tools that we need in the subsequent chapters. We start with introducing some basic notation in Section 2.1. In Section 2.2 we state our problems formally.

Section 2.3 is dedicated to the uncrossing property. We present the most relevant examples of uncrossable cycle families and give a slightly stronger, but equivalent definition of the uncrossing property. We show how to implement the necessary oracle access for these cycle families in Section 2.4. Most of Sections 2.3 and 2.4 is based on joint work with Hanjo Thiele and Jens Vygen [79].

In Section 2.5 we introduce more notation that is related to graphs embedded in (orientable) surfaces. We review some easy and mostly well-known implications of Euler’s formula on those graphs in Section 2.5.2.

In Section 2.6 then we present uncrossing techniques for cycles in graphs that are embedded in an orientable surface. Most of these results are joint work with Hanjo Thiele and Jens Vygen [79].

In Section 2.7 we review relations between the edge and the vertex versions of both our main problems. The reductions we give are quite simple, but to the best of our knowledge they have not been stated in this generality before.

Section 2.8 contains an NP-hardness proof for the planar VERTEX-DISJOINT CYCLE PACKING PROBLEM in a relatively general setting.

Finally, we validate the lower bounds for the integrality gaps and the Erdős–Pósa ratios from Tables 1.1 and 1.2 in Section 2.9. Most of the examples are very easy. We also give a new lower bound for the vertex-disjoint Erdős–Pósa ratio for general uncrossable cycle families.

2.1 Basic notation

Throughout this thesis, G will denote a graph and \mathcal{C} will be a family of cycles in G . Usually, \mathcal{C} will be uncrossable. If not stated otherwise, all graphs that we consider will be undirected and contain no loops (i.e., edges where both endpoints coincide).

A *cycle* in G is a connected two-regular subgraph $C = (V(C), E(C))$ of G . We will sometimes also use the notion of a cycle for the edge set $E(C)$. If the underlying graph is embedded in a surface, it will sometimes be convenient to identify the cycle with its embedding in the surface. This will not lead to ambiguities.

Given an edge set $X \subseteq E(G)$ and weights or costs $w: E(G) \rightarrow \mathbb{R}$ we write $w(X)$ for $\sum_{x \in X} w(x)$. By $G - X$ we denote the subgraph of G that arises from G by deleting all edges in X . Similarly, if $X \subseteq V(G)$ is a set of vertices, $G - X$ is constructed from G by deleting all vertices in X and edges incident to X .

If \mathcal{C} is a family of cycles in G and X is a set of vertices (or edges) of G , then $\mathcal{C}[X]$ denotes the set of all cycles $C \in \mathcal{C}$ with $V(C) \subseteq X$ (or $E(C) \subseteq X$, respectively). Similarly, we denote by $V(\mathcal{C})$ and $E(\mathcal{C})$ the set of all vertices or edges that are part of a cycle in \mathcal{C} , respectively. The subgraph of G consisting of those vertices and edges is denoted by $G[\mathcal{C}] := (V(\mathcal{C}), E(\mathcal{C}))$. A *connected component* of \mathcal{C} is a maximal subset $\mathcal{C}' \subseteq \mathcal{C}$ such that $G[\mathcal{C}']$ is connected. Equivalently, $\mathcal{C}' = \mathcal{C}[X]$, where $X \subseteq V(G)$ is a connected component of $G[\mathcal{C}]$. We call \mathcal{C} *connected* if it has only one connected component.

2.2 Problem definition

In this section we define our main problems formally.

Definition 2.1 (cycle packing). Let G be a graph and \mathcal{C} a family of cycles in G . A *vertex-disjoint cycle packing* for \mathcal{C} is a subset $\mathcal{S} \subseteq \mathcal{C}$ of pairwise vertex-disjoint cycles. An *edge-disjoint cycle packing* for \mathcal{C} is defined analogously. We denote the maximum size of a vertex- or edge-disjoint cycle packing for \mathcal{C} by $\nu_v(\mathcal{C})$ and $\nu_e(\mathcal{C})$, respectively.

Definition 2.2 (cycle transversal). Let G be a graph and \mathcal{C} a family of cycles in G . A *vertex cycle transversal* for \mathcal{C} is a subset $T \subseteq V(G)$ such that $T \cap V(C) \neq \emptyset$ for all $C \in \mathcal{C}$. Again, an *edge cycle transversal* for \mathcal{C} is defined analogously. We denote the minimum size of a vertex or edge cycle transversal for \mathcal{C} by $\tau_v(\mathcal{C})$ and $\tau_e(\mathcal{C})$, respectively.

We will usually consider the case of vertex-disjoint cycle packings and vertex cycle transversals; all of our theorems allow for an edge version which can be proven similarly or even deduced from the vertex version by a simple reduction. Therefore, we use the notion of *cycle packings* and *cycle transversals* for the vertex versions.

Definition 2.3. The CYCLE PACKING PROBLEM is defined as follows.

Input: A graph G , a family \mathcal{C} of cycles in G (usually given by an oracle)

Task: Find a subset $\mathcal{S} \subseteq \mathcal{C}$ such that the cycles in \mathcal{S} are pairwise vertex-disjoint (in the VERTEX-DISJOINT CYCLE PACKING PROBLEM) or edge-disjoint (in the EDGE-DISJOINT CYCLE PACKING PROBLEM), respectively, such that $|\mathcal{S}|$ is maximum.

Definition 2.4. The CYCLE TRANSVERSAL PROBLEM is defined as follows.

Input: A graph G , a family \mathcal{C} of cycles in G (usually given by an oracle)

Task: Find a vertex set $T \subseteq V(G)$ (in the VERTEX CYCLE TRANSVERSAL PROBLEM) or an edge set $T \subseteq E(G)$ (in the EDGE CYCLE TRANSVERSAL PROBLEM) such that any $C \in \mathcal{C}$ contains an element of T , minimizing $|T|$.

Since the size of \mathcal{C} can be exponential in the size of G we do not want to list all cycles in \mathcal{C} in the input. Instead, we will assume \mathcal{C} to be accessed through a certain oracle. Our algorithms will require one of the following two oracles:

Definition 2.5 (weight oracle, support oracle). A *weight oracle*, for a family \mathcal{C} of cycles in a graph G , takes as input nonnegative weights for the edges of G ; it outputs a cycle in \mathcal{C} that has minimum total weight.

A *support oracle*, for a family \mathcal{C} of cycles in a graph G , takes as input a subset $X \subseteq E(G)$ of edges and outputs the union of the edge sets of all cycles $C \in \mathcal{C}$ that do not contain any edge of X .

Both of these oracles exist for all examples of uncrossable cycle families given in this thesis. Implementations and basic observations about the oracles can be found in Section 2.4.

2.3 Uncrossable cycle families

2.3.1 Examples for uncrossable cycle families

In this section we study the uncrossing property from Definition 1.1. First, we give some examples of cycle families that are always uncrossable.

Definition 2.6 (examples of uncrossable families). Let G be a graph. We define \mathcal{C}_{all} to denote the set of all cycles in G and \mathcal{C}_{odd} the set of all odd cycles in G (i.e., cycles with an odd number of edges). Given an orientation \vec{G} of G , let $\vec{\mathcal{C}}_{\text{all}}$ denote the set of all directed cycles in \vec{G} . For a given set $D \subseteq E(G)$ of demand edges, let $\mathcal{C}_D^{\geq 1}$ denote the set of cycles containing at least one edge from D and $\mathcal{C}_D^=1$ the set of cycles containing exactly one edge from D . Cycles in the latter family $\mathcal{C}_D^=1$ are also called *D-cycles*.

For several of these families, the following was already noted by [41]:

Proposition 2.7. *Let G be a graph and $D \subseteq E(G)$. The families \mathcal{C}_{all} , $\vec{\mathcal{C}}_{\text{all}}$, \mathcal{C}_{odd} , $\mathcal{C}_D^=1$, and $\mathcal{C}_D^{\geq 1}$ are all uncrossable.*

Proof. Let \mathcal{C} denote one of these five families of cycles. Let $C_1, C_2 \in \mathcal{C}$ and P_2 a path in C_2 such that P_2 shares only its endpoints with C_1 . Denote these endpoints by v and w . These partition C_1 into two v - w -paths P_1' and P_1'' .

For $\vec{\mathcal{C}}_{\text{all}}$ we take $P_1 \in \{P_1', P_1''\}$ such that $P_1 + P_2$ is a directed cycle. For \mathcal{C}_{all} , we can take $P_1 \in \{P_1', P_1''\}$ such that $C_1 - P_1 \neq C_2 - P_2$. In both cases, $(C_1 - P_1) + (C_2 - P_2)$ is Eulerian and contains some edge exactly once. In particular, $(C_1 - P_1) + (C_2 - P_2)$ contains some cycle.

For \mathcal{C}_{odd} , we take $P_1 \in \{P_1', P_1''\}$ such that $P_1 + P_2$ contains an odd number of edges (note that P_1' and P_1'' have different parity since C_1 is odd). Then also $(C_1 - P_1) + (C_2 - P_2)$ contains an odd number of edges and can be partitioned into cycles, at least one of which is odd.

For $\mathcal{C}_D^=1$, we take $P_1 \in \{P_1', P_1''\}$ such that $P_1 + P_2$ (and hence also $(C_1 - P_1) + (C_2 - P_2)$) contains exactly one edge of D (note that exactly one of P_1' and P_1'' contains an edge of D). Again, $(C_1 - P_1) + (C_2 - P_2)$ can be partitioned into cycles, one of which is a D -cycle.

For $\mathcal{C}_D^{\geq 1}$, choose $P_1 \in \{P_1', P_1''\}$ such that the subsets of edges in D differ for P_1 and P_2 and also for $C_1 - P_1$ and $C_2 - P_2$. Again, $(C_1 - P_1) + (C_2 - P_2)$ can be partitioned into cycles, at least one of which contains an edge in D and another edge. \square

In contrast, the set of even cycles is not uncrossable (except in special cases such as bipartite graphs). In directed graphs, neither D -cycles nor odd cycles nor cycles hitting D are uncrossable in general.

Another example of an uncrossable family that has been considered (e.g., by [54]) is the set of D -odd cycles: a cycle is D -odd if it contains an odd number of edges from D . However, both the CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM for D -odd cycles reduce to the corresponding problem for odd cycles by subdividing all edges in $E(G) \setminus D$; so we will not consider D -odd cycles any further.

Goemans and Williamson [41] also considered the set of cycles in an undirected graph G that contain at least one vertex from a subset $S \subseteq V(G)$. However, if we let D be the set of all edges with at least one endpoint in S , then this family is $\mathcal{C}_D^{\geq 1}$, so we do not need to consider this separately.

Another important example of an uncrossable cycle family is the set of all shortest cycles w.r.t. given edge lengths from a given uncrossable cycle family:

Proposition 2.8. *Let G be a graph and \mathcal{C} an uncrossable family of cycles in G . Let $l: E(G) \rightarrow \mathbb{Z}_{\geq 0}$. Define the cycle family $\mathcal{C}[l] := \{C \in \mathcal{C} : l(C) = \min_{C' \in \mathcal{C}} l(C')\}$. The cycle family $\mathcal{C}[l]$ is uncrossable.*

Proof. Let $C_1, C_2 \in \mathcal{C}[l]$ and P_2 a path in C_2 that shares only its endpoints, v and w , with C_1 . Since \mathcal{C} is uncrossable, choose a v - w -path P_1 in C_1 such that $C_* := P_1 + P_2 \in \mathcal{C}$ and $(C_1 - P_1) + (C_2 - P_2)$ contains a cycle $C'_* \in \mathcal{C}$. Now by definition of $\mathcal{C}[l]$ we have

$$l(C_*) + l(C'_*) \geq 2 \min_{C' \in \mathcal{C}} l(C') = l(C_1) + l(C_2) \geq l(C_*) + l(C'_*).$$

Thus, the above inequality is an equality and both C_* and C'_* are in $\mathcal{C}[l]$. □

Packing shortest cycles (i.e., $\mathcal{C}_{\text{all}}[\mathbb{1}_{E(G)}]$, where $\mathbb{1}_{E(G)}$ denotes the function that maps all edges to 1) has been considered for example by [71]. Note that the construction from Proposition 2.8 allows us to also model the DISJOINT SHORTEST PATHS PROBLEM by setting the values of l on the demand edges accordingly. Furthermore, Proposition 2.8 shows that the families of all cycles in $\mathcal{C}_D^{\geq 1}$ that contain a minimum number (or a minimum odd number) of D -edges are uncrossable.

Berman and Yaroslavtsev [16] considered another uncrossable family of cycles to model the planar STEINER FOREST PROBLEM as an (uncrossable) CYCLE TRANSVERSAL PROBLEM in planar graphs. There are probably more uncrossable cycle families of interest.

2.3.2 A stronger uncrossing property

Part of our algorithms are based on a slightly stronger uncrossing property, which we review in this section.

Definition 2.9 (strongly uncrossable). A family \mathcal{C} of cycles in a graph is called *strongly uncrossable* if the following property holds.

Let $C_1, C_2 \in \mathcal{C}$, and let v and w be two vertices that belong to both cycles C_1 and C_2 . Then there are v - w -paths P_1 in C_1 and P_2 in C_2 such that both $P_1 + P_2$ and $(C_1 - P_1) + (C_2 - P_2)$ contain a cycle in \mathcal{C} .

It is obvious that strongly uncrossable families of cycles are uncrossable. We will now show the nontrivial fact that the two definitions are in fact equivalent. To this end, we call an edge set *good* if it contains a cycle in \mathcal{C} . The following lemma is the core of the proof:

Lemma 2.10. *Let \mathcal{C} be an uncrossable family of cycles in a graph and $C_1, C_2 \in \mathcal{C}$. Let v and w be two vertices, P_1 a v - w -path in C_1 , and P_2 a v - w -path in C_2 . Then $P_1 + P_2$ or $P_1 + (C_2 - P_2)$ is good.*

Proof. By induction on the number of edges in P_1 .

Case 1: If v and w are the only common vertices of P_1 and C_2 , then (since \mathcal{C} is uncrossable) $P_1 + P_2 \in \mathcal{C}$ or $P_1 + (C_2 - P_2) \in \mathcal{C}$.

Case 2: Otherwise (by possibly swapping P_2 and $C_2 - P_2$) we may assume that P_1 and P_2 have a common inner vertex x . Let x be the first such vertex when traveling along P_2 from v . Let P'_1 and P''_1 be the v - x -subpath and the x - w -subpath of P_1 , and let P'_2 and P''_2 be the v - x -subpath and the x - w -subpath of P_2 . (See Figure 2.1.) Then P'_2 and P''_1 have no inner vertex in common.

Suppose $P_1 + P_2$ is not good. First we apply the induction hypothesis to C_1 , C_2 , P'_1 , and P'_2 . Since $P_1 + P_2$ is not good, $P'_1 + P'_2$ is not good either, and hence the induction hypothesis implies that $A = P'_1 + (C_2 - P'_2)$ is good.

Since A is good, it contains a cycle $C_3 \in \mathcal{C}$. If C_3 does not contain any edge of P'_2 , we conclude that $P_1 + (C_2 - P_2)$ is good, as required. Otherwise (by the choice of x), C_3 contains P'_2 entirely.

Now we apply the induction hypothesis to C_1 , C_3 , P'_1 , and P'_2 . Since $P_1 + P_2$ is not good, $P'_1 + P'_2$ is not good either, and hence the induction hypothesis implies that $B = P'_1 + (C_3 - P'_2)$ is good. Since $B \subseteq P'_1 + (A - P'_2) = P_1 + (C_2 - P_2)$, we conclude that $P_1 + (C_2 - P_2)$ is good. \square




Figure 2.1: Illustrating the proof of Lemma 2.10 (Case 2). Here P_1 is the blue horizontal v - w -path, and P_2 is the red curved path.

This implies:

Theorem 2.11. *Any uncrossable family of cycles is strongly uncrossable.*

Proof. Let \mathcal{C} be an uncrossable family of cycles and $C_1, C_2 \in \mathcal{C}$. Let v and w be two vertices that belong to both cycles C_1 and C_2 , and let P_1 be a v - w -path in C_1 and P_2 a v - w -path in C_2 .

If $P_1 + P_2$ and $(C_1 - P_1) + (C_2 - P_2)$ are good, we are done. Suppose, without loss of generality, $P_1 + P_2$ is not good. Then we apply Lemma 2.10 (once to C_1, C_2, P_1, P_2 , and once to C_2, C_1, P_2, P_1) and obtain that $P_1 + (C_2 - P_2)$ is good and $P_2 + (C_1 - P_1)$ is good. Again, we are done. \square

2.4 Oracles

This section is dedicated to implementing the weight and support oracle from Definition 2.5. For all uncrossable cycles mentioned in Section 2.3.1 both oracles can be implemented in polynomial time. We start with the weight oracle:

Proposition 2.12. *Let G be a graph and $D \subseteq E(G)$. For each of the families \mathcal{C}_{all} , $\overrightarrow{\mathcal{C}}_{\text{all}}$, \mathcal{C}_{odd} , $\mathcal{C}_D^{\geq 1}$, and $\mathcal{C}_D^{\equiv 1}$ there is a polynomial-time algorithm that implements the weight oracle.*

Proof. For the weight oracle, we are given nonnegative weights of the edges of G . Finding a minimum-weight cycle in \mathcal{C}_{all} , $\overrightarrow{\mathcal{C}}_{\text{all}}$, $\mathcal{C}_D^{\geq 1}$ or $\mathcal{C}_D^{\equiv 1}$ is easy, for example by applying Dijkstra's shortest path algorithm for finding a minimum-weight path from s to t in $G - e$ or $G - D$ for each edge $e = \{t, s\}$ or $e = (t, s)$ in $E(G)$ or D , respectively. Finding an odd cycle of minimum total weight reduces to weighted matching (cf. Section 29.11e of [80]). \square

Proposition 2.13. *Let G be a graph and \mathcal{C} an uncrossable family of cycles in G , given by a weight oracle. Let $l: E(G) \rightarrow \mathbb{Z}_{\geq 0}$. Then there exists a polynomial-time algorithm that implements the weight oracle for $\mathcal{C}[l]$.*

Proof. Again, assume we are given nonnegative edge weights $w: E(G) \rightarrow \mathbb{R}_{\geq 0}$. Let $C^* \in \mathcal{C}$ be a cycle that lexicographically minimizes $(l(C^*), w(C^*))$. We need to find a cycle $C \in \mathcal{C}$ with $(l(C), w(C)) = (l(C^*), w(C^*))$. For this, we choose a large constant

$$c > |V(G)| \cdot \max_{e \in E(G)} w(e) \geq \max_{C \in \mathcal{C}} w(C).$$

Define $w': E(G) \rightarrow \mathbb{R}_{\geq 0}$ as $w'(e) := w(e) + c \cdot l(e)$ for all $e \in E(G)$. By the weight oracle for \mathcal{C} , we can find a cycle $C \in \mathcal{C}$ minimizing $w'(C)$. If $l(C) \neq l(C^*)$, we have $l(C) \geq l(C^*) + 1$ due to integrality of l and thus

$$w'(C) = w(C) + c \cdot l(C) \geq c \cdot l(C^*) + c > c \cdot l(C^*) + w(C^*) = w'(C^*).$$

This contradicts the choice of C . Therefore, $l(C) = l(C^*)$ and by choice of C also $w(C) = w(C^*)$. \square

Note that the weight oracle also allows for vertex weights instead of edge weights:

Proposition 2.14. *Let G be a graph and \mathcal{C} an uncrossable family of cycles in G , given by a weight oracle. Let $w: V(G) \rightarrow \mathbb{R}_{\geq 0}$. Then we can find a cycle in \mathcal{C} minimizing $w(V(C))$ in polynomial time.*

Proof. Define $w': E(G) \rightarrow \mathbb{R}_{\geq 0}$ by $w'(\{s, t\}) := w(s) + w(t)$. Using our weight oracle, we can find a cycle $C \in \mathcal{C}$ that minimizes $w'(E(C)) = 2w(V(C))$. \square

Next, we implement the support oracle.

Proposition 2.15. *Let G be a graph and $D \subseteq E(G)$. Let $l: E(G) \rightarrow \mathbb{Z}_{\geq 0}$. For each of the families $\mathcal{C}_{\text{all}}[l]$, $\overrightarrow{\mathcal{C}}_{\text{all}}[l]$, $\mathcal{C}_{\text{odd}}[l]$, $\mathcal{C}_D^{\geq 1}[l]$, and $\mathcal{C}_D^{\equiv 1}[l]$ there is a polynomial-time algorithm that implements the support oracle.*

Proof. Let \mathcal{C} be one of these five families. Let $l^* := \min_{C \in \mathcal{C}} l(C)$. It suffices to show how to compute l^* and the set E^* of edges that belong to at least one cycle in \mathcal{C} : Given a set $X \subseteq E(G)$ we can first check whether the deletion of X increases the value of l^* . If this is the case, we return the empty set, otherwise, we apply the computation of E^* to $G - X$ and $\mathcal{C}[G - X]$.

For both $\mathcal{C} = \mathcal{C}_{\text{all}}[l]$ and $\mathcal{C} = \overrightarrow{\mathcal{C}}_{\text{all}}[l]$ we can compute a minimum-length cycle containing e for any $e \in E(G)$ by Dijkstra's shortest path algorithm. This yields both l^* and E^* .

For $\mathcal{C} = \mathcal{C}_{\text{odd}}[l]$ we can compute a minimum-length odd cycle containing e for any $E(G)$, by a reduction to weighted matching as in Proposition 2.12 (cf. Section 29.11e of [80]).

Finally, consider the cases $\mathcal{C} = \mathcal{C}_D^{\geq 1}[l]$ and $\mathcal{C} = \mathcal{C}_D^{\leq 1}[l]$. Let $d = \{s, t\} \in D$. By the same argument as in the previous cases, it clearly suffices to compute the set E_d of all edges on a shortest s - t -path in $G - d$ or $G - D$, respectively. Thus, w.l.o.g. we only consider the case $\mathcal{C}_D^{\geq 1}$ and set $G' := G - d$. Let $G'_{>0}$ arise from G' by contracting all edges of length zero. By Dijkstra's algorithm we can compute the distances from s to all vertices in $G'_{>0}$ and thus also the set E'_d of all edges that are on a shortest s - t -walk in $G'_{>0}$. In $G'_{>0}$ we have that $l > 0$ and therefore E'_d coincides with the set of all edges that are on a shortest s - t -path. Clearly, $E_d \cap E(G'_{>0}) = E'_d$.

Now consider a connected component X of $\{e \in E(G') : l(e) = 0\}$. For any $e \in E(X)$ it holds that $e \in E_d$ if and only if there exist two edges $e_1 = \{s_1, t_1\}, e_2 = \{s_2, t_2\} \in \delta_{G'}(X)$ such that e_1 and e_2 lie on a common shortest s - t -path in $G'_{>0}$ and there exists a t_1 - s_2 -path in X containing e . For any possibilities of e_1 and e_2 , the former can be checked easily by Dijkstra's algorithm in $G'_{>0}$, while the latter can be checked on the block structure of X . \square

Note that in particular the cycle families $\mathcal{C}_{\text{all}}, \overrightarrow{\mathcal{C}}_{\text{all}}, \mathcal{C}_{\text{odd}}, \mathcal{C}_D^{\geq 1}$, and $\mathcal{C}_D^{\leq 1}$ allow for a polynomial-time support oracle because we can set l to zero. The existence of a black-box support oracle for $\mathcal{C}[l]$ that only needs a support (or weight) oracle for \mathcal{C} remains as an open question. Furthermore, we do not know whether the existence of a polynomial-time support oracle implies the existence of a polynomial-time weight oracle or vice versa.

Once we have a weight oracle or support oracle, we can do simple operations. For instance, to apply the uncrossing property algorithmically, we also need to uncross two cycles. To this end, we note:

Proposition 2.16. *Given a graph G and a weight oracle or support oracle for a family \mathcal{C} of cycles in G , we can decide by one oracle call whether a given edge set C contains the edge set of a cycle in \mathcal{C} .*

Proof. If we have a support oracle, we set $X := E(G) \setminus E(C)$. Calling the support oracle for X and checking whether the result is nonempty does the job.

If we have a weight oracle, set the weight of every edge outside C to 1 and the weight of every edge of C to 0. Calling the weight oracle for these weights will produce a cycle with total weight 0 if and only if C contains the edge set of a cycle in \mathcal{C} . \square

This also yields a *membership oracle* by applying Proposition 2.16 to the edge set of a cycle.

2.5 Embedded graphs

Throughout this thesis, G will usually be a graph that is embedded in an orientable surface Σ of constant genus g , the sphere if G is planar. Note that for fixed g there is a linear-time algorithm that embeds a given graph in an orientable surface of genus g or decides that there is no such embedding [63]. We will often identify edges, vertices, paths or cycles in G with their embedding in Σ . Sometimes it will be convenient to fix a point ∞ on Σ that is outside the embedding of all vertices and edges. The open connected components that result from removing all vertices and edges from Σ are called the *faces* of G . The face containing ∞ is called *infinite*; all other faces are *finite*.

There are two types of cycles in G : We call a simple and closed curve q on Σ *separating* if $\Sigma \setminus q$ is disconnected, otherwise q is *non-separating*. Since cycles in G define simple and closed curves on Σ this notion extends to cycles in G . Clearly, if $g = 0$, i.e. G is embedded in the sphere, then all cycles are separating. For a separating cycle C in G we call the connected components of $\Sigma \setminus C$ the *sides* of C . The side that contains ∞ is called the *exterior* of C , the other side is the *interior* of C . We denote the interior of C by $\text{int}(C)$.

We say that a separating cycle C_1 *contains* another separating cycle C_2 and write $C_2 \subseteq_{\infty} C_1$ if the interior of C_1 contains the interior of C_2 , i.e. $\text{int}(C_2) \subseteq \text{int}(C_1)$. The minimal cycles with respect to this partial order among a cycle family \mathcal{C} are called the *face-minimal* cycles of \mathcal{C} , which we usually denote by $\mathcal{C}_{\min} \subseteq \mathcal{C}$.

Similarly, we say that a side S of a cycle C_1 *contains* a cycle C_2 if S contains a side of C_2 .

2.5.1 Laminar cycle families

We say that two separating cycles C_1 and C_2 *cross* if their interiors are not disjoint and none of the two interiors contains the other. A multi-set of separating cycles is *laminar* if no pair of its cycles crosses. Equivalently, for any $C_1, C_2 \in \mathcal{L}$ one of the sides of C_1 contains C_2 .

Given a laminar family \mathcal{L} of (separating) cycles in a graph G that is embedded in an orientable surface of constant genus g we call a cycle $C \in \mathcal{L}$ *one-sided* if it has a \subseteq -minimal side among all cycles in \mathcal{L} , otherwise C is *two-sided*. Note that if $|\mathcal{L}| > 1$ then only one of the sides of C can be minimal, we call this the *one-sided side* of C (cf. Figure 2.2). Note that the one-sided side does not have to be finite; there can be at most one cycle in \mathcal{L} for which the infinite side is one-sided. This is the difference between the one-sided and the face-minimal cycles of \mathcal{L} . If $|\mathcal{L}| = 1$ then both sides of the cycle in \mathcal{L} are called one-sided.

A laminar family \mathcal{L} with only two one-sided sides is called a *chain*. Note that in this case we can choose a side S_C for each $C \in \mathcal{L}$ such that the family $\{S_C : C \in \mathcal{L}\}$ actually forms a chain w.r.t. \subseteq .

We say two cycles $C, C' \in \mathcal{L}$ are \mathcal{L} -*homotopic* if the set of face-minimal (or one-sided) cycles of \mathcal{L} that C and C' contain coincide. See Figure 2.2 for an example. This clearly defines an equivalence relation. We call the equivalence classes w.r.t. this relation the \mathcal{L} -*homotopy classes*. Clearly, each such class defines a chain, and the number of \mathcal{L} -homotopy classes is bounded by $2|\mathcal{L}_{\min}| - 1$, where \mathcal{L}_{\min} is the set of face-minimal cycles of \mathcal{L} .

Since we are interested in cycles that are vertex-disjoint, the notion of the *neighbourhood* of a cycle $C \in \mathcal{L}$ will be useful: This is the set $\mathcal{N}_{\mathcal{L}}(C) := \{C' \in \mathcal{L} : V(C) \cap V(C') \neq \emptyset\}$ of cycles that share a vertex with C .




Figure 2.2: A laminar family \mathcal{L} in the plane with seven cycles. The one-sided sides are filled; note that the one-sided cycles b, e, f and g are also face-minimal in \mathcal{L} , while a is the single one-sided cycle with one-sided exterior. The cycles c and d are \mathcal{L} -homotopic because both cycles contain exactly the one-sided sides of e, f and g in their interior.

2.5.2 Euler's formula

A main tool in almost every result of this thesis will be Euler's formula. It states:

Fact 2.17. Let $G = (V, E)$ be a planar graph embedded in the sphere. Let \mathcal{F} denote the set of faces of G . If G is connected, then

$$|V| + |\mathcal{F}| = |E| + 2.$$

There exists a more general version for graphs embedded in bounded-genus surfaces:

Fact 2.18. Let $G = (V, E)$ be a graph embedded in an orientable surface of genus g . Let \mathcal{F} denote the set of faces of G . If the embedding is *cellular*, i.e. the boundary of every face is connected then

$$|V| + |\mathcal{F}| + 2g = |E| + 2.$$

A very useful consequence of this is that planar (or bounded-genus) graphs are sparse if they have no parallel edges (that are embedded similarly).

Definition 2.19 (homotopic edges). Let G be a (directed or undirected) graph embedded in an orientable surface of genus g . Two parallel edges e and e' of G are called *homotopic* if they bound an area homeomorphic to the 2-dimensional disk that contains no vertices of G .

Definition 2.20 (degree of a face). Let G be a graph embedded in an orientable surface of genus g . Let F be a face of G and $E_F \subseteq E(G)$ the set of edges on the boundary of F . Let $E_{FF} \subseteq E_F$ be the set of edges that are adjacent to no other face than F , i.e. for which F is “on both sides” of the edge. We define the *degree* of F as $\deg(F) := |E_F| + |E_{FF}|$.

This means that the degree of F is the length of the closed walk (or the collection of closed walks if the boundary is disconnected) of going once around the boundary of F .

Lemma 2.21. *Let $G = (V, E)$ be a graph embedded in an orientable surface of genus g . If G contains no homotopic edges, then $|E| \leq 3|V| + 6g - 6$.*

Proof. We can assume w.l.o.g. that G is connected because adding edges to G only strengthens the lemma’s statement. By a similar argument we can assume w.l.o.g. that the embedding of G is cellular. Let \mathcal{F} denote the set of faces of G . Since G contains no homotopic edges, $\deg(F) \geq 3$ for all $F \in \mathcal{F}$. Also, we have $2|E| = \sum_{F \in \mathcal{F}} \deg(F) \geq 3|\mathcal{F}|$. Plugging this into Euler’s formula yields

$$|E| + 2 = |V| + |\mathcal{F}| + 2g \leq |V| + \frac{2}{3}|E| + 2g,$$

finishing our proof. □

As a direct consequence, there always exists a low-degree vertex:

Lemma 2.22. *Let $G = (V, E)$ be a graph embedded in an orientable surface of genus g . If G contains no homotopic edges then there exists a vertex $v \in V$ with $|\delta_G(v)| \leq 6g + 5$.*

Proof. Clearly, $\sum_{v \in V} |\delta_G(v)| = 2|E|$. Using this and the bound in Lemma 2.21 we can upper bound the minimum degree in G by

$$\min_{v \in V} |\delta_G(v)| \leq \left\lfloor \frac{2|E|}{|V|} \right\rfloor \leq \left\lfloor \frac{6|V| + 12g - 12}{|V|} \right\rfloor = 6 + \left\lfloor \frac{12}{|V|}(g - 1) \right\rfloor \leq 6g + 5 \quad \square$$

Lemma 2.23. *Let $G = (V, E)$ be a directed graph embedded in an orientable surface of genus g . If G contains no homotopic edges then there exists a vertex $v \in V$ with $|\delta_G^+(v)| \leq 6g + 5$.*

Proof. Ignoring the orientation of all edges can only make pairs of edges homotopic that were oriented in opposite directions; therefore merging such pairs and applying Lemma 2.22 proves the assertion. □

For planar graphs the bound from Lemma 2.22 can be sharpened by distinguishing between faces of degree > 3 . The following has been proven already by [24] and [60]:

Lemma 2.24. *Let $G = (V, E)$ be a connected graph, embedded in the sphere without homotopic edges. Then there exists a vertex $v \in V$ such that*

1. $|\delta_G(v)| \leq 3$ or
2. $|\delta_G(v)| = 4$ and v is incident to a face of degree 3 or
3. $|\delta_G(v)| = 5$ and v is incident to 4 faces of degree 3.

Proof. Let \mathcal{F} denote the set of faces of G . Again, $\deg(F) \geq 3$ for any $F \in \mathcal{F}$ because G has no homotopic edges. For any edge $e \in E$ let F_e^1 and F_e^2 denote the two faces of G adjacent to e (possibly $F_e^1 = F_e^2$ if e is a bridge). We define the weight of e as

$$w(e) := 1 + \frac{1}{4} |\{i \in \{1, 2\} : \deg(F_e^i) \geq 4\}|.$$

Note that

$$\sum_{e \in E} w(e) = |E| + \frac{1}{4} \sum_{F \in \mathcal{F} : \deg(F) \geq 4} \deg(F) \leq |E| + \sum_{F \in \mathcal{F}} \deg(F) - 3.$$

Since we can triangulate any face $F \in \mathcal{F}$ by adding $\deg(F) - 3$ edges inside F without creating homotopic edges we can use Lemma 2.21 to deduct $\sum_{e \in E} w(e) \leq 3|V| - 6$. Thus, we can choose $v \in V$ such that $\sum_{e \in \delta_G(v)} w(e) < 6$. An easy case distinction shows that v has the desired properties. \square

Remark 2.25. An easy consequence of this lemma is that the Erdős–Pósa ratio for the cycle family \mathcal{C}_{all} in a planar graph G is at most 3 (as observed by [24, 60]): We can directly construct a feasible cycle packing and transversal of suitable sizes. First, assume w.l.o.g. that each vertex of G has degree at least 3 by removing bridges and replacing paths whose internal vertices have degree 2 by single edges. Then, Lemma 2.24 can be applied to the planar dual of G . It immediately yields a face F of G such that all cycles that touch the boundary of F can be hit by a set T of at most 3 vertices of F . Adding F to our cycle packing solution, T to our transversal solution and recursing on $G - T$ yields the desired solutions.

2.6 Uncrossing

As the name suggests, an uncrossable family of cycles allows for uncrossing. We first review the simpler situation in planar graphs.

2.6.1 Uncrossing in planar graphs

A folklore result states that in planar graphs, every uncrossable cycle family has an optimum cycle packing that is laminar. While in the vertex-disjoint case actually any cycle packing must be laminar, this is a non-trivial statement for edge-disjoint cycle packing. In fact, an even stronger statement holds: For uncrossable cycle families in a planar graph there always exist optimum solutions to the cycle packing LPs (1.1) and (1.3) where the *support*, i.e., the set of all cycles with non-zero LP value, forms a laminar family of cycles. This is essentially shown in Lemma 4.2 of [41]; for completeness we briefly re-prove it here:

Proposition 2.26. *Let G be a planar graph and \mathcal{C} an uncrossable family of cycles in G given by a weight oracle. Given an explicit multi-subset $\mathcal{F} \subseteq \mathcal{C}$, one can compute in polynomial time a laminar multi-subset $\mathcal{L} \subseteq \mathcal{C}$ with $|\mathcal{L}| = |\mathcal{F}|$ and such that for every vertex v and every edge e , the number of cycles that contain v (or e) is no more in \mathcal{L} than in \mathcal{F} .*

Proof. We initialize $\mathcal{L} = \emptyset$ and iterate the following until $\mathcal{F} = \emptyset$:

Pick a cycle C_1 from \mathcal{F} . While there is a cycle $C_2 \in \mathcal{F}$ that crosses C_1 , we can find a path P_2 on C_2 inside the interior of C_1 that shares only its endpoints with C_1 . Since $C_1, C_2 \in \mathcal{F} \subseteq \mathcal{C}$

and \mathcal{C} is uncrossable, we can replace C_1 by a cycle C'_1 (consisting of P_2 and a part of C_1) with $\text{int}(C'_1) \subsetneq \text{int}(C_1)$ and C_2 by another cycle C'_2 such that no vertex or edge is contained more often in C'_1 and C'_2 than in C_1 and C_2 . By Proposition 2.16 we can find such cycles using our weight oracle.

Iterate with C'_1 in the role of C_1 , note that its interior contains fewer faces. Thus, after linearly many steps, the resulting cycle C_1 does not cross any cycle in \mathcal{F} . We then remove this cycle from \mathcal{F} and add it to \mathcal{L} . Throughout, we maintain the invariant that a cycle in \mathcal{L} does not cross any cycle in $\mathcal{L} \cup \mathcal{F}$. \square

2.6.2 Uncrossing in bounded-genus graphs

We also want to uncross cycles in a bounded-genus graph, which is much more difficult. In the special case of D -cycles, Huang et al. [48] showed how to uncross as much as possible. They implicitly use the fact that D -cycles are strongly uncrossable (cf. Definition 2.9). However, Theorem 2.11 states that any uncrossable cycle family is even strongly uncrossable. Therefore, we can generalize their proof to the case of uncrossable cycle families.

In order to do so, we first need a more general notion of crossing cycles that works also for non-separating cycles and quantifies the number of crossings of two cycles.

Definition 2.27 (crossings, uncrossed). Let C_1 and C_2 be two cycles in a graph G embedded in an orientable surface. A *set of crossings* of C_1 and C_2 is a set X of vertices of G such that for every $\varepsilon > 0$ there are simple closed curves \tilde{C}_1 and \tilde{C}_2 in an ε -environment of the embeddings of C_1 and C_2 , respectively, such that $\tilde{C}_1 \cap \tilde{C}_2 = X$ and $|X|$ is minimum. We say that C_1 and C_2 cross $|X|$ times.

We call a multi-set of cycles *uncrossed* if no pair of its cycles crosses more than once.

Note that the set X of crossings is not unique in general (if C_1 and C_2 share a path incident to a crossing, any vertex of that path could be chosen as crossing). We remark that the definition of [48] is slightly different but equivalent.

Definition 2.27 is consistent with the previous definition of crossing cycles: a pair of separating cycles crosses if and only if the number of their crossings is at least one (in fact at least two). We note:

Proposition 2.28. *If a set of separating cycles is uncrossed, then it is laminar.*

Proof. Every pair of separating cycles crosses an even number of times. Hence, if a set of separating cycles is uncrossed, no pair of its cycles cross. \square

However, for non-separating cycles there is no notion of laminarity. Nevertheless we can uncross cycles in an uncrossable family if a pair of cycles crosses at least twice. The following generalizes Proposition 2.26; its proof follows closely [48].

Lemma 2.29. *Let $G = (V, E)$ be a graph embedded in a fixed orientable surface, and let \mathcal{C} be an uncrossable family in G given by a weight oracle. Given an explicit multi-subset $\mathcal{F} \subseteq \mathcal{C}$, one can compute in polynomial time an uncrossed multi-subset $\mathcal{L} \subseteq \mathcal{C}$ with $|\mathcal{L}| = |\mathcal{F}|$ and such that for every vertex v and every edge e , the number of cycles that contain v (or e) is no more in \mathcal{L} than in \mathcal{F} .*

Proof. First, we can assume without loss of generality that all cycles in \mathcal{F} are pairwise edge-disjoint; this can be achieved by replacing each edge by sufficiently many parallel edges and can be done without increasing the number of crossings. This makes the set X of crossings in Definition 2.27 unique. Therefore, for a vertex $v \in V$ we can define $\text{cr}(v)$ to be the number of pairs of cycles in \mathcal{F} that cross in v .

While \mathcal{F} is not already uncrossed, take $C_1, C_2 \in \mathcal{F}$ and two vertices v, w of G such that C_1 and C_2 cross in v and in w . By Theorem 2.11 and Definition 2.9, there are v - w -paths P_1 in C_1 and P_2 in C_2 such that $P_1 + P_2$ contains a cycle $C'_1 \in \mathcal{C}$ and $(C_1 - P_1) + (C_2 - P_2)$ contains a cycle $C'_2 \in \mathcal{C}$. We can find such cycles using Proposition 2.16. We claim that replacing C_1 and C_2 by C'_1 and C'_2 decreases $(\sum_{C \in \mathcal{F}} |E(C)|, \sum_{x \in V} \text{cr}(x))$ lexicographically. Since $\sum_{C \in \mathcal{F}} |E(C)|$ can decrease at most $|\mathcal{F}| \cdot |E|$ times and $\sum_{x \in V} \text{cr}(x)$ can decrease at most $|\mathcal{F}|^2 |V|$ times while $\sum_{C \in \mathcal{F}} |E(C)|$ is constant, this will complete the proof.

If at least one edge of C_1 or C_2 is neither part of C'_1 nor of C'_2 , then replacing C_1 and C_2 by C'_1 and C'_2 decreases $\sum_{C \in \mathcal{F}} |E(C)|$.

Otherwise, $\sum_{C \in \mathcal{F}} |E(C)|$ remains constant, and replacing C_1 and C_2 by C'_1 and C'_2 does not increase $\text{cr}(x)$ for any $x \in V \setminus \{v, w\}$. Let now $x \in \{v, w\}$. Since C_1 and C_2 crossed in x , the new cycles C'_1 and C'_2 do not cross in x . Also, for every other cycle $C \in \mathcal{C} \setminus \{C_1, C_2\}$ one can verify by an easy case distinction (cf. Figure 2.3) that it does not cross more cycles in $\{C'_1, C'_2\}$ at x than in $\{C_1, C_2\}$. This means that replacing C_1 and C_2 by C'_1 and C'_2 decreases the total number of crossings by at least two.



Figure 2.3: On the left we see C_1 and C_2 crossing in x , together with three possible other cycles in \mathcal{F} that go through x and might cross C'_1 or C'_2 at x . However, on the right one can see that none of the three cycles crosses more cycles in x after replacing C_1 and C_2 by C'_1 and C'_2 .

Overall, iterating those replacements yields a multi-set \mathcal{L} as demanded after at most $O(|\mathcal{F}|^3 |V| |E|)$ iterations. \square

2.6.3 Uncrossing an LP solution

We first recall the well-known fact that both the cycle packing LPs (1.1) and (1.3) and the cycle transversal LPs (1.2) and (1.4) can be solved in polynomial time if the underlying cycle family \mathcal{C} allows for a weight oracle. For the cycle transversal LPs this is an immediate consequence of the Ellipsoid Method because the weight oracle solves the corresponding separation problems (for the vertex transversal LP see Proposition 2.14).

Despite the fact that the cycle packing LPs can have exponentially many variables we can apply standard methods to solve them in polynomial time. In particular, the support will be

polynomially bounded:

Proposition 2.30. *Given a graph G and a weight oracle for an uncrossable family \mathcal{C} of cycles in G , one can compute in polynomial time an optimum solution x to the linear program (1.1) or (1.3), given by an explicit list of all pairs (C, x_C) with $C \in \mathcal{C}$ and $x_C > 0$. \square*

Proof. We focus on (1.1); the proof for (1.3) is analogous. First solve the dual LP (1.2) by the Ellipsoid Method. To solve the primal LP, we can ignore all $C \in \mathcal{C}$ except those returned by the separation oracle while solving (1.2). Restricting the primal LP to these variables (whose number is bounded by a polynomial in $|V(G)|$), we can solve it in polynomial time. \square

Lemma 2.31. *Let $\varepsilon > 0$ be a fixed constant. Given a graph $G = (V, E)$ embedded in a fixed orientable surface and a weight oracle for an uncrossable family \mathcal{C} of cycles in G , one can compute in polynomial time a feasible solution x with uncrossed support to the linear program (1.1) or (1.3), and such that $\sum_{C \in \mathcal{C}} x_C$ is at least $(1 - \varepsilon)$ times the LP value.*

Proof. First compute an optimum solution x^* to the LP by applying Proposition 2.30. Let $\text{LP} := \sum_{C \in \mathcal{C}} x_C^*$ and $K := |\{C \in \mathcal{C} : x_C^* > 0\}|$; note that K is bounded by a polynomial in $|V|$. Now define $y_C := \lfloor \frac{K}{\varepsilon \cdot \text{LP}} x_C^* \rfloor$ and consider the multi-set \mathcal{F} that contains y_C copies of C for every $C \in \mathcal{C}$. Observe that \mathcal{F} contains at most $\frac{K}{\varepsilon}$ cycles (counting multiplicities). Apply Lemma 2.29 to \mathcal{F} and obtain an uncrossed multi-subset \mathcal{L} of \mathcal{C} . Finally, if \mathcal{L} contains z_C copies of C , then set $x_C := \frac{\varepsilon \cdot \text{LP}}{K} z_C$ for all $C \in \mathcal{C}$.

Obviously $\frac{\varepsilon \cdot \text{LP}}{K} y$ and hence $\frac{\varepsilon \cdot \text{LP}}{K} z = x$ is a feasible LP solution. Moreover,

$$\sum_{C \in \mathcal{C}} x_C = \frac{\varepsilon \cdot \text{LP}}{K} |\mathcal{L}| = \frac{\varepsilon \cdot \text{LP}}{K} |\mathcal{F}| = \frac{\varepsilon \cdot \text{LP}}{K} \sum_{C \in \mathcal{C}} \left\lfloor \frac{K}{\varepsilon \cdot \text{LP}} x_C^* \right\rfloor \geq \sum_{C \in \mathcal{C}} x_C^* - \varepsilon \cdot \text{LP} = (1 - \varepsilon) \text{LP}.$$

\square

It is an open question whether an optimum LP solution with uncrossed support can be computed in polynomial time. However, the $(1 - \varepsilon)$ factor will not be relevant for our main results. As far as existence is concerned, we immediately note:

Corollary 2.32. *Let G be a graph and \mathcal{C} an uncrossable family of cycles in G . Then the linear programs (1.1) and (1.3) have an optimum solution with uncrossed support. \square*

For planar graphs, we can remove the $(1 - \varepsilon)$ factor:

Theorem 2.33. *Given a planar graph $G = (V, E)$ embedded in the sphere and a weight oracle for an uncrossable family \mathcal{C} of cycles in G , one can compute in polynomial time an optimum solution x^* to the linear program (1.1) or (1.3) such that the support of x^* is laminar.*

Proof. Again we argue for (1.1); the proof for (1.3) is completely analogous. To allow for efficient uncrossing, we want to minimize $\sum_{C \in \mathcal{C}} x_C |C|$ among all optimum solutions to (1.1). This can be done by first computing the optimum value OPT of (1.1) and then solving the LP

$$\min \left\{ \sum_{C \in \mathcal{C}} x_C |C| : \sum_{C \in \mathcal{C}: v \in C} x_C \leq 1 \ (v \in V), \sum_{C \in \mathcal{C}} x_C = \text{OPT}, x_C \geq 0 \ (C \in \mathcal{C}) \right\}. \quad (2.1)$$

Similarly to the proof of Proposition 2.30, we do this by first solving the dual LP

$$\max \left\{ z \text{ OPT} - \sum_{v \in V} y_v : z \leq \sum_{v \in C} (y_v + 1) \ (C \in \mathcal{C}), \ y_v \geq 0 \ (v \in V) \right\}, \quad (2.2)$$

whose separation problem again reduces to calling the weight oracle, by the Ellipsoid Method, and keeping only primal variables corresponding to cycles returned by the separation oracle.

We now describe how we uncross the support of an optimum solution x to (2.1) (which is obviously also an optimum solution to (1.1)). First we describe an uncrossing operation of a pair of cycles, C_1 and C_2 , that cross. Take two vertices v, w of G such that C_1 and C_2 cross in v and in w . By Theorem 2.11 and Definition 2.9, there are v - w -paths P_1 in C_1 and P_2 in C_2 such that $P_1 + P_2$ contains a cycle $C'_1 \in \mathcal{C}$ and $(C_1 - P_1) + (C_2 - P_2)$ contains a cycle $C'_2 \in \mathcal{C}$. Again we can find such cycles using Proposition 2.16. If C'_1 and C'_2 still cross (we know they cross less often than C_1 and C_2), we apply the same step again to this pair of cycles. Let \bar{C}_1 and \bar{C}_2 denote the final outcome. These two cycles do not cross. Now we would set $\delta := \min\{x_{C_1}, x_{C_2}\}$ and change the LP solution by subtracting δ from x_{C_1} and x_{C_2} and adding δ to $x_{\bar{C}_1}$ and $x_{\bar{C}_2}$. Note that x remains an optimum LP solution. One of C_1 and C_2 vanishes from the support of x , but in general not both.

This operation cannot decrease $\sum_{C \in \mathcal{C}} x_C |C|$ because $\sum_{C \in \mathcal{C}} x_C$ remains constant and x was an optimum solution to (2.1). Hence every edge is contained the same number of times in \bar{C}_1 and \bar{C}_2 as in C_1 and C_2 . This is a key property that we will exploit now.

This single step makes some progress, but we need to be very careful to obtain a polynomial bound on the number of steps. Since \bar{C}_1 and \bar{C}_2 do not cross, $\text{int}(\bar{C}_1)$ and $\text{int}(\bar{C}_2)$ are either disjoint or one set is a subset of the other. In the first case, $\{\text{int}(\bar{C}_1), \text{int}(\bar{C}_2)\} = \{\text{int}(C_1) \setminus \text{int}(C_2), \text{int}(C_2) \setminus \text{int}(C_1)\}$. In the second case, $\{\text{int}(\bar{C}_1), \text{int}(\bar{C}_2)\} = \{\text{int}(C_1) \cap \text{int}(C_2), \text{int}(C_1) \cup \text{int}(C_2)\}$.

Therefore we can apply Karzanov's uncrossing algorithm [50] for uncrossing set systems and terminate with laminar support after polynomially many uncrossing steps. \square

2.7 Relations between the edge and vertex versions

Throughout this thesis we will usually work with the VERTEX-DISJOINT CYCLE PACKING PROBLEM and the VERTEX CYCLE TRANSVERSAL PROBLEM instead of the corresponding edge versions. All of our main results on the vertex variants of those problems will also allow for an edge version which can be proven in a similar way, or sometimes even easier. However, it will often also be possible to deduce the edge version directly from the vertex version via relatively straightforward reductions. We show the main reductions that we use in this section.

We start with the CYCLE TRANSVERSAL PROBLEM. Conveniently, our main result on this problem (Theorem 6.1) will consider the weighted version, for which we can give a simple reduction:

Proposition 2.34. *Let G be a graph with edge weights $w: E(G) \rightarrow \mathbb{R}_{\geq 0}$ and \mathcal{C} a family of cycles in G . Let G' arise from G by replacing each edge $e \in E(G)$ by a path P_e of length 2. Denote the degree-2-vertex of P_e by v_e . Let \mathcal{C}' be the family of cycles in G' with edge set $\bigcup_{e \in E(\mathcal{C})} E(P_e)$ for some $C \in \mathcal{C}$. Define vertex weights $w': V(G') \rightarrow \mathbb{R}_{\geq 0}$ by $w'(v_e) := w(e)$ for $e \in E(G)$ and $w'(v) := \sum_{e \in \delta_G(v)} w(e)$ for $v \in V(G)$. Then the following holds:*

1. For any feasible solution $y = (y_e)_{e \in E(G)}$ to the weighted edge cycle transversal LP (1.6) for \mathcal{C} there exists a feasible solution $y' = (y'_v)_{v \in V(G')}$ to the weighted vertex cycle transversal LP for \mathcal{C}' with the same LP value. If y is integral then y' can be chosen integral.
2. For any feasible solution $y' = (y'_v)_{v \in V(G')}$ to the weighted vertex cycle transversal LP (1.5) for \mathcal{C}' there exists a feasible solution $y = (y_e)_{e \in E(G)}$ to the weighted edge cycle transversal LP for \mathcal{C} with the same LP value. If y' is integral then y can be chosen integral.

Proof. Item 1 is simple: Given our solution y , we just set $y'_{v_e} := y_e$ for $e \in E(G)$ and $y'_v := 0$ for $v \in V(G)$. For item 2 we need a little bit more effort. Let y' be a feasible solution to the LP (1.5) for \mathcal{C}' . For an edge $e = \{v, w\} \in E(G)$ we set $y_e := y'_{v_e} + y'_v + y'_w$. This is a feasible solution to the LP (1.6) for \mathcal{C} : Let $C \in \mathcal{C}$ and $C' \in \mathcal{C}'$ be the corresponding cycle in G' . Then

$$\sum_{e \in E(C)} y_e = \sum_{e \in E(G)} y'_{v_e} + 2 \sum_{v \in V(C')} y'_v \geq \sum_{v \in V(G')} y'_v.$$

But the LP value of y is not larger than the LP value of y' because

$$\sum_{e \in E(G)} w(e)y_e = \sum_{e \in E(G)} w(e)y'_{v_e} + \sum_{v \in V(G')} \sum_{e \in \delta_G(v)} w(e)y'_v = \sum_{v \in V(G')} w'(v)y'_v \quad \square$$

In particular, this implies:

Corollary 2.35. *Let Σ be a surface and α_v and α_e denote the integrality gaps of the LPs (1.5) and (1.6) for uncrossable cycle families in graphs that can be embedded in Σ . Then $\alpha_e \leq \alpha_v$.*

Proof. Consider a graph G , embedded in Σ , with edge weights $w: E(G) \rightarrow \mathbb{R}_{\geq 0}$ and an uncrossable family \mathcal{C} of cycles in G . Clearly, the graph G' constructed in Proposition 2.34 can still be embedded in Σ and the family \mathcal{C}' is also uncrossable. By Proposition 2.34 the integrality gap of (1.6) for \mathcal{C} is the same as the integrality gap of (1.5) for \mathcal{C}' . This concludes the proof. \square

It is not known whether such a reduction also exists for the unweighted case. For the CYCLE PACKING PROBLEM it seems to be more difficult to give a direct reduction of the edge-disjoint to the vertex-disjoint version, preserving uncrossability of the cycle family and the oracles. Currently, such a reduction only exists for the D -CYCLE PACKING PROBLEM [62]. Instead, we give a reduction for laminar cycle families. This will be sufficient for all of our LP-based approaches because our results usually bound the laminar cycle packing integrality gap.

Proposition 2.36. *Let G be a graph, embedded in an orientable surface Σ of genus g . Let \mathcal{L} be a laminar family of separating cycles in G . Then we can compute a graph G' that can also be embedded in Σ , together with a laminar family \mathcal{L}' of separating cycles in G' and a bijection $f: \mathcal{L} \rightarrow \mathcal{L}'$, such that the following holds:*

1. For any edge $e \in E(G)$ where the cycles in $\mathcal{L}_e \subseteq \mathcal{L}$ meet (i.e., $e \in E(C)$ for any $C \in \mathcal{L}_e$) there exists a vertex $v \in V(G')$ such that $v \in V(C)$ for any $C \in f(\mathcal{L}_e)$.
2. For any vertex $v \in V(G')$ where the cycles in $\mathcal{L}'_v \subseteq \mathcal{L}'$ meet (i.e., $v \in V(C)$ for any $C \in \mathcal{L}'_v$) there exists an edge $e \in E(G)$ such that $e \in E(C)$ for any $C \in f^{-1}(\mathcal{L}'_v)$.

Proof. Define $V(G') := E(G)$. For any path $P = e_1e_2$ of length two in a cycle $C \in \mathcal{L}$ we add the edge $e_C^P := \{e_1, e_2\}$ to $E(G')$. For any $C \in \mathcal{L}$ let $f(C)$ be the cycle consisting of all edges $e_C^P \in E(G')$ for any path P of length two inside C .

Since \mathcal{L} is laminar, G' can be embedded in Σ such that $\mathcal{L}' := \{f(C) : C \in \mathcal{L}\}$ defines a laminar family of separating cycles, as shown in Figure 2.4. Properties 1 and 2 hold by construction. \square




Figure 2.4: Example for the construction of G' and \mathcal{L}' : The left picture shows four cycles in \mathcal{L} containing the vertex v . The six edges incident to v correspond to vertices of G' , as shown in the right picture. Since \mathcal{L} is laminar, the paths of length two in cycles of \mathcal{L} can be embedded in Σ as edges of G' . Cycles in \mathcal{L} share an edge if and only if the corresponding cycles in \mathcal{L}' share a vertex.

Due to uncrossing the reduction from Proposition 2.36 is all we need to relate the integrality gaps of both the vertex-disjoint and the edge-disjoint cycle packing LP to the laminar cycle packing integrality gap.

Proposition 2.37. *Let α^* be the laminar cycle packing integrality gap, i.e. the integrality gap of the LP (1.1) for laminar cycle families \mathcal{C} (in planar graphs). Let α_v and α_e be the integrality gaps of the vertex-disjoint and edge-disjoint cycle packing LPs (1.1) and (1.3), respectively, for uncrossable cycle families \mathcal{C} in planar graphs. Then $\alpha_v \leq \alpha^*$ and $\alpha_e \leq \alpha^*$.*

Proof. The first inequality $\alpha_v \leq \alpha^*$ follows directly from the fact that any uncrossable cycle family in a planar graph allows for an optimum LP solution with laminar support (Theorem 2.33).

Now consider the edge-disjoint cycle packing LP (1.3). Again, by Theorem 2.33 there exists an optimum LP solution x with laminar support \mathcal{L} . By Proposition 2.36 we can find another laminar family \mathcal{L}' of cycles in a planar graph with a bijection $f: \mathcal{L} \rightarrow \mathcal{L}'$ such that the edge packing constraints from the LP (1.3) on \mathcal{L} translate to vertex packing constraints on \mathcal{L}' . In particular, x induces a feasible solution to the vertex-disjoint cycle packing LP (1.1) on \mathcal{L}' , and any vertex-disjoint subset $\mathcal{S}' \subseteq \mathcal{L}'$ induces an edge-disjoint subset $\mathcal{S} \subseteq \mathcal{L}$. This proves $\alpha_e \leq \alpha^*$. \square

We will show later that α^* is bounded by a constant: Our best upper bound is $\frac{20+\sqrt{130}}{9} < 3.5$; the proof can be found in Section 5.

2.8 NP-hardness

In this section we give a brief NP-hardness proof for a large variety of (vertex-disjoint) cycle packing problems in planar graphs. For many cycle families such a result is already known, as described in Section 1.3. For the family $\overrightarrow{\mathcal{C}}_{\text{all}}$ however, this seems to be the first NP-hardness proof for planar graphs.

We will reduce the STABLE SET PROBLEM in certain planar graphs to the VERTEX-DISJOINT CYCLE PACKING PROBLEM. In the STABLE SET PROBLEM, the task is to compute a *stable set*, i.e. a set of pairwise non-adjacent vertices, of maximum cardinality in a given graph. Garey and Johnson [35] proved:

Proposition 2.38 (Garey, Johnson [35]). *The STABLE SET PROBLEM is NP-hard, even when restricted to planar graphs with maximum degree 3.*

By a simple observation, replacing an edge by a path of length 3 increases the maximum size of a stable set by exactly one. In particular, the problem still remains NP-hard if we add the restriction that all degree-3-vertices are “far apart”. The following extension of Proposition 2.38 follows from a result by Murphy [64]:

Proposition 2.39 (Murphy [64]). *Let $k \in \mathbb{Z}_{>0}$. The STABLE SET PROBLEM is NP-hard, even when restricted to planar graphs G with maximum degree 3 such that any edge in G is contained in some path of length k whose internal vertices have degree 2 in G .*

This enables us to prove the following hardness result:

Theorem 2.40. *Let \mathcal{I} be a family of instances of the VERTEX-DISJOINT CYCLE PACKING PROBLEM, i.e. \mathcal{I} consists of pairs (G, \mathcal{C}) of a graph G and a family \mathcal{C} of cycles in G . Assume \mathcal{I} has the following property: Given a planar graph G and a set \mathcal{C} of pairwise edge-disjoint cycles in G , then \mathcal{I} contains an instance (G', \mathcal{C}') such that G' arises from G by possibly subdividing some edges and all cycles in G' corresponding to cycles in \mathcal{C} are contained in \mathcal{C}' . Then the VERTEX-DISJOINT CYCLE PACKING PROBLEM restricted to instances in \mathcal{I} is NP-hard.*

Proof. Let G be a planar graph with maximum degree 3 such that any edge in G is contained in some path of length 4 whose internal vertices have degree 2 in G . We reduce the STABLE SET PROBLEM in G to the VERTEX-DISJOINT CYCLE PACKING PROBLEM in an instance of \mathcal{I} , which proves the theorem because of Proposition 2.39.

W.l.o.g. we assume that any vertex in G has degree at least 2. We construct a planar graph \hat{G} with a family $\hat{\mathcal{C}} = \{C_v : v \in V(G)\}$ of pairwise edge-disjoint cycles in \hat{G} such that for any $v, w \in V(G)$ we have that C_v and C_w intersect in exactly one vertex if $\{v, w\} \in E(G)$ and are vertex-disjoint otherwise. This can be achieved as follows: Set $V(\hat{G}) := E(G)$. For any $v \in V(G)$ let $E_v \subseteq V(\hat{G})$ denote the vertices corresponding to the edges in G that are adjacent to v . If $|E_v| = 2$ we add two parallel edges e_1, e_2 between those vertices to \hat{G} and define $E(C_v) := \{e_1, e_2\}$. Otherwise, if $|E_v| = 3$ we add the edges of a complete graph on E_v to \hat{G} and let C_v consist of these three edges.

We claim that any maximum set of pairwise vertex-disjoint cycles in \hat{G} is a subset of $\hat{\mathcal{C}}$. Assume this is false, and let \mathcal{S} denote a maximum-cardinality set of pairwise vertex-disjoint cycles in \hat{G} with a cycle $C \in \mathcal{S} \setminus \hat{\mathcal{C}}$. Since any edge in G is contained in a path of length 4 whose internal vertices have degree 2, no vertices of degree 3 in G are adjacent. Therefore, $C \notin \hat{\mathcal{C}}$

implies that C contains an edge of some cycle $C_v \in \hat{\mathcal{C}}$ of length 2. Let $v_1v_2v_3v_4v_5$ denote a path in G of length 4 that contains v such that v_2, v_3 and v_4 have degree 2 in G . Since $C \notin \hat{\mathcal{C}}$, we know that C contains exactly one edge from C_{v_2}, C_{v_3} and C_{v_4} . Therefore, replacing C by C_{v_2} and C_{v_4} yields a larger set of pairwise vertex-disjoint cycles in \hat{G} , a contradiction.

Let $(G', \mathcal{C}') \in \mathcal{I}$ such that G' arises from \hat{G} by possibly subdividing some edges and \mathcal{C}' includes all cycles in G' corresponding to cycles in $\hat{\mathcal{C}}$. If we find a maximum-cardinality subset $\mathcal{S}' \subseteq \mathcal{C}'$ of pairwise vertex-disjoint cycles then the above claim and the fact that G' arises from \hat{G} only by subdividing edges imply that \mathcal{S} induces a set $\hat{\mathcal{S}} \subseteq \hat{\mathcal{C}}$ of pairwise vertex-disjoint cycles. By the construction of \hat{G} now the set $\{v : C_v \in \hat{\mathcal{S}}\}$ is a stable set in G . On the other hand, any stable set $S \subseteq V(G)$ in G induces a cycle packing in \mathcal{C}' by adding the cycle in G' corresponding to the cycle C_v in \hat{G} for each $v \in S$. \square

The setting in Theorem 2.40 is quite general: It is easy to see that all cycle families introduced in Section 2.3.1 have the property needed to apply Theorem 2.40. Similarly, also the CYCLE PACKING PROBLEM for even-length cycles, cycles of some given length or combinations of the mentioned properties is NP-hard (in planar graphs). The most interesting application that has not been stated before seems to be the set $\overrightarrow{\mathcal{C}}_{\text{all}}$ of directed cycles in a digraph. We state:

Corollary 2.41. *The VERTEX-DISJOINT CYCLE PACKING PROBLEM for the family $\overrightarrow{\mathcal{C}}_{\text{all}}$ of all directed cycles in a planar digraph is NP-hard.*

Remark 2.42. Note that a similar construction as in the proof of Theorem 2.40 can also be used to derive NP-hardness results for the EDGE-DISJOINT CYCLE PACKING PROBLEM in various settings. In this case, we have to make sure that two cycles $C_v, C_w \in \hat{\mathcal{C}}$ intersect in exactly one *edge* if $\{v, w\} \in E(G)$ and are *edge*-disjoint otherwise. Also, for proving the claim we might now have to delete two cycles from \mathcal{S} before we can add the C_{v_i} for even i , but the claim still holds if we assume that any edge in G is part of a path of length 6 whose internal vertices have degree 2.

However, note that for this construction the property that is demanded for \mathcal{I} in Theorem 2.40 does not suffice any more because the cycles in $\hat{\mathcal{C}}$ are not edge-disjoint. Still, for many cycle families, NP-hardness of the corresponding EDGE-DISJOINT CYCLE PACKING PROBLEM can be deduced similarly as in Theorem 2.40. On the other hand, e.g. for the family $\overrightarrow{\mathcal{C}}_{\text{all}}$ the construction does not work (and indeed the corresponding CYCLE PACKING PROBLEM is solvable in polynomial time, cf. Table 1.1).

2.9 Lower bounds

In this section we give examples attaining (or approaching) the lower bounds presented in the Tables 1.1 and 1.2. All are quite simple, and some of them have been published before.

2.9.1 Edge-disjoint packing and edge transversal

We start with the edge versions of our problems. For the family \mathcal{C}_{all} the following example is due to Král'; it is published in [60]: Let $n \in \mathbb{N}$. The *n-wheel graph* is defined as the following planar graph on vertex set $\{1, \dots, n\} \times \{1, \dots, 8\}$. For each $1 \leq i \leq n$ the eight vertices of the form (i, j) form a cycle C_i . For any even $1 \leq i < n$, we join the cycles C_i and C_{i+1} by

the four edges $\{(i, j), (i + 1, j)\}$ with even j ; for any odd $1 \leq i < n$ we instead add the edges $\{(i, j), (i + 1, j)\}$ with odd j . The graph can be embedded planarly such that the interiors of the C_i form a chain, see Figure 2.5.




Figure 2.5: Example for the wheel graph for $n = 5$.

It is easy to see that the n -wheel graph has packing number n , while the edge transversal number is $4n - 3$. There also exists a fractional packing of value $2n - 1$ (pack each face with a factor of $\frac{1}{2}$) and a fractional edge transversal of value $2n - \frac{2}{3}$ (set $y_e := \frac{1}{6}$ for all edges).

Next, consider the family \mathcal{C}_{odd} . Here, the worst case example is given by a complete graph on 4 vertices, which we denote as K_4 . Any edge-disjoint (odd) cycle packing in a K_4 contains at most a single cycle, while we can pack 2 odd cycles fractionally (again, pack each face with a factor of $\frac{1}{2}$). The same example also yields the lower bound on the Erdős–Pósa ratio.

Finally, consider the family \mathcal{C}_D^{-1} of D -cycles. For this family, we first consider the instance where G is a K_4 and D forms a perfect matching on the four vertices of G . Similarly to the previous example this shows that both the Erdős–Pósa ratio and the integrality gap of (1.3) is at least 2. The lower bound for the integrality gap of (1.4) is more difficult. Huang et al. [47] pointed out that a lower bound by Cheriyan et al. [25] for the integrality gap of the tree augmentation LP also yields a lower bound of $\frac{3}{2}$ for planar edge D -cycle transversal.

2.9.2 Vertex-disjoint packing and vertex transversal

Now we consider the vertex versions of our problems. For the family \mathcal{C}_{all} we have two lower bound examples: First, the graph with 2 vertices and 3 parallel edges allows to pack only one cycle vertex-disjointly, but $\frac{3}{2}$ cycles fractionally (pack each cycle with a factor of $\frac{1}{2}$). The other two lower bounds are attained on a K_4 : An optimum integral transversal consists of 2 vertices, while setting $y_v := \frac{1}{3}$ for all vertices v yields a fractional transversal of size $\frac{4}{3}$. This shows a lower bound of $\frac{3}{2}$ for the cycle transversal LP (1.2), while the lower bound for the Erdős–Pósa ratio follows from the fact that the packing number of the K_4 is 1.

Next, we look at the family $\overrightarrow{\mathcal{C}}_{\text{all}}$ of directed cycles in a digraph. Consider the orientation of the octahedron shown in Figure 2.6. Any two directed cycles in this example share a vertex, so the packing number is 1, while the transversal number is 2. Packing each of the four directed cycles of length 3 with a value of $\frac{1}{2}$ also yields a fractional packing of value 2. For the corresponding CYCLE TRANSVERSAL PROBLEM we consider the complete directed graph on four vertices with a pair of parallel edges, oriented in opposite directions, between any two

vertices. Any feasible cycle transversal in this graph has ≥ 3 vertices, while using each vertex with value $\frac{1}{2}$ yields a fractional transversal of value 2.

For the family \mathcal{C}_{odd} we can subdivide some of the edges of an octahedron such that a similar argumentation as for $\overrightarrow{\mathcal{C}}_{\text{all}}$ yields lower bounds of 2 for both the Erdős–Pósa ratio and the integrality gap of (1.1), see Figure 2.6. For the integrality gap of the cycle transversal LP (1.2) we can again just consider the K_4 , as already for \mathcal{C}_{all} .




Figure 2.6: Lower bound examples for the Erdős–Pósa ratio and the integrality gap of (1.1) for directed cycles (left) and odd cycles (right). On the right the only edges of the octahedron that are not subdivided are marked in red, so a cycle is odd if and only if it contains an odd number of red edges. In both examples, any cycle of our cycle family contains at least three of the large colored vertices (blue and red), and only four of those cycles contain exactly 3 large colored vertices. Any two of them share a vertex, but packing all of them with a value of $\frac{1}{2}$ yields a feasible packing of value 2. The red vertices show optimum transversals.

Finally, for the family $\mathcal{C}_D^{\leq 1}$ of D -cycles, Middendorf and Pfeiffer’s reduction [62] shows that the integrality gap of (1.1) is not smaller than for the edge-disjoint version (1.3). In particular, this also bounds the Erdős–Pósa ratio from below by 2. For the D -cycle vertex transversal LP we again use the K_4 with two parallel edges, one D -edge and one non- D -edge, between any pair of vertices. The transversal number in this example is 3, while the fractional transversal number is 2.

Note that we do not know any example showing that the Erdős–Pósa ratio is larger than 2 for the examples considered so far. For the (uncrossable) family $\mathcal{C}_{D \geq 1}$ however this ratio is ≥ 4 :

Proposition 2.43. *The vertex-disjoint Erdős–Pósa ratio for the cycle family $\mathcal{C}_D^{\geq 1}$ in planar graphs is at least 4.*

Proof. Let G be the n -wheel graph from Section 2.9.1 for some $n \in \mathbb{N}$. Note that any vertex of G has degree ≤ 3 . Construct G' from G by replacing each degree-3-vertex $v \in V(G)$ by a triangle (a complete graph on three vertices) and connecting each of the three edges in $\delta_G(v)$ to a different vertex of this triangle. See Figure 2.7. Let $D \subseteq E(G')$ be the set of edges corresponding to the edges of G , i.e. all edges except the triangle edges for the vertices $v \in V(G)$. The cycle family $\mathcal{C}_D^{\geq 1}$ in G' behaves very similarly to the cycle family \mathcal{C}_{all} in G . In particular, it is easy to see that vertex-disjoint cycle packings in the former family correspond

to edge-disjoint packings in the latter, and both packing numbers are n (cf. Section 2.9.1). On the other hand, any vertex transversal $T' \subseteq V(G')$ for $\mathcal{C}_D^{\geq 1}$ induces an edge transversal $T \subseteq E(G)$ for \mathcal{C}_{all} by adding all D -edges incident to T . Thus, the vertex transversal number for $\mathcal{C}_D^{\geq 1}$ is at least $4n - 3$. This concludes the proof. \square




Figure 2.7: Example of the modification of the wheel graph in Proposition 2.43. A vertex of degree three in G (left) is replaced by a triangle in G' (right). The D -edges in G' are marked in red.

The same example also shows that the integrality gap for the LP (1.2) is at least 2 for $\mathcal{C}_D^{\geq 1}$, by a similar argument as for the edge transversal LP of \mathcal{C}_{all} . This lower bound of 2 has already been observed by [79] by considering a different example.

Chapter 3

Combinatorial approximation algorithms for planar cycle packing

In this chapter we devise several constant-factor approximation algorithms for the (uncrossable) CYCLE PACKING PROBLEM in planar graphs. All of these algorithms are combinatorial and do not compare to the (fractional) optimum solutions of the cycle packing LPs but only to optimum integral solutions. As shown in Tables 1.1 and 1.2 they yield the best-known constant-factor approximations for several of the uncrossable cycle families of interest. The results from Sections 3.1 and 3.2 are joint work with Hanjo Thiele and Jens Vygen [79].

A common technique that is used in all of our algorithms in this chapter is a polynomial-time approximation scheme (PTAS) for the CYCLE PACKING PROBLEM in planar graphs where all cycles in our family have disjoint interior. In particular this is the case for the cycles in \mathcal{C} with minimal interior, i.e. the face-minimal cycles of \mathcal{C} . The underlying principle of the PTAS is known as Baker's technique [11].

3.1 A PTAS for face-minimal cycles

In this section we consider the VERTEX-DISJOINT CYCLE PACKING PROBLEM in planar graphs on only the face-minimal cycles among an uncrossable family \mathcal{C} of cycles in a planar graph G . The following easy consequence of the uncrossing property has been noted by Goemans and Williamson [41]:

Proposition 3.1 ([41]). *Let G be a planar graph, embedded in the sphere, and \mathcal{C} an uncrossable family of cycles in G . If $G = (V(\mathcal{C}), E(\mathcal{C}))$ then any face-minimal cycle in \mathcal{C} is given by the boundary of a finite face of G .*

Proof. Let $C \in \mathcal{C}$ be face-minimal. Assume C is not given by the boundary of a finite face of G , i.e. let $e \in E(G)$ be embedded in the interior of C . Choose $C' \in \mathcal{C}$ such that $e \in E(C')$. Since C is face-minimal, C' contains also some edge outside the interior of C , so we can pick a subpath P' of C' that is embedded in the interior of C and shares exactly its endpoints with C . By the uncrossing property we can choose a subpath P of C such that $P + P' \in \mathcal{C}$, contradicting face-minimality of C . \square

Due to this proposition, for the CYCLE PACKING PROBLEM on the face-minimal cycles of \mathcal{C} all constraints can be formulated locally. In those cases, Baker [11] devised an approximation scheme that works for a large variety of problems. It uses the notion of k -outerplanar graphs:

Definition 3.2 (k -outerplanar). Let G be a planar graph, embedded in the sphere. We assign level 1 to all vertices that are on the boundary of the infinite face. For $i \geq 1$, we assign level $i + 1$ to all vertices that are on the boundary of the infinite face after removing all vertices of level at most i . G is called k -outerplanar if all vertices have level at most k .

Using the well-known facts that any k -outerplanar graph has treewidth at most $3k - 1$ [18] and tree decompositions with constant tree-width can be found efficiently [76, 10, 17], Baker's technique [11] can be described in a simpler form (the dependence of the runtime on k is worse, but we will apply it only for fixed $k = \lceil \frac{1}{\epsilon} \rceil$ anyway).

Lemma 3.3. *Let $k \in \mathbb{N}$ be a fixed constant. Then there is a linear-time algorithm that, given a k -outerplanar graph G and a family \mathcal{C}_{\min} of face-minimal cycles, computes a maximum vertex-disjoint subset of \mathcal{C}_{\min} .*

Proof. Since \mathcal{C}_{\min} is given explicitly, we can delete all vertices and edges that do not belong to any cycle in \mathcal{C}_{\min} ; then every cycle in \mathcal{C}_{\min} bounds a finite face by Proposition 3.1. We extend G to another planar graph G' by adding a new vertex inside each such face of G and connecting it to each vertex on the boundary of that face via a single edge. Let V denote the set of vertices of G and V' denote the set of new vertices (corresponding to the cycles in \mathcal{C}_{\min}). It is easy to see that G' is still $2k$ -outerplanar and therefore has treewidth at most $6k - 1$ [18]. Thus, using Bodlaender's algorithm [17], we can compute a rooted tree decomposition of width at most $6k - 1$ for G' in linear time such that each leaf bag contains exactly one vertex and each other bag B fulfills one of the following properties:

1. B has exactly two children in the tree decomposition, and they contain exactly the same set of vertices as B .
2. B has exactly one child in the tree decomposition, and the symmetric difference of B and its child contains exactly one element.

The algorithm to compute a vertex-disjoint subset \mathcal{S} of \mathcal{C}_{\min} is given by a dynamic program that specifies for each vertex in V' whether the cycle corresponding to that face is in \mathcal{S} or not. A candidate for a bag B now specifies for each such vertex that is contained in some bag in the subtree rooted at B whether the corresponding face is in \mathcal{S} or not (obeying the rule that $\mathcal{S} \subseteq \mathcal{C}_{\min}$ is a family of pairwise vertex-disjoint cycles). The label of that candidate indicates for each vertex in $B \cap V'$ whether the corresponding face has been chosen to be in \mathcal{S} or not and for each vertex in $B \cap V$ whether the vertex belongs to some cycle that has already been chosen to be in \mathcal{S} . Now for each of the at most 2^{6k} possible labels at some bag B we will compute a candidate maximizing the number of cycles chosen to be in \mathcal{S} . For each of the above types of bags, it is easy to see that this can be done in constant time (for fixed k) if the optimum candidates for the children of B are already computed. Therefore, going through the tree decomposition in reverse topological order and computing optimum candidates for all labels yields the result. \square

Continuing like Baker [11], we obtain:

Lemma 3.4. *For every fixed $\varepsilon > 0$ there is a linear-time algorithm that, given a planar graph G and a family \mathcal{C}_{\min} of face-minimal cycles in G , computes a vertex-disjoint subset of \mathcal{C}_{\min} with at least $(1 - \varepsilon)\nu_v(\mathcal{C}_{\min})$ elements.*

Proof. Let $k := \lceil \frac{1}{\varepsilon} \rceil$. We again remove all vertices and edges that do not belong to any cycle in \mathcal{C}_{\min} ; then every cycle in \mathcal{C}_{\min} bounds a finite face. For each integer $i > 0$ let V_i be the set of vertices of level i in G as in Definition 3.2, and let $V_i := \emptyset$ for $i \leq 0$. For each integer i define $G_i := G \left[\bigcup_{j=i}^{i+k-1} V_j \right]$. Since each of these subgraphs G_i is k -outerplanar by definition, we can find an optimum solution \mathcal{S}_i in G_i (i.e., a maximum vertex-disjoint subset of $\mathcal{C}_{\min}[G_i]$) in linear time by Lemma 3.3. This defines feasible solutions $\mathcal{S}_i^+ := \bigcup_{j \in \mathbb{Z}} \mathcal{S}_{i+jk}$ for $i = 1, \dots, k$.

Because the vertices on the boundary of a face of G differ by at most one in their level, each finite face of G is contained in at least $k - 1$ of the subgraphs G_i . Therefore, the best solution among the \mathcal{S}_i^+ ($i = 1, \dots, k$) has at least $\frac{k-1}{k}$ times the size of an optimum solution. This proves the lemma. \square

3.2 A $(3 + \varepsilon)$ -approximation

In this section we give a combinatorial $(3 + \varepsilon)$ -approximation algorithm for the uncrossable CYCLE PACKING PROBLEM:

Theorem 3.5. *For any fixed $\varepsilon > 0$, there is a polynomial-time $(3 + \varepsilon)$ -approximation algorithm for each of the following problems. Given a planar graph G and a support oracle for an uncrossable family \mathcal{C} of cycles in G ,*

- (a) *find a maximum-cardinality vertex-disjoint subset of \mathcal{C} ;*
- (b) *find a maximum-cardinality edge-disjoint subset of \mathcal{C} .*

We start with explaining the algorithm for the vertex-disjoint case (a) in Section 3.2.1. The modifications needed to the algorithm to work for the edge-disjoint case are described in Section 3.2.2.

3.2.1 The algorithm

Our main algorithm works as follows.

Algorithm 1: $(3 + \varepsilon)$ -approximation for uncrossable cycle packing

Input: A planar graph G , a support oracle for an uncrossable family \mathcal{C} of cycles in G

Output: A set $\mathcal{S} \subseteq \mathcal{C}$ of pairwise vertex-disjoint cycles

- 1 Compute an embedding of G in the sphere
 - 2 Compute the set \mathcal{C}_{\min} of face-minimal cycles of \mathcal{C}
 - 3 Find a set $\mathcal{S}_1 \subseteq \mathcal{C}_{\min}$ of pairwise vertex-disjoint cycles using Lemma 3.4
 - 4 Remove $V(\mathcal{S}_1)$ from G and recurse on the remaining graph to obtain a vertex-disjoint subset $\mathcal{S}_2 \subseteq \mathcal{C}[G - V(\mathcal{S}_1)]$
 - 5 Output $\mathcal{S}_1 \cup \mathcal{S}_2$
-

Step 1 can be done in linear time [26]. For step 1 we only need to call the support oracle once:

Proposition 3.6. *Given a planar graph G and a support oracle for an uncrossable family \mathcal{C} of cycles in G , we can compute the set \mathcal{C}_{\min} of face-minimal cycles in \mathcal{C} in linear time.*

Proof. Call the support oracle with $X = \emptyset$ and remove all edges that the oracle does not return. Now every edge belongs to a cycle in \mathcal{C} . Then, by Proposition 3.1 every face-minimal cycle bounds a finite face of the induced embedding. So calling the membership oracle (Proposition 2.16) for each of the cycles bounding a finite face yields \mathcal{C}_{\min} . \square

For step 2 we use the linear-time $(1 + \varepsilon)$ -approximation algorithm from Section 3.1 for a fixed $0 < \varepsilon < \frac{1}{3}$. Since the support oracle also allows to remove a set X of edges and the family $\mathcal{C}[G - V(\mathcal{S}_1)]$ is still uncrossable the recursion in step 3 works as well. Thus, the overall algorithm takes only quadratic time. Finally, we prove the approximation guarantee.

Lemma 3.7. *Algorithm 1 computes a vertex-disjoint subset of \mathcal{C} with at least $(\frac{1}{3} - \varepsilon)\nu_v(\mathcal{C})$ elements.*

Proof. Let $G = (V, E)$ denote the given planar graph. We use induction on $|V|$. Let \mathcal{S}^* be a maximum-cardinality vertex-disjoint subset of \mathcal{C} . Clearly \mathcal{S}^* is laminar. We may assume that every cycle with minimal interior in \mathcal{S}^* is also face-minimal in \mathcal{C} , otherwise we could replace it by a smaller cycle in its interior.

By combining Proposition 3.6 and Lemma 3.4 we can find a vertex-disjoint subset $\mathcal{S}_1 \subseteq \mathcal{C}_{\min}$ with $|\mathcal{S}_1| \geq (1 - \varepsilon)|\mathcal{S}_{\min}^*|$ in polynomial time. Moreover, $\mathcal{S}^* \cup \mathcal{S}_1$ is laminar by Proposition 3.1. Consider a tree representation T of $\mathcal{S}^* \cup \mathcal{S}_1$: an arborescence whose arcs represent the cycles in $\mathcal{S}^* \cup \mathcal{S}_1$ so that the arc representing C_1 is reachable from the arc representing C_2 if and only if $C_1 \subseteq_{\infty} C_2$. For $C \in \mathcal{S}_1$ let $\text{parent}(C)$ be the tail of the arc representing C . For a vertex $v \in T$ let $\delta_{\mathcal{S}^*}^+(v)$ denote the set of arcs that represent a cycle in \mathcal{S}^* and leave v in T , and $\delta_{\mathcal{S}^*}^-(v)$ denotes the (empty or one-element) set of arcs that represent a cycle in \mathcal{S}^* and enter v in T . See Figure 3.1 for an example.

Let $\mathcal{B} \subseteq \mathcal{S}^*$ be the set of cycles in \mathcal{S}^* that have a vertex in common with at least one cycle in \mathcal{S}_1 (these are the cycles from \mathcal{S}^* that do not exist anymore when we recurse on $G - V(\mathcal{S}_1)$). Note that the arc representing some $B \in \mathcal{B}$ must be incident to $\text{parent}(C)$ for some $C \in \mathcal{S}_1$. Hence

$$\begin{aligned} |\mathcal{B}| &\leq \left| \bigcup_{C \in \mathcal{S}_1} \delta_{\mathcal{S}^*}^-(\text{parent}(C)) \cup \delta_{\mathcal{S}^*}^+(\text{parent}(C)) \right| \\ &\leq 2|\mathcal{S}_1| + \sum_{v \in T} \max \{0, |\delta_{\mathcal{S}^*}^+(v)| - 1\} \\ &= 2|\mathcal{S}_1| + |\mathcal{S}_{\min}^*| - 1 \\ &\leq 2|\mathcal{S}_1| + \frac{1}{1 - \varepsilon} |\mathcal{S}_1| - 1 \\ &\leq \frac{1}{\frac{1}{3} - \varepsilon} |\mathcal{S}_1|, \end{aligned}$$

where we used in the equation that the \subseteq_{∞} -minimal cycles in \mathcal{S}^* are face-minimal in \mathcal{C} .

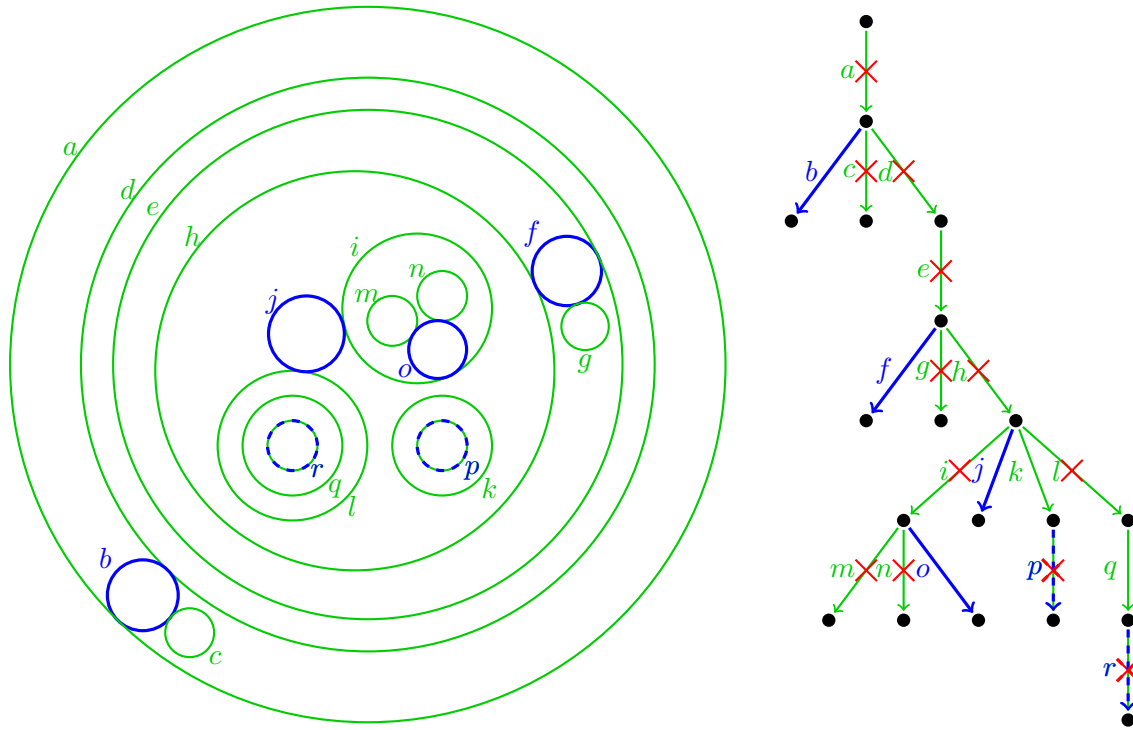


Figure 3.1: The left-hand side shows an example of cycles in \mathcal{S}^* (green) and \mathcal{S}_1 (blue and bold). The cycles p and r are contained both in \mathcal{S}_1 and \mathcal{S}^* . The right-hand side shows the tree representation of $\mathcal{S}^* \cup \mathcal{S}_1$. The cycles in \mathcal{B} are crossed out in red; these cycles are “eliminated” by \mathcal{S}_1 . For the cycles corresponding to the first four levels of the tree, the bound for $|\mathcal{B}|$ in the proof is tight.

By the induction hypothesis, recursing on $G - V(\mathcal{S}_1)$ yields a set \mathcal{S}_2 of vertex-disjoint cycles with $|\mathcal{S}_2| \geq (\frac{1}{3} - \varepsilon)|\mathcal{S}^* \setminus \mathcal{B}|$. Hence

$$\begin{aligned}
 |\mathcal{S}_1 \cup \mathcal{S}_2| &= |\mathcal{S}_1| + |\mathcal{S}_2| \\
 &\geq \left(\frac{1}{3} - \varepsilon\right) |\mathcal{B}| + \left(\frac{1}{3} - \varepsilon\right) |\mathcal{S}^* \setminus \mathcal{B}| \\
 &= \left(\frac{1}{3} - \varepsilon\right) |\mathcal{S}^*|. \quad \square
 \end{aligned}$$

This completes the proof of Theorem 3.5 (a).

3.2.2 The edge-disjoint case

We now prove part (b) of Theorem 3.5. Since both the algorithm and the analysis are very similar to the vertex-disjoint case, we only discuss the differences. Algorithm 1 still works as in the vertex-disjoint case, with two modifications:

First, in Step 2 we want to find an edge-disjoint subset \mathcal{S}_1 of \mathcal{C}_{\min} . This can be done similarly to Lemma 3.3 and Lemma 3.4; however, they can be simplified since, after deleting

redundant edges, edge-disjoint packings of face-minimal cycles in G correspond to stable sets in (a subgraph of) the planar dual of G . Therefore we can directly use Baker's approximation scheme for the stable set problem in planar graphs [11].

Second, in Step 3 we only delete the edges of \mathcal{S}_1 instead of the vertices and recurse on $\mathcal{C}[G - E(\mathcal{S}_1)]$.

In order to adapt the analysis of the algorithm for the vertex-disjoint case to the edge-disjoint case, we need to ensure that $\mathcal{S}^* \cup \mathcal{S}_1$ is a laminar family of cycles. This is not trivial anymore since edge-disjoint cycles can cross. However, \mathcal{S}^* can be chosen to be laminar by Proposition 2.26; then by Proposition 3.1 also $\mathcal{S}^* \cup \mathcal{S}_1$ is laminar.

If we then choose \mathcal{B} to be the set of cycles in \mathcal{S}^* that share an edge with at least one cycle in \mathcal{S}_1 , the analysis from Section 3.2.1 proves Theorem 3.5 (b).

3.3 A $(2 + \varepsilon)$ -approximation for \mathcal{C}_{all}

In this section we devise a simple linear-time $(2 + \varepsilon)$ -approximation algorithm for any $\varepsilon > 0$ for the planar VERTEX-DISJOINT CYCLE PACKING PROBLEM for the family \mathcal{C}_{all} of all cycles in the underlying graph G . Actually, our algorithm will consist of only a single iteration of Algorithm 1 from Section 3.2, i.e. we only need to apply the approximation scheme for face-minimal cycles from Section 3.1 once. For this to work, however, we first need to make our graph 2-vertex-connected:

Lemma 3.8. *Let G be a planar graph, embedded in the sphere. In linear time, we can compute a graph G' such that all connected components of G' are 2-vertex-connected and the VERTEX-DISJOINT CYCLE PACKING PROBLEM for the set of all cycles in G' is equivalent to the VERTEX-DISJOINT CYCLE PACKING PROBLEM for the set of all cycles in G .*

Proof. We can clearly remove bridges of G without changing \mathcal{C}_{all} . Since also the connected components of G are independent we may assume that G is 2-edge-connected. Let G_1 and G_2 be two blocks (maximal 2-vertex-connected components) of G with a common vertex $v \in V(G_1) \cap V(G_2)$. Add a new vertex v' with an edge $\{v, v'\}$ and for each $i \in \{1, 2\}$ pick an edge $\{v, w\} \in E(G_i)$ to replace by $\{v', w\}$ without destroying planarity. Call the result G_v . Since we added the edge $\{v, v'\}$ any vertex-disjoint cycle packing in G induces a vertex-disjoint cycle packing in G_v of the same size. However, any cycle in G_v that contains the vertex v' also contains v because $|\delta_{G_v}(v') \setminus \{\{v, v'\}\}| = 2$ and v is the only common vertex in the blocks G_1 and G_2 . Therefore, vertex-disjoint cycle packings in G_v also induce vertex-disjoint cycle packings in G of the same size. Doing the above modification until the graph is 2-vertex-connected yields the desired graph G' . \square

The following structural result is very similar to an observation by Caprara, Panconesi and Rizzi [20] which they used to derive a $(2 + \varepsilon)$ -approximation for the edge-disjoint case.

Lemma 3.9. *Let G be a 2-vertex-connected planar graph, embedded in the sphere. Let \mathcal{S} be a set of pairwise vertex-disjoint cycles in G . Then there is a set \mathcal{S}' of pairwise vertex-disjoint face-minimal cycles in G with $|\mathcal{S}| \leq 2|\mathcal{S}'|$.*

Proof. Since \mathcal{S} is pairwise vertex-disjoint, \mathcal{S} is laminar. As in the proof of Lemma 3.7 we consider a tree representation T of \mathcal{S} , i.e. an arborescence whose arcs correspond to cycles

in \mathcal{S} where the arc corresponding to C_1 is reachable from the arc corresponding to C_2 if and only if $C_1 \subseteq_\infty C_2$. Let r denote the root of T . For any $e \in E(T)$ let $C_e \in \mathcal{S}$ denote the cycle corresponding to e . For any $v \in V(T) \setminus \{r\}$ let C_v be the boundary of a face of G which is in $\text{int}(C_e)$ for the unique edge $e \in \delta_T^-(v)$ but outside $\text{int}(C_e)$ for each $e \in \delta_T^+(v)$. Since G is 2-vertex-connected, C_v is a (face-minimal) cycle.

Let T' denote the undirected version of T , i.e. T' arises from T by replacing each directed edge (v, w) by the undirected edge $\{v, w\}$. Since T' is bipartite we can find a stable set $S \subseteq V(T') \setminus \{r\}$ with $|S| \geq \frac{|V(T')|-1}{2} = \frac{|S|}{2}$. Then $\mathcal{S}' := \{C_v : v \in S\}$ defines a cycle packing as wished: For $v, w \in S$ let P_{vw} denote the unique v - w -path in T' and let $\mathcal{C}_P := \{C_e : e \in E(P)\} \cup \{C_v, C_w\}$. It is easy to see that \mathcal{C}_P is a chain and C_v and C_w are the one-sided cycles in \mathcal{C}_P . Thus, vertex-disjointness of the cycles in \mathcal{S} and $|\mathcal{C}_P \cap \mathcal{S}| \geq 2$ shows vertex-disjointness of C_v and C_w . \square

As a direct consequence, we get:

Theorem 3.10. *For any fixed $\varepsilon > 0$, there is a linear-time $(2 + \varepsilon)$ -approximation algorithm for the problem of finding a maximum-cardinality set of pairwise vertex-disjoint cycles in a planar graph G .*

Proof. We can embed G in the sphere in linear time [26]. By Lemma 3.8 we can assume w.l.o.g. that G is 2-vertex-connected. This means that the boundary of any face of G is a cycle, in particular the set \mathcal{C}_{\min} of face-minimal cycles of \mathcal{C}_{all} is given exactly by all finite faces of G . Using Lemma 3.4 we can find a $(1 + \frac{\varepsilon}{2})$ -approximation for a maximum-cardinality vertex-disjoint subset of \mathcal{C}_{\min} in linear time. By Lemma 3.9 this defines a $(2 + \varepsilon)$ -approximation for a maximum-cardinality vertex-disjoint subset of \mathcal{C}_{all} . \square

3.4 A PTAS for laminar families

In this section we extend the PTAS for the CYCLE PACKING PROBLEM on face-minimal cycles from Lemma 3.4 to the case of arbitrary laminar families. Note that this is particularly interesting as the natural cycle packing LPs always admit an optimum solution with laminar support. Running our PTAS on this support will yield an $(\alpha^* + \varepsilon)$ -approximation for the CYCLE PACKING PROBLEM for uncrossable cycle families in planar graphs for any $\varepsilon > 0$, where α^* is the laminar cycle packing integrality gap (cf. Corollary 3.14).

We start with the case where all cycles that are not face-minimal are relatively large. We can deal with the smaller cycles differently. As in Section 3.2 we denote the set of face-minimal cycles of a cycle family \mathcal{L} by \mathcal{L}_{\min} .

Lemma 3.11. *Let $k \in \mathbb{N}$ with $k > 4$. There exists a polynomial-time $\left(\frac{k}{k-4}\right)^3$ -approximation algorithm for the VERTEX-DISJOINT CYCLE PACKING PROBLEM in laminar families \mathcal{L} of cycles in planar graphs G where for all cycles C in $\mathcal{L}' := \mathcal{L} \setminus \mathcal{L}_{\min}$ there exist k^2 pairwise vertex-disjoint cycles in \mathcal{L}_{\min} that are in the interior of C and contain no vertex of C .*

Proof. Let \mathcal{B}_1 be the set of maximal cycles of each \mathcal{L}' -homotopy class, i.e. for any $C \in \mathcal{L}'$ we add the cycle in \mathcal{L}' that is \mathcal{L}' -homotopic to C with the largest interior to \mathcal{B}_1 . For any $B \in \mathcal{B}_1$ let $\mathcal{B}'_1(B) \subseteq \mathcal{L}'$ be the set of all $C \in \mathcal{L}'$ that are \mathcal{L}' -homotopic to B with $V(C) \cap V(B) \neq \emptyset$ and define $\mathcal{B}'_1 := \bigcup_{B \in \mathcal{B}_1} \mathcal{B}'_1(B)$. See Figure 3.2 for an example.

We will restrict ourselves to finding a cycle packing among $\mathcal{L} \setminus \mathcal{B}'_1$. By removing \mathcal{B}'_1 we get the nice property that any cycles in $\mathcal{L}^- := \mathcal{L}' \setminus \mathcal{B}'_1$ that share a vertex are \mathcal{L}^- -homotopic: Assume $C_1, C_2 \in \mathcal{L}^-$ are not \mathcal{L}^- -homotopic. W.l.o.g. $C_2 \not\subseteq_\infty C_1$. Choose $B \in \mathcal{B}_1$ that is \mathcal{L}' -homotopic to C_1 . Then C_1 and C_2 are contained in different sides of B , but B and C_1 are vertex-disjoint. Thus, C_1 and C_2 are vertex-disjoint.

Let $\mathcal{F} \subseteq \mathcal{L}^-$ be any maximal subset of pairwise vertex-disjoint cycles. This implies that any $C \in \mathcal{L}^-$ contains a vertex of some cycle in \mathcal{F} . Let $\mathcal{B}_2 \subseteq \mathcal{F}$ contain the $k-1$ cycles with largest interior of every \mathcal{F} -homotopy class (and all cycles of each \mathcal{F} -homotopy class of size $< k$). Define $\mathcal{B}'_2 := \{C \in \mathcal{L}^- : V(C) \cap V(\mathcal{B}_2) \neq \emptyset\}$. See Figure 3.2.

We will only consider cycle packings among $\mathcal{L} \setminus (\mathcal{B}'_1 \cup \mathcal{B}'_2)$. To see that this is no problem, consider an optimum solution $\mathcal{S}^* \subseteq \mathcal{L}$, i.e. a maximum-cardinality subset of pairwise vertex-disjoint cycles. Let $\mathcal{S}^*_- := \mathcal{S}^* \setminus (\mathcal{B}'_1 \cup \mathcal{B}'_2)$. For each $B \in \mathcal{B}_1$ the cycles in $\mathcal{B}'_1(B)$ form a chain (in which B is one-sided) and intersect $V(B)$, so $|\mathcal{S}^* \cap \mathcal{B}'_1(B)| \leq 1$. Also, for each $B \in \mathcal{B}_2$ the cycles in \mathcal{L}^- that intersect $V(B)$ form a chain, so $|\mathcal{S}^* \cap \mathcal{B}'_2| \leq 2|\mathcal{B}_2| \leq 2(k-1)|\mathcal{B}_1|$. Thus, $|\mathcal{S}^* \setminus \mathcal{S}^*_-| \leq (2k-1)|\mathcal{B}_1|$. Furthermore, no cycles in \mathcal{B}_1 are \mathcal{B}_1 -homotopic, so only the root and the leaves in a tree representation for \mathcal{B}_1 can have out-degree at most one. This implies that $|\mathcal{B}_1| \leq 2|\mathcal{B}_{1\min}|$ and we conclude $|\mathcal{S}^* \setminus \mathcal{S}^*_-| \leq 4k|\mathcal{B}_{1\min}|$. By the assumption in our Lemma, each cycle $B \in \mathcal{B}_{1\min}$ contains at least k^2 pairwise vertex-disjoint cycles that are also vertex-disjoint to B . Thus, $|\mathcal{S}^*_-| \geq k^2|\mathcal{B}_{1\min}|$. Together this yields $|\mathcal{S}^*| \leq \frac{k}{k-4}|\mathcal{S}^*_-|$.

As in the proof of Lemma 3.7 let T be a tree representation for the laminar family \mathcal{F} . For any $e \in E(T)$ let V_e be the set of vertices that are embedded on the cycle corresponding to e or in its interior, but outside the interior of all cycles corresponding to other edges reachable from e in T . Let $r \in V(T)$ be the root of T . To simplify our arguments in the following, we add a new root r' and a path $r'e_1v_1e_2 \dots v_{k-1}e_kr$ to T . Furthermore, we define $V_{e_i} := \emptyset$ for $i = 1, \dots, k-1$ and let V_{e_k} be the set of all vertices outside the interior of all cycles in \mathcal{F} .

For any $e \in E(T)$ we let $G_e := G \left[\bigcup_{f \in E(T) : \text{dist}_T(e,f) < k} V_f \right]$, where $\text{dist}_T(e, f)$ denotes the length of the unique directed path in T starting with e and ending with f (or ∞ if no such path exists). Also, define $\mathcal{L}_e := \mathcal{L}[G_e] \setminus (\mathcal{B}'_1 \cup \mathcal{B}'_2)$. This makes sure that $\mathcal{L}_e \cap \mathcal{F}$ is a chain: If not, let $C, C' \in \mathcal{L}_e \cap \mathcal{F}$ have disjoint interior. Then G_e contains a cycle $F \in \mathcal{F}$ that contains both C and C' . By the definition of \mathcal{B}_2 we have cycles $B_1, \dots, B_{k-1} \in \mathcal{B}_2$ that are \mathcal{F} -homotopic to C with $\text{int}(C) \subsetneq \text{int}(B_1) \subsetneq \dots \subsetneq \text{int}(B_{k-1}) \subsetneq \text{int}(F)$. In particular, this defines a chain of $k+1$ cycles in $\mathcal{F} \cap \mathcal{L}_e$, contradicting the definition of G_e . So indeed $\mathcal{L}_e \cap \mathcal{F}$ is a chain (cf. Figure 3.2).

Therefore, $|\mathcal{L}_e \cap \mathcal{F}| \leq k$. Let $\mathcal{S}_e^* \subseteq \mathcal{L}_e$ be a maximum-cardinality set of pairwise vertex-disjoint cycles. By the construction of \mathcal{F} and \mathcal{B}'_2 , any cycle $C \in \mathcal{S}_e^* \cap \mathcal{L}'$ shares a vertex with some cycle $C' \in \mathcal{F} \cap \mathcal{L}_e$. Furthermore, in this case C and C' are \mathcal{L}^- -homotopic due to the construction of \mathcal{B}'_1 . In particular, for each $C' \in \mathcal{F} \cap \mathcal{L}_e$ there can be at most two such cycles $C \in \mathcal{S}_e^* \cap \mathcal{L}'$ and therefore $|\mathcal{S}_e^* \cap \mathcal{L}'| \leq 2k$.

Thus, we can compute a set $\mathcal{S}_e \subseteq \mathcal{L}_e$ of pairwise vertex-disjoint cycles with $|\mathcal{S}_e^*| \leq \frac{k}{k-1}|\mathcal{S}_e|$ by guessing all cycles in $\mathcal{S}_e^* \cap \mathcal{L}'$ and using Lemma 3.4 on the remaining face-minimal cycles after deleting the vertices of $\mathcal{S}_e^* \cap \mathcal{L}'$. As in Lemma 3.4 we can define feasible solutions

$$\mathcal{S}_i^+ := \bigcup_{e \in E(T) : \text{dist}_T(e_0, e) \equiv i \pmod k} \mathcal{S}_e$$

for $i = 1, \dots, k$. Since any cycle in $\mathcal{L} \setminus (\mathcal{B}'_1 \cup \mathcal{B}'_2)$ appears in at least $k-1$ of the \mathcal{L}_e , there

exists an $i \in \{1, \dots, k\}$ such that

$$|\mathcal{S}^*| \leq \frac{k}{k-4} |\mathcal{S}^*| \leq \left(\frac{k}{k-4}\right)^2 \sum_{e \in E(T) : \text{dist}_T(e_0, e) \equiv i \pmod k} |\mathcal{S}_e^*| \leq \left(\frac{k}{k-4}\right)^3 |\mathcal{S}_i^+|. \quad \square$$

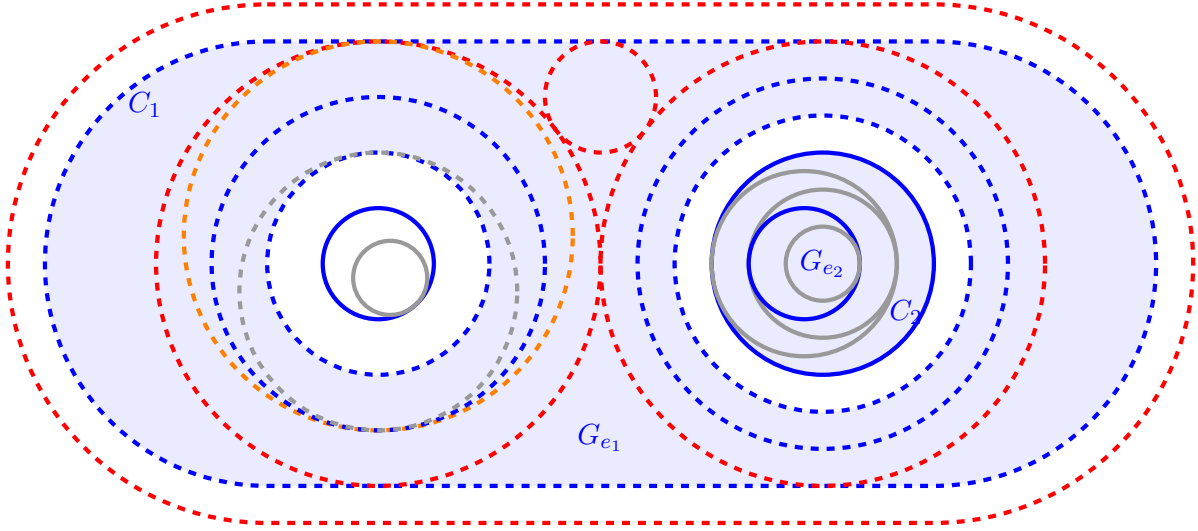


Figure 3.2: Example for the laminar family \mathcal{L}' in the proof of Lemma 3.11. The cycles in \mathcal{B}_1 are drawn in red and dashed, the only cycle in $\mathcal{B}'_1 \setminus \mathcal{B}_1$ is drawn in orange and dashed. Any two cycles in $\mathcal{L}' \setminus \mathcal{B}'_1$ that share a vertex are also homotopic.

Among those cycles we choose a maximal pairwise vertex-disjoint set \mathcal{F} , which is drawn in blue. For each \mathcal{F} -homotopy class the $k-1$ largest cycles are added to \mathcal{B}_2 (dashed blue cycles); for simplicity, we consider the case $k=3$ here although the lemma asserts $k>4$. The set \mathcal{B}'_2 (dashed gray and blue cycles) consists of the remaining cycles intersecting $V(\mathcal{B}_2)$. We only consider solutions that contain no cycles from $\mathcal{B}'_1 \cup \mathcal{B}'_2$, which are the dashed cycles in this example.

The filled blue areas show the graphs G_{e_1} and G_{e_2} , where e_1 and e_2 are edges of T corresponding to the cycles C_1 and C_2 . Note that due to neglecting cycles in $\mathcal{B}'_1 \cup \mathcal{B}'_2$ the cycles in $\mathcal{L}_e \cap \mathcal{F}$ form a chain for all $e \in E(T)$.

Now we can show the (vertex-disjoint version of the) main theorem of this section:

Theorem 3.12. *There is a Polynomial Time Approximation Scheme (PTAS) for the VERTEX-DISJOINT CYCLE PACKING PROBLEM on laminar families of cycles in planar graphs.*

Proof. Let $\varepsilon > 0$ be a constant. Choose $k \in \mathbb{N}$ minimal such that $\left(\frac{k}{k-4}\right)^3 < 1 + \varepsilon$. Let G be a planar graph and \mathcal{L} a laminar family of cycles in G . Let $\mathcal{L}' := \mathcal{L} \setminus \mathcal{L}_{\min}$. We do an induction on $|\mathcal{L}'|$.

We call a cycle $C \in \mathcal{L}'$ “large” if the interior of C contains the interior of k^2 pairwise vertex-disjoint cycles that are also vertex-disjoint to C ; otherwise we call C “small”. For each $C \in \mathcal{L}'$ we can check in polynomial time by complete enumeration whether C is large or

small. In the case where all cycles in \mathcal{L}' are large we can simply apply Lemma 3.11 to get a $(1 + \varepsilon)$ -approximation.

Now consider the other case and pick a small cycle $C \in \mathcal{L}'$ with minimal interior. In particular, the interior of C contains only face-minimal cycles of \mathcal{L} . Let $\mathcal{L}_C \subseteq \mathcal{L}_{\min}$ denote the set of these cycles and $\mathcal{L}_C^- \subseteq \mathcal{L}_C$ the set of cycles that are also vertex-disjoint to C . Since C is small we can compute a maximum-cardinality set $\mathcal{S}_C^- \subseteq \mathcal{L}_C^-$ of pairwise vertex-disjoint cycles among \mathcal{L}_C^- in polynomial time by complete enumeration. Similarly, we can decide whether there exist $|\mathcal{S}_C^-| + 1$ pairwise vertex-disjoint cycles among \mathcal{L}_C . We distinguish two cases:

Case 1: There exists a set $\mathcal{S}_C \subseteq \mathcal{L}_C$ of pairwise vertex-disjoint cycles with $|\mathcal{S}_C| > |\mathcal{S}_C^-|$. In this case there always exists an optimum cycle packing for \mathcal{L} that does not contain C : If the cycle packing $\mathcal{S} \subseteq \mathcal{L}$ contains C then we can replace C and all at most $|\mathcal{S}_C^-|$ cycles contained in C by \mathcal{S}_C without decreasing the size of \mathcal{S} .

Therefore, any $(1 + \varepsilon)$ -approximation for the VERTEX-DISJOINT CYCLE PACKING PROBLEM in $\mathcal{L} \setminus \{C\}$ yields a $(1 + \varepsilon)$ -approximation for \mathcal{L} and we can continue on $\mathcal{L} \setminus \{C\}$, which decreases the size of \mathcal{L}' .

Case 2: The maximum number of pairwise vertex-disjoint cycles in \mathcal{L}_C is $|\mathcal{S}_C^-|$. Then, there exists an optimum cycle packing \mathcal{S}^* for \mathcal{L} with $\mathcal{S}^* \cap \mathcal{L}_C = \mathcal{S}_C^-$ because for any cycle packing \mathcal{S} for \mathcal{L} we can feasibly replace $\mathcal{S} \cap \mathcal{L}_C$ by \mathcal{S}_C^- without decreasing the size of \mathcal{S} .

Thus, we can recurse on $\mathcal{L} \setminus \mathcal{L}_C$, which decreases the size of \mathcal{L}' . Any $(1 + \varepsilon)$ -approximation for $\mathcal{L} \setminus \mathcal{L}_C$ yields a $(1 + \varepsilon)$ -approximation for \mathcal{L} by adding the cycles in \mathcal{S}_C^- to the solution. \square

Clearly, there is also an edge-disjoint variant of Theorem 3.12. We can deduce it directly from the reduction in Proposition 2.36:

Theorem 3.13. *There is a Polynomial Time Approximation Scheme (PTAS) for the EDGE-DISJOINT CYCLE PACKING PROBLEM for laminar families of cycles in planar graphs.* \square

Due to the fact that uncrossable cycle families allow for uncrossing (cf. Section 2.6) this also gives approximation guarantees for the CYCLE PACKING PROBLEM for uncrossable cycle families in planar graphs:

Corollary 3.14. *For any fixed $\varepsilon > 0$ there is a polynomial-time $(\alpha^* + \varepsilon)$ -approximation algorithm for the (vertex- or edge-disjoint) CYCLE PACKING PROBLEM for uncrossable cycle families in planar graphs, given by a weight oracle, where α^* denotes the laminar cycle packing integrality gap.*

Proof. We first compute an optimum solution x to the cycle packing LP (1.1) or (1.3), respectively, with laminar support. This can be done by Theorem 2.33. Let \mathcal{L} denote the support of our LP solution. Clearly, x is still an optimum solution to our LP after replacing the uncrossable cycle family \mathcal{C} by \mathcal{L} . By the definition of α^* there exists a cycle packing $\mathcal{S}^* \subseteq \mathcal{L}$ with $\sum_{C \in \mathcal{L}} x_C \leq \alpha^* |\mathcal{S}^*|$. By applying Theorem 3.12 or Theorem 3.13 to \mathcal{L} we can compute a solution $\mathcal{S} \subseteq \mathcal{L}$ of pairwise disjoint cycles such that $|\mathcal{S}^*| \leq (1 + \frac{\varepsilon}{\alpha^*}) |\mathcal{S}|$. This concludes the proof. \square

Note that there is currently no example for an uncrossable cycle family known where the best known approximation ratio for the CYCLE PACKING PROBLEM is worse than the corresponding integrality gap. However, on the other hand we know of no example showing $\alpha^* > 2$.

Chapter 4

A simple LP-based approximation

In this chapter we analyze a simple LP-based greedy rounding algorithm for the (uncrossable) CYCLE PACKING PROBLEM. We will start with planar graphs, where we achieve a 5-approximation (Theorem 4.3). In Section 4.1 we also note that a previous result by Garg, Kumar and Sebő [37] on the fully planar EDGE-DISJOINT PATHS PROBLEM carries over to the general uncrossable EDGE-DISJOINT CYCLE PACKING PROBLEM. Both of these results also bound the corresponding integrality gaps. They are not the state of the art any more since we will improve on both bounds in Chapter 5. However, the approaches in this section are simple and allow for extensions to cycle packing in bounded-genus graphs (cf. Section 4.3) and weighted cycle packing (Section 4.4). All results from this chapter are joint work with Hanjo Thiele and Jens Vygen [79].

4.1 Edge-disjoint packing in planar graphs

Garg, Kumar and Sebő [37] recently gave a 4-approximation algorithm for EDGE-DISJOINT PATHS PROBLEM in fully planar instances, which is equivalent to the planar EDGE-DISJOINT D -CYCLE PACKING PROBLEM. In fact, after computing an LP solution x with laminar support \mathcal{L} they do not need any specific properties of D -cycles any more, but only use the laminar structure of \mathcal{L} :

First, they round x to a half-integral solution, losing at most half of the total LP value. For this they observe that all cycles that contain an edge e form a chain. (Note that this property is false for vertex-disjoint cycle packing: a vertex can belong to many cycles with disjoint interior.) They divide this chain into two chains, say $L_1(e)$ and $L_2(e)$ such that any cycles in $L_1(e)$ have disjoint interior to cycles in $L_2(e)$. Therefore the LP

$$\max \left\{ \sum_{C \in \mathcal{C}^{>0}} x_C : \sum_{C \in L_i(e)} x_C \leq 1 \ (e \in E, i = 1, 2), \ x_C \geq 0 \ (C \in \mathcal{C}) \right\} \quad (4.1)$$

has an integral optimum solution, computable in polynomial time. This solution is given by a subset $\mathcal{C}_{1/2} \subseteq \mathcal{L}$, such that setting $x_C = \frac{1}{2}$ for all $C \in \mathcal{C}_{1/2}$ and $x_C = 0$ for all other cycles C constitutes a feasible LP solution with $|\mathcal{C}_{1/2}| \geq \sum_{C \in \mathcal{C}} x_C$. Then they exploit an observation of [33] that the conflict graph (with vertex set $\mathcal{C}_{1/2}$ and edges between cycles that share an edge) is planar. Therefore, using (an algorithmic version of) the four-color theorem one can find a subset $\mathcal{C}_1 \subseteq \mathcal{C}_{1/2}$ of pairwise edge-disjoint cycles with $|\mathcal{C}_1| \geq \frac{1}{4}|\mathcal{C}_{1/2}|$.

Clearly, replacing the computation of x by our Theorem 2.33 yields a 4-approximation for general uncrossable edge-disjoint cycle packing:

Theorem 4.1. *Given a planar graph G and a weight oracle for an uncrossable family \mathcal{C} of cycles in G , we can find an edge-disjoint subset of \mathcal{C} whose cardinality is at least $\frac{1}{4}$ the value of the LP (1.3) in polynomial time. \square*

4.2 Vertex-disjoint packing in planar graphs

In this section now we show our approach for the (more difficult) vertex-disjoint case. The core idea is still to exploit the laminar structure of the support of an optimum LP solution and to round that fractional solution to an integral one. We first only present the version for planar graphs; however our approach for separating cycles in more complex orientable surfaces works similarly.

Algorithm 2: Greedy Rounding for Cycle Packing

Input: A planar graph G , a weight oracle for an uncrossable family \mathcal{C} of cycles in G
Output: A set $\mathcal{S} \subseteq \mathcal{C}$ of pairwise vertex-disjoint cycles

- 1 Compute an embedding of G in the sphere
- 2 Compute an optimum solution x to the LP (1.1) with laminar support
- 3 **while** $x \neq 0$ **do**
- 4 Let $\mathcal{L}_x := \{C \in \mathcal{C} : x_C > 0\}$ be the support of x
- 5 Pick a cycle $F^* \in \mathcal{L}_x$
- 6 Add F^* to the solution \mathcal{S}
- 7 Set $x_C := 0$ for all $C \in \mathcal{N}_{\mathcal{L}_x}(F^*)$
- 8 **end**
- 9 Output \mathcal{S}

Step 1 works in polynomial time [26]. For step 2 we can apply Theorem 2.33. Afterwards, in each iteration we add a cycle F^* in the support to our solution and remove all neighboring cycles. Recall that $\mathcal{N}_{\mathcal{L}_x}(F^*)$ denotes the set of all cycles in the support \mathcal{L}_x that share a vertex with F^* . In order to show that Algorithm 2 is a constant-factor approximation, we only need to bound the decrease of our LP solution in each iteration by a constant because each iteration increases the solution size by exactly one. To this end, we show a new structure lemma, which we call *Efficient Cycle Lemma*:

Lemma 4.2 (Efficient Cycle Lemma). *Let G be a planar graph embedded in the sphere, and let \mathcal{L} be a non-empty laminar set of cycles in G . Then there exists a cycle $F^* \in \mathcal{L}$ and a vertex set $W \subseteq V(F^*)$ such that $|W| \leq 5$ and every $C \in \mathcal{N}_{\mathcal{L}}(F^*)$ contains a vertex from W .*

We will prove the Efficient Cycle Lemma in the following sections. As an easy consequence, Algorithm 2 yields a 5-approximation:

Theorem 4.3. *Given a planar graph G and a weight oracle for an uncrossable family \mathcal{C} of cycles in G , we can find a vertex-disjoint subset of \mathcal{C} whose cardinality is at least $\frac{1}{5}$ the value of the LP (1.1) in polynomial time.*

Proof. We apply Algorithm 2. Clearly, the output \mathcal{S} is a set of pairwise vertex-disjoint cycles because whenever we add a cycle F^* to \mathcal{S} Step 7 makes sure that all remaining cycles in consideration are vertex-disjoint to F^* . For Step 5 we pick a cycle $F^* \in \mathcal{L}_x$ minimizing $x(\mathcal{N}_{\mathcal{L}_x}(F^*))$. By the Efficient Cycle Lemma 4.2 there always exists a cycle $F \in \mathcal{L}_x$ such that all cycles in $\mathcal{N}_{\mathcal{L}_x}(F)$ can be covered by a vertex set $W \subseteq V(F)$ of size at most 5. In particular,

$$x(\mathcal{N}_{\mathcal{L}_x}(F)) \leq \sum_{w \in W} \sum_{C \in \mathcal{L}_x: w \in V(C)} x_C \leq |W| \leq 5$$

and therefore also $x(\mathcal{N}_{\mathcal{L}_x}(F^*)) \leq 5$. This means that in each iteration of our algorithm we add a cycle to \mathcal{S} and remove cycles with a total LP value of at most 5 from \mathcal{L}_x . This finishes the proof. \square

In Section 4.3 we explain the differences in the bounded-genus case. Here we obtain an LP solution with uncrossed support, and we proceed exactly as above for the set of separating cycles. The only difference is that the constant 5 in Lemma 4.2 will depend on the genus. For the non-separating cycles we proceed exactly as in [48].

Combining our Efficient Cycle Lemma with the fractional local ratio method also yields constant-factor approximations for the weighted version of the CYCLE PACKING PROBLEM. We describe this result in Section 4.4.

4.2.1 Nice paths

The next three sections are devoted to the Efficient Cycle Lemma 4.2. In fact, we immediately prove the generalization for the bounded-genus case which results in a larger, but still constant size of W . So let G be a graph embedded in an orientable surface Σ of genus g , the sphere if G is planar. Let also \mathcal{L} be a laminar family of separating cycles in G . We assume $|\mathcal{L}| \geq 2$ since otherwise Lemma 4.2 is trivial. Thus, each cycle has at most one one-sided side.

We will see that it is always possible to choose a one-sided cycle F^* in Lemma 4.2. We will construct a graph G' that can also be embedded in Σ . The vertices of G' are the one-sided cycles, and the edges represent conflicts. In the end, F^* will correspond to a low-degree vertex in G' . Since many one-sided cycles can share a vertex, our graph will not contain an edge for all pairs of such cycles. Moreover, we have to take conflicts to two-sided cycles into account. To embed the edges of G' , we will use *nice paths*:

Definition 4.4 (nice path). Let S be a side of a cycle $C \in \mathcal{L}$ and x a point on the embedding of C that does not lie on the embedding of any other cycle contained in S . A *nice path* for (x, S) is a continuous path P on Σ such that

1. P starts in x and ends in a point on the embedding of some one-sided cycle C' that is contained in S
2. Except for the start point x , P is contained in S
3. P intersects any cycle in \mathcal{L} in at most one point

Remark 4.5. If (x, S) and (x, S') are two pairs as in the above definition such that S and S' are disjoint (or equivalently $S \neq S'$), it is easy to check that the concatenation of nice paths for (x, S) and (x, S') is again a nice path in both directions.

Lemma 4.6. *Let T be a finite set of pairs (x, S) as in Definition 4.4. Then there are nice paths $(P_t)_{t \in T}$ for the $t \in T$ such that all the P_t are pairwise disjoint, except for possibly coinciding start points $(x, S), (x, S') \in T$.*

Proof. Let $X := \{x : (x, S) \in T \text{ for some } S\}$ be the set of all start points. We construct the paths one by one, ensuring in addition that the nice path P_t for $t = (x, S) \in T$ contains no point in $X \setminus \{x\}$. Assume that there are already nice paths P_t with the desired properties for all $t \in T' \subset T$, and let $t_0 \in T \setminus T'$. We show that there is a nice path for $t_0 = (x_0, S_0) \in T$ that (except possibly at x_0) avoids X and all previously constructed paths.

We prove this by induction on the number of cycles of \mathcal{L} that have a side strictly contained in S_0 . If S_0 does not strictly contain a side of any cycle in \mathcal{L} , x_0 is already a feasible endpoint and we can take a trivial path as P_{t_0} . Otherwise, consider the connected components that arise from S_0 after deleting the embeddings of the cycles in \mathcal{L} and the previously constructed paths P_t ($t \in T'$). Let A be a connected component such that x_0 is on the boundary of A .

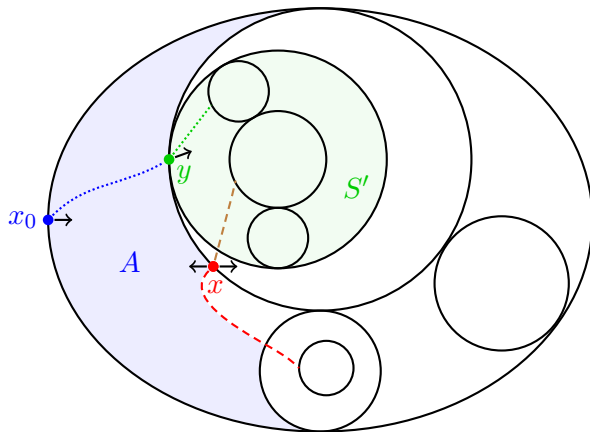


Figure 4.1: Example for nice paths. The set X contains two points x and x_0 . For the point x nice paths to both incident sides are demanded and have already been found (dashed). Note that the concatenation is again a nice path in both directions. Now we ask for a nice path for (x_0, S_0) , where S_0 is the interior of the largest cycle in the figure. We can connect x_0 feasibly to all points on the boundary of the blue area A . We choose a point y on a cycle with a side $S' \subseteq S_0$ (and choose that cycle so that S' is minimal) and complete an x_0 - y -path inside A (blue, dotted) to a nice path by a nice path for (y, S') (green, dotted), which exists by induction.

If two previously constructed paths have a common point, their concatenation is also a nice path by Remark 4.5. Any nice path P_t can touch the boundary of S_0 at most once. Moreover S_0 strictly contains a side of a cycle in \mathcal{L} . Hence the boundary of A contains a point y that is neither on the boundary of S_0 nor on a previously constructed path P_t ($t \in T'$) nor in X (cf. Figure 4.1).

Therefore y must lie on a cycle $C \in \mathcal{L}$ with a side S' strictly contained in S_0 , such that S' does not contain any other cycle that touches y . Then the induction hypothesis implies that there is a nice path P' for (y, S') that is disjoint to each of the P_t ($t \in T'$) and to X . By prepending an x_0 - y -path whose interior is inside A we get the desired nice path for t_0 . \square

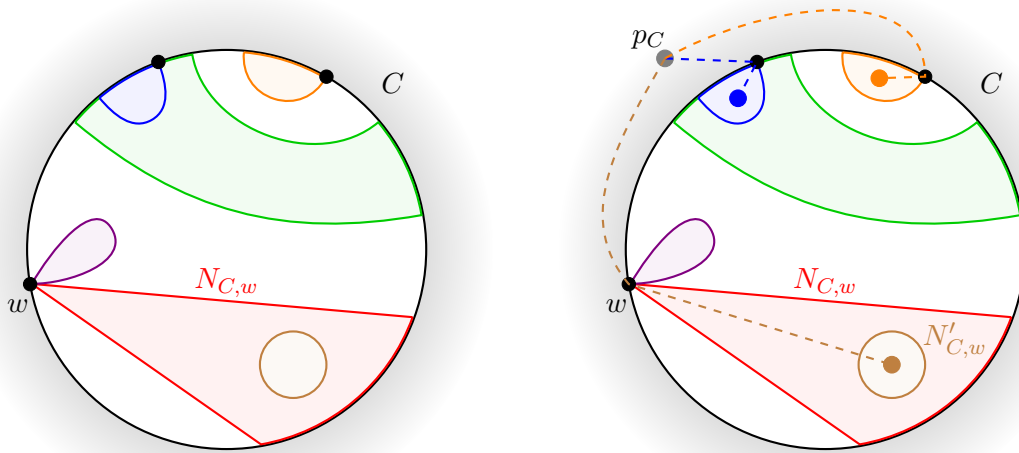


Figure 4.2: The left picture shows a one-sided cycle C and all cycles that share a vertex with C . Note that the one-sided side of C is drawn as the outer face in this example. The set $\mathcal{N}^*(C)$ contains all the colored cycles except the brown (which is not a neighbour) and the green one (which is not a \subseteq_C -minimal neighbour). The black vertices mark the set $W(C)$. For the specified vertex w the neighbour $N_{C,w}$ is given by the red cycle.

In the right picture the dashed lines represent embeddings of the edges that are added for C to G' . Note that the red cycle is two-sided and the corresponding edge ends inside a one-sided cycle $N'_{C,w}$ inside $N_{C,w}$.

4.2.2 Proof of the Efficient Cycle Lemma

Now we can prove the more general version of the Efficient Cycle Lemma 4.2 which considers G to be embedded in a bounded-genus surface. We can also enforce F^* to be one-sided.

Lemma 4.7. *Let G be a graph embedded in an orientable surface Σ of genus g , the sphere if G is planar. Let \mathcal{L} be a non-empty laminar set of separating cycles in G . Then there exists a one-sided cycle $F^* \in \mathcal{L}$ and a vertex set $W \subseteq V(F^*)$ such that $|W| \leq 6g + 5$ and every $C \in \mathcal{N}_{\mathcal{L}}(F^*)$ contains a vertex from W .*

Proof. W.l.o.g. $|\mathcal{L}| \geq 2$. Let $\mathcal{L}_1 \subseteq \mathcal{L}$ denote the set of one-sided cycles in \mathcal{L} . Let $C \in \mathcal{L}_1$ and let S be the one-sided side of C . We say that a cycle $C_1 \in \mathcal{L}$ contains another cycle $C_2 \in \mathcal{L}$ w.r.t. C and write $C_2 \subseteq_C C_1$ if the side of C_1 that is disjoint to S contains a side of C_1 .

For a one-sided cycle $C \in \mathcal{L}_1$ let $\mathcal{N}^*(C) \subseteq \mathcal{N}_{\mathcal{L}}(C)$ be the set of \subseteq_C -minimal cycles in $\mathcal{N}_{\mathcal{L}}(C)$. Let $W(C)$ be a minimal set such that $W(C) \cap V(N) \neq \emptyset$ for all $N \in \mathcal{N}^*(C)$. Now it suffices to find a one-sided cycle C such that $|W(C)| \leq 6g + 5$: Let N be some cycle that shares a vertex with C . Pick $N' \in \mathcal{N}^*(C)$ such that $N' \subseteq_C N$. Now there is some $w \in V(N') \cap W(C)$ and since $N' \subseteq_C N$, we also have $w \in V(N) \cap V(C)$.

In order to find a cycle $C \in \mathcal{L}_1$ with $|W(C)| \leq 6g + 5$, we first construct an embedding in Σ of a directed graph G' with vertex set \mathcal{L}_1 . For each $C \in \mathcal{L}_1$ we choose a point p_C in the one-sided side of C . Now for each $C \in \mathcal{L}_1$ with one-sided side S and $w \in W(C)$ let $N_{C,w}$ be

the cycle in $\mathcal{N}^*(C)$ that comes in anti-clockwise direction first after C in w . We can find a nice path for (w, S') from w to the boundary of some one-sided cycle $N'_{C,w}$, where S' is the side of $N_{C,w}$ that is disjoint to S . We add the edge $(C, N'_{C,w})$ to G' and get an embedding of that edge by prepending a p_C - w -path inside S and appending a path to $p_{N'_{C,w}}$ inside the one-sided side of $N'_{C,w}$ to the nice path (cf. Figure 4.2).

By Lemma 4.6 the nice paths can be drawn disjointly (except for coinciding start points), therefore the embeddings of the edges in G' can intersect only in their endpoints and in points $w \in W(C) \cap W(C')$. However, in this case the path that was added for (C, w) is continued inside a cycle that comes in anti-clockwise direction before C' in w , while the path that was added for (C', w) is continued inside a cycle that comes in anti-clockwise direction before C in w (cf. Figure 4.3). Therefore the embeddings of edges in G' can only 'touch' but not 'cross'; in particular they induce an embedding of G' .




Figure 4.3: Three one-sided cycles, C_1 to C_3 , and three two-sided cycles, N_1 to N_3 , meet in the point w . Note that w can be in $W(C_i)$ for all $1 \leq i \leq 3$. The dotted lines represent the constructed embeddings of edges added to G' .

Note that the out-degree of any vertex $C \in \mathcal{L}_1$ in this digraph G' is exactly $|W(C)|$. We claim that G' contains no pair of homotopic edges. To prove this, assume there are $C_1, C_2 \in \mathcal{L}_1$ such that homotopic edges from C_1 to C_2 were added for (C_1, w) and (C_1, w') . Since both edges point to p_{C_2} , we know that $N_{C_1,w}$ and $N_{C_1,w'}$ must be ordered by \subseteq_{C_1} . Furthermore, both $N_{C_1,w}$ and $N_{C_1,w'}$ are in $\mathcal{N}(C_1)$ and thus $N_{C_1,w} = N_{C_1,w'}$ (cf. Figure 4.4). However, due to the minimality of $W(C)$ there must be other \subseteq_{C_1} -minimal cycles C_3 at w and C_4 at w' that come after $N_{C_1,w}$ in anti-clockwise order (cf. Figure 4.4). The cycles C_3 and C_4 are not necessarily one-sided, but both sides contain a one-sided cycle and hence a vertex of G' .

Thus, G' is a directed graph embedded in Σ and without pairs of homotopic edges. By Lemma 2.23 there is some $F^* \in \mathcal{L}_1$ with out-degree at most $6g + 5$ in G' ; therefore $|W(F^*)| \leq 6g + 5$. \square




Figure 4.4: If there are two parallel edges in G' , say from C_1 to C_2 , added for (C_1, w) and (C_1, w') , then the embedding of these two edges does not bound an area without any cycle in \mathcal{L} : The case on the left, in which the edges are continued by nice paths in different cycles, cannot happen because one of them must contain the other and is therefore not in $\mathcal{N}^*(C_1)$. In the other case, due to the minimality of $W(C_1)$, there must be other \subseteq_{C_1} -minimal neighbours C_3 at w and C_4 at w' as shown.

4.2.3 Tightness of the Efficient Cycle Lemma

The constant 5 in the Efficient Cycle Lemma for planar graphs is best possible: let G be a truncated dodecahedron, where the cycles in \mathcal{L} are given exactly by the decagonal faces, i.e., every cycle in \mathcal{L} is one-sided and touches exactly five neighbours, but there is no vertex contained in more than two of these cycles. Then the constructed graph G' is an icosahedron; all vertices have degree 5.

However, the postulation that $W \subseteq V(F^*)$ in the Efficient Cycle Lemma is not necessary for the proof of Theorem 4.3. If this postulation is relaxed, the best lower bound for the constant in the planar version of the Efficient Cycle Lemma that we know is 4:

Lemma 4.8. *There exists a planar graph G and a laminar set \mathcal{L} of cycles in G such that for every $C \in \mathcal{L}$ and every $X \subseteq V(G)$ with $|X| \leq 3$ there is a cycle in \mathcal{L} that shares a vertex with C but contains no vertex of X .*

Proof. Let G' be the planar graph that arises from replacing each face of a cube by a 7×7 grid. Let G be a planar graph that consists of edge-disjoint cycles $(C_v)_{v \in V(G')}$, each bounding a face on the sphere, such that C_v and C_w share exactly one vertex if and only if $\{v, w\} \in E(G')$. Let $\mathcal{L} := \{C_v : v \in V(G')\}$.

It is clear that except for the cycles corresponding to the corners of the cube, each cycle in \mathcal{L} has exactly four neighbours that are pairwise vertex-disjoint. Thus it is not possible to cover those with three vertices. For each of the corner cycles we add three two-sided neighbours as in Figure 4.5 to \mathcal{L} (and G). Now there are seven cycles sharing a vertex with a corner cycle, including the corner cycle itself, which means it is not possible to cover all of them by three vertices in G as each such vertex covers only two of them. A similar argument for the added two-sided cycles concludes the proof. \square




Figure 4.5: The black dashed lines show one of the six 7×7 grids G' is built of. For each grid vertex, G contains one one-sided cycle in \mathcal{L} , as shown by the blue cycles. For each corner cycle C of the grid we add a two-sided cycle (red) that meets C in a vertex that is not contained in any other neighbour of C . Since C is adjacent to exactly three of those grids (three sides of the cube), it shares a vertex with exactly three one-sided and three two-sided cycles. Therefore it is not possible to cover C and all its neighbours with three vertices.

Closing the gap between the lower bound 4 (this example) and the upper bound 5 in the planar Efficient Cycle Lemma remains an open problem.

4.3 Packing cycles in bounded-genus graphs

In this section we generalize Theorem 4.1 and 4.3 to the case where G is not planar, but embedded in an orientable surface of constant genus g . Note that if all (or most) of the support of our LP solution consists of separating cycles then the proof of Theorem 4.3 still works by using the more general version of the Efficient Cycle Lemma 4.7. For the non-separating cycles we use the following lemma by Huang et al. [48]:

Lemma 4.9 ([48]). *Let \mathcal{L} be an uncrossed family of cycles in a graph G , embedded in an orientable surface of genus g . If all cycles in \mathcal{L} are freely homotopic, then \mathcal{L} can be ordered $\mathcal{L} = \{C_1, \dots, C_k\}$ such that for each $z \in V(G) \cup E(G)$ the cycles containing z are given as an interval $\{C_i, C_{i+1}, \dots, C_j\}$ or $\mathcal{L} \setminus \{C_i, \dots, C_j\}$ for some $1 \leq i \leq j \leq k$.*

In particular, for such cycle families we get a similar statement as our Efficient Cycle Lemma.

Corollary 4.10. *Let \mathcal{L} be an uncrossed family of cycles in a graph G , embedded in an orientable surface of genus g . Assume all cycles in \mathcal{L} are freely homotopic. Let x be a feasible solution to the LP (1.1) on the cycle family \mathcal{L} . Then for any $C \in \mathcal{L}$ we have $x(\mathcal{N}_{\mathcal{L}}(C)) \leq 2$.*

Proof. Let $\mathcal{L} = \{C_1, \dots, C_k\}$ be an ordering as in Lemma 4.9. W.l.o.g. $C = C_1$. Let $i^+ \leq k$ maximum such that there exists a vertex $v \in V(C_1) \cap \dots \cap V(C_{i^+})$. Similarly, choose $i^- \geq 0$ minimum such that there exists a vertex $v \in V(C_{i^-}) \cap \dots \cap V(C_{k+1})$, where we denote $k+1 := 1$. By the property from Lemma 4.9 we have $\mathcal{N}_{\mathcal{L}}(C) = \{C_1, \dots, C_{i^+}\} \cup \{C_{i^-}, \dots, C_k\}$ and therefore $x(\mathcal{N}_{\mathcal{L}}(C)) \leq 2$. \square

This enables us to prove a bounded-genus version of Theorem 4.3. The proof follows closely the approach for the EDGE-DISJOINT PATHS PROBLEM in [48], so we only sketch it here.

Theorem 4.11. *Given a graph G embedded in a fixed orientable surface of genus g and a weight oracle for an uncrossable family \mathcal{C} of cycles in G , we can find a*

(a) *vertex-disjoint subset of \mathcal{C} whose cardinality is $\Omega\left(\frac{1}{g^2}\right)$ times the value of the LP (1.1)*

(b) *edge-disjoint subset of \mathcal{C} whose cardinality is $\Omega\left(\frac{1}{g^2}\right)$ times the value of the LP (1.3)*

in polynomial time.

Proof. Let us start with the vertex-disjoint case (a). First apply Lemma 2.31 to obtain a near-optimal solution x to the LP (1.1) with uncrossed support. Let \mathcal{C}^{sep} denote the set of the separating cycles in \mathcal{C} . If $x(\mathcal{C}^{\text{sep}}) \geq \frac{1}{2}x(\mathcal{C})$, then set $x_C := 0$ for all $C \in \mathcal{C} \setminus \mathcal{C}^{\text{sep}}$ and proceed with Algorithm 2. In each iteration we choose a cycle F^* as guaranteed by Lemma 4.7 (note that the support of x on \mathcal{C}^{sep} is then laminar by Proposition 2.28). Similar to Theorem 4.3, we end up with a vertex-disjoint set of at least $\frac{1}{6g+5} \sum_{C \in \mathcal{C}^{\text{sep}}} x_C$ cycles.

Otherwise partition the non-separating cycles in the support of x into free homotopy classes $\mathcal{C}_1, \dots, \mathcal{C}_k$. Greene [44] showed that k is in $O(g^2 \log g)$. Very recently, Aougab and Gaster [7] improved this bound to $k \leq O(g^2)$. Take the class \mathcal{C}_i with the largest LP value $x(\mathcal{C}_i)$ and ignore other cycles. By Corollary 4.10, applying Algorithm 2 to \mathcal{C}_i yields a vertex-disjoint set $\mathcal{S}_i \subseteq \mathcal{C}_i$ with $|\mathcal{S}_i| \geq \frac{1}{2}x(\mathcal{C}_i)$, independently of the choice of F^* in Step 5.

For the edge-disjoint version (b) of Theorem 4.11, we solve the edge-disjoint cycle packing LP (1.3) instead of (1.1). For the non-separating cycles, we can again use Aougab and Gaster's [7] bound and otherwise follow [48] exactly. For the separating cycles the reduction from Proposition 2.36 translates our edge-disjoint cycle packing problem into a vertex-disjoint one. \square

4.4 Weighted Cycle Packing

A natural generalization of the CYCLE PACKING PROBLEM is to allow for weights on the cycles of \mathcal{C} and try to maximize the total weight of the cycles in the solution:

Definition 4.12. The MAXIMUM-WEIGHT CYCLE PACKING PROBLEM is defined as follows.

Input: A graph G , a family \mathcal{C} of cycles in G with weights $w: \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ (usually given by an oracle)

Task: Find a subset $\mathcal{S} \subseteq \mathcal{C}$ such that the cycles in \mathcal{S} are pairwise vertex-disjoint (in the MAXIMUM-WEIGHT VERTEX-DISJOINT CYCLE PACKING PROBLEM) or edge-disjoint (in the MAXIMUM-WEIGHT EDGE-DISJOINT CYCLE PACKING PROBLEM), respectively, such that $w(\mathcal{S})$ is maximum.

It turns out that our greedy rounding methods from Theorem 4.1, Theorem 4.3 and Theorem 4.11 still work in this more general setting as soon as an optimum LP solution with laminar (or uncrossed) support is achieved. This is due to the fractional local ratio method. We give the details in Section 4.4.1.

The more difficult part for using the results from Section 4.4.1 is to incorporate the weights in the oracles as well as solving the LP and uncrossing the solution without changing the total weight. However, for some uncrossable families, in particular the family of D -cycles, there is a natural way of doing this. We describe our main applications in Section 4.4.2.

4.4.1 Rounding the LP with the Local Ratio Method

As in the unit-weight case also the MAXIMUM-WEIGHT CYCLE PACKING PROBLEM admits natural LP relaxations: The edge-disjoint and vertex-disjoint versions are given by:

$$\max \left\{ \sum_{C \in \mathcal{C}} w(C)x_C : \sum_{C \in \mathcal{C}: e \in C} x_C \leq 1 \ (e \in E(G)), \ x_C \geq 0 \ (C \in \mathcal{C}) \right\} \quad (4.2)$$

$$\max \left\{ \sum_{C \in \mathcal{C}} w(C)x_C : \sum_{C \in \mathcal{C}: v \in C} x_C \leq 1 \ (v \in V(G)), \ x_C \geq 0 \ (C \in \mathcal{C}) \right\} \quad (4.3)$$

In this section we will show how to round a feasible solution x for one of these LPs to an integral solution, preserving a constant fraction of the total weight. For the vertex-disjoint approach we use the fractional local ratio method in a similar way as Chan and Lau [21] did for the hypergraph matching problem. The following result is implicit in their paper:

Theorem 4.13 ([21]). *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a hypergraph and $x: \mathcal{E} \rightarrow \mathbb{R}_{\geq 0}$ such that $x(\{e \in \mathcal{E} : v \in e\}) \leq 1$ for all $v \in \mathcal{V}$. For any $e \in \mathcal{E}$ define $N[e] := \{e' \in \mathcal{E} : e \cap e' \neq \emptyset\}$. Let further $k \in \mathbb{N}$ and $\mathcal{E} = \{e_1, \dots, e_m\}$ such that $x(N[e_i] \cap \{e_i, e_{i+1}, \dots, e_m\}) \leq k$ for all $i = 1, \dots, m$.*

Then for any given edge weights $w: \mathcal{E} \rightarrow \mathbb{R}_{\geq 0}$ one can find a set $M \subseteq \mathcal{E}$ of pairwise disjoint hyperedges such that $w(M) \geq \frac{1}{k} \sum_{e \in \mathcal{E}} w(e)x(e)$ in polynomial time.

The algorithm works by (i) finding the hyperedge e_i with smallest index such that $w(e_i)$ and $x(e_i)$ are both positive, (ii) modifying the weights by subtracting $w(e_i)$ from the weight of all edges in $N[e_i]$, (iii) applying the algorithm recursively to the resulting instance, and (iv) adding e_i to the resulting solution if this is feasible. The approximation guarantee follows easily from the induction hypothesis and the fact that the returned solution contains at least one element of $N[e_i]$. This has been called the fractional local ratio method; see also [58] for details.

Since the VERTEX-DISJOINT CYCLE PACKING PROBLEM can be formulated as a hypergraph matching problem where each hyperedge corresponds to a cycle in \mathcal{C} , this can be used to obtain constant-factor approximation algorithms for the MAXIMUM-WEIGHT VERTEX-DISJOINT CYCLE PACKING PROBLEM. Once we have a solution x to the corresponding LP (4.2) the order of the hyperedges that is needed in Theorem 4.13 can be established with the Efficient Cycle Lemma 4.7.

Theorem 4.14. *Let G be a planar graph, embedded in the sphere. Let x be a feasible solution with laminar support \mathcal{L} to the LP (4.2) for a family \mathcal{C} of cycles in G . Then we can find a vertex-disjoint subset $\mathcal{S} \subseteq \mathcal{L}$ with $w(\mathcal{S}) \geq \frac{1}{5} \sum_{C \in \mathcal{L}} w(C)x_C$ in polynomial time.*

Proof. We can view the laminar LP solution as a fractional hypergraph matching in a hypergraph $\mathcal{G} := (\mathcal{V}, \mathcal{E})$: Let $\mathcal{V} := V(G)$ and $\mathcal{E} := \{V(C) : C \in \mathcal{L}\}$. For $l = 1, \dots, |\mathcal{L}|$ we know from Lemma 4.2 that there exists some $e_l \in \mathcal{E}$ such that $N[e_l] \setminus \{e_1, \dots, e_{l-1}\}$ can be covered by

at most five vertices and thus $x(N[e_l] \setminus \{e_1, \dots, e_{l-1}\}) \leq 5$. Since we can find such an e_l in polynomial time by complete enumeration, we find an order $\mathcal{E} = \{e_1, \dots, e_{|\mathcal{L}|}\}$ as in the setting of Theorem 4.13. Therefore we can apply Theorem 4.13 which yields a solution as desired. \square

For the edge-disjoint case it turns out that the methods of Garg, Kumar and Sebő (see Section 4.1) still work in the weighted case:

Theorem 4.15. *Let G be a planar graph, embedded in the sphere. Let x be a feasible solution with laminar support \mathcal{L} to the LP (4.3) for a family \mathcal{C} of cycles in G . Then we can find an edge-disjoint subset $\mathcal{S} \subseteq \mathcal{L}$ with $w(\mathcal{S}) \geq \frac{1}{4} \sum_{C \in \mathcal{L}} w(C)x_C$ in polynomial time.*

Proof. We use the same algorithm as sketched in Section 4.1. The weighted version of LP (4.1) is still integral, therefore we get a half-integral solution to (4.3) of at least half the value of x . Since the conflict graph is planar and hence four-colorable, we can partition all cycles in the support of that half-integral LP solution into four edge-disjoint families, one of which must have enough weight. \square

Similarly, we can also generalize our methods for bounded-genus instances from Section 4.3 to the weighted case:

Theorem 4.16. *Let G be a graph, embedded in a fixed orientable surface of genus g , the sphere if G is planar. Let x be a feasible solution with uncrossed support \mathcal{L} to the LP (1.1) or (1.3) for a family \mathcal{C} of cycles in G . Then we can find an integral solution to the LP, given by a set $\mathcal{S} \subseteq \mathcal{L}$ with $w(\mathcal{S}) \geq \Omega\left(\frac{1}{g^2}\right) \sum_{C \in \mathcal{L}} w(C)x_C$ in polynomial time.*

Proof. If the separating cycles contribute more than half to the LP value of x , we get an $O(g)$ -approximation by replacing the planar Efficient Cycle Lemma 4.2 in Theorem 4.14 by the bounded-genus version 4.7.

Now consider the case that the non-separating cycles contribute more to the LP value of x . Pick the free homotopy class \mathcal{C}^* that carries most of the (weighted) LP value, which is at least an $\Omega\left(\frac{1}{g^2}\right)$ fraction [7]. Ignore all other cycles and use the proof of Theorem 4.14 on \mathcal{C}^* , but replace the planar Efficient Cycle Lemma 4.2 by Corollary 4.10. \square

4.4.2 Applications of Weighted Cycle Packing

For the methods from Section 4.4.1 to achieve good approximations of the MAXIMUM-WEIGHT CYCLE PACKING PROBLEM we need to first compute an (almost) optimum LP solution with laminar or uncrossed support. A crucial step for this is uncrossing. In the weighted setting our uncrossing methods still work if single uncrossing steps do not lose any weight:

Definition 4.17. Let G be a graph and \mathcal{C} an uncrossable family of cycles with weights $w: \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$. We call (\mathcal{C}, w) *weight-uncrossable* if for any cycles $C_1, C_2 \in \mathcal{C}$ and $C'_1, C'_2 \in \mathcal{C}$ with $E(C'_1) + E(C'_2) \subseteq E(C_1) + E(C_2)$ we have $w(C'_1) + w(C'_2) \geq w(C_1) + w(C_2)$.

Note that this definition is a bit stronger than what we actually need, however all relevant applications that we know of have this property. We note:

Proposition 4.18. *Let G be a graph and $D \subseteq E(G)$ a set of demand edges with weights $w: D \rightarrow \mathbb{R}_{\geq 0}$. Then the set \mathcal{C} of D -cycles, together with weights given by $w(C) := w(E(C) \cap D)$ for $C \in \mathcal{C}$ is weight-uncrossable.*

Proof. We observe that in the setting of Definition 4.17 we have $(E(C_1) \cap D) + (E(C_2) \cap D) = (E(C'_1) \cap D) + (E(C'_2) \cap D)$. This finishes the proof. \square

Proposition 4.19. *Let G be a graph and \mathcal{C} an uncrossable family of cycles in G . Let $l: E(G) \rightarrow \mathbb{Z}_{>0}$ and $w: E(G) \rightarrow \mathbb{R}_{\geq 0}$. Then the cycle family $\mathcal{C}[l]$ as defined in Proposition 2.8, together with weights given by $w(C) := w(E(C))$ for $C \in \mathcal{C}[l]$, is weight-uncrossable.*

Proof. In the setting of Definition 4.17 we have $l(E(C'_1)) + l(E(C'_2)) \leq l(E(C_1)) + l(E(C_2))$. By the definition of $\mathcal{C}[l]$ this implies equality. Since $l > 0$ we conclude $E(C'_1) + E(C'_2) = E(C_1) + E(C_2)$, finishing the proof. \square

Since we do not lose any weight when doing uncrossing steps the machinery from Section 2.6 allows us to transform an optimum solution to one of the LPs (4.3) and (4.2) into a near-optimum one with uncrossed support.

It remains to solve the LPs (4.3) and (4.2). For the family $\mathcal{C}[l]$ with weights as given in Proposition 4.19 we remark that the separation oracle for the duals of (4.3) and (4.2) reduces to finding a cycle $C \in \mathcal{C}[l]$ with $y(C) < w(C)$ for given weights $y, w: E(G) \rightarrow \mathbb{R}_{\geq 0}$. This can be done by finding a negative cycle in \mathcal{C} w.r.t. the weights $y - w$, which is possible similar to Proposition 2.13. Thus, the Theorems in Section 4.4.1 yield constant-factor approximation algorithms for the corresponding MAXIMUM-WEIGHT CYCLE PACKING PROBLEM.

However, the more interesting application is given by Proposition 4.18: In the MAXIMUM-WEIGHT DISJOINT PATHS PROBLEM we are given a graph G and a set $D \subseteq E(G)$ of demand edges with weights $w: D \rightarrow \mathbb{R}_{\geq 0}$ and want to find a maximum-weight set of demands together with pairwise disjoint paths for these demands. The problem is clearly equivalent to the MAXIMUM-WEIGHT D -CYCLE PACKING PROBLEM with weights as defined in Proposition 4.18. Note that for this particular problem there exist equivalent formulations of the LPs (4.3) and (4.2) of polynomial size: We can replace the variables for all cycle $C \in \mathcal{C}$ containing a demand edge $d = \{s, t\} \in D$ by an s - t -flow of value at most 1 in $G - D$. Thus, we can solve the LPs directly via the Ellipsoid Method.

As a consequence, we note:

Theorem 4.20. *Let G be a planar graph embedded in the sphere and $D \subseteq E(G)$ a set of demand edges with weights $w: D \rightarrow \mathbb{R}_{\geq 0}$. Let \mathcal{C} be the set of D -cycles in G with weights as defined in Proposition 4.18. Then we can find a*

- (a) *vertex-disjoint subset of \mathcal{C} whose weight is $\frac{1}{5}$ times the value of the LP (4.2)*
- (b) *edge-disjoint subset of \mathcal{C} whose weight is $\frac{1}{4}$ times the value of the LP (4.3)*

in polynomial time.

Proof. We can solve each of the two LPs in polynomial time due to their equivalent polynomial-size formulations. As in Theorem 2.33 we can even compute optimum solutions with laminar support because also the weighted version of the corresponding LP (2.1) allows for a polynomial-size formulation. Afterwards, we apply Theorem 4.14 or Theorem 4.15, respectively. \square

Theorem 4.21. *Let G be a graph embedded in a fixed orientable surface of genus g and $D \subseteq E(G)$ a set of demand edges with weights $w: D \rightarrow \mathbb{R}_{\geq 0}$. Let \mathcal{C} be the set of D -cycles in G with weights as defined in Proposition 4.18. Then we can find a*

(a) *vertex-disjoint subset of \mathcal{C} whose weight is $\Omega\left(\frac{1}{g^2}\right)$ times the value of the LP (4.2)*

(b) *edge-disjoint subset of \mathcal{C} whose weight is $\Omega\left(\frac{1}{g^2}\right)$ times the value of the LP (4.3)*

in polynomial time.

Proof. After solving the (polynomial-size version of the) LP we apply the uncrossing procedure from Lemma 2.31. Afterwards, Theorem 4.16 yields the result. \square

Previously, only very little was known about the MAXIMUM-WEIGHT DISJOINT PATHS PROBLEM: In general graphs we still have the $O(\sqrt{n})$ -approximation from [22]. For the edge-disjoint version we have constant-factor approximations for very limited settings: Chekuri, Mydlarz and Shepherd [23] gave a 4-approximation algorithm for the case when $G - D$ results from a tree by duplicating edges. If $G - D$ is outerplanar, Naves, Shepherd and Xia [65] obtained a 224-approximation.

Chapter 5

Bounding the integrality gap for planar cycle packing

In this chapter we improve on the bounds of 4 and 5 for the edge-disjoint and vertex-disjoint CYCLE PACKING PROBLEM, respectively, from Section 4:

Theorem 5.1. *The laminar cycle packing integrality gap is at most $\frac{20+\sqrt{130}}{9} < 3.5$, i.e. for any fractional solution to the LP (1.1) for a laminar cycle family \mathcal{C} in a planar graph there exists an integral solution with at least $\frac{9}{20+\sqrt{130}}$ times the LP value.*

Due to Proposition 2.37 this upper bound extends to the integrality gaps of the LPs (1.1) and (1.3) for uncrossable cycle families in planar graphs. Combining these bounds with an upper bound on the integrality gaps of the dual LPs (1.2) and (1.4) of 2.4 by Berman and Yaroslavtsev [16] directly yields the following bound on the Erdős–Pósa ratio:

Corollary 5.2. *Let G be a planar graph and \mathcal{C} an uncrossable family of cycles in G . Define $\alpha := 2.4 \cdot \frac{20+\sqrt{130}}{9} < 8.38$. Then $\tau_e(\mathcal{C}) \leq \alpha \nu_e(\mathcal{C})$ and $\tau_v(\mathcal{C}) \leq \alpha \nu_v(\mathcal{C})$.*

This improves the upper bound on the Erdős–Pósa ratio for many interesting examples of uncrossable cycle families in planar graphs, see Tables 1.1 and 1.2. The best known lower bound is 4 for both the edge and vertex version of our problems.

5.1 A new rounding algorithm

In this section we explain our main algorithm, which is only a slight modification of Algorithm 2. We first only analyze the easiest variant of the algorithm. This already yields an upper bound of 4 on the integrality gap for the cycle packing LP, equalizing the best known upper bounds for edge-disjoint and vertex-disjoint cycle packing. In Section 5.3 we will analyze a more refined version of the algorithm, which yields an upper bound of below 3.5.

Note that we again only consider the vertex-disjoint version of the CYCLE PACKING PROBLEM. Due to the reduction in Proposition 2.36 the same methods also work for edge-disjoint packing.

Definition 5.3. Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Let \mathcal{L}_1 be the set of one-sided cycles in \mathcal{L} . Recall that for any $C \in \mathcal{L}$ the set $\mathcal{N}_{\mathcal{L}}(C)$

denotes the set of “neighbours” of C , i.e. cycles in \mathcal{L} that contain a vertex of C . Define $\mathcal{N}_{\mathcal{L}}^1(C) := \mathcal{N}_{\mathcal{L}}(C) \cap \mathcal{L}_1$ to be the set of one-sided “neighbours” of C .

Our modified algorithm works as follows.

Algorithm 3: Structured Rounding for Cycle Packing

Input: A planar graph G , a weight oracle for an uncrossable family \mathcal{C} of cycles in G
Output: A set $\mathcal{S} \subseteq \mathcal{C}$ of pairwise vertex-disjoint cycles

- 1 Compute an embedding of G in the sphere
- 2 Compute an optimum solution x to the LP (1.1) with laminar support
- 3 **while** $x \neq 0$ **do**
- 4 Modify x to make it structured (see Definition 5.5)
- 5 Let $\mathcal{L}_x := \{C \in \mathcal{C} : x_C > 0\}$ be the support of x
- 6 Pick a non-empty subset $\mathcal{F}^* \subseteq \mathcal{L}_x$ of pairwise vertex-disjoint cycles
- 7 Add all cycles in \mathcal{F}^* to the solution \mathcal{S}
- 8 Set $x_C := 0$ for all $C \in \bigcup_{C' \in \mathcal{F}^*} \mathcal{N}_{\mathcal{L}_x}(C')$
- 9 **end**
- 10 Output \mathcal{S}

Note that there are only two modifications to Algorithm 2: First, we start each iteration with the new Step 4 which modifies our LP solution. This is needed for our analysis to work, however this step will always maintain the property that x is a feasible solution to our LP and we will never lose any LP value in this step. We explain the notion of *structuredness* and the algorithm for Step 4 in Definition 5.5 and Lemma 5.6.

The second modification compared to Algorithm 2 is that the Steps 6, 7 and 8 allow for a family $\mathcal{F}^* \subseteq \mathcal{L}_x$ of pairwise vertex-disjoint cycles instead of using a single cycle. However, it is still clear that our new algorithm produces a feasible solution to the VERTEX-DISJOINT CYCLE PACKING PROBLEM because the cycles in \mathcal{F}^* are pairwise vertex-disjoint and Step 8 removes all cycles that are in conflict with any of the cycles in \mathcal{F}^* .

By similar arguments as for Algorithm 2 (cf. the proof of Theorem 4.3) we can see that the algorithm achieves approximation guarantee α if we can always find a set \mathcal{F}^* in Step 6 such that $x(\bigcup_{C \in \mathcal{F}^*} \mathcal{N}_{\mathcal{L}_x}(C)) \leq \alpha |\mathcal{F}^*|$. Thus, in order to bound the laminar cycle packing integrality gap we only need to analyze the minimum values of $\frac{x(\bigcup_{C \in \mathcal{F}^*} \mathcal{N}_{\mathcal{L}_x}(C))}{|\mathcal{F}^*|}$ that we can achieve with different choices of \mathcal{F}^* .

Let us now explain in more detail what Step 4 does. To do so, we first introduce the notion of *redundant cycles*, which uses the notion of \mathcal{L} -homotopy in a slightly different way than it is defined in Section 2.5.1: Here, we say two cycles C, C' in a laminar family \mathcal{L} are \mathcal{L} -homotopic if there exist sides S of C and S' of C' such that the set of one-sided cycles that S and S' contain coincide. This is equivalent to the definition in Section 2.5.1 if the point ∞ is chosen within a one-sided side in \mathcal{L} , which we will assume throughout the following sections. Note that in this case Definition 5.4 does not depend on which one-sided side the point ∞ lies in.

Definition 5.4 (redundant cycle). Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Assume the point ∞ lies in a one-sided side of a cycle in \mathcal{L} . We call a two-sided cycle $C \in \mathcal{L}$ *redundant* if it is \mathcal{L} -homotopic to a one-sided cycle in \mathcal{L} (cf. Figure 5.1).

Definition 5.5 (structured LP solution). Let $x \in \mathbb{R}^{\mathcal{C}}$ be a solution to the LP (1.1). We call it *structured* if the support of x is laminar and each connected component \mathcal{L} of the support of x contains no redundant cycles.

Lemma 5.6. Let $x \in \mathbb{R}^{\mathcal{C}}$ be a feasible solution to the LP (1.1) with laminar support \mathcal{L} . Then we can compute a structured solution $x' \in \mathbb{R}^{\mathcal{L}}$ to (1.1) with $x(\mathcal{L}) = x'(\mathcal{L})$ in polynomial time in the size of \mathcal{L} .

Proof. We can consider the connected components of \mathcal{L} separately, so we assume w.l.o.g. that \mathcal{L} is connected. Assume that x is not structured. Let $C \in \mathcal{L}$ be redundant with a side S that contains only one cycle $C' \neq C$ in \mathcal{L} . This can be achieved by choosing C and S minimal. Since \mathcal{L} is connected, $x_C + x_{C'} \leq 1$ holds. Thus, we can shift the LP value from C to C' , i.e. set $x'_{C'} := x_C + x_{C'}$ and $x'_C := 0$, removing C from the support. This does not affect feasibility of the LP solution since it increases the LP value only on vertices strictly inside S , which are contained in no other cycles than C' due to minimality of S . See Figure 5.1.

Applying this reduction at most $|\mathcal{L}|$ times results in a solution as desired. \square




Figure 5.1: The left picture shows a possible laminar support \mathcal{L} of a feasible LP solution, consisting of six cycles. Dashed cycles have LP value $\frac{1}{3}$, while the others have LP value $\frac{2}{3}$. The cycles b and e are redundant because they have a side containing only one other cycle; also a is redundant because it is \mathcal{L} -homotopic to f .

In the proof of Lemma 5.6 we would pick the interiors of b and e as S in the first and second step, increasing the LP value of d and f and removing b and e from the support. This yields a support as in the right image. The cycle a is still redundant, however in the laminar family given by its connected component it is one-sided and therefore not redundant. Thus, the solution is structured.

Next, we analyze the ratio $\frac{x(\bigcup_{C \in \mathcal{F}^*} \mathcal{N}_{\mathcal{L}_x}(C))}{|\mathcal{F}^*|}$ that we can achieve. In this section we only consider the simple case where \mathcal{F}^* consists of a single one-sided cycle, as in Section 4. This will already allow us to improve on Theorem 4.3 by using the following Structure Lemma. The proof of this Structure Lemma can be found in Section 5.2.

Lemma 5.7. Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Assume that the point ∞ lies in a one-sided side of a cycle in \mathcal{L} and that \mathcal{L} contains no redundant cycles. Let \mathcal{L}_1 be the set of one-sided cycles in \mathcal{L} . Then there is a multi-subset $M^* \subseteq V(G)$ with $|M^*| \leq 3|\mathcal{L}_1|$ such that for any $C \in \mathcal{L}$ at least $|\mathcal{N}_{\mathcal{L}}^1(C) \setminus \{C\}|$ vertices in M^* are in $V(C)$.

As a consequence we get:

Lemma 5.8. *Let x be a structured solution to the vertex-disjoint cycle packing LP (1.1). Let \mathcal{L} be a connected component of the support of x and \mathcal{L}_1 the set of one-sided cycles in \mathcal{L} . Then*

$$\sum_{C \in \mathcal{L}_1} x(\mathcal{N}_{\mathcal{L}}(C) \setminus \{C\}) \leq 3|\mathcal{L}_1|$$

Proof. By the Structure Lemma 5.7, choose $M^* \subseteq V(G)$ with $|M^*| \leq 3|\mathcal{L}_1|$ such that for any $C \in \mathcal{L}$ we have that M^* contains at least $|\mathcal{N}_{\mathcal{L}}^1(C) \setminus \{C\}|$ vertices in $V(C)$. We get

$$\begin{aligned} & \sum_{C \in \mathcal{L}_1} x(\mathcal{N}_{\mathcal{L}}(C) \setminus \{C\}) \\ &= \sum_{C \in \mathcal{L}} x_C \cdot |\mathcal{N}_{\mathcal{L}}^1(C) \setminus \{C\}| \\ &\leq \sum_{C \in \mathcal{L}} x_C \cdot |\{v \in M^* : v \in V(C)\}| \\ &\leq \sum_{v \in M^*} \sum_{C \in \mathcal{L} : v \in C} x_C \\ &\leq |M^*| \\ &\leq 3|\mathcal{L}_1| \end{aligned} \quad \square$$

Since the LP value of each single cycle itself is bounded by 1 this immediately yields an upper bound of 4 for the laminar cycle packing integrality gap:

Theorem 5.9. *Let x be a feasible solution to the LP (1.1) for a laminar family \mathcal{C} of cycles in a planar graph G . Then we can compute a set $\mathcal{S} \subseteq \mathcal{C}$ of pairwise vertex-disjoint cycles with $x(\mathcal{C}) \leq 4|\mathcal{S}|$ in polynomial time.*

Proof. We apply the rounding procedure from Algorithm 3. By applying Lemma 5.6 in each iteration we can assume x to be structured. Let $\mathcal{L}_x := \{C \in \mathcal{C} : x_C > 0\}$ be the (laminar) support of x . We proceed on each connected component of \mathcal{L}_x individually, so we may assume \mathcal{L}_x to be connected.

Let $\mathcal{L}_1 \subseteq \mathcal{L}_x$ be the set of one-sided cycles in the support. In each step of our greedy rounding algorithm we add a one-sided cycle $F^* \in \mathcal{L}_x$ to our solution and set x on all cycles containing a vertex of F^* to 0, removing them from the support of x .

Lemma 5.8 implies

$$\sum_{C \in \mathcal{L}_1} x(\mathcal{N}_{\mathcal{L}_x}(C)) \leq 3|\mathcal{L}_1| + \sum_{C \in \mathcal{L}_1} x_C \leq 4|\mathcal{L}_1|$$

So there exists a one-sided cycle F^* where removing $\mathcal{N}_{\mathcal{L}_x}(F^*)$ decreases x by at most 4.

After the first iteration we again apply Lemma 5.6 and split the support of x into connected components. Iterating this procedure until $x = 0$ yields a solution as desired.

In each step we can find F^* by picking a one-sided cycle in \mathcal{L}_x minimizing $x(\mathcal{N}_{\mathcal{L}_x}(F^*))$, so our algorithm works in polynomial time. \square

The greedy rounding algorithm described above also allows for adding several cycles at once to the solution. We will exploit this in Section 5.3 to decrease the upper bound for the laminar cycle packing integrality gap to below 3.5.

5.2 Proof of the Structure Lemma

Before we prove the Structure Lemma 5.7 for general laminar cycle families we first briefly consider the case where no two-sided cycles exist. In this case we start by constructing another planar graph G' on vertex set $\mathcal{L}_1 = \mathcal{L}$ as follows:

For any vertex $v \in V(G)$ let $\mathcal{L}_v \subseteq \mathcal{L}$ be the set of cycles containing v . Since all cycles are one-sided, there is a natural cyclic order $\mathcal{L}_v = \{C_1 =: C_{k+1}, C_2, \dots, C_k\}$ on \mathcal{L}_v . Then we add for any $i = 1, \dots, k$ the edge $\{C_i, C_{i+1}\}$ with its obvious planar embedding to G' . Finally, we identify homotopic edges in G' (i.e. parallel edges bounding an area homeomorphic to the disk). See Figure 5.2.

Now from G' we can construct our multi-set M^* : For each $e = \{C_1, C_2\} \in E(G')$ we add an arbitrary vertex in $V(C_1) \cap V(C_2)$ to M^* ; furthermore, for each vertex $v \in V(G)$ that is contained in $k > 3$ cycles we add $k - 3$ copies of v to M^* . Since in this case v lies inside a face of G' with exactly k edges on its boundary we can construct another planar graph G^* from G' by triangulating each such face F with $k - 3$ edges inside F (cf. Figure 5.2).

This yields a planar graph G^* on vertex set \mathcal{L}_1 with $|M^*|$ edges and no homotopic edges. According to Lemma 2.21 Euler's formula implies $|M^*| = |E(G^*)| \leq 3|V(G^*)| - 6 = 3|\mathcal{L}_1| - 6$.

Let now $C \in \mathcal{L}$ and $\mathcal{B} \subseteq \mathcal{N}_{\mathcal{L}}^1(C) \setminus \{C\}$ be the set of all neighbours of C that are not connected to C in G' . By construction of G' this means that for any vertex $v \in V(C)$ that is contained in k cycles at most $k - 3$ of them can be in \mathcal{B} . But we added $k - 3$ copies of v to M^* . This proves that $V(C)$ contains at least $|\mathcal{B}| + |\delta_{G'}(C)| \geq |\mathcal{N}_{\mathcal{L}}^1(C) \setminus \{C\}|$ vertices of M^* .




Figure 5.2: Example for the case that no two-sided cycles exist: The coloured cycles are the elements of \mathcal{L}_1 . The vertices of G' are drawn as nodes inside the one-sided sides. The edges of G' are drawn as thick dashed lines. Since five cycles meet in the vertex v we would add v twice to M^* , in addition to the vertices in M^* corresponding to edges of G' . This is possible while keeping $|M^*| \leq 3|\mathcal{L}_1| - 6$ because we can triangulate the face of G' that v lies in with two additional edges; as indicated by the dotted lines.

Next, we have to consider also two-sided cycles. However, we do not know how to extend the relatively easy construction of G' and G^* to this more general case. Instead, we will use the notion of *incidences*:

Definition 5.10. Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Assume that the point ∞ lies in a one-sided side of a cycle in \mathcal{L} . Let $\mathcal{L}_1 \subseteq \mathcal{L}$ be the set

of one-sided cycles. A *neighbour pair* is a pair $(\{C, N\}, v)$ of a set of two cycles $C, N \in \mathcal{L}$ that are not \mathcal{L} -homotopic and a vertex $v \in V(C) \cap V(N)$. We call two neighbour pairs $(\{C, N\}, v)$ and $(\{C, N\}, v')$ *homotopic* if there exist v - v' -paths P in C and P' in N such that $P + P'$ bounds an area that contains all one-sided sides in \mathcal{L} .

It is easy to see that homotopy defines an equivalence relation on neighbour pairs between two cycles $C, N \in \mathcal{L}$. An equivalence class of such neighbour pairs is called an *incidence* between C and N (cf. Figure 5.3). The *vertex set* $V(I)$ of an incidence I between C and N is the set of all v with $(\{C, N\}, v) \in I$. We also denote I by $I = (\{C, N\}, V(I))$. For a cycle $C \in \mathcal{L}$ let $\mathcal{I}_{\mathcal{L}}^1(C)$ be the set of all incidences between C and one-sided cycles in $\mathcal{N}_{\mathcal{L}}^1(C)$.

Let now I be an incidence between $C \in \mathcal{L}$ and $N \in \mathcal{N}_{\mathcal{L}}(C)$. Let S_C be a side of C and S_N a side of N such that S_C and S_N are disjoint. We call an incidence $I' = (\{C', N'\}, V(I'))$ a *sub-incidence* of I if C' is inside S_C , N' is inside S_N and $V(I') \subseteq V(I)$. We call I *minimal* if any sub-incidence of I is equal to I . We call I *crossing* if $V(I) = \{v\}$ for some $v \in V(G)$ and there exist cycles $C_1, C_2 \in \mathcal{L}$ that also contain v with sides S_1 and S_2 such that S_C, S_1, S_N, S_2 are all disjoint and are ordered in this way around v . Such incidences are also called *v -incidences*. If I is not crossing we call it *non-crossing* (cf. Figure 5.3).

Extending the idea of including the edges of G' in M^* , in order to prove Lemma 5.7 we will construct a set of incidences instead of a set of vertices.

Definition 5.11. Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Let $\mathcal{L}_1 \subseteq \mathcal{L}$ be the set of one-sided cycles. Let M be a multi-set of incidences in \mathcal{L} . We say that an element $I \in M$ *hits* a cycle $C \in \mathcal{L}$ if $V(I) \subseteq V(C)$. We call a cycle $C \in \mathcal{L}$ *M -good* if at least $|\mathcal{I}_{\mathcal{L}}^1(C)|$ elements of M hit C . We call M *good* if all cycles in \mathcal{L} are M -good and $|M| \leq 3|\mathcal{L}_1| - 6$. Furthermore, we call M *structured* if the following properties hold:

1. M contains every non-crossing incidence between one-sided cycles.
2. For each $C \in \mathcal{L}_1$ and $v \in V(C)$ there exist at least as many v -incidences in M as there are v -incidences between C and $\mathcal{N}_{\mathcal{L}}^1(C)$ in \mathcal{L} .

This notion of structured incidence sets is inspired from the construction of M^* in the case $\mathcal{L} = \mathcal{L}_1$: The edges in G' correspond to non-crossing incidences between one-sided cycles, which are included in M by property 1. Property 2 makes sure that vertices in which many one-sided cycles meet are included in M . In particular, we get the following as a direct consequence of the above definition:

Lemma 5.12. *Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Let M be a structured multi-set of incidences in \mathcal{L} . Then every one-sided cycle is M -good.*

Lemma 5.13. *Let \mathcal{L} be a laminar family of at least two cycles in a planar graph G , embedded in the sphere. Let \mathcal{L}_1 be the set of one-sided cycles in \mathcal{L} . Then there exists a good and structured multi-set M of incidences in \mathcal{L} .*

Proof. We use an induction on $|\mathcal{L}|$. We start with the case $|\mathcal{L}_1| = 2$, which is trivial because all cycles are homotopic and there are no incidences. In particular, $M := \emptyset$ yields a good and structured set.

Let us now assume $|\mathcal{L}| = |\mathcal{L}_1| = 3$. In this case there is at most one incidence between any two cycles, thus let M be the set of all incidences, which is a good and structured set.




Figure 5.3: An example of a part of the embedding of \mathcal{L} . One-sided cycles are filled, two-sided cycles are drawn dashed. The cycle C^* has exactly one non-crossing incidence to each of the five one-sided cycles; the neighbour pairs $(\{C^*, C_6\}, v)$ and $(\{C^*, C_6\}, w)$ are not homotopic and therefore yield two incidences: $(\{C^*, C_6\}, \{v\})$ is a v -incidence, while $(\{C^*, C_6\}, \{w\})$ is non-crossing.

The lower half of the picture corresponds to the side S_1 which contains no two-sided cycles; the arrows on C^* indicate the orientation of C^*_{\rightarrow} . For each one-sided cycle C_i , $1 \leq i \leq 5$, the dotted arrow starting at C_i represents the element $f_j(I) \in M'_j$ that replaces the non-crossing incidence I between C_i and C^* in M_j , where S_j is the side containing C_i . For example, $f_1(\{C^*, C_4\}, \{v, w\}) = (\{C_3, C_4\}, \{w\})$ because C_3 is the “first” cycle in S_2 (w.r.t. C^*_{\rightarrow}) touching C_4 , while $f_2(\{C^*, C_2\}, \{v, w\}) = (\{C_4, C_2\}, \{v, w\})$ because $(\{C_5, C_2\}, \{w\})$ is crossing.

For proving that C_6 is M -good we need to consider consecutive incidences in A w.r.t. the ordering inherited by C^*_{\rightarrow} with the same image under g , like $a_1 = (\{C_6, C_5\}, \{w\})$ and $a_2 = (\{C_6, C_4\}, \{w\})$. However in this case $f_1(\{C^*, C_4\}, \{v, w\})$ must hit C_6 .

Let now $|\mathcal{L}| > 3$. We first assume $\mathcal{L} \neq \mathcal{L}_1$; we will show how to deal with the other case at the end. Let $C^* \in \mathcal{L}$ be two-sided with a side S_1 that contains no other two-sided cycles. Let S_2 be the other side of C^* . For $i = 1, 2$ let $\mathcal{L}_{S_i} \subseteq \mathcal{L}$ be the set of all cycles with a side inside S_i . In particular, $C^* \in \mathcal{L}_{S_1} \cap \mathcal{L}_{S_2}$ is one-sided in both families. Since both families are strictly smaller than \mathcal{L} , by induction hypothesis we can find good and structured sets M_i for \mathcal{L}_{S_i} for $i = 1, 2$.

Let C^*_{\rightarrow} be an arbitrary (Eulerian) orientation of C^* . Assume $I \in \mathcal{I}_{\mathcal{L}_{S_i}}^1(C^*)$ to be a non-crossing incidence between C^* and a cycle $N \in \mathcal{L}_1 \cap \mathcal{L}_{S_i}$ for some $i \in \{1, 2\}$. Let \mathcal{I}_I be the set of minimal sub-incidences of I in \mathcal{L} between N and cycles in $\mathcal{L}_{S_{3-i}}$. The orientation of C^*_{\rightarrow} induces a natural linear order on \mathcal{I}_I . Let $f_1(I)$ be a first element in this order. Let $f_2(I)$ be a first non-crossing incidence in this order if it exists and $f_2(I) = f_1(I)$ otherwise (cf. Figure 5.3). Let M'_i arise from M_i by replacing each non-crossing incidence $I \in \mathcal{I}_{\mathcal{L}_{S_i}}(C^*)$ by $f_i(I)$. Since $V(f_i(I)) \subseteq V(I)$ holds for any I , M'_i is also good for \mathcal{L}_{S_i} .

Let $M := M'_1 + M'_2$. We first show that M is structured. To show property 1 of Definition 5.11, let $C_1, C_2 \in \mathcal{L}_1$ and $I = (\{C_1, C_2\}, V(I))$ be a non-crossing incidence. If C_1 and C_2 are on the same side S_i of C^* then $I \in M_i \setminus \mathcal{I}_{\mathcal{L}_{S_i}}^1(C^*) \subseteq M$. Otherwise, w.l.o.g. C_i is inside S_i for $i = 1, 2$. First assume that there is an $i \in \{1, 2\}$ such that C_i is the only one-sided

cycle in S_i . Then S_{3-i} contains at least two one-sided cycles because otherwise all cycles are \mathcal{L} -homotopic. Thus, there is a non-crossing incidence $I_{3-i} \in \mathcal{I}_{\mathcal{L}_{S_{3-i}}}^1(C^*)$ such that I is a sub-incidence of I_{3-i} . Also, the fact that \mathcal{L}_{S_i} is a chain implies $|\mathcal{I}_{I_{3-i}}| = 1$ and therefore $f_{3-i}(I_{3-i}) = I \in M$.

Now we assume that both S_i contain at least two one-sided cycles. Again, there exist non-crossing incidences $I_i \in \mathcal{I}_{\mathcal{L}_{S_i}}^1(C^*)$ such that I is a sub-incidence of I_i for each $i = 1, 2$. Since M_i is structured, $I_i \in M_i$. Let $C'_1 \in \mathcal{L}_{S_1}$ and $C'_2 \in \mathcal{L}_{S_2}$ such that $f_i(I_i)$ is an incidence between C_i and C'_{3-i} for $i = 1, 2$. If neither $C_1 = C'_1$ nor $C_2 = C'_2$ holds then all four of those cycles must meet in one vertex due to minimality of $f_i(I_i)$ in the order on \mathcal{I}_{I_i} (cf. Figure 5.4). But then both $f_1(I_1)$ and $f_2(I_2)$ are crossing, contradicting the definition of f_2 . Thus, we get $f_1(I_1) = I$ or $f_2(I_2) = I$.

Next, we show property 2 of Definition 5.11. Let $C \in \mathcal{L}_1 \cap \mathcal{L}_{S_i}$ for some $i \in \{1, 2\}$ and $v \in V(C)$. If all one-sided cycles that contain v are in the same \mathcal{L}_{S_i} then M_i already contains enough v -incidences. So we only have to consider the case $v \in V(C^*)$ such that both $\mathcal{L}_{S_1} \cap \mathcal{L}_1$ and $\mathcal{L}_{S_2} \cap \mathcal{L}_1$ contain a cycle that contains v . Let $\mathcal{N} \subseteq \mathcal{L}_1$ be the set of one-sided cycles with a v -incidence to C .

Case 1: There is a $j \in \{1, 2\}$ such that the cycles in \mathcal{L}_{S_j} that contain v form a chain. In this case there is a natural correspondence between v -incidences between one-sided cycles in $\mathcal{L}_{S_{3-j}}$ and v -incidences between one-sided cycles in \mathcal{L} (cf. Figure 5.4). Thus, M_{3-j} already contains $|\mathcal{N}|$ v -incidences.

Case 2: For $j = 1, 2$ the cycles in \mathcal{L}_{S_j} that contain v do not form a chain. Since S_1 contains no two-sided cycles except C^* , this implies that by definition of f_1 there is a non-crossing incidence I of C^* in M_1 such that $f_1(I)$ is a v -incidence, as shown in Figure 5.4. Clearly, M_i contains at least $|\mathcal{N} \cap \mathcal{L}_{S_i}|$ v -incidences. Also, applying property 2 of Definition 5.11 for C^* and v in $\mathcal{L}_{S_{3-i}}$ yields that the number of v -incidences in M_{3-i} is at least $|\mathcal{N} \cap \mathcal{L}_{S_{3-i}}| - 2$ because C^* can have at most two non-crossing incidences to cycles in \mathcal{N} at v . However, if the second lower bound is tight then also the incidence between C^* and C in \mathcal{L}_{S_i} is crossing (see Figure 5.4) and M_1 contains even $|\mathcal{N} \cap \mathcal{L}_{S_i}| + 1$ v -incidences. So in total M contains at least $|\mathcal{N}|$ v -incidences, concluding the proof that M is structured.

It is left to show that M is good. First, we have

$$|M| \leq 3(|\mathcal{L}_1 \cap \mathcal{L}_{S_1}| + 1) - 6 + 3(|\mathcal{L}_1 \cap \mathcal{L}_{S_2}| + 1) - 6 \leq 3|\mathcal{L}_1| - 6$$

Also all one-sided cycles are M -good since M is structured. Let now $C \in \mathcal{L}$ be two-sided. By choice of C^* , $C \in \mathcal{L}_{S_2}$. Let A be the set of incidences between C and one-sided cycles inside S_1 in \mathcal{L} , and let B be the set of incidences between C and C^* in the family \mathcal{L}_{S_2} . Note that $B = \emptyset$ if $C = C^*$ or C is homotopic to C^* . By definition we get $|\mathcal{I}_{\mathcal{L}}(C)| - |\mathcal{I}_{\mathcal{L}_{S_2}}(C)| = |A| - |B|$. Also, A is naturally ordered cyclically by C^* .

We first assume $B \neq \emptyset$. Then there is a natural map $g: A \rightarrow B$ such that each $a \in A$ is a sub-incidence of $g(a)$; furthermore, for each $b \in B$ the cyclic order on A induces a linear order on $g^{-1}(\{b\})$. Let now $b \in B$ and $a_1, a_2 \in A$ be consecutive elements in the linear order on $g^{-1}(\{b\})$. Then there is an incidence $I_2 \in \mathcal{I}_{\mathcal{L}_{S_1}}(C^*)$ such that a_2 is a sub-incidence of I_2 . Then by the definition of f_1 we have that $f_1(I_2)$ hits C (cf. Figure 5.3).




Figure 5.4: Three examples for the proof that M is structured: The left image shows a situation with a non-crossing incidence I between C_1 and C_2 . If for each $i = 1, 2$ the incidence between C_i and C^* is mapped by f_i to an incidence between C_i and $C'_{3-i} \neq C_{3-i}$ then each C'_i must come “before” C_i w.r.t. C^* . So all cycles meet in a point v , contradicting the choice of f_2 .

The other two images show the different situations for property 2: The middle image shows case 1 where all cycles on one side (in this case S_2) form a chain. Then it is easy to see that M_1 already contains enough v -incidences.

The right image shows the remaining case 2. Let us assume that C is inside S_2 . M_2 already contains $|\mathcal{N} \cap \mathcal{L}_{S_2}|$ v -incidences and M_1 contains at least $|\mathcal{N} \cap \mathcal{L}_{S_1}| - 2$ v -incidences because C_1 and C_3 are the only cycles that might be in \mathcal{N} without a v -incidence to C^* . However, if both C_1 and C_3 are in \mathcal{N} then C cannot be among the non-crossing neighbours (C_4 and C_7) of C^* at v . Thus, M_2 contains an additional v -incidence because $(\{C^*, C\}, \{v\})$ is crossing.

In the other case $B = \emptyset$ one can see similarly that for any consecutive elements $a_1, a_2 \in A$ in the cyclic order on A there is an incidence $I_2 \in \mathcal{I}_{\mathcal{L}_{S_1}}(C^*)$ that a_2 is a sub-incidence of such that $f_1(I_2)$ hits C . Thus, in both cases the number of incidences in M'_1 that hit C is at least $|A| - |B|$ and

$$\begin{aligned}
 |\{I \in M : V(I) \subseteq V(C)\}| &\geq |\{I \in M'_1 : V(I) \subseteq V(C)\}| + |\{I \in M'_2 : V(I) \subseteq V(C)\}| \\
 &\geq |\{I \in M'_1 : V(I) \subseteq V(C)\}| + |\mathcal{I}_{\mathcal{L}_{S_2}}(C)| \\
 &= |\{I \in M'_1 : V(I) \subseteq V(C)\}| + |\mathcal{I}_{\mathcal{L}}(C)| - (|A| - |B|) \\
 &\geq |\mathcal{I}_{\mathcal{L}}(C)|
 \end{aligned}$$

We are left with the case that $|\mathcal{L}| > 3$, but $\mathcal{L} = \mathcal{L}_1$. In this case, however, we can just add a new two-sided cycle C^* with at least two one-sided sides on each side both to G and to \mathcal{L} . This only strengthens the statement we prove because any good and structured set for the constructed instance is also good and structured for the original instance. Also, in the induction step we can still apply the induction hypothesis to \mathcal{L}_{S_1} and \mathcal{L}_{S_2} because both S_1 and S_2 contain at least two one-sided cycles. \square

As a direct consequence of this lemma we get the Structure Lemma 5.7, which we restate here:

Lemma 5.7. *Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Assume that the point ∞ lies in a one-sided side of a cycle in \mathcal{L} and that \mathcal{L} contains no redundant cycles. Let \mathcal{L}_1 be the set of one-sided cycles in \mathcal{L} . Then there is a multi-subset*

$M^* \subseteq V(G)$ with $|M^*| \leq 3|\mathcal{L}_1|$ such that for any $C \in \mathcal{L}$ at least $|\mathcal{N}_{\mathcal{L}}^1(C) \setminus \{C\}|$ vertices in M^* are in $V(C)$.

Proof. W.l.o.g. $|\mathcal{L}| \geq 2$. By Lemma 5.13 let M be a good multi-set of incidences in \mathcal{L} . Note that Lemma 5.13 allows us to choose M also structured, but we do not need this property here. Let M^* arise from adding one element of $V(I)$ for every $I \in M$. In particular, $|M^*| = |M| \leq 3|\mathcal{L}_1| - 6$.

Let $C \in \mathcal{L}$. Since C is not redundant, it is not \mathcal{L} -homotopic to any cycle in \mathcal{L}_1 (except possibly itself) and therefore $|\mathcal{I}_{\mathcal{L}}^1(C)| \geq |\mathcal{N}_{\mathcal{L}}^1(C) \setminus \{C\}|$. But the number of vertices in M^* that hit $V(C)$ is at least

$$|\{I \in M : V(I) \subseteq V(C)\}| \geq |\mathcal{I}_{\mathcal{L}}^1(C)|.$$

Together this concludes the proof. \square

5.3 Improving the bounds below 3.5

In this section we improve the bound on the laminar cycle packing integrality gap to $\frac{20+\sqrt{130}}{9} < 3.5$. We still use Algorithm 3 from Section 5.1 for the improved approximation guarantee, but add further possibilities for the set \mathcal{F}^* of cycles that are added to our solution during a single iteration. Note that the one-cardinality set chosen in Theorem 5.9 already gives a good approximation guarantee if the average LP value on one-sided cycles is small. On the other hand, if the average LP value on one-sided cycles is large then we will find a large set of pairwise vertex-disjoint cycles with relatively small neighbourhood which we can take as \mathcal{F}^* . For analyzing the case of \mathcal{F}^* containing more than one cycle we use the following Lemma:

Lemma 5.14. *Let \mathcal{L} be a laminar family of cycles in a planar graph G , embedded in the sphere. Let \mathcal{L}_1 be the set of one-sided cycles in \mathcal{L} . Let $\mathcal{F} \subseteq \mathcal{L}_1$. Then there is a set $M \subseteq V(\mathcal{F})$ with $|M| \leq |\mathcal{F}| + |\mathcal{L}_1|$ such that each cycle in $\bigcup_{C \in \mathcal{F}} \mathcal{N}_{\mathcal{L}}(C)$ contains a vertex from M .*

Proof. W.l.o.g. $|\mathcal{L}| > 1$. Let $\mathcal{B}_{\text{int}} \subseteq \mathcal{L}$ be the set of cycles C such that there is a cycle $C' \in \mathcal{F}$ with $C' \subseteq_{\infty} C$ and $V(C) \cap V(C') \neq \emptyset$. In particular, $\mathcal{F} \subseteq \mathcal{B}_{\text{int}}$. Let $f: \mathcal{B}_{\text{int}} \rightarrow \mathcal{F}$ such that each $C \in \mathcal{B}_{\text{int}}$ contains $f(C)$ and shares a vertex with $f(C)$. Then for any $C \in \mathcal{F}$ all cycles in $f^{-1}(C)$ must form a chain and therefore meet in some vertex $v_C \in V(C)$. Thus, $M_{\text{int}} := \{v_C : C \in \mathcal{F}\}$ hits all cycles in \mathcal{B}_{int} .

Let now $\mathcal{B}_{\text{ext}} \subseteq \mathcal{L} \setminus \mathcal{B}_{\text{int}}$ be the set of all cycles in $\mathcal{L} \setminus \mathcal{B}_{\text{int}}$ that share a vertex with any cycle in \mathcal{F} . We show by induction on $|\mathcal{L}_1|$ that we can hit all cycles in \mathcal{B}_{ext} with some $M_{\text{ext}} \subseteq V(\mathcal{F})$ with $|M_{\text{ext}}| \leq |\mathcal{L}_1|$: For $|\mathcal{L}_1| = 0$ this is trivial. Otherwise, let $C_1 \in \mathcal{B}_{\text{ext}}$ with minimal interior. Choose $C_2 \in \mathcal{F}$ with some vertex $v \in V(C_1) \cap V(C_2)$ (cf. Figure 5.5). Construct another laminar family \mathcal{L}' by deleting all cycles inside C_1 and all cycles in $\mathcal{L} \setminus \mathcal{F}$ that contain v . Since C_1 contains some one-sided cycle, \mathcal{L}' contains strictly fewer one-sided cycles and we can use the induction hypothesis on \mathcal{L}' . Also the deletion of cycles inside C_1 does (except for C_1 itself) not change \mathcal{B}_{ext} because $C_1 \notin \mathcal{B}_{\text{int}}$ and C_1 was minimal. Thus, the induction hypothesis gives us a set $M'_{\text{ext}} \subseteq V(\mathcal{F})$ that hits all cycles in $\mathcal{B}_{\text{ext}} \cap \mathcal{L}'$ with $|M'_{\text{ext}}| \leq |\mathcal{L}_1| - 1$, so $M_{\text{ext}} := M'_{\text{ext}} \cup \{v\}$ has the desired properties.

Setting $M := M_{\text{int}} \cup M_{\text{ext}}$ yields a set as desired in the Lemma. \square

In the following, we assume $x \in \mathbb{R}^{\mathcal{L}_x}$ to be a structured solution to the LP (1.1) with support \mathcal{L}_x , which is a laminar family of cycles in a planar graph G , embedded in the sphere.




Figure 5.5: The cycles in \mathcal{F} are drawn in green, the cycles in \mathcal{B}_{ext} in blue. Note that the orange cycle is not in \mathcal{B}_{ext} because it is in \mathcal{B}_{int} . The inside of C_1 contains no other cycles in \mathcal{B}_{ext} and C_1 meets $C_2 \in \mathcal{F}$ in v . We then add v to our set M_{ext} and recurse on the laminar family \mathcal{L}' that is constructed by removing all cycles inside C_1 and all cycles that contain v and are not in \mathcal{F} . These cycles are drawn dashed. This step decreases the number of one-sided cycles.

As in Theorem 5.9 we can assume \mathcal{L}_x to be connected. Let $\mathcal{L}_1 \subseteq \mathcal{L}_x$ be the set of one-sided cycles in \mathcal{L}_x . For each $0 \leq \alpha < 1$ we define $\mathcal{L}_1^{>\alpha} \subseteq \mathcal{L}_1$ to be the set of one-sided cycles with LP value $> \alpha$ and set $r_\alpha := \frac{|\mathcal{L}_1^{>\alpha}|}{|\mathcal{L}_1|}$.

We will now give three possible choices for \mathcal{F}^* in Algorithm 3. The first possibility is to choose a one-cardinality subset of \mathcal{L}_1 as \mathcal{F}^* , as in Theorem 5.9. By Lemma 5.8 we directly get:

Lemma 5.15. *There exists a cycle $F^* \in \mathcal{L}_1$ with $x(\mathcal{N}_{\mathcal{L}_x}(F^*)) \leq 3 + \frac{x(\mathcal{L}_1)}{|\mathcal{L}_1|}$.*

As a second possibility we define $\mathcal{F}_\alpha^* := \mathcal{L}_1^{>\alpha}$ for any $\alpha \geq \frac{1}{2}$. Note that the cycles in \mathcal{F}_α^* are pairwise vertex-disjoint because x is a feasible LP solution.

Lemma 5.16. *For any $\alpha \geq \frac{1}{2}$ with $r_\alpha > 0$ we have*

$$\frac{x\left(\bigcup_{C \in \mathcal{F}_\alpha^*} \mathcal{N}_{\mathcal{L}_x}(C)\right)}{|\mathcal{F}_\alpha^*|} \leq 1 + \frac{1 - \alpha}{r_\alpha}$$

Proof. Let $\alpha \geq \frac{1}{2}$. By Lemma 5.14 there is a set $M \subseteq V(\mathcal{F}_\alpha^*)$ with $|M| \leq |\mathcal{F}_\alpha^*| + |\mathcal{L}_1|$ such that each cycle in $\bigcup_{C \in \mathcal{F}_\alpha^*} \mathcal{N}_{\mathcal{L}_x}(C)$ contains a vertex of M . Now

$$\begin{aligned} & x\left(\bigcup_{C \in \mathcal{F}_\alpha^*} \mathcal{N}_{\mathcal{L}_x}(C)\right) \\ & \leq \sum_{v \in M} x(\{C \in \mathcal{L}_x \setminus \mathcal{F}_\alpha^* : v \in V(C)\}) + x(\mathcal{F}_\alpha^*) \\ & \leq |\mathcal{F}_\alpha^*| + (1 - \alpha)|\mathcal{L}_1| \end{aligned}$$

holds, where the last inequality follows from the fact that there are $|\mathcal{F}_\alpha^*|$ vertices in M covering \mathcal{F}_α^* , and on the other vertices we only have to count the LP value of cycles not in \mathcal{F}_α^* . \square

As a third possibility we take a look at the sets $\mathcal{L}_1^{>\alpha}$ with $\frac{1}{4} \leq \alpha < \frac{1}{2}$. For these we know that at most three cycles in $\mathcal{L}_1^{>\alpha}$ can share a vertex. Let G' be the *conflict graph* for the cycles

in $\mathcal{L}_1^{>\alpha}$; i.e. G' is the graph on vertex set $\mathcal{L}_1^{>\alpha}$ such that two cycles in $\mathcal{L}_1^{>\alpha}$ are connected by an edge in G' if and only if they share a vertex in G . Since each vertex is contained in at most three cycles of $\mathcal{L}_1^{>\alpha}$, G' is planar. Furthermore, the cycles in $\mathcal{L}_1^{>1-\alpha} \subseteq \mathcal{L}_1^{>\alpha}$ correspond to isolated vertices in G' . By the Four Colour Theorem [8] we can partition $V(G') - \mathcal{L}_1^{>1-\alpha}$ into four stable sets. The largest of those, together with $\mathcal{L}_1^{>1-\alpha}$, yields a stable set in G' of size at least $|\mathcal{L}_1^{>1-\alpha}| + \frac{1}{4}(|\mathcal{L}_1^{>\alpha}| - |\mathcal{L}_1^{>1-\alpha}|)$. We let \mathcal{F}_α^* be the set of cycles in $\mathcal{L}_1^{>\alpha}$ corresponding to such a stable set in G' .

Lemma 5.17. *For any $\frac{1}{4} \leq \alpha < \frac{1}{2}$ with $r_\alpha > 0$ we have*

$$\frac{x\left(\bigcup_{C' \in \mathcal{F}_\alpha^*} \mathcal{N}_{\mathcal{L}_x}(C')\right)}{|\mathcal{F}_\alpha^*|} \leq 1 + \frac{4(1-\alpha)}{r_\alpha + 3r_{1-\alpha}}$$

Proof. Let $\frac{1}{4} \leq \alpha < \frac{1}{2}$. By a similar argument as in Lemma 5.16 we get

$$\frac{x\left(\bigcup_{C' \in \mathcal{F}_\alpha^*} \mathcal{N}_{\mathcal{L}_x}(C')\right)}{|\mathcal{F}_\alpha^*|} \leq 1 + \frac{(1-\alpha)|\mathcal{L}_1|}{|\mathcal{F}_\alpha^*|}$$

Inserting the bound $|\mathcal{F}_\alpha^*| \geq |\mathcal{L}_1^{>1-\alpha}| + \frac{1}{4}(|\mathcal{L}_1^{>\alpha}| - |\mathcal{L}_1^{>1-\alpha}|)$ yields the lemma. \square

One of these possibilities for \mathcal{F}^* will be sufficient to prove the desired upper bound of $\frac{20+\sqrt{130}}{9}$ for the integrality gap:

Lemma 5.18. *There exists a set $\mathcal{F}^* \subseteq \mathcal{L}_1$ with $\frac{x(\bigcup_{C' \in \mathcal{F}^*} \mathcal{N}_{\mathcal{L}_x}(C'))}{|\mathcal{F}^*|} \leq \frac{20+\sqrt{130}}{9}$*

Proof. Define $\beta := \frac{20+\sqrt{130}}{9}$. We will either pick one of the sets \mathcal{F}_α^* for $\frac{1}{4} \leq \alpha < 1$ from Lemma 5.16 or Lemma 5.17 or we will use $\mathcal{F}^* := \{F^*\}$ with the cycle F^* from Lemma 5.15. Assume none of these sets \mathcal{F}^* fulfills the above inequality. Then Lemma 5.16 implies

$$\begin{aligned} 1 + \frac{1-\alpha}{r_\alpha} &> \beta \\ \Leftrightarrow r_\alpha &< \frac{1-\alpha}{\beta-1} \end{aligned}$$

for all $\alpha \geq \frac{1}{2}$. For $\frac{1}{4} \leq \alpha < \frac{1}{2}$, Lemma 5.17 yields

$$\begin{aligned} 1 + \frac{4(1-\alpha)}{r_\alpha + 3r_{1-\alpha}} &> \beta \\ \Leftrightarrow r_\alpha + 3r_{1-\alpha} &< \frac{4(1-\alpha)}{\beta-1} \end{aligned}$$

Furthermore, we have

$$x(\mathcal{L}_1) = \sum_{C \in \mathcal{L}_1} \int_0^1 \mathbb{1}_{x_C > \alpha} d\alpha = \int_0^1 \sum_{C \in \mathcal{L}_1} \mathbb{1}_{x_C > \alpha} d\alpha = |\mathcal{L}_1| \int_0^1 r_\alpha d\alpha$$

Thus, Lemma 5.15 implies

$$\beta < 3 + \int_0^1 r_\alpha d\alpha \tag{5.1}$$

For any threshold $0 < \delta \leq \frac{1}{6}$ that can be chosen later we can bound this as follows, using the fact that the r_α are non-increasing:

$$\begin{aligned}
\beta &< 3 + \int_0^1 r_\alpha d\alpha \\
&= 3 + \int_0^{\frac{1}{2}-\delta} r_\alpha d\alpha + \int_{\frac{1}{2}-\delta}^{\frac{1}{2}+3\delta} r_\alpha d\alpha + \int_{\frac{1}{2}+3\delta}^1 r_\alpha d\alpha \\
&\leq 3 + \int_0^{\frac{1}{2}-\delta} r_\alpha d\alpha + \int_{\frac{1}{2}-\delta}^{\frac{1}{2}} r_\alpha d\alpha + 3 \int_{\frac{1}{2}}^{\frac{1}{2}+\delta} r_\alpha d\alpha + \int_{\frac{1}{2}+3\delta}^1 r_\alpha d\alpha \\
&= 3 + \int_0^{\frac{1}{2}-\delta} r_\alpha d\alpha + \int_{\frac{1}{2}-\delta}^{\frac{1}{2}} r_\alpha + 3r_{1-\alpha} d\alpha + \int_{\frac{1}{2}+3\delta}^1 r_\alpha d\alpha \\
&\leq 3 + \int_0^{\frac{1}{2}-\delta} 1 d\alpha + \int_{\frac{1}{2}-\delta}^{\frac{1}{2}} \frac{4(1-\alpha)}{\beta-1} d\alpha + \int_{\frac{1}{2}+3\delta}^1 \frac{1-\alpha}{\beta-1} d\alpha \\
&= \frac{7}{2} - \delta + \frac{2\delta(\delta+1)}{\beta-1} + \frac{(1-6\delta)^2}{8(\beta-1)}
\end{aligned}$$

This term attains its minimum of $\frac{\beta^2-94\beta+90}{26(1-\beta)}$ at $\delta := \frac{2\beta-3}{26} < \frac{1}{6}$, so we get

$$\beta < \frac{\beta^2 - 94\beta + 90}{26(1 - \beta)}$$

This is a contradiction for $\beta = \frac{20+\sqrt{130}}{9}$. \square

As an immediate consequence we can prove Theorem 5.1, improving on the easy bound from Theorem 5.9. We restate Theorem 5.1 here for convenience:

Theorem 5.1. *The laminar cycle packing integrality gap is at most $\frac{20+\sqrt{130}}{9} < 3.5$, i.e. for any fractional solution to the LP (1.1) for a laminar cycle family \mathcal{C} in a planar graph there exists an integral solution with at least $\frac{9}{20+\sqrt{130}}$ times the LP value.*

Proof. Let x be a solution to the LP (1.1) with laminar support \mathcal{L}_x . As in Theorem 5.9 we apply Algorithm 3. In contrast to the procedure in Theorem 5.9 however we use a set \mathcal{F}^* as guaranteed by Lemma 5.18 in step 6 of the algorithm instead of a single one-sided cycle. Thus, in each step we increase the number of cycles with LP value 1 by $|\mathcal{F}^*|$ while decreasing the LP value on \mathcal{L}_x by at most $\frac{20+\sqrt{130}}{9}|\mathcal{F}^*|$. Therefore we arrive at an integral solution to the LP with at least $\frac{9}{20+\sqrt{130}}$ the LP value. \square

Note that a set \mathcal{F}^* as in Lemma 5.18 can be found in polynomial time in $|\mathcal{L}_x|$: For the cycle guaranteed by Lemma 5.15 we can try all one-sided cycles. For the sets \mathcal{F}_α^* used in Lemma 5.16, note that there are only linearly many different sets $\mathcal{L}_1^{>\alpha}$ to consider. The sets \mathcal{F}_α^* used in Lemma 5.17 are constructed from $\mathcal{L}_1^{>\alpha}$ by applying the Four Colour Theorem, which can also be done in polynomial time [75]. Therefore, Algorithm 3 can be carried out in polynomial time if \mathcal{C} is given by a weight oracle. Furthermore, Proposition 2.36 allows us to extend our vertex-disjoint packing results also to edge-disjoint packing. We note:

Theorem 5.19. *Given a planar graph G and a weight oracle for an uncrossable family \mathcal{C} of cycles in G , we can find a*

(a) *vertex-disjoint subset of \mathcal{C} whose cardinality is $\frac{9}{20+\sqrt{130}}$ times the value of the LP (1.1)*

(b) *edge-disjoint subset of \mathcal{C} whose cardinality is $\frac{9}{20+\sqrt{130}}$ times the value of the LP (1.3)*

in polynomial time.

Chapter 6

A constant Erdős–Pósa ratio in bounded-genus graphs

In this chapter we present a constant upper bound on the Erdős–Pósa ratio of several cycle families in graphs that can be embedded in an orientable surface of genus g . As in Corollary 5.2 our bound results from multiplying the integrality gaps of the cycle packing LP and the cycle transversal LP. For the former we gave an upper bound of $O(g^2)$ for uncrossable cycle families in Section 4.3. So in this section we focus on the cycle transversal LP. Actually, our methods will all work for the weighted version of the problem, see the LPs (1.5) and (1.6). Since the WEIGHTED EDGE CYCLE TRANSVERSAL PROBLEM can be reduced to the WEIGHTED VERTEX CYCLE TRANSVERSAL PROBLEM by Proposition 2.34 and the reduction even preserves bounds on the integrality gap we only consider the vertex version (1.5) of our problem.

Recently, Sun [83] gave an (LP-based) $O(g)$ -approximation for the WEIGHTED VERTEX CYCLE TRANSVERSAL PROBLEM for the family $\mathcal{C} := \overrightarrow{\mathcal{C}}_{\text{all}}$ of all directed cycles in an orientation of a graph G , embedded in an orientable surface of genus g . Their approach works in two steps:

First, they show how to eliminate all *facial* cycles of \mathcal{C} , i.e., cycles that bound an area homeomorphic to the disk. To this end, they use the primal-dual framework by Goemans and Williamson [40] in a similar way as Goemans and Williamson [41] as well as Berman and Yaroslavtsev [16] did for the planar CYCLE TRANSVERSAL PROBLEM. Sun implicitly uses that the facial cycles of \mathcal{C} are uncrossable, which is actually false (cf. Figure 6.1). However, we show that this step can still be carried out for all uncrossable cycle families by a more careful “uncrossing” argument. The details are given in Section 6.2.

The second step of [83] is to iteratively eliminate a fraction of cycles such that the remaining graph can be embedded in a surface of smaller genus. We will use similar techniques to bound the integrality gap of the weighted cycle transversal LP (1.5) for many other uncrossable cycle families by $O(g)$. However, this result does not work for all uncrossable families. We capture the properties we need in the notion of *subgraph-defined* uncrossable cycle families (cf. Definition 6.16).

Theorem 6.1. *Given a graph G with vertex weights $w: V(G) \rightarrow \mathbb{R}_{\geq 0}$, embedded in an orientable surface of genus g , and a subgraph-defined uncrossable family \mathcal{C} of cycles in G , then the integrality gap of the LP (1.5) is in $O(g)$.*

Throughout this chapter we will denote the optimum value of the LP (1.5) for some given

cycle family \mathcal{C} by $\text{OPT}_{\text{LP}}(\mathcal{C})$. As already noted before, the main application of Theorem 6.1 is to combine it with Theorem 4.11 to yield an upper bound for the Erdős–Pósa ratio:

Corollary 6.2. *Given a graph G , embedded in an orientable surface of genus g , and a subgraph-defined uncrossable family \mathcal{C} of cycles in G , then there exists a cycle transversal $T \subseteq V(G)$ and a set $\mathcal{S} \subseteq \mathcal{C}$ of pairwise vertex-disjoint cycles such that $|T| \leq O(g^3)|\mathcal{S}|$.*

We also note that the methods used in Theorem 6.1 are constructive and a corresponding cycle transversal can be found in polynomial time.

This chapter is organized as follows. First, we introduce some useful tools from topology as well as some topological notation in Section 6.1. Section 6.2 presents our results on facial cycles and in Section 6.3 we prove Theorem 6.1. Note that while the main result of [83] builds on methods to remove facial cycles, this is not the case for our main result. In particular, Sections 6.2 and 6.3 are independent.

6.1 Topological facts we need

For our results on bounded-genus graphs we need a few results from topology, which we summarize in this section. First, we need to extend our notion of crossings between cycles from Definition 2.27 to curves:

Definition 6.3 (crossing). Let Σ be an orientable surface of genus g and q_1, q_2 two curves on Σ that intersect in a finite set $X := q_1 \cap q_2$ of points on Σ . For each $x \in X$ we say that q_1 and q_2 *cross* in x if for any sufficiently small neighbourhood $U \subseteq \Sigma$ of x the curve q_2 intersects more than one connected component of $U \setminus q_1$. If $X' \subseteq X$ is the set of points where q_1 and q_2 cross then we say that q_1 and q_2 cross $|X'|$ times.

Definition 6.4. Let G be a graph embedded in an orientable surface of genus g . We call a simple, closed and non-separating curve q on Σ *G -avoiding* if $q \cap V(G) = \emptyset$ and any $e \in E(G)$ is either embedded disjointly to q or intersects q in exactly one point in which e and q cross.

As in [83] we note that locally even a non-separating curve on Σ has two “sides”. In the following sections will use this fact more intuitively rather than formally correct. Therefore, we only state the following proposition informally:

Proposition 6.5 ([83]). *Let G be a graph embedded in an orientable surface of constant genus g and q a G -avoiding curve on Σ . We can partition a small neighbourhood around q into a “left” and a “right” part. For each edge $e \in E(G)$ that crosses q we can denote e by $e = \{a, b\}$ such that traversing e from a to b enters the embedding of q “from the left” and leaves the embedding of q “to the right”. See Figure 6.2 for an image.*

The following is explained e.g. in [9]:

Definition 6.6 (surgery). Let Σ be a surface and q a simple, closed and non-separating curve on Σ . The surface obtained from *applying surgery to Σ along q* is constructed as follows: Delete a small neighbourhood of q on Σ from Σ such that the boundary of the resulting surface consists of exactly two cycles. Glue a disk along each of those cycles to the result.

Proposition 6.7 ([9]). *Let Σ be an orientable surface of genus g and q a simple, closed and non-separating curve on Σ . The surface obtained after applying surgery along q has genus $g' < g$.*

The following observation goes back to Youngs [88]. We give a brief proof using the surgery operation from Definition 6.6.

Proposition 6.8 ([88]). *Let G be a graph embedded in an orientable surface of genus g and q a G -avoiding curve on Σ . Let $E_q \subseteq E(G)$ be the set of edges crossed by q . For any $e \in E_q$ let $e = \{a_e, b_e\}$ such that traversing e from a to b enters the embedding of q “from the left” and leaves the embedding of q “to the right” according to Proposition 6.5. Let $A := \{a_e : e \in E_q\}$ and $B := \{b_e : e \in E_q\}$. If there is no A - B -path in $G - E_q$ then each connected component of G can be embedded in an orientable surface of genus at most $g - 1$.*

Proof. W.l.o.g. G is connected. We show the existence of a simple, closed and non-separating curve q^* on Σ that does not intersect any edge or vertex of G . Since G can then still be embedded in the result of applying surgery along q^* this will finish the proof due to Proposition 6.7.

If $E_q = \emptyset$ then setting $q^* = q$ suffices, so assume $E_q \neq \emptyset$. Let G_A be the subgraph of $G - E_q$ that is connected to A and G_B the subgraph of $G - E_q$ that is connected to B .

Deleting the embeddings of q and G from Σ must yield a connected component F such that the boundary of F is not connected: Since q is non-separating, there exists some curve q' on $\Sigma \setminus q$ that starts in A and ends in B . The boundary of the connected component F that q' traverses before first reaching the embedding of some edge or vertex of G_B contains vertices of both G_A and G_B and is thus disconnected. Let q_F be the closed curve on Σ corresponding to traversing one connected component of the boundary of F . Since G is connected, q_F is non-separating. Slightly shifting q_F into F yields q^* . \square

We also need the following:

Proposition 6.9. *Let G be a connected graph embedded in an orientable surface Σ of genus $g \geq 1$. We can find a G -avoiding curve on Σ in polynomial time.*

Proof. As in the proof of Proposition 6.8, if the boundary of a face of G is disconnected then we can find a G -avoiding curve in Σ that does not intersect any edge or vertex of G . Otherwise, we consider the dual graph G^* of G w.r.t. the embedding of G . By a result by Thomassen [85] one can find a non-separating cycle in G^* in polynomial time. The embedding of this cycle in G^* corresponds to a G -avoiding curve. \square

6.2 Eliminating facial cycles

In this section we show how to eliminate all *facial* cycles from an uncrossable cycle family. These are the cycles whose embedding is null-homotopic:

Definition 6.10 (facial). Let G be a graph, embedded in a fixed orientable surface Σ of genus g . A separating cycle C in G is called *facial* if one of its sides is homeomorphic to the 2-dimensional disk. Clearly, this side is unique if $g \geq 1$. In this case we call it the *facial side* of C .

We use the primal-dual framework by Goemans and Williamson [40]. They proved:

Theorem 6.11 ([40]). *Let \mathcal{C} be a family of cycles in a graph G with vertex weights $w: V(G) \rightarrow \mathbb{R}_{\geq 0}$. Let $\gamma > 0$. Assume there exists an oracle that takes as input a subgraph G' of G such that $\mathcal{C}[G'] \neq \emptyset$ and outputs a non-empty set $\mathcal{M} \subseteq \mathcal{C}[G']$ such that for any minimal transversal T for $\mathcal{C}[G']$ the following holds:*

$$\sum_{C \in \mathcal{M}} |C \cap T| \leq \gamma |\mathcal{M}|$$

Then there exists a transversal T^ for \mathcal{C} such that $w(T^*) \leq \gamma \cdot OPT_{LP(\mathcal{C})}$.*

In the case that G is planar, Goemans and Williamson [41] proved that for any uncrossable family \mathcal{C} of cycles in G the oracle that returns all face-minimal cycles in $\mathcal{C}[G']$ fulfills this requirement with $\gamma = 3$. To show this, they use an uncrossing procedure to find a *laminar* family of *witness cycles*. Sun [83] claimed that the same approach also works for the facial cycles in the family $\vec{\mathcal{C}}_{\text{all}}$ of a graph G , embedded in an orientable surface of constant genus. However, the approach implicitly uses that the set of facial cycles in $\vec{\mathcal{C}}_{\text{all}}$ is uncrossable, which is actually not the case; a counterexample is given in Figure 6.1.




Figure 6.1: Example showing that the facial cycles among the (uncrossable) family of directed cycles in a digraph are not uncrossable anymore. It is embedded on the torus; the opposite edges of the square are identified.

Both colored cycles (green and violet) are facial. The path P_2 shares only its endpoints with the violet cycle, but cannot be extended to a facial cycle by using only violet edges.

We will see in the following that the full uncrossing property for the facial cycles in \mathcal{C} is not needed. First, we introduce the notion of witness cycles.

Definition 6.12 (witness cycle). *Let \mathcal{C} be a family of cycles in a graph G and $T \subseteq V(G)$ a minimal cycle transversal for \mathcal{C} . A *witness cycle* for some vertex $v \in T$ is a cycle $C_v \in \mathcal{C}$ with $V(C) \cap T = \{v\}$. A *family of witness cycles* for T consists of a witness cycle C_v for each $v \in T$.*

Lemma 6.13. *Let \mathcal{C} be an uncrossable family of cycles in a graph G , embedded in an orientable surface Σ of genus $g \geq 1$. Let $\mathcal{C}_f \subseteq \mathcal{C}$ denote the facial cycles in \mathcal{C} and let T be a minimal transversal for \mathcal{C}_f . Then there exists a laminar family $\{C_v : v \in T\} \subseteq \mathcal{C}_f$ of facial witness cycles for T .*

Proof. Let \mathcal{F} denote the set of faces of G , i.e. the connected components of Σ after deleting the embedding of G . Since T is a minimal transversal, there exists a facial witness cycle for each $v \in T$ and thus there exists a family $\mathcal{W} = \{C_v : v \in T\} \subseteq \mathcal{C}_f$ of witness cycles for T .

Among all families of facial witness cycles we choose \mathcal{W} such that

$$\phi(\mathcal{W}) := \sum_{F \in \mathcal{F}} |\{C \in \mathcal{W} : \text{the facial side of } C \text{ contains } F\}|$$

is minimum. We show that \mathcal{W} is laminar.

Assume there are $v_1, v_2 \in T$ such that the facial sides S_1, S_2 of $C_1 := C_{v_1}$ and $C_2 := C_{v_2}$ cross, i.e. $S_1 \setminus S_2 \neq \emptyset$ and $S_2 \setminus S_1 \neq \emptyset$. This implies that C_i must contain an edge that is (except for its endpoints) embedded inside S_{3-i} for each $i = 1, 2$. Let \mathcal{P}_i denote the set of subpaths of C_i that share exactly their endpoints with C_{3-i} and are embedded inside S_{3-i} otherwise. By the above observation and because of $g \geq 1$, both of the \mathcal{P}_i are non-empty.

First, assume that for some $i \in \{1, 2\}$ there exists a path $P \in \mathcal{P}_i$ such that $v_i \notin V(P)$. By the uncrossing property for \mathcal{C} there exists a subpath P' of C_{3-i} such that $C' := P + P' \in \mathcal{C}$. Since C_{3-i} is facial also C' is facial and the facial side of C' is strictly smaller than the facial side of C_{3-i} . Also, $v_{3-i} \in V(P')$ because T is a transversal for \mathcal{C}_f and $T \cap V(C_j) = \{v_j\}$ for $j = 1, 2$. But then replacing C_{3-i} by C' in \mathcal{W} yields another family of witness cycles for T and reduces ϕ , contradicting minimality of \mathcal{W} .

Therefore, for $i = 1, 2$ all paths in \mathcal{P}_i contain v_i . Since $v_i \notin V(C_{3-i})$ there can be at most one such path in \mathcal{P}_i . Thus, \mathcal{P}_i contains exactly one element, which we call P_i . Since $(C_1 - P_1) + (C_2 - P_2)$ contains no element of T , it also contains no cycle in \mathcal{C} (as an edge set). Applying the uncrossing property to the subpath P_2 of C_2 and C_1 therefore yields $C'_1 := P_1 + (C_2 - P_2) \in \mathcal{C}$ and $C'_2 := P_2 + (C_1 - P_1) \in \mathcal{C}$. Again the fact that C_1 and C_2 are facial implies $C'_1, C'_2 \in \mathcal{C}_f$ and furthermore $\phi(\{C'_1, C'_2\}) < \phi(\{C_1, C_2\})$. Also, C'_i is a witness cycle for v_i for $i = 1, 2$. Thus, replacing C_1 and C_2 by C'_1 and C'_2 in \mathcal{W} yields another contradiction to the minimality of \mathcal{W} . This proves that \mathcal{W} is already laminar. \square

Lemma 6.14. *Let \mathcal{C} be a non-empty uncrossable family of cycles in a graph G , embedded in an orientable surface of genus $g \geq 1$. Let $\mathcal{C}_f \subseteq \mathcal{C}$ denote the set of facial cycles in \mathcal{C} and let $\mathcal{M} \subseteq \mathcal{C}_f$ be the set of facial cycles with minimal facial side. Let $T \subseteq V(G)$ be a minimal transversal for \mathcal{C}_f . Then*

$$\sum_{C \in \mathcal{M}} |V(C) \cap T| \leq (15 + 6g)|\mathcal{M}|$$

Furthermore, if $|\mathcal{M}| \leq 2$ then

$$\sum_{C \in \mathcal{M}} |V(C) \cap T| \leq 8|\mathcal{M}|$$

Proof. By Lemma 6.13, let $\mathcal{W} \subseteq \mathcal{C}_f$ be a laminar family of witness cycles for T . Let $T' := T \cap V(\mathcal{M})$ and $\mathcal{W}' := \{W \in \mathcal{W} : V(W) \cap T' \neq \emptyset\}$. By a similar proof as for Proposition 3.1 we can assume that the cycles in \mathcal{M} are given by boundaries of (not necessarily finite) faces of G . In particular, $\mathcal{W}' \cup \mathcal{M}$ is still laminar, and the facial sides of cycles in \mathcal{M} are pairwise disjoint.

Define a function $f: \mathcal{W}' \rightarrow 2^{\mathcal{M}}$ such that for any $C \in \mathcal{M}$ and $W \in \mathcal{W}'$ we have that $C \in f(W)$ if and only if the facial side of W contains the facial side of C . Since \mathcal{W}' is laminar, the family $\mathcal{L} := \{f(W) : W \in \mathcal{W}'\}$ of sets is also laminar. In particular, $|\mathcal{L}| \leq 2|\mathcal{M}|$.

Assume there are three different cycles $W_1, W_2, W_3 \in \mathcal{W}'$ with $f(W_1) = f(W_2) = f(W_3)$. For $i = 1, 2, 3$ let S_i denote the facial side of W_i . Since $g \geq 1$, each W_i has a unique facial side and the S_i form a laminar family (of sets). Also, $f(W_i) \neq \emptyset$ for each $i = 1, 2, 3$, so the S_i are not

disjoint and therefore w.l.o.g. $S_1 \subsetneq S_2 \subsetneq S_3$. Choose $v \in T' \cap V(W_2)$ and $C \in \mathcal{M}$ with $v \in V(C)$. Then $C \in f(W_3) \setminus f(W_1)$, a contradiction. Thus, $|\mathcal{W}'| \leq \sum_{L \in \mathcal{L}} |f^{-1}(L)| \leq 2|\mathcal{L}| \leq 4|\mathcal{M}|$.

Now define the graph G' on vertex set $\mathcal{M} \cup T'$ by adding an edge between any $C \in \mathcal{M}$ and $v \in T' \cap V(C)$. G' can be embedded in Σ by embedding the vertex corresponding to $C \in \mathcal{M}$ somewhere in the facial side of C and embedding all edges incident to C inside this facial side. Since G' contains no parallel edges, according to Lemma 2.21 Euler's formula implies

$$\sum_{C \in \mathcal{M}} |V(C) \cap T| = \sum_{C \in \mathcal{M}} |V(C) \cap T'| = |E(G')| \leq 3|V(G')| + 6g = 3|\mathcal{M}| + 3|T'| + 6g.$$

The first assertion of the lemma now follows from $|T'| = |\mathcal{W}'| \leq 4|\mathcal{M}|$. For the second assertion, note that

$$\sum_{C \in \mathcal{M}} |V(C) \cap T| = \sum_{C \in \mathcal{M}} |V(C) \cap T'| \leq |\mathcal{M}| \cdot |T'| = |\mathcal{M}| \cdot |\mathcal{W}'| \leq 4|\mathcal{M}|^2. \quad \square$$

As an immediate consequence we get the following result:

Lemma 6.15. *Let \mathcal{C} be an uncrossable family of cycles in a graph G with node weights $w: V(G) \rightarrow \mathbb{R}_{\geq 0}$, embedded in a fixed orientable surface Σ of genus $g \geq 1$. There exists a set $T \subseteq V(G)$ with $w(T) \leq (15 + 6g) \text{OPT}_{\text{LP}(\mathcal{C})}$ such that $G - T$ contains no facial cycles in \mathcal{C} . If there exist two points $p, p' \in \Sigma$ on the surface such that the facial side of any facial cycle in \mathcal{C} contains p or p' , then we can bound $w(T)$ by $8 \text{OPT}_{\text{LP}(\mathcal{C})}$.*

Proof. We apply Theorem 6.11 to the family \mathcal{C}_f of facial cycles in \mathcal{C} and the oracle that, given a subgraph G' of G , outputs the set of all facial cycles with minimal facial side in $\mathcal{C}_f[G']$. By Lemma 6.14 this fulfills the requirements from Theorem 6.11 with $\gamma := 15 + 6g$. If any facial side of a cycle in \mathcal{C} contains p or p' then the oracle will always output at most 2 cycles because their facial sides must be disjoint. Thus, in this case we can use Theorem 6.11 with $\gamma := 8$. The proof is finished by the observation that $\text{OPT}_{\text{LP}(\mathcal{C}_f)} \leq \text{OPT}_{\text{LP}(\mathcal{C})}$. \square

We will see in the following Section 6.3 how this result can be used to reduce the problem of bounding the integrality gap of the LP (1.5) for uncrossable cycle families to the case where the family contains no facial cycles. See Theorem 6.25.

6.3 Bounding the integrality gap

In this section we give a proof of Theorem 6.1. Some tools and the general approach are quite similar to [83]. We first introduce the notion of *subgraph-defined* uncrossable cycle families.

Definition 6.16 (subgraph-defined). An uncrossable family \mathcal{C} in an undirected graph G is called *subgraph-defined* if there is a subset $S \subseteq E(G)$ of edges such that the following two properties hold:

1. For any $C \in \mathcal{C}$ at most one edge of C does not belong to S .
2. Given a cycle $C \in \mathcal{C}$ and two points $v, w \in V(C)$ with a v - w -path P in (V, S) , then C contains a v - w -path Q such that $P + Q$ contains a cycle in \mathcal{C} (as an edge set).

Here are the most important examples that have this property. Most of the results from this section only apply for subgraph-defined uncrossable cycle families.

Proposition 6.17. *Given an undirected graph G and a set $D \subseteq E(G)$ of demand edges, the uncrossable cycle families \mathcal{C}_{all} , \mathcal{C}_{odd} , $\mathcal{C}_D^{\geq 1}$ and $\mathcal{C}_D^=1$ are all subgraph-defined.*

Proof. We set $S := E(G)$ for the examples \mathcal{C}_{all} , \mathcal{C}_{odd} and $\mathcal{C}_D^{\geq 1}$ and $S := E(G) \setminus D$ for $\mathcal{C}_D^=1$. \square

On the other hand, it is easy to see that the set $\mathcal{C}_{\text{all}}[l]$ of all shortest cycles w.r.t. edge lengths l is not subgraph-defined. A counterexample is given e.g. by a graph on three vertices with a pair of parallel edges between any two vertices and $l \equiv 1$.

Throughout this section let $G = (V, E)$ denote a graph with vertex weights $w: V \rightarrow \mathbb{R}_{\geq 0}$, embedded in a fixed orientable surface Σ of genus g . Furthermore, let \mathcal{C} denote a subgraph-defined uncrossable family of cycles in G if not stated otherwise. We let $S \subseteq E$ be an edge set as guaranteed by Definition 6.16.

We first prove an easy consequence of the Definition 6.16:

Lemma 6.18. *If $G = G[\mathcal{C}]$ is connected then also the graph (V, S) is connected.*

Proof. Assume not, and let $X \subsetneq V$ be a connected component of (V, S) . Let $e \in \delta(X)$ be an outgoing edge of X . Let $C \in \mathcal{C}$ such that $e \in E(C)$. Since $e \notin S$, Definition 6.16 implies $E(C) \setminus \{e\} \in S$, contradicting the choice of e . \square

The following three lemmata establish tools we will use to bound the weight of our transversal for \mathcal{C} . A similar technique has been used by Sun [83] for hitting directed cycles.

Definition 6.19. Let $x \in \mathbb{R}_{\geq 0}^V$ be a feasible solution to the LP (1.5). For a path P in G , define its distance w.r.t. x to be $d_x(P) := x(V(P))$. For $A, B \subseteq V$ define $d_x^G(A, B)$ to be the minimum distance w.r.t. x of an A - B -path in G . If B is not reachable from A in G , we set $d_x^G(A, B) := \infty$. For $a, b \in V$ we write $d_x^G(a, b)$ in short for $d_x^G(\{a\}, \{b\})$.

Note that by including the x values of both endpoints of P in d_x^G this does not define a metric on V because $d_x^G(v, v) > 0$ for some $v \in V$. However, this definition makes the following proofs easier.

Lemma 6.20. *Let $x \in \mathbb{R}_{\geq 0}^V$ be a feasible solution to the LP (1.5) and G' a subgraph of G . Let $\alpha > 0$ and $A, B \subseteq V(G')$ such that $d_x^{G'}(A, B) \geq \alpha$. Then there exists a set $X \subseteq V(G')$ such that $w(X) \leq \frac{1}{\alpha} \cdot \sum_{v \in V} w(v)x(v)$ and B is not reachable from A in $G' - X$.*

Proof. Choose $p \in [0, \alpha]$ uniformly at random. Define

$$X_p := \{v \in V(G') : p \leq d_x^{G'}(A, \{v\}) \leq p + x(v)\}$$

By definition, in expectation $w(X_p)$ can be bounded by $\sum_{v \in V(G')} w(v) \frac{x(v)}{\alpha}$. Thus, we can choose $p^* \in [0, \alpha]$ such that $w(X_{p^*}) \leq \frac{1}{\alpha} \cdot \sum_{v \in V} w(v)x(v)$. Let $X := X_{p^*}$.

It is left to show that B is not reachable from A in $G - X$. Let $P = v_1 v_2 \dots v_k$ be some A - B -path in G' . Since $d_x^{G'}(A, B) \geq \alpha \geq p^*$, choose $1 \leq i^* \leq k$ minimal such that $p^* \leq d_x^{G'}(A, \{v_{i^*}\})$. By minimality of i^* , also $d_x^{G'}(A, \{v_{i^*}\}) \leq p^* + x(v_{i^*})$. Thus, $P \cap X \neq \emptyset$, finishing the proof. \square

Lemma 6.21. *Let $x \in \mathbb{R}_{\geq 0}^V$ be a feasible solution to the LP (1.5) and G' a subgraph of G . Let G_S denote the graph $G'[S]$. Let $0 < \alpha < \frac{1}{2}$ and $A \subseteq V(G')$ be connected in G_S such that $x(A) \leq \alpha$. Then there is a set $X \subseteq V(G')$ such that $w(X) \leq \frac{4}{1-2\alpha} \cdot \sum_{v \in V} w(v)x(v)$ and no cycle in $\mathcal{C}[G' - X]$ contains a vertex of A .*

Proof. Define $\beta := \frac{1-2\alpha}{4} > 0$. Let $B := \{v \in V(G') : d_x^{G_S}(A, \{v\}) \geq \beta\}$. By Lemma 6.20, we can find a set $X \subseteq V(G_S)$ such that $w(X) \leq \frac{1}{\beta} \cdot \sum_{v \in V} w(v)x(v)$ and B is not reachable from A in $G_S - X$. It suffices to show that any cycle in $\mathcal{C}[G']$ intersects B .

Assume this is wrong, and let $C \in \mathcal{C}[G' - B]$ be a cycle of minimum LP value $x(V(C))$. Choose a minimal subpath P of C such that $d_x(P) \geq \frac{1}{2}x(V(C))$ and let v and w be its endpoints. This ensures $d_x(P') - x(v) - x(w) \leq \frac{1}{2}x(V(C))$ for any of the two v - w -paths P' in C . Let P_v and P_w be paths in $G_S - B$ of distance less than β w.r.t. x such that P_v connects v to A and P_w connects w to A . By concatenating P_v , some path inside $G_S[A]$ and P_w we get a v - w -path P^* in $G_S - B$ of distance less than $2\beta + \alpha = \frac{1}{2} \leq \frac{1}{2}x(C)$ w.r.t. x . Using Property 2 from Definition 6.16 for P^* and C yields a cycle $C^* \in \mathcal{C}[G' - B]$ with $x(V(C^*)) < x(V(C))$, contradicting minimality of C . \square

Lemma 6.22. *Let $x \in \mathbb{R}_{\geq 0}^V$ be a feasible solution to the LP (1.5) and G' a subgraph of G . Let G_S denote the graph $G'[S]$. Let $\alpha > 0$ and $A, B \subseteq V(G')$ such that no edge $\{v, w\} \in E(G')$ fulfills $d_x^{G_S}(A, \{v\}) \leq \alpha$ and $d_x^{G_S}(B, \{w\}) \leq \alpha$. Then there exists a set $X \subseteq V(G')$ such that $w(X) \leq \frac{2}{\alpha} \cdot \sum_{v \in V} w(v)x(v)$ and $G' - X$ contains no A - B -path with at most one edge outside S .*

Proof. For $M \in \{A, B\}$ define $M' := \{v \in V(G') : d_x^{G_S}(M, \{v\}) \geq \alpha\}$. By Lemma 6.20 we can choose $X_M \subseteq V(G_S)$ with $w(X_M) \leq \frac{1}{\alpha} \cdot \sum_{v \in V(G_S)} w(v)x(v)$ such that M' is not reachable from M in $G_S - X_M$. Define $X := X_A \cup X_B$. It suffices to show that $G' - X$ contains no A - B -path with at most one edge outside S .

Assume this is wrong, and let P be an A - B -path in $G' - X$ with an edge $e = \{v, w\} \in E(P)$ such that $E(P) \setminus \{e\} \subseteq S$. W.l.o.g. P contains an A - v -path and a B - w -path in G_S . However, by the assumption in the lemma's statement $v \in A'$ or $w \in B'$ holds, contradicting that P avoids X . \square

With these tools, we can finally show how to eliminate enough cycles from \mathcal{C} such that the remaining graph has lower genus. Applying this result iteratively then yields Theorem 6.1.

Lemma 6.23. *Assume the genus g of our surface Σ is at least 1. Then, there exists a set $X \subseteq V$ with $w(X) \leq 24OPT_{LP(\mathcal{C})}$ such that each connected component of $\mathcal{C}[G - X]$ can be embedded in an orientable surface of genus at most $g - 1$.*

Proof. By Proposition 6.9 let q be a G -avoiding curve on Σ . Let $E_q \subseteq E(G)$ be the set of edges that q crosses. The embedding of q induces a cyclic order on E_q . Enumerate the edges in E_q according to this order by $E_q = \{e_1, \dots, e_k\}$. For $i = 1, \dots, k$ let $e_i = \{a_i, b_i\}$ such that traversing e from a to b enters the embedding of q “from the left” and leaves the embedding of q “to the right” according to Proposition 6.5. Let $A := \{a_1, \dots, a_k\}$ and $B := \{b_1, \dots, b_k\}$. Let $G' := G - E_q$ and $G_S := G'[S]$.

Let $x \in \mathbb{R}_{\geq 0}^V$ be an optimum solution to the LP (1.5) and set $\alpha := \frac{1}{4}$. We say that vertices $a \in A$ and $b \in B$ are *close* if there exists an edge $e = \{v, w\} \in E(G')$ such that $d_x^{G_S}(a, v) \leq \alpha$ and $d_x^{G_S}(b, w) \leq \alpha$. There are two cases:

Case 1: There is no pair of close vertices $a \in A, b \in B$. In this case we apply Lemma 6.22 to find a set $X \subseteq V$ with $w(X) \leq \frac{2}{\alpha} \cdot \sum_{v \in V} w(v)x(v) = 8\text{OPT}_{\text{LP}(\mathcal{C})}$ such that $G' - X$ contains no A - B -path with at most one edge outside S . Define $G_X := G[\mathcal{C}[G - X]]$. We show that there is no A - B -path in $G_X - E_q$, which finishes the proof for this case due to Proposition 6.8. Assume there is an A - B -path in $G_X - E_q$, and let P be such a path containing a minimum number of edges outside S ; clearly this number is at least 2. Let $e \in E(P) \setminus S$. Minimality of P implies that one of the endpoints of e , say v , is not reachable from $A \cup B$ in $G_X[S]$. By definition of G_X we can choose a cycle $C \in \mathcal{C}[G_X]$ containing e . Since replacing e in P by $C - e$ yields an A - B -path with fewer edges in S we know that C contains some edge from E_q . But then C yields an $(A \cup B)$ - v -path in $G_X[S]$, a contradiction.

Case 2: There exist close vertices $a \in A, b \in B$. In this case, choose $1 \leq i^* \leq j^* \leq k$ such that a_{i^*} and b_{j^*} or a_{j^*} and b_{i^*} are close, but for any $i, j \in \{i^* + 1, \dots, j^* - 1\}$ neither a_i and b_j nor a_j and b_i are close. W.l.o.g. a_{i^*} and b_{j^*} are close. Thus, we can choose an edge $e = \{v, w\} \in E(G')$ and an a_{i^*} - v path P_1 and a w - b_{j^*} -path P_2 in G_S , both of distance at most α w.r.t. x . By applying Lemma 6.21 to $V(P_1)$ and $V(P_2)$ we can find sets $X_1, X_2 \subseteq V(G)$ such that setting $X := X_1 \cup X_2$ yields

$$w(X) \leq \frac{8}{1 - 2\alpha} \cdot \sum_{v \in V} w(v)x(v) = 16\text{OPT}_{\text{LP}(\mathcal{C})}$$

and no cycle in $\mathcal{C}[G - X]$ contains any vertex of P_1 or P_2 . Define $G_X := G[\mathcal{C}[G - X]]$.

Define the curve q' by concatenating the embeddings of P_1 , e , P_2 , part of the embedding of e_{j^*} , the part of q that crosses $e_{i^*+1}, \dots, e_{j^*-1}$ and part of the embedding of e_{i^*} (see Figure 6.2). Now q' is a G_X -avoiding curve (q' is non-separating because it is homotopic to a curve that crosses q exactly once) and by minimality of i^* and j^* we know that restarting this proof for G_X and q' brings us into case 1 (note that distances w.r.t. x can only increase by deleting vertices or edges). Thus, in total we get a vertex set of weight at most $16\text{OPT}_{\text{LP}(\mathcal{C})} + 8\text{OPT}_{\text{LP}(\mathcal{C})}$ with the desired properties. \square

Now we can give a proof of Theorem 6.1, which we restate here for convenience:

Theorem 6.1. *Given a graph G with vertex weights $w: V(G) \rightarrow \mathbb{R}_{\geq 0}$, embedded in an orientable surface of genus g , and a subgraph-defined uncrossable family \mathcal{C} of cycles in G , then the integrality gap of the LP (1.5) is in $O(g)$.*

Proof. We show by induction on $g \in \mathbb{Z}_{\geq 0}$ that the integrality gap of the LP (1.5) is at most $(24g + 2.4)$. For $g = 0$ this has been proven by Berman and Yaroslavtsev [16]. Now assume $g > 0$ and the assertion holds for smaller g . By Lemma 6.23, there exists a set $X_1 \subseteq V(G)$ with $w(X_1) \leq 24\text{OPT}_{\text{LP}(\mathcal{C})}$ such that each connected component of $G_{X_1} := G[\mathcal{C}[G - X_1]]$ can be embedded in an orientable surface of genus at most $g - 1$. Thus, we can apply the induction hypothesis to each connected component of G_{X_1} . Together this yields a set $X_2 \subseteq V(G')$ such that $G' - X_2$ contains no cycles in \mathcal{C} , and $w(X_2) \leq (24(g - 1) + 2.4)\text{OPT}_{\text{LP}(\mathcal{C}[G - X_1])} \leq (24(g - 1) + 2.4)\text{OPT}_{\text{LP}(\mathcal{C})}$. Thus, $X := X_1 \cup X_2$ is a feasible integral solution to (1.5) of at most $24g + 2.4$ times the LP value. \square




Figure 6.2: The curve q is the horizontal line in the middle, going once around the torus (note that the blue arrows as well as the red arrows are identified). Clearly, q is G -avoiding and each edge $e_i = \{a_i, b_i\}$ crosses q in the same direction. Note that the edges e_1, \dots, e_5 are crossed by q in this order.

If in case 2 of the proof of Lemma 6.23, P_1 and P_2 are short paths w.r.t. x between $a_{i^*} = a_2$ and v and $b_{j^*} = b_5$ and w , respectively, we can eliminate all cycles in \mathcal{C} intersecting a vertex of P_1 or P_2 . Then, the curve q' as depicted in violet crosses only $\{e_{i^*+1}, \dots, e_{j^*-1}\}$ and brings us to case 1.

Note that all steps in the proof can be carried out in polynomial time if \mathcal{C} has a support oracle and S is given in the input. If \mathcal{C} also has a weight oracle we can also solve the LP (1.5). Thus, we note:

Theorem 6.24. *Let G be a graph with vertex weights $w: V(G) \rightarrow \mathbb{R}_{\geq 0}$, embedded in an orientable surface of genus g . For any subgraph-defined uncrossable cycle family that has both a support and a weight oracle and where the set S from Definition 6.16 is known we can compute a transversal for \mathcal{C} of weight at most $O(g)$ times the LP value in polynomial time.*

Proving an analogous result to Theorem 6.1 or 6.24 for more or even all uncrossable cycle families remains an open problem. However, Lemma 6.15 makes a first step in that direction. Similar to [83], we note:

Theorem 6.25. *Assume that for any uncrossable cycle family \mathcal{C} in a graph G with vertex weights $w: V(G) \rightarrow \mathbb{R}_{\geq 0}$, embedded in a fixed orientable surface Σ of genus g , the following holds: If \mathcal{C} contains no facial cycles, we can compute a vertex set $X \subseteq V(G)$ of weight at most $\alpha \text{OPT}_{\text{LP}(\mathcal{C})}$ such that there exists a simple, closed and non-separating curve q on Σ that is disjoint to the embedding of any edge or vertex of $\mathcal{C}[G - X]$.*

Then the integrality gap of the LP (1.5) is at most $O(\alpha g)$.

Proof. We show by induction on $g \in \mathbb{Z}_{\geq 0}$ that the integrality gap of the LP (1.5) is at most $(\alpha + 8)g + 1$ if \mathcal{C} contains no facial cycles. This time, the case $g = 0$ implies $\mathcal{C} = \emptyset$ and is therefore trivial. Let now $g > 0$. By the assumption let $X_1 \subseteq V(G)$ with $w(X_1) \leq \alpha \text{OPT}_{\text{LP}(\mathcal{C})}$ and q a closed, simple and non-separating curve on Σ that is disjoint to the embedding of any

edge or vertex of $G_{X_1} := G[\mathcal{C}[G - X_1]]$. Let Σ' arise from applying surgery to Σ along q . The embedding of G on Σ induces an embedding of G_{X_1} on Σ' . Note that the facial side of any facial cycle in $\mathcal{C}[G_{X_1}]$ must contain one of the two disks glued to Σ' when doing surgery. In particular, Lemma 6.15 allows us to remove all facial cycles in $\mathcal{C}[G_{X_1}]$ by deleting a set X_2 of weight at most $8\text{OPT}_{\text{LP}(\mathcal{C})}$. By Proposition 6.7 we can apply the induction hypothesis to G_{X_1} to get a set X_3 of weight at most $((\alpha + 8)(g - 1) + 1)\text{OPT}_{\text{LP}(\mathcal{C})}$ such that $X_1 \cup X_2 \cup X_3$ is a feasible integral solution to our LP, finishing our induction.

Finally, consider the case that \mathcal{C} contains also facial cycles. We first apply Lemma 6.15 to find a vertex set X_0 with $w(X_0) \leq (15 + 6g)\text{OPT}_{\text{LP}(\mathcal{C})}$ such that $\mathcal{C}[G - X_0]$ contains no facial cycles. Applying the above induction to $\mathcal{C}[G - X_0]$ shows that we can extend X_0 to a transversal for \mathcal{C} of weight at most $O(\alpha g)\text{OPT}_{\text{LP}(\mathcal{C})}$. \square

Chapter 7

The Erdős–Pósa ratio of odd cycles in planar graphs

In this section we consider a specific example of an uncrossable family in a graph G : The set of all odd cycles in G . As we will see in the following, these are much more structured than arbitrary uncrossable cycle families. For example, the edge versions of the corresponding packing and transversal problem in planar graphs are essentially solved (cf. Table 1.1) and it is known that the Erdős–Pósa ratio is exactly 2 [55]. For the vertex version in planar graphs Král', Sereni and Stacho [54] already proved an upper bound of 6 on the Erdős–Pósa ratio for this particular cycle family, which is better than our upper bound for arbitrary uncrossable cycle families from Corollary 5.2. We will combine methods from [54] with new ideas to improve this upper bound to 4:

Theorem 7.1. *Let G be a planar graph and \mathcal{C} the family of odd cycles in G . Then $\tau_v(\mathcal{C}) \leq 4\nu_v(\mathcal{C})$.*

This is joint work with Luise Puhlmann [69]. Note that this matches the best known lower bound for the Erdős–Pósa ratio of arbitrary uncrossable cycle families in planar graphs (cf. Table 1.2). For this particular cycle family however there is no example showing that the Erdős–Pósa ratio exceeds 2.

7.1 Properties of odd cycles

Recall that an odd cycle transversal for a graph G is a set of vertices that intersects the vertex set of each odd cycle in G . For this particular cycle family we can give another equivalent definition which is very useful for us. For a planar graph G , embedded in the sphere, let $\mathcal{F}(G)$ be the set of (not necessarily finite) faces of G . For a face $F \in \mathcal{F}(G)$ let $V(F)$ be the set of vertices on the boundary of F . The edge set $E(F)$ is the set of boundary edges of F , excluding bridges. Therefore, $E(F)$ is Eulerian. We call F *odd* if $|E(F)|$ is odd. In that case $E(F)$ contains an odd cycle. We define $\mathcal{F}_{\text{odd}}(G) \subseteq \mathcal{F}(G)$ to be the set of odd faces of G .

Proposition 7.2. *Let G be a planar graph, embedded in the sphere, and C a cycle in G . Let $\mathcal{F}(C) \subseteq \mathcal{F}(G)$ be the set of faces of G in the interior of C . C is odd if and only if $\mathcal{F}(C)$ contains an odd number of odd faces.*

Proof. This follows directly from the fact that $E(C)$ can be written as the symmetric difference

$$E(C) = \bigoplus_{F \in \mathcal{F}(C)} E(F). \quad \square$$

As in [54], we define the *vertex-face incidence graph*:

Definition 7.3 (Vertex-face incidence graph). Given a planar graph G with a fixed planar embedding, its *vertex-face incidence graph* $\text{VF}(G)$ is the planar graph on the vertex set $\mathcal{F}(G) \cup V(G)$ with the edge set $E(\text{VF}(G))$ being $\{\{F, v\} : F \in \mathcal{F}(G), v \in V(F)\}$. We embed $\text{VF}(G)$ such that each edge $\{v, F\}$ is embedded inside F , see Figure 7.1 for an example.




Figure 7.1: The figure shows a graph G where $V(G)$ are the filled vertices and $E(G)$ are the solid edges. Its vertex-face incidence graph $\text{VF}(G)$ is the graph on all vertices (filled and empty) with the dotted edges. When choosing $T \subseteq V(G)$ to be the filled red vertices, $\text{VF}(G)[\mathcal{F}(G) \cup T]$ decomposes into three connected components, drawn in red, green, and blue, respectively. Hence T is an \mathcal{F} -transversal when choosing \mathcal{F} to be the faces of G that are filled in red. Note that \mathcal{F} are exactly the odd faces of G ; thus T is an odd cycle transversal for G by Proposition 7.5.

The notion of the vertex-face incidence graph allows us to define the following important concept:

Definition 7.4 (\mathcal{F} -Transversal). Let G be a planar graph, embedded in the plane. Let $\mathcal{F} \subseteq \mathcal{F}(G)$ have even cardinality $|\mathcal{F}|$. A set $T \subseteq V(G)$ is an \mathcal{F} -transversal if each connected component of $\text{VF}(G)[\mathcal{F}(G) \cup T]$ contains an even number of elements of \mathcal{F} . We define $\tau_{\text{odd}}(\mathcal{F})$ to be the size of a minimum-cardinality \mathcal{F} -transversal. For an example see Figure 7.1.

Note that Definition 7.4 does not require the faces of \mathcal{F} to be odd. However, the notion of \mathcal{F} -transversals coincides with the notion of odd cycle transversals for $\mathcal{F} := \mathcal{F}_{\text{odd}}(G)$:

Proposition 7.5. *Let G be a planar graph with a fixed planar embedding. A set $T \subseteq V(G)$ is an odd cycle transversal for G if and only if T is an $\mathcal{F}_{\text{odd}}(G)$ -transversal according to Definition 7.4.*

Proof. First we show “ \Rightarrow ”. Let $X \subseteq \mathcal{F}(G) \cup T$ be a connected component of $\text{VF}(G)[\mathcal{F}(G) \cup T]$ containing an odd number of elements of $\mathcal{F}_{\text{odd}}(G)$. Then the symmetric difference $C := \bigoplus_{F \in X \cap \mathcal{F}(G)} E(F)$ is odd and Eulerian and therefore contains some odd cycle. However, if any edge $e \in E(G)$ is incident to some vertex in T , then the two faces adjacent to e are connected in $\text{VF}(G)[\mathcal{F}(G) \cup T]$, so $e \notin E(C)$. Thus, $V(C) \cap T = \emptyset$ and T is no odd cycle transversal.

Next we show “ \Leftarrow ”. Assume there exists an odd cycle C in G with $V(C) \cap T = \emptyset$. Let $\mathcal{F}(C) \subseteq \mathcal{F}(G)$ be the set of faces inside C (w.r.t. the planar embedding of G). Clearly, no face in $\mathcal{F}(C)$ can be connected to a face outside $\mathcal{F}(C)$ in $\text{VF}(G)[\mathcal{F}(G) \cup T]$. But $\mathcal{F}(C)$ contains an odd number of elements of $\mathcal{F}_{\text{odd}}(G)$ because $C = \bigoplus_{F \in \mathcal{F}(C)} E(F)$ is odd. Thus, one of the connected components of $\text{VF}(G)[\mathcal{F}(C) \cup T]$ must also contain an odd number of elements of $\mathcal{F}_{\text{odd}}(G)$. \square

Our definition of \mathcal{F} -transversals is closely related to the well-known notion of T -joins:

Definition 7.6 (T -join). Given a graph G and a set $T \subseteq V(G)$, a T -join is a set $J \subseteq E(G)$ such that for any vertex $v \in V(G)$, $|\delta(v) \cap J|$ is odd if and only if $v \in T$.

Clearly a T -join in G exists if and only if each connected component of G contains an even number of elements of T . In particular, the definition of \mathcal{F} -transversals from Definition 7.4 can be reformulated as $\text{VF}(G)[\mathcal{F}(G) \cup T]$ containing an \mathcal{F} -join.

Note that in the edge-disjoint case the size of a minimum edge odd cycle transversal is actually equal to the size of a minimum $\mathcal{F}_{\text{odd}}(G)$ -join in the planar dual of G , while odd cycles in G correspond to $\mathcal{F}_{\text{odd}}(G)$ -cuts in the planar dual. As Fiorini et al.[33] pointed out, a relatively straightforward application of Seymour’s [81] theorem on T -joins yields $\tau_e(\mathcal{C}) \leq 2\nu_e(\mathcal{C})$ in this case, which was first proved by Král’, Sereni and Stacho [54]. The vertex-disjoint case is much harder because many edge-disjoint $\mathcal{F}_{\text{odd}}(G)$ -cuts in $\text{VF}(G)$ can correspond to odd cycles that meet in the same vertex. However, this happens only if many odd faces are nearby. If odd faces are further apart, then Král’, Sereni and Stacho [54] obtained a strong Erdős–Pósa-like bound, which we will also use as a black box (cf. Lemma 7.9).

Using the fact that for $T, T' \subseteq V(G)$, a T -join J and a T' -join J' , the symmetric difference $J \oplus J'$ is a $T \oplus T'$ -join, we can show the following property of \mathcal{F} -transversals:

Proposition 7.7. *Given a planar graph G with a fixed planar embedding, let $\mathcal{F}, \mathcal{F}' \subseteq \mathcal{F}(G)$ have even cardinality. Let $T \subseteq V(G)$ be an \mathcal{F} -transversal and $T' \subseteq V(G)$ be an \mathcal{F}' -transversal. Then $T \cup T'$ is an $\mathcal{F} \oplus \mathcal{F}'$ -transversal.*

Proof. Let J be an \mathcal{F} -join contained in $\text{VF}(G)[\mathcal{F}(G) \cup T]$ and J' be an \mathcal{F}' -join contained in $\text{VF}(G)[\mathcal{F}(G) \cup T']$. Then $J \oplus J'$ is contained in $\text{VF}(G)[\mathcal{F}(G) \cup (T \cup T')]$. As $J \oplus J'$ is an $\mathcal{F} \oplus \mathcal{F}'$ -join, this implies that $T \cup T'$ is an $\mathcal{F} \oplus \mathcal{F}'$ -transversal. \square

Similar to \mathcal{F} -transversals, we define the packing number with respect to some face set:

Definition 7.8. Given a planar graph G embedded in the sphere and a set $\mathcal{F} \subseteq \mathcal{F}(G)$ of its faces, define $\nu(\mathcal{F})$ to be the maximum number of pairwise vertex-disjoint faces of \mathcal{F} .

7.2 Proof of the Main Theorem

In this section we prove Theorem 7.1. Throughout this section let G denote a planar graph, embedded in the sphere, and let \mathcal{C} be the family of odd cycles in G .

7.2.1 High-level outline

The outline of our proof roughly follows the proof of $\tau_v(\mathcal{C}) \leq 6\nu_v(\mathcal{C})$ by Král', Sereni and Stacho [54]. They use the notion of *clouds* of a planar graph G , which are the maximal subsets of $\mathcal{F}_{\text{odd}}(G)$ that are connected in $\text{VF}(G)[\mathcal{F}_{\text{odd}}(G) \cup V(G)]$. For any cloud, they compute a vertex set T and a set \mathcal{P} of vertex-disjoint odd faces of the cloud such that $|T| \leq 6|\mathcal{P}| - 2$ and the whole cloud is contained in a single face of $G - T$. They mark such a face as “deadly”. After applying this to each cloud, they use the following lemma to finish the proof:

Lemma 7.9 ([54]). *Let G be a planar graph, embedded in the sphere. Let $\mathcal{F}_{\text{dead}} \subseteq \mathcal{F}(G)$ be a set of faces marked as deadly, including all odd faces of G . Then $\tau_v(\mathcal{C}) \leq \nu_{\text{dead}}(G) + 2|\mathcal{F}_{\text{dead}}|$, where $\nu_{\text{dead}}(G)$ is the maximum number of pairwise vertex-disjoint odd cycles that are vertex-disjoint to all deadly faces in G .*

For our approach it will be convenient to generalize the notion of clouds:

Definition 7.10. Let G be a planar graph, embedded in the sphere, and $\mathcal{F}^* \subseteq \mathcal{F}(G)$ with $\mathcal{F}_{\text{odd}}(G) \subseteq \mathcal{F}^*$ define a set of *special faces* of G . An \mathcal{F}^* -*cloud* of G is a maximal set $\mathcal{W} \subseteq \mathcal{F}^*$ that is connected in $\text{VF}(G)[\mathcal{F}^* \cup V(G)]$, i.e., \mathcal{W} can be written as $\mathcal{W} = X \cap \mathcal{F}^*$, where X is a connected component of $\text{VF}(G)[\mathcal{F}^* \cup V(G)]$.

Instead of finding any packing of odd cycles in G , we will work with packings with a special structure:

Definition 7.11. Let G be a planar graph, embedded in the sphere. Let $\mathcal{F}^* \subseteq \mathcal{F}(G)$ with $\mathcal{F}_{\text{odd}}(G) \subseteq \mathcal{F}^*$ be a set of special faces of G . A set \mathcal{P} is called a *special packing* for G if each $C \in \mathcal{P}$ is either a special face $C \in \mathcal{F}^*$ or C is an odd cycle in G with $V(C) \cap F = \emptyset$ for all $F \in \mathcal{F}^*$, such that any $C, C' \in \mathcal{P}$ are vertex-disjoint, i.e., $V(C) \cap V(C') = \emptyset$.

Note that for $\mathcal{F}^* = \mathcal{F}_{\text{odd}}(G)$ any special packing for G can be transformed into an odd cycle packing (i.e., a set of pairwise vertex-disjoint odd cycles) of the same size by replacing each $F \in \mathcal{F}^* = \mathcal{F}_{\text{odd}}(G)$ by some odd cycle in $E(F)$.

We will show that we can always find a special packing for G and an odd cycle transversal of at most four times the size. Our algorithm works as follows: Observe that in the case where each \mathcal{F}^* -cloud consists of a single face, Lemma 7.9 directly yields a special packing and an odd cycle transversal of at most twice the size. By a simple observation, we still get an odd cycle transversal of at most four times the size of a special packing in the case where any \mathcal{F}^* -cloud \mathcal{W} has packing number $\nu(\mathcal{W}) = 1$. In the case of an \mathcal{F}^* -cloud \mathcal{W} with packing number $\nu(\mathcal{W}) > 1$ however, we will find a special face in \mathcal{W} with few neighbours, “merge” it with all its neighbours (and possibly one additional face) and recurse on the smaller instance. We prove the existence of such a special face in \mathcal{W} in Section 7.3.




Figure 7.2: The left picture shows the faces F_1 and F_2 , which meet at the vertex v , before merging. The edges incident to v are ordered counterclockwise. Note that $j_1 = 3$, while for j_2 there are two options: $j_2 = 6$ and $j_2 = 8$.

The right picture shows the graph obtained after merging F_1 and F_2 along v for the choice $j_2 = 8$. The large face F corresponds to the face set $\{F_1, F_2\}$.

7.2.2 Detailed proof

We first define a simple way to merge a set of faces into one face, without modifying the structure of the graph.

Definition 7.12. Let G be a planar graph, embedded in the sphere. Let F_1, F_2 be faces of G and $v \in V(F_1) \cap V(F_2)$. A graph obtained after *merging* the faces F_1 and F_2 along v can be constructed as follows: Let $e_1, \dots, e_k \in E(G)$ be the edges incident to v , ordered counterclockwise around v (for an example, see Figure 7.2). Let $j_1, j_2 \in \{1, \dots, k\}$ such that F_i lies between e_{j_i} and e_{j_i+1} for $i = 1, 2$ (where we set $k + 1 := 1$). W.l.o.g. $j_1 < j_2$. Now split the vertex v into two vertices v' and v'' , connect the edges $e_{j_1+1}, \dots, e_{j_2}$ to v' and the remaining edges from $\delta_G(v)$ to v'' .

Let G' be the graph obtained by merging the faces F_1 and F_2 along v . It is clear that G' contains exactly one face less than G : All faces except for F_1 and F_2 remain unchanged, and G' contains one other face F with $E(F) = E(F_1) \oplus E(F_2)$. We say that F corresponds to the face set $\{F_1, F_2\}$. Also, all vertices except for v remain unchanged, and we say that both v' and v'' in $V(G')$ correspond to $v \in V(G)$. See Figure 7.2.

Given a set $\mathcal{F} \subseteq \mathcal{F}(G)$ of faces that are connected in $\text{VF}(G)[\mathcal{F} \cup V(G)]$, we can define a graph obtained after merging all faces in \mathcal{F} as follows: Start by setting $\mathcal{F}' := \mathcal{F}$. While $|\mathcal{F}'| > 1$, choose two faces $F_1, F_2 \in \mathcal{F}'$ and merge them via Definition 7.12. Delete both F_1 and F_2 from \mathcal{F}' and add the new face in the constructed graph that corresponds to $\{F_1, F_2\}$ to \mathcal{F}' . In the graph resulting from all these merging operations, all faces except for the faces in \mathcal{F} remain unchanged and the faces from \mathcal{F} are replaced by exactly one additional face that corresponds to \mathcal{F} .

Lemma 7.13. Let G be a planar graph, embedded in the sphere, and $\mathcal{F} \subseteq \mathcal{F}(G)$ a subset of its faces that is connected in $\text{VF}(G)[\mathcal{F} \cup V(G)]$. Let G' be a graph arising from G when merging all faces of \mathcal{F} and let T be an odd cycle transversal for G' . Let $T_G \subseteq V(G)$ be constructed by replacing each $v \in T$ by the vertex in G corresponding to v . Then there exists a subset $\mathcal{F}' \subseteq \mathcal{F}$ of even cardinality such that for any \mathcal{F}' -transversal T' , $T_G \cup T'$ defines an odd cycle transversal for G .

Proof. Given a connected component X of $\text{VF}(G)[\mathcal{F}(G) \cup T_G]$ with $X \cap \mathcal{F} = \emptyset$, clearly also $X \cap V(\mathcal{F}) = \emptyset$. Thus, X is also a connected component of $\text{VF}(G')[\mathcal{F}(G') \cup T]$ and therefore $|X \cap \mathcal{F}_{\text{odd}}(G)|$ is already even.

Thus, we can construct a set $\mathcal{F}' \subseteq \mathcal{F}$ by adding exactly one element from each connected component X of $\text{VF}(G)[\mathcal{F}(G) \cup T_G]$ with $|X \cap \mathcal{F}_{\text{odd}}(G)|$ odd. In particular, each connected component of $\text{VF}(G)[\mathcal{F}(G) \cup T_G]$ contains an even number of elements of $\mathcal{F}_{\text{odd}}(G) \oplus \mathcal{F}'$ and therefore T_G is an $\mathcal{F}_{\text{odd}}(G) \oplus \mathcal{F}'$ -transversal.

Now let $T' \subseteq V(G)$ be an \mathcal{F}' -transversal. According to Proposition 7.7, $T_G \cup T'$ is an $\mathcal{F}_{\text{odd}}(G)$ -transversal because $\mathcal{F}_{\text{odd}}(G) \oplus \mathcal{F}' \oplus \mathcal{F}' = \mathcal{F}_{\text{odd}}(G)$. By Proposition 7.5, $T_G \cup T'$ hence is an odd cycle transversal for G . \square \square

The following two lemmata show structural properties of \mathcal{F}^* -clouds. We defer their proofs to Section 7.3.

Lemma 7.18. *Let G be a planar graph, embedded in the plane. Let $\mathcal{F} \subseteq \mathcal{F}(G)$ with $\nu(\mathcal{F}) = 1$. Then for any $\mathcal{F}' \subseteq \mathcal{F}$ of even cardinality we have $\tau_{\text{odd}}(\mathcal{F}') \leq 2$.*

Lemma 7.19. *Let G be a planar graph, embedded in the plane, and $\mathcal{F}^* \subseteq \mathcal{F}(G)$ with $\mathcal{F}_{\text{odd}}(G) \subseteq \mathcal{F}^*$ a set of special faces. Let \mathcal{W} be an \mathcal{F}^* -cloud of G with $\nu(\mathcal{W}) > 1$. Then there exists a set $\mathcal{F} \subseteq \mathcal{W}$ such that*

- (i) \mathcal{F} is connected in $\text{VF}(G)[\mathcal{F} \cup V(G)]$ and
- (ii) $\nu(\mathcal{F}) > 1$ and
- (iii) there is a face $F \in \mathcal{F}$ such that all $F' \in \mathcal{W}$ where $V(F) \cap V(F') \neq \emptyset$ are also contained in \mathcal{F} and
- (iv) for any even-cardinality subset $\mathcal{F}' \subseteq \mathcal{F}$, we have $\tau_{\text{odd}}(\mathcal{F}') \leq 4$.

Note that item (iv) is stronger than just postulating $\tau_{\text{odd}}(\mathcal{F}) \leq 4$ since τ_{odd} is not monotone. We are now ready to show Theorem 7.1. As explained in Section 7.2.1, we will show an even stronger statement by requiring a special packing:

Theorem 7.14. *Let G be a planar graph, embedded in the plane. Let $\mathcal{F}^* \subseteq \mathcal{F}(G)$ with $\mathcal{F}_{\text{odd}}(G) \subseteq \mathcal{F}^*$ be a set of special faces. There exists a special packing \mathcal{P}^* for G and an odd cycle transversal T^* for G such that $|T^*| \leq 4|\mathcal{P}^*|$.*

Proof. We prove the theorem by induction on the number of faces $|\mathcal{F}(G)|$. Let \mathfrak{W} be the set of \mathcal{F}^* -clouds of G .

As the base case of our induction we consider the case where $\nu(\mathcal{W}) = 1$ for all clouds $\mathcal{W} \in \mathfrak{W}$. Let G' arise from G by merging the faces of \mathcal{W} for each \mathcal{F}^* -cloud $\mathcal{W} \in \mathfrak{W}$. Then, G' contains exactly one face per \mathcal{F}^* -cloud $\mathcal{W} \in \mathfrak{W}$ and one face per non-special face of G . We mark the faces corresponding to the \mathcal{F}^* -clouds as “deadly”.

By applying Lemma 7.9, we can find an odd cycle transversal T for G' and a set \mathcal{P} of pairwise vertex-disjoint odd cycles of G' such that $|T| \leq |\mathcal{P}| + 2|\mathfrak{W}|$ and no cycle in \mathcal{P} contains a vertex of a deadly face of G' . In particular, \mathcal{P} defines a special packing for G' and thus also for G . By adding exactly one face of each $\mathcal{W} \in \mathfrak{W}$ to \mathcal{P} we get a special packing \mathcal{P}^* for G with $|\mathcal{P}^*| = |\mathcal{P}| + |\mathfrak{W}|$.

Let $T_G \subseteq V(G)$ arise by replacing each $v \in T$ by the vertex of $V(G)$ corresponding to v . By applying Lemma 7.13 and Lemma 7.18 to each \mathcal{F}^* -cloud $\mathcal{W} \in \mathfrak{W}$, we can find a set $T_{\mathcal{W}}$ with $|T_{\mathcal{W}}| \leq 2$ for each $\mathcal{W} \in \mathfrak{W}$ such that $T^* := T_G \cup \bigcup_{\mathcal{W} \in \mathfrak{W}} T_{\mathcal{W}}$ is an odd cycle transversal for G with

$$|T^*| \leq |T| + 2|\mathfrak{W}| \leq |\mathcal{P}| + 4|\mathfrak{W}| \leq 4|\mathcal{P}^*|.$$

This concludes the case where $\nu(\mathcal{W}) = 1$ for all $\mathcal{W} \in \mathfrak{W}$.

Next, assume that there exists an \mathcal{F}^* -cloud $\mathcal{W} \in \mathfrak{W}$ with $\nu(\mathcal{W}) > 1$. Choose a subset $\mathcal{F} \subseteq \mathcal{W}$ as given by Lemma 7.19. Let G' arise from G by merging the faces in \mathcal{F} . Note that G' has fewer faces than G because $\nu(\mathcal{F}) > 1$ implies $|\mathcal{F}| > 1$. The set $\mathcal{F}^* \subseteq \mathcal{F}(G')$ of special faces of G' can be constructed from \mathcal{F}^* by deleting the faces in \mathcal{F} and replacing them by the face $F_{\mathcal{F}}$ of G' that corresponds to \mathcal{F} .

By the induction hypothesis, we can find an odd cycle transversal T and a special packing \mathcal{P} for G' such that $|T| \leq 4|\mathcal{P}|$. Let $T_G \subseteq V(G)$ arise by replacing each $v \in T$ by the vertex of G corresponding to v . By the properties of \mathcal{F} and Lemma 7.13 we can find a set $T_{\mathcal{F}} \subseteq V(G)$ with $|T_{\mathcal{F}}| \leq 4$ such that $T^* := T_G \cup T_{\mathcal{F}}$ is an odd cycle transversal for G . We will show how to find a special packing \mathcal{P}^* for G with $|\mathcal{P}^*| = |\mathcal{P}| + 1$. This will finish the proof.

Clearly, $\mathcal{P} \setminus \{F_{\mathcal{F}}\}$ defines a special packing for G . We consider two cases:

Case 1: $F_{\mathcal{F}} \in \mathcal{P}$. Since $\nu(\mathcal{F}) \geq 2$, we can replace $F_{\mathcal{F}}$ by two vertex-disjoint deadly faces of \mathcal{F} to get a special packing for G .

Case 2: $F_{\mathcal{F}} \notin \mathcal{P}$. In this case \mathcal{P} already defines a special packing for G . Also, by Lemma 7.19 (iii) there is a face $F \in \mathcal{F}$ such that all special faces F' with $V(F) \cap V(F') \neq \emptyset$ are also contained in \mathcal{F} . In particular, $\mathcal{P} \cup \{F\}$ is a special packing for G . \square

As noted above, for $\mathcal{F}^* := \mathcal{F}_{\text{odd}}(G)$ any special packing for G can be transformed into an odd cycle packing of the same size. Furthermore, all steps in the proof of Theorem 7.14 are constructive and can be carried out in polynomial time, including embedding G in the plane ([26]). Thus we get Theorem 7.1 as a corollary:

Corollary 7.15. *Let G be an undirected planar graph. Then there exists an odd cycle transversal T and a set \mathcal{P} of pairwise vertex-disjoint odd cycles in G with $|T| \leq 4|\mathcal{P}|$. T and \mathcal{P} can be computed in polynomial time.*

7.3 Structural properties of clouds

The goal of this section is to prove the statements about clouds that we used in our proof of Theorem 7.14. We first define a “reduced” planar version of the conflict graph on the faces of an \mathcal{F}^* -cloud. Note that this construction is quite similar to the graph used in the simple case of our structure lemma in Section 5.2.

Definition 7.16 (Reduced conflict graph). Given a planar graph G with a fixed planar embedding and a set $\mathcal{F} \subseteq \mathcal{F}(G)$ of faces of G , we define the *reduced conflict graph* R for \mathcal{F} as a planar graph on $V(R) := \mathcal{F}$ as follows: For a vertex $v \in V(G)$ let $\mathcal{F}_v \subseteq \mathcal{F}$ be the set of faces containing v and $k := |\mathcal{F}_v|$. The planar embedding of $VF(G)$ induces a cyclic order on $\delta_{VF(G)}(v) = \{\{v, F\} : F \in \mathcal{F}_v\}$. Enumerate the faces in \mathcal{F}_v according to this order as $F_1 =: F_{k+1}, F_2, \dots, F_k$. If $k \geq 2$, add the edges $\{F_i, F_{i+1}\}$ for any $i = 1, \dots, k$ to R with the

obvious planar embedding. Finally, we identify homotopic edges in R (i.e., parallel edges that bound a face of R). For an example see Figure 7.3.




Figure 7.3: Example for a situation as in Definition 7.16, showing the surroundings of the light blue face of a graph G : The colored faces are the faces in $\mathcal{F} \subseteq \mathcal{F}(G)$. The relevant vertices of G are drawn in black. The gray dotted lines are the edges in the vertex-face incidence graph $\text{VF}(G)$ that are incident to \mathcal{F} and the black vertices. Then the reduced conflict graph R for \mathcal{F} is constructed on the colored vertices (each representing a face in \mathcal{F}). The edges of R are the thick dashed lines (black and light and dark green). Note that R depends on the embedding of $\text{VF}(G)$ since there are two possibilities for embedding the edge between v_2 and the orange face. When we choose the light blue face as $F \in \mathcal{F}$ in Lemma 7.17, we get $T_1 = \{v_1, v_2\}$. Hence \mathcal{F}_1 consists of the light blue, the violet, the orange, the dark blue, and the magenta face. \mathcal{F}_2 then consists of the red and the yellow face; we get $T_2 = \{v_3, v_4\}$. The edges in $\delta_R(F)$ that ultimately carry a charge of $\frac{1}{2}$ are colored in light green; the edges with a charge of 1 are colored in dark green: v_1 adds a charge of $\frac{1}{2}$ to each of e_1 and e_2 , v_2 adds a charge of $\frac{1}{2}$ to each of e_2 and e_3 , and v_3 and v_4 add a charge of 1 to e_4 and e_5 , respectively.

The following lemma helps us formalize some intuition about the reduced conflict graph, see also Figure 7.3.

Lemma 7.17. *Let G be a planar graph with a fixed planar embedding, $\mathcal{F} \subseteq \mathcal{F}(G)$ a set of faces of G , and R the reduced conflict graph for \mathcal{F} . Given a face $F \in \mathcal{F}$, there is a subset of vertices $T \subseteq V(F)$ with $|T| \leq |\delta_R(F)|$ such that $T \cap V(F') \neq \emptyset$ for all $F' \in \mathcal{F}$ with $V(F') \cap V(F) \neq \emptyset$.*

Proof. We will define non-negative charges $c(e)$ on the edges $e \in \delta_R(F)$ as follows. Let $T_1 \subseteq V(F)$ be the set of vertices of F that are contained in at least three faces in \mathcal{F} . When constructing the reduced conflict graph R for \mathcal{F} , each $v \in T_1$ produces a face B_v of R that is incident to F in R . For every such face B_v , add a charge of $\frac{1}{2}$ to two different edges in $\delta_R(F)$ that are boundary edges of B_v . Because each edge is only incident to at most two faces, $c(e) \leq 1$ for all $e \in \delta_R(F)$. See also Figure 7.3.

Define $\mathcal{F}_1 := \{F' \in \mathcal{F} : V(F') \cap T_1 \neq \emptyset\}$ and $\mathcal{F}_2 := \{F' \in \mathcal{F} : V(F') \cap V(F) \neq \emptyset\} \setminus \mathcal{F}_1$. Construct $T_2 \subseteq V(F)$ by adding exactly one vertex $v' \in V(F') \cap V(F)$ for each $F' \in \mathcal{F}_2$. Now consider a face $F' \in \mathcal{F}_2$, and the corresponding $v' \in V(F') \cap V(F) \cap T_2$. Since $v' \notin T_1$, F and F' are the only faces in \mathcal{F} containing v' . Hence R contains an edge $e_{v'} = \{F, F'\} \in \delta_R(F)$. Note that $e_{v'}$ is not incident to any face B_v for any $v \in T_1$ because otherwise we would have $v \in V(F')$ for this $v \in T_1$. Therefore $e_{v'}$ has not been charged at all so far. Now v' adds a charge of 1 to $e_{v'}$. In doing so, we still guarantee $c(e) \leq 1$ for all $e \in \delta_R(F)$.

Ultimately we obtain $|\delta_R(F)| \geq \sum_{e \in \delta_R(F)} c(e) = |T_1| + |T_2|$ because every vertex in T_1 and every vertex in T_2 has added a total charge of 1 to the edges in $\delta_R(F)$. But $T := T_1 \cup T_2$ is a vertex set as required in the lemma. \square

We can now prove Lemma 7.18 which we restate here for convenience. Recall that for any set \mathcal{F} of (not necessarily odd) faces $\tau_{\text{odd}}(\mathcal{F})$ denotes the size of a minimum \mathcal{F} -transversal (cf. Definition 7.4).

Lemma 7.18. *Let G be a planar graph, embedded in the plane. Let $\mathcal{F} \subseteq \mathcal{F}(G)$ with $\nu(\mathcal{F}) = 1$. Then for any $\mathcal{F}' \subseteq \mathcal{F}$ of even cardinality we have $\tau_{\text{odd}}(\mathcal{F}') \leq 2$.*

Proof. First, observe that $\nu(\mathcal{F}) = 1$ implies that for any $F, F' \in \mathcal{F}$ we have $V(F) \cap V(F') \neq \emptyset$. Thus, the statement clearly holds if $|\mathcal{F}| \leq 4$: Given an even-cardinality subset $\mathcal{F}' \subseteq \mathcal{F}$, partition \mathcal{F}' into at most two pairs. For each pair, choosing a vertex where the two paired up faces meet yields an \mathcal{F}' -transversal.

So we assume $|\mathcal{F}| \geq 5$ from now on. In this case, we will show something stronger: Call a set $T \subseteq V(\mathcal{F})$ *good* if $\text{VF}(G)[T \cup \mathcal{F}]$ is connected. In particular, a good set is an \mathcal{F}' -transversal for any even-cardinality $\mathcal{F}' \subseteq \mathcal{F}$. We will show the existence of a good set T with $|T| \leq 2$.

Let $G' := \text{VF}(G)[\mathcal{F} \cup V(G)]$ with a planar embedding as in Definition 7.3. If there is no vertex $v \in V(G)$ where more than 3 faces of \mathcal{F} meet, then we can embed the conflict graph of \mathcal{F} , i.e. the complete graph on \mathcal{F} , planarly, which is a contradiction. Thus, choose $v \in V(G)$ with four different faces $F_1, \dots, F_4 \in \mathcal{F}$ such that the edges $\{v, F_1\}, \{v, F_2\}, \{v, F_3\}, \{v, F_4\} \in \delta_{G'}(v)$ are arranged in this order around v (in the planar embedding of G').

Case 1: For some $i \in \{1, 2\}$, F_i and F_{i+2} meet in another vertex $v' \neq v$. W.l.o.g. $i = 2$. We consider the cycle $C = vF_2v'F_4v$ in G' . Both connected components of $\mathbb{R}^2 \setminus C$ contain a face in \mathcal{F} , namely F_1 and F_3 , respectively (cf. Figure 7.4). Thus, any other cycle in \mathcal{F} must contain v or v' in order to intersect the vertex sets of both F_1 and F_3 . In particular, $T := \{v, v'\}$ is good.

Case 2: $V(F_1) \cap V(F_3) = \{v\}$ and $V(F_2) \cap V(F_4) = \{v\}$. We can assume the existence of a face $F \in \mathcal{F}$ with $v \notin V(F)$ because otherwise $T := \{v\}$ is good. For $i = 1, \dots, 4$ let $v_i \in V(G)$ such that F meets F_i in v_i . As in Case 1, the cycle $C_1 = vF_1v_1Fv_3F_3v$ in G' has the property that both connected components of $\mathbb{R}^2 \setminus C_1$ contain a face in \mathcal{F} ,




Figure 7.4: The left part shows the situation of Case 1 in Lemma 7.18 where F_2 and F_4 meet in v' . The cycle C in $\text{VF}(G)$ is drawn dashed. The right part shows the more complicated Case 2. The embeddings of the cycles C_1 (dashed lines) and C_2 (dotted lines) both separate the faces $F_1^*(v_3)$ and $F_2^*(v_4)$ from each other.

namely F_2 and F_4 , respectively. Analogously, F_1 and F_3 are part of different connected components of $\mathbb{R}^2 \setminus C_2$ for the cycle $C_2 = vF_2v_2Fv_4F_4v$ in G' . See Figure 7.4. Let $T_1 := V(C_1) \cap V(G) = \{v, v_1, v_3\}$ and $T_2 := V(C_2) \cap V(G) = \{v, v_2, v_4\}$.

We will show that some two-element subset T of T_1 or T_2 is good. Assume this is false, i.e., for any $i \in \{1, 2\}$ and any $t \in T_i$ there is $F_i^*(t) \in \mathcal{F}$ such that $V(F_i^*(t)) \cap T_i = \{t\}$. W.l.o.g. $F_1^*(v) = F_2$ and $F_2^*(v) = F_1$.

Now we consider $F_1^*(v_3)$. We know that $F_2 = F_1^*(v)$ and $F_1^*(v_3)$ share a vertex, but $v, v_1 \notin V(F_1^*(v_3))$ and $v_1, v_3 \notin V(F_2)$. Hence F_2 and $F_1^*(v_3)$ must lie in the same connected component of $\mathbb{R}^2 \setminus C_1$. Also, $F_1^*(v_3)$ lies in the same connected component of $\mathbb{R}^2 \setminus C_2$ as F_3 because $v_3 \in V(F_1^*(v_3)) \cap V(F_3)$, but $v_3 \notin C_2$.

Analogously, $F_2^*(v_4)$ and F_4 lie in the same connected component of $\mathbb{R}^2 \setminus C_1$, and $F_2^*(v_4)$ and F_1 lie in the same connected component of $\mathbb{R}^2 \setminus C_2$.

Since C_2 separates F_1 from F_3 , and C_1 separates F_2 from F_4 , this implies that $F_2^*(v_4)$ lies in the other connected component of both $\mathbb{R}^2 \setminus C_1$ and $\mathbb{R}^2 \setminus C_2$ than $F_1^*(v_3)$. In particular, a vertex where $F_1^*(v_3)$ and $F_2^*(v_4)$ meet must lie in $V(C_1) \cap V(C_2) \cap V(G) = T_1 \cap T_2 = \{v\}$, a contradiction. \square

We finally prove Lemma 7.19. To this end, we want to find a face $F \in \mathcal{W}$ such that for any even-cardinality subset \mathcal{F}' of the set of F and all its adjacent faces in \mathcal{W} we can find a small \mathcal{F}' -transversal. We will see that a vertex in the reduced conflict graph with small degree and many adjacent faces of degree 3 as guaranteed by Lemma 2.24 can serve as F .

Lemma 7.19. *Let G be a planar graph, embedded in the plane, and $\mathcal{F}^* \subseteq \mathcal{F}(G)$ with $\mathcal{F}_{\text{odd}}(G) \subseteq \mathcal{F}^*$ a set of special faces. Let \mathcal{W} be an \mathcal{F}^* -cloud of G with $\nu(\mathcal{W}) > 1$. Then there exists a set $\mathcal{F} \subseteq \mathcal{W}$ such that*

- (i) \mathcal{F} is connected in $\text{VF}(G)[\mathcal{F} \cup V(G)]$ and
- (ii) $\nu(\mathcal{F}) > 1$ and
- (iii) there is a face $F \in \mathcal{F}$ such that all $F' \in \mathcal{W}$ where $V(F) \cap V(F') \neq \emptyset$ are also contained in \mathcal{F} and
- (iv) for any even-cardinality subset $\mathcal{F}' \subseteq \mathcal{F}$, we have $\tau_{\text{odd}}(\mathcal{F}') \leq 4$.

Proof. Let R be the reduced conflict graph for \mathcal{W} . As R is planar and has no homotopic edges, we distinguish two cases by Lemma 2.24:

Case 1: There is $F \in V(R)$ with $|\delta_R(F)| \leq 4$. Let $\mathcal{N} \subseteq \mathcal{W}$ consist of all faces of \mathcal{W} that have at least one common vertex with F .

If $\nu(\mathcal{N}) > 1$, then $\mathcal{F} := \mathcal{N}$ has the demanded properties: Items (i), (ii) and (iii) hold by construction. Moreover, as $|\delta_R(F)| \leq 4$, we can find a set $T \subseteq V(F)$ with $|T| \leq 4$ such that $T \cap V(F') \neq \emptyset$ for all $F' \in \mathcal{F}$ by Lemma 7.17. Thus, the graph $\text{VF}(G)[\mathcal{F} \cup T]$ is connected. In particular, T is an \mathcal{F}' -transversal for any $\mathcal{F}' \subseteq \mathcal{F}$ where $|\mathcal{F}'|$ is even and therefore $\tau_{\text{odd}}(\mathcal{F}') \leq 4$.

If $\mathcal{N} = \mathcal{W}$ then $\nu(\mathcal{N}) = \nu(\mathcal{W}) > 1$ and we are in the above case with which we already dealt.

So let us now consider the case where $\nu(\mathcal{N}) = 1$ and thus in particular $\mathcal{N} \neq \mathcal{W}$. We take a face $F^* \in \mathcal{W} \setminus \mathcal{N}$ that has a common vertex with some $F' \in \mathcal{N}$ (say $v' \in V(F') \cap V(F^*)$). Define $\mathcal{F} := \mathcal{N} \cup \{F^*\}$. Note that $\nu(\mathcal{F}) > 1$ because $V(F) \cap V(F^*) = \emptyset$. Furthermore, (i) and (iii) hold by construction. It remains to prove (iv):

Let $\mathcal{F}' \subseteq \mathcal{F}$ have even cardinality. If $F^* \notin \mathcal{F}'$, we have $\mathcal{F}' \subseteq \mathcal{N}$. Since $\nu(\mathcal{N}) = 1$, we have $\tau_{\text{odd}}(\mathcal{F}') \leq 2$ by Lemma 7.18.

If $F^* \in \mathcal{F}'$, consider $\mathcal{F}_1 := \mathcal{F}' \oplus \{F^*, F'\}$ and $\mathcal{F}_2 := \{F^*, F'\}$. Since $F^* \notin \mathcal{F}_1$, there is an \mathcal{F}_1 -transversal $T_1 \subseteq V(G)$ with $|T_1| \leq 2$ (again by Lemma 7.18). Moreover $\{v'\}$ is an \mathcal{F}_2 -transversal. Hence $T_1 \cup \{v'\}$ is an $\mathcal{F}_1 \oplus \mathcal{F}_2$ -transversal by Proposition 7.7. But $\mathcal{F}_1 \oplus \mathcal{F}_2 = \mathcal{F}' \oplus \{F^*, F'\} \oplus \{F^*, F'\} = \mathcal{F}'$ and $|T_1 \cup \{v'\}| \leq 3$.

Case 2: There is $F \in V(R)$ with $|\delta_R(F)| = 5$ and F has at least four incident faces in R with degree 3. As in the first case, we consider the set $\mathcal{N} \subseteq \mathcal{W}$ of faces that have at least one common vertex with F . If $\nu(\mathcal{N}) = 1$, we can proceed in exactly the same way as in the situation in Case 1 where $\nu(\mathcal{N}) = 1$ and thus in particular $\mathcal{N} \neq \mathcal{W}$.

Now suppose that we have $\nu(\mathcal{N}) > 1$. We show that then $\mathcal{F} := \mathcal{N}$ has the demanded properties: Items (i), (ii), and (iii) already hold by construction.

Let \mathcal{B} be the set of faces $F' \in \mathcal{F} \setminus \{F\}$ for which there is no edge $\{F, F'\}$ in $E(R)$. First consider the case that $\mathcal{B} \neq \emptyset$. In R , each $F' \in \mathcal{B}$ is incident to a face B of R with $F \in V(B)$ and $\deg(B) \geq 4$. Since there is at most one face with degree ≥ 4 in R that is incident to F , all $F' \in \mathcal{B}$ are incident to the same face B of R .

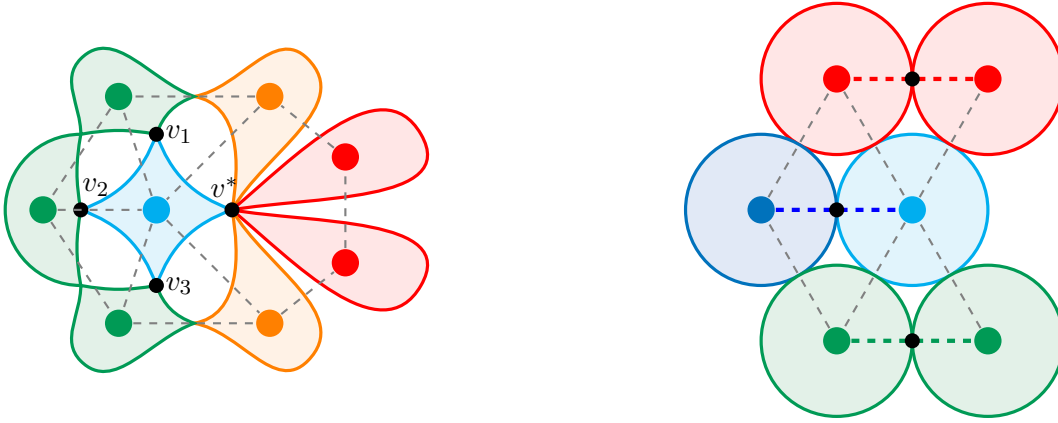


Figure 7.5: Example for two possible situations in Case 2 of the proof of Lemma 7.19. The colored faces are the faces in \mathcal{N} , the light blue face in the middle is F . The dashed lines represent the reduced conflict graph R .

On the left hand side, we see the case where $\mathcal{B} \neq \emptyset$: \mathcal{B} consists of the red faces that clearly contain v^* . F and the orange faces also contain v^* . The only 3 faces that do not contain v^* are the green faces, but they contain v_1 , v_2 , and v_3 , respectively, so $\{v_1, v_2, v_3, v^*\}$ is an \mathcal{F}' -transversal for any \mathcal{F}' .

On the right hand side, we see the case where $\mathcal{B} = \emptyset$ and $|\mathcal{F}'| = |\mathcal{F}| = 6$. R contains a perfect matching on \mathcal{F} as indicated by the thick, colored lines. This induces an \mathcal{F}' -transversal of size 3, indicated by the black vertices.

Choose $v^* \in V(G)$ such that the boundary edges of B are the edges added for v^* in Definition 7.16. In particular, all faces of \mathcal{B} contain v^* . Furthermore, at least three faces of $\mathcal{F} \setminus \mathcal{B}$ (including F) contain v^* . See also Figure 7.5.

Since $|\mathcal{F} \setminus \mathcal{B}| \leq 6$, there are at most 3 faces in \mathcal{F} that do not contain v^* . Hence we can find $v_1, v_2, v_3 \in V(F)$ such that $V(F') \cap \{v_1, v_2, v_3, v^*\} \neq \emptyset$ for each $F' \in \mathcal{F}$. In particular, $\text{VF}(G)[\mathcal{F} \cup \{v_1, v_2, v_3, v^*\}]$ is connected and thus $\tau_{\text{odd}}(\mathcal{F}') \leq 4$ for any subset $\mathcal{F}' \subseteq \mathcal{F}$ of even cardinality.

If however $\mathcal{B} = \emptyset$, there is an edge $\{F, F'\} \in E(R)$ for all faces $F' \in \mathcal{F} \setminus \{F\}$. Then $|\mathcal{F}| \leq 6$. We distinguish again two cases for even subsets $\mathcal{F}' \subseteq \mathcal{F}$:

1. $|\mathcal{F}'| \leq 4$. For each $F' \in \mathcal{F}'$, choose $v_{F'} \in V(F) \cap V(F')$. Then $\{v_{F'} : F' \in \mathcal{F}'\}$ defines an \mathcal{F}' -transversal and $\tau_{\text{odd}}(\mathcal{F}') \leq |\mathcal{F}'| \leq 4$.
2. $|\mathcal{F}'| = 6$. In particular we have $\mathcal{F}' = \mathcal{F}$. Since F has at least 4 incident faces in R with degree 3, R contains a perfect matching M for \mathcal{F} (cf. Figure 7.5). Adding a vertex in $V(F_1) \cap V(F_2)$ for each $\{F_1, F_2\} \in M$ yields an \mathcal{F} -transversal of size 3.

This finishes the proof. □

Chapter 8

Maximum k -systems on the torus

In this chapter we study k -systems on compact surfaces. Given a number $k \in \mathbb{N}$ and a compact surface Σ , a k -system on Σ is a collection of non-trivial and simple closed curves on Σ such that any two are non-homotopic and intersect at most k times. The notion of k -systems is not directly related to the CYCLE PACKING PROBLEM or the CYCLE TRANSVERSAL PROBLEM; however, a crucial step in the constant-factor approximation for the CYCLE PACKING PROBLEM in bounded-genus graphs (cf. Theorem 4.11) is to bound the number of homotopy classes in the support of an uncrossed LP solution. Since the LP solution is uncrossed, the number of homotopy classes can be bounded by the maximum size of a 1-system on the given surface, which is $O(g^2)$ due to a bound by Aougab and Gaster [7].

Despite its fundamental and elementary nature, determining the maximum size of a k -system on Σ , which we denote by $N(\Sigma, k)$, has remained an open problem, even in the case that Σ is the torus \mathbb{T}^2 . In this chapter we determine $N(\mathbb{T}^2, k)$ for every large enough $k \in \mathbb{N}$. In Appendix A we use computer assistance to compute (\mathbb{T}^2, k) also for small k , completely resolving the question about the maximum size of k -systems on the torus. Both results are joint work with Igor Balla, Marek Filakovský, Bartłomiej Kielak, and Daniel Král' [14]. Note that Kriepke and Schymura [56] derived almost the same result as the one presented in this chapter independently and slightly after our result was first published on arXiv.

The maximum size of a k -system on Σ has been the subject of an intensive line of research for various surfaces Σ and values k [2, 3, 4, 5, 7, 44, 49, 61, 68], also see [67] for results concerning a punctured plane. We remark that, as discussed in [44, 68], particularly the case $k = 1$ enjoys having interesting relations including those to the systolic curve complex [77] and Dehn surgery [12]. A priori, it is not clear whether $N(\Sigma, k)$ is even finite; to this end, Juvan, Malnič and Mohar [49] showed that $N(\Sigma, k)$ is finite for every compact surface Σ and every $k \in \mathbb{N}$. When Σ is the closed orientable surface of genus g , Greene [44] showed that $N(\Sigma, k) \leq O(g^{k+1} \log g)$ for any fixed $k \in \mathbb{N}$, improving Przytycki's bound from [68]. If furthermore $k = 1$ then Aougab and Gaster [7] showed $N(\Sigma, 1) = O(g^2)$.

In this chapter we consider the torus \mathbb{T}^2 as our surface Σ ; this arguably simplest case turns out to have surprising connections to number theory, which are presented below.

Juvan, Malnič and Mohar [49] showed that $k + 1 \leq N(\mathbb{T}^2, k) \leq 2k + 3$ and noted that the upper bound can be improved to $\frac{3}{2}k + O(1)$. The connection between this problem and number theory was pointed out by Agol [1], who observed that the size of a k -system on the torus is at most one more than the smallest prime larger than k . This implies that $N(\mathbb{T}^2, k)$ is at most

$(1 + o(1))k$ and specifically, using the bound on the size of prime gaps by Baker, Harman and Pintz [13], $N(\mathbb{T}^2, k)$ is at most $k + O(k^{21/40})$. Cramér [29] showed that a positive resolution of the Riemann hypothesis would yield a bound on prime gaps implying that $N(\mathbb{T}^2, k)$ is at most $k + O(\sqrt{k} \log k)$ and formulated a stronger number-theoretic conjecture that would imply an upper bound of $k + O(\log^2 k)$; we refer to [43, 66] for further discussion including the suspicion that Cramér's error term should actually be $O(\log^{2+\varepsilon} k)$.

Very recently, Aougab and Gaster [6] used combinatorial and geometric arguments in conjunction with estimates from analytic number theory to show that $N(\mathbb{T}^2, k)$, i.e., the maximum size of a k -system on the torus, is at most $k + O(\sqrt{k} \log k)$ (note that this matches Cramér's bound, which is conditioned on a positive resolution of the Riemann hypothesis).

Our main result determines the maximum size of a k -system for every $k \in \mathbb{N}$. Aougab and Gaster also noted that they are not aware of any k -system on the torus whose size exceeds $k + 6$, and our main result shows that there is indeed no k -system whose size exceeds $k + 6$.

Theorem 8.1. *Let K_0 be the set containing the 59 integers listed in Table 8.1. For every $k \in \mathbb{N} \setminus K_0$, it holds that*

$$N(\mathbb{T}^2, k) = \begin{cases} k + 4 & \text{if } k \bmod 6 = 2, \\ k + 3 & \text{if } k \bmod 6 \in \{1, 3, 5\}, \text{ and} \\ k + 2 & \text{otherwise.} \end{cases}$$

The values of $N(\mathbb{T}^2, k)$ for $k \in K_0$ are given in Table 8.1.

Note that there are only four values of k such that $N(\mathbb{T}^2, k) = k + 6$, namely $k \in \{24, 48, 120, 168\}$, and only 13 values such that $N(\mathbb{T}^2, k) = k + 5$.

The proof of Theorem 8.1 is computer assisted. In this chapter we only prove the weaker Theorem 8.16 which can be proven without computer assistance. Theorem 8.16 asserts that $N(\mathbb{T}^2, k) \leq k + 4$ for every sufficiently large $k \in \mathbb{N}$, which implies that $N(\mathbb{T}^2, k) \leq k + O(1)$ for all k . The proof of Theorem 8.1 can be found in Appendix A. The source code used to obtain the computer assisted parts of the proof can be found in Appendix B.

8.1 Overview of the proof

We now provide a brief overview of the proofs of Theorem 8.1 and Theorem 8.16. They are based on the analysis of the geometric and number-theoretic structure of a subset of \mathbb{Z}^2 that describes the homotopy classes of curves contained in a k -system; we introduce the relevant notation and basic properties in Section 8.2. We will call those subsets of \mathbb{Z}^2 corresponding to k -systems on the torus k -nice throughout this chapter. In Section 8.3, we show that the area of the convex hull of any k -nice set is at most $\frac{\pi}{2}k$ and show that the size of any k -nice set with height h is at most $Ck + O(h)$ for a constant $C \in (0, 1)$ (Lemma 8.5); the height of a subset is the smallest h such that it is contained in a strip with width $2h$ centered around the x -axis. This result relies, implicitly, on the fact that the density of coprime points in \mathbb{Z}^2 is asymptotically equal to the Euler product $\prod_p (1 - p^{-2})$ over all primes p , which is equal to $\frac{1}{\zeta(2)} = \frac{6}{\pi^2}$ and, crucially, is strictly smaller than $\frac{2}{\pi}$.

In Section 8.4, we focus on analyzing the height of k -nice sets and the size of k -nice sets with specific height. First, we show that any k -nice set can be modified to a k -nice set of the same

k	1	2	19	23	24	25	33	34	37	47
$N(\mathbb{T}^2, k)$	3	4	23	27	30	30	37	38	42	51
$N(\mathbb{T}^2, k) - k$	+2	+2	+4	+4	+6	+5	+4	+4	+5	+4
“pattern”	+3	+4	+3	+3	+2	+3	+3	+2	+3	+3
k	48	49	53	54	55	61	62	63	64	76
$N(\mathbb{T}^2, k)$	54	54	57	59	60	65	67	67	68	80
$N(\mathbb{T}^2, k) - k$	+6	+5	+4	+5	+5	+4	+5	+4	+4	+4
“pattern”	+2	+3	+3	+2	+3	+3	+4	+3	+2	+2
k	83	84	85	89	90	94	113	114	115	118
$N(\mathbb{T}^2, k)$	87	89	89	93	94	98	117	119	119	122
$N(\mathbb{T}^2, k) - k$	+4	+5	+4	+4	+4	+4	+4	+5	+4	+4
“pattern”	+3	+2	+3	+3	+2	+2	+3	+2	+3	+2
k	119	120	121	124	127	139	141	142	143	144
$N(\mathbb{T}^2, k)$	123	126	126	128	132	143	145	147	147	149
$N(\mathbb{T}^2, k) - k$	+4	+6	+5	+4	+5	+4	+4	+5	+4	+5
“pattern”	+3	+2	+3	+2	+3	+3	+3	+2	+3	+2
k	145	154	167	168	169	174	184	204	208	214
$N(\mathbb{T}^2, k)$	149	158	171	174	174	178	188	208	212	217
$N(\mathbb{T}^2, k) - k$	+4	+4	+4	+6	+5	+4	+4	+4	+4	+3
“pattern”	+3	+2	+3	+2	+3	+2	+2	+2	+2	+2
k	234	244	264	274	294	304	324	354	384	
$N(\mathbb{T}^2, k)$	238	247	268	277	297	307	327	357	387	
$N(\mathbb{T}^2, k) - k$	+4	+3	+4	+3	+3	+3	+3	+3	+3	
“pattern”	+2	+2	+2	+2	+2	+2	+2	+2	+2	

Table 8.1: The values of $N(\mathbb{T}^2, k)$ for $k \in K_0$. The values “pattern” are the additive constants based on $k \bmod 6$ given as in Theorem 8.1, which determines the values of $N(\mathbb{T}^2, k)$ for all sufficiently large k .

size with height at most $\sqrt{2k}$ (Lemma 8.6). The size of a k -nice set with height h is analyzed by a suitable linear program, which yields that every k -nice set with height h has at most $\gamma_h k + \beta_h$ elements for some $\gamma_h \in (0, 1)$ and some constant β_h (Lemma 8.7 and 8.9) whenever $h \geq 4$. In particular, if k is sufficiently large, the size of k -nice set of height $h \geq 4$ is either at most $Ck + O(\sqrt{k}) < k$ by Lemma 8.5 and 8.6 or at most $\gamma_h k + \beta_h < k$ by Lemma 8.7 and 8.9. This line of reasoning is refined using computer assistance to eventually yield that for every $k \geq 1892$, there exists a k -nice set with maximum size that has height at most three (Theorem A.4).

To complete the proof of Theorem 8.1, we determine the maximum size of k -nice sets with height at most three in Section 8.5 and we determine the maximum size of k -nice sets for $k \in \{3, \dots, 1891\}$ with computer assistance in Section A.2. Without computer assistance, it is possible to show a weaker version of Theorem A.4 where 1892 is replaced with a larger constant k_0 (see Lemma 8.10); this yields a computer-free proof that every k -nice set for $k \geq k_0$ has size at most $k + 4$ and so the general upper bound $k + O(1)$ on the size of a k -nice set, which is the main result of this chapter.

8.2 Nice sets and k -systems

We now recall basic results concerning closed simple curves in the torus, which we use later; we refer to e.g. [82] for a more detailed exposition. We view the torus as $\mathbb{R}^2/\mathbb{Z}^2$ and let $C_{m,n}$ for $(m, n) \in \mathbb{Z}^2 \setminus \{(0, 0)\}$, be the closed curve in the torus parameterized as $(m \cdot t \bmod 1, n \cdot t \bmod 1)$ for $t \in [0, 1]$ (cf. Figure 8.1). Every non-trivial, i.e. not null-homotopic, closed curve in the torus is freely homotopic to $C_{m,n}$ for some $(m, n) \in \mathbb{Z}^2 \setminus \{(0, 0)\}$; if the curve is non-self-intersecting, then m and n are coprime, i.e., $\gcd(m, n) = 1$. Note that $(2, 0)$ is not a coprime pair as $\gcd(2, 0) = 2$.




Figure 8.1: The left shows three curves on the torus which form a (maximum) 1-system. It is easy to see that they are homotopic to the curves $C_{1,0}$, $C_{0,1}$ and $C_{1,1}$ as drawn on the right in the universal cover \mathbb{R}^2 of the torus.

Consider coprime pairs $(m, n), (m', n') \in \mathbb{Z}^2 \setminus \{(0, 0)\}$ of integers. The minimum number of crossings of closed simple curves freely homotopic to $C_{m,n}$ and $C_{m',n'}$ is equal to $|mn' - m'n|$, and this minimum is attained by the curves $C_{m,n}$ and $C_{m',n'}$ themselves. This leads us to the following definition: for an integer $k \in \mathbb{N}$, we say that a set $Q \subseteq \mathbb{Z}^2$ is k -nice if

1. Q contains coprime pairs only,

2. Q does not contain both (m, n) and $(-m, -n)$ for any $(m, n) \in \mathbb{Z}^2$, and
3. $|mn' - m'n| \leq k$ for all (m, n) and (m', n') contained in Q .

Note that property 2 already implies $(0, 0) \notin Q$. Since (m, n) and $(-m, -n)$ represent opposite orientations of the same curve, any k -system of simple closed curves on the torus can be represented by a k -nice set $Q \subseteq \mathbb{Z}^2$, and vice versa. Hence, to prove Theorem 8.1, it suffices to determine the maximum size of a k -nice set for every $k \in \mathbb{N}$.

Let $A \in \mathbb{Z}^{2 \times 2}$ be a unimodular matrix, i.e., A has integer entries and $|\det A| = 1$. Observe that if $Q \subseteq \mathbb{Z}^2$ is k -nice, then the set

$$AQ = \{Ax, x \in Q\}$$

is also k -nice. Indeed, since A is unimodular, the inverse A^{-1} is also unimodular and so, if Ax were not coprime, i.e., both coordinates of Ax were divisible by an integer $s > 1$, then both coordinates of $A^{-1}Ax = x$ would also be divisible by s . Likewise, if $x \in Q$ and $y \in Q$, then $|\det(x|y)| = |x_1y_2 - x_2y_1| \leq k$ and so

$$|(Ax)_1(Ay)_2 - (Ax)_2(Ay)_1| = |\det(Ax|Ay)| = |\det A| \cdot |\det(x|y)| = |\det(x|y)| \leq k.$$

Note that the application of the matrix A to the elements of the set $Q \subseteq \mathbb{Z}^2$ corresponds to a reparameterization of the torus. Also note that negating both coordinates of an element of Q does not change what curve it corresponds to. Hence, we say that two subsets Q and Q' of \mathbb{Z}^2 are *equivalent* if there exists an integer unimodular matrix A such that Q' can be obtained from AQ by negating a subset of its elements. Since the operations of multiplying each element by A and negating a subset of elements commute, and A^{-1} is also unimodular, it is not hard to see that being equivalent is indeed an equivalence relation on k -nice sets.

We say that $Q \subseteq \mathbb{Z}^2$ is *x -non-negative* if $m \geq 0$ for all $(m, n) \in Q$, and it is *y -non-negative* if $n \geq 0$ for all $(m, n) \in Q$. The *height* of a set $Q \subseteq \mathbb{Z}^2$ is the maximum $z \in \mathbb{N}$ such that the set Q contains (m, n) with $|n| = z$, and the *width* is the maximum $z \in \mathbb{N}$ such that the set Q contains (m, n) with $|m| = z$.

We conclude this section with the following observation on inclusion-wise maximal k -nice y -non-negative sets.

Proposition 8.2. *Let Q be an inclusion-wise maximal k -nice y -non-negative subset of \mathbb{Z}^2 . If $(m, n) \in \mathbb{Z}^2$ is contained in the convex hull of Q and the integers m and n are coprime, then (m, n) is contained in Q .*

Proof. Fix an inclusion-wise maximal k -nice y -non-negative set $Q \subseteq \mathbb{Z}^2$ and let $(m, n) \in \mathbb{Z}^2$ be contained in the convex hull of Q such that m and n are coprime. Note that $n \geq 0$ as Q is y -non-negative. Also note that since Q is k -nice, it can contain at most one element of the form $(a, 0)$ (the value of a is either -1 or $+1$).

Suppose that a point $(m, n) \in \mathbb{Z}^2$ with m and n being coprime is not contained in Q ; note that $(-m, -n)$ is also not contained in Q as Q is y -non-negative. The inclusion-wise maximality of Q implies that there exists $(m', n') \in Q$ such that $|mn' - m'n| > k$. If $mn' - m'n > k$, then there must exist $(m'', n'') \in Q$ such that $m''n' - m'n'' > k$ (as (m, n) is contained in the convex hull of Q and any linear function on a convex set is maximized at a boundary point), which is impossible. Similarly, if $mn' - m'n < -k$, then there must exist $(m'', n'') \in Q$ such that $m''n' - m'n'' < -k$, which is also impossible. We conclude that (m, n) is contained in Q . \square




Figure 8.2: Visualization of a maximum 24-nice set with 30 elements.

In particular, any maximal k -nice set is given by the set of all coprime integer pairs within some convex area in \mathbb{R}^2 . An example for a k -nice set is given in Figure 8.2. This also shows one of the four examples where $N(\mathbb{T}, k) = k + 6$.

8.3 Nice sets with large height

In this section, we show that, when h is sufficiently large, the size of a k -nice set with height h does not exceed $\gamma k + O(h)$ for some $\gamma \in (0, 1)$. We start with the following auxiliary geometric result, which is a geometric analogue of the upper bound that we wish to prove.

Lemma 8.3. *Let $k \in \mathbb{N}$. If $S \subseteq \mathbb{R}^2$ is a convex set such that $y \geq 0$ for every $(x, y) \in S$ and $|xy' - yx'| \leq k$ for all $(x, y), (x', y') \in S$, then the area of S is at most $\frac{\pi k}{2}$.*

Proof. It is enough to prove the lemma for $k = 1$. Indeed, if $k > 1$, consider the set $S' = \left\{ \left(\frac{x}{\sqrt{k}}, \frac{y}{\sqrt{k}} \right) : (x, y) \in S \right\}$. The set S' satisfies the assumption of the lemma for $k = 1$ and its area is the area of S divided by k .

Fix a convex set $S \subseteq \mathbb{R}^2$ such that $y \geq 0$ for every $(x, y) \in S$ and $|xy' - yx'| \leq 1$ for all $(x, y), (x', y') \in S$. We may assume that S is closed since the closure of S also satisfies the assumption of the lemma. In addition, we may assume that S is bounded: if the area of any convex bounded subset of S is at most $\frac{\pi}{2}$, then the area of S is at most $\frac{\pi}{2}$. Finally, we may assume that $\max\{y : (x, y) \in S\} = 1$ (note that maximum exists as S is bounded and closed): if $\max\{y : (x, y) \in S\} = z \neq 1$, then we can consider $S' = \left\{ \left(xz, \frac{y}{z} \right) : (x, y) \in S \right\}$ instead of S , which has the same area as S and also satisfies the assumptions of the lemma.

We next define two auxiliary functions $f^-, f^+ : [0, 1] \rightarrow \mathbb{R}$ as follows:

$$\begin{aligned} f^-(y) &= \min\{x : (x, y) \in S\}, \\ f^+(y) &= \max\{x : (x, y) \in S\}. \end{aligned}$$

For $y \in [0, 1]$, let $y' = \sqrt{1 - y^2}$; since the points $(f^+(y), y)$ and $(f^-(y'), y')$ are contained in S , we get

$$f^+(y) - \frac{y}{\sqrt{1 - y^2}} f^-(\sqrt{1 - y^2}) \leq \frac{1}{\sqrt{1 - y^2}}.$$

It follows that

$$\int_0^1 f^+(y) dy - \int_0^1 \frac{y}{\sqrt{1 - y^2}} f^-(\sqrt{1 - y^2}) dy \leq \int_0^1 \frac{1}{\sqrt{1 - y^2}} dy = \frac{\pi}{2}. \quad (8.1)$$

We next obtain by substituting that

$$\int_0^1 \frac{y}{\sqrt{1-y^2}} f^-\left(\sqrt{1-y^2}\right) dy = -1 \int_1^0 f^-(y) dy = \int_0^1 f^-(y) dy. \quad (8.2)$$

We combine (8.1) and (8.2) to conclude that the area of S is

$$\int_0^1 f^+(y) - f^-(y) dy = \int_0^1 f^+(y) dy - \int_0^1 \frac{y}{\sqrt{1-y^2}} f^-\left(\sqrt{1-y^2}\right) dy \leq \frac{\pi}{2},$$

which finishes the proof. \square

We next define quantities ρ_ℓ and α_ℓ for every $\ell \in \mathbb{N}$. The values for $\ell \in \{1, \dots, 20\}$ can be found in Table 8.2 in Section 8.4. Appendix B contains the source code of a python script for computing those values (Section B.3). They are defined as follows:

$$\rho_\ell = \prod_{\text{primes } p, p|\ell} \left(1 - \frac{1}{p}\right) \text{ and}$$

$$\alpha_\ell = \max_{1 \leq a \leq b \leq 2\ell} |\{z, a \leq z \leq b \text{ and } \gcd(z, \ell) = 1\}| - \rho_\ell(b - a + 1).$$

The numbers ρ_ℓ and α_ℓ are chosen such that the following Proposition holds:

Proposition 8.4. *Let X be a set of n consecutive integers and let $\ell \in \mathbb{N}$. The number of $x \in X$ that are coprime with ℓ is at most $\rho_\ell n + \alpha_\ell$.*

Proof. We denote the number of $x \in X$ that are coprime with ℓ by $\phi_\ell(X)$. First, note the well-known fact that $\phi_\ell(\{1, \dots, \ell\}) = \phi(\ell) = \ell\rho_\ell$, where $\phi(\ell)$ denotes Euler's totient function. Since any $a \in \mathbb{N}$ is coprime with ℓ if and only if $a + \ell$ is coprime with ℓ , a straightforward induction also shows $\phi_\ell(\{a, \dots, a + \ell - 1\}) = \phi_\ell(\{1, \dots, \ell\}) = \ell\rho_\ell$.

Now consider a set $X = \{a, \dots, a + n - 1\}$ of n consecutive integers. By the above fact we can assume w.l.o.g. that $1 \leq a \leq \ell$. Choose $r, s \in \mathbb{Z}_{\geq 0}$ such that $n = r\ell + s$ and $s < \ell$. Define $X_1 := \{a, \dots, a + s - 1\}$ and $X_2 := X \setminus X_1$. X_2 is a set of $r\ell$ consecutive integers and therefore $\phi_\ell(X_2) = r\ell\rho_\ell$. For X_1 we know that $1 \leq a \leq a + s - 1 \leq 2\ell$ and thus $\phi_\ell(X_1) \leq \rho_\ell s + \alpha_\ell$ by the definition of α_ℓ . We conclude $\phi_\ell(X) \leq \phi_\ell(X_1) + \phi_\ell(X_2) = \rho_\ell n + \alpha_\ell$. \square

We are now ready to prove the main lemma of this section.

Lemma 8.5. *For every $h \geq 41020$ and every $k \geq h$, the maximum size of a k -nice set with height h is at most*

$$\frac{3264\pi}{10255} \cdot k + \frac{4946}{3675} \cdot h + 1.$$

Proof. Set $h_0 = 41020$. Consider a k -nice set Q with height $h \geq h_0$ for some $k \geq h$; without loss of generality, we may assume that $(1, 0) \in Q$ (we use that $h \leq k$) and Q is y -non-negative (by replacing any element (x, y) with $y < 0$ with the element $(-x, -y)$). Let \widehat{Q} denote the convex hull of Q , let s_i and t_i be the minimum and maximum real such that the points (s_i, i) and (t_i, i) are contained in \widehat{Q} for $i = 0, \dots, h$, and let $\ell_i = t_i - s_i$. Note that $\ell_0 = 0$. Finally,

for $i = 1, \dots, h$, let P_i be the set of all $p \in \{2, 3, 5, 7\}$ such that $p|i$ (for instance, $P_{132} = \{2, 3\}$) and let p_i be the product of the elements contained in P_i ; if $P_i = \emptyset$, we set $p_i = 1$.

Since \widehat{Q} is convex, the sequence ℓ_0, \dots, ℓ_h is concave and in particular unimodal, i.e., the values of ℓ_i 's first increase and then decrease; let $m \in \{0, \dots, h\}$ be such that ℓ_m is the maximum element of this sequence. Lemma 8.3 and the convexity of \widehat{Q} imply that

$$\sum_{i=1}^h \ell_i = \frac{\ell_h}{2} + \sum_{i=1}^h \frac{\ell_{i-1} + \ell_i}{2} \leq \frac{\pi}{2} \cdot k + \frac{\ell_m}{2}, \quad (8.3)$$

as the sum in the middle expression is a lower bound on the area of \widehat{Q} . On the other hand, note that \widehat{Q} contains the triangles with corners $(s_m, m), (t_m, m), (s_0, 0)$ and $(s_m, m), (t_m, m), (s_h, h)$, which have a combined area of $\frac{h\ell_m}{2}$. Thus, the area of \widehat{Q} is at least $\frac{h\ell_m}{2}$, which yields using Lemma 8.3 that

$$\ell_m \leq \frac{2}{h} \cdot \frac{\pi}{2} \cdot k \leq \frac{\pi}{h_0} \cdot k. \quad (8.4)$$

Let H be the smallest multiple of $210 = 2 \cdot 3 \cdot 5 \cdot 7$ larger than h and set $\ell_{h+1} = \dots = \ell_H = 0$; note that the sequence ℓ_0, \dots, ℓ_H is unimodal. Define I_a for $a \in \{1, \dots, 210\}$ as the set of all $i \in \{1, \dots, H\}$ such that $i = a \pmod{210}$; note that $|I_a| = \frac{H}{210}$. We next show that the following holds for any $a, b \in \{1, \dots, 210\}$:

$$\left| \sum_{i \in I_a} \ell_i - \sum_{i \in I_b} \ell_i \right| \leq \ell_m. \quad (8.5)$$

By symmetry, we may assume that $a < b$ (the inequality (8.5) trivially holds if $a = b$). Observe that there exists A such that $\ell_{210A+a} \geq \ell_i$ for all $i \in I_a \cup I_b$ or there exists B such that $\ell_{210B+b} \geq \ell_i$ for all $i \in I_a \cup I_b$. The two cases are completely analogous and so we analyze the former case only. Observe that $\ell_{210j+a} \leq \ell_{210j+b}$ for $j \in \{0, \dots, A-1\}$ and $\ell_{210(j+1)+a} \leq \ell_{210j+b}$ for $j \in \{A, \dots, H/210 - 2\}$. We obtain that

$$\sum_{i \in I_a} \ell_i \leq \ell_{210A+a} + \sum_{j=0, j \neq A}^{H/210-1} \ell_{210j+a} \leq \ell_{210A+a} + \sum_{j=0}^{H/210-2} \ell_{210j+b} \leq \ell_m + \sum_{i \in I_b} \ell_i.$$

Likewise, it holds that $\ell_{210j+b} \leq \ell_{210(j+1)+a}$ for $j \in \{0, \dots, A-1\}$ and $\ell_{210j+b} \leq \ell_{210j+a}$ for $j \in \{A, \dots, H/210 - 1\}$, and we obtain that

$$\sum_{i \in I_b} \ell_i = \sum_{j=0}^{H/210-1} \ell_{210j+b} \leq \sum_{j=1}^A \ell_{210j+a} + \sum_{j=A}^{H/210-1} \ell_{210j+a} \leq \ell_{210A+a} + \sum_{i \in I_a} \ell_i \leq \ell_m + \sum_{i \in I_a} \ell_i.$$

Hence, the inequality (8.5) follows. Using (8.3) and (8.5), we obtain that the following holds for every $a \in \{1, \dots, 210\}$:

$$\sum_{i \in I_a} \ell_i \leq \ell_m + \frac{1}{210} \sum_{i=1}^H \ell_i \leq \frac{\pi}{2 \cdot 210} \cdot k + \frac{3\ell_m}{2},$$

which yields using (8.4) that

$$\sum_{i \in I_a} \ell_i \leq \frac{\pi}{420} \cdot k + \frac{3\pi}{2h_0} \cdot k = \frac{\pi}{2} \cdot \left(\frac{1}{210} + \frac{3}{h_0} \right) \cdot k \quad (8.6)$$

Finally, we obtain using (8.6) the following:

$$\begin{aligned} \sum_{i=1}^h \ell_i \prod_{p \in P_i} \left(1 - \frac{1}{p} \right) &= \sum_{a=1}^{210} \left(\sum_{i \in I_a} \ell_i \right) \prod_{p \in P_a} \left(1 - \frac{1}{p} \right) \\ &\leq \frac{\pi}{2} \cdot \left(\frac{1}{210} + \frac{3}{h_0} \right) \cdot k \cdot \sum_{a=1}^{210} \prod_{p \in P_a} \left(1 - \frac{1}{p} \right) \\ &= \frac{\pi}{2} \cdot \left(\frac{1}{210} + \frac{3}{h_0} \right) \cdot k \cdot 210 \cdot \prod_{p=2,3,5,7} \left(1 - \frac{1}{p^2} \right) = \frac{3264\pi}{10255} \cdot k. \end{aligned} \quad (8.7)$$

For $i = 1, \dots, h$, let Q_i be the set of the elements of Q with their second coordinate equal to i . Proposition 8.4 yields that

$$|Q_i| \leq \rho_{p_i}(\ell_i + 1) + \alpha_{p_i} = \rho_{p_i} + \alpha_{p_i} + \ell_i \prod_{p \in P_i} \left(1 - \frac{1}{p} \right) \quad (8.8)$$

since at most $\rho_{p_i}(\ell_i + 1) + \alpha_{p_i}$ integers among $\lceil s_i \rceil, \dots, \lfloor t_i \rfloor$ are coprime with p_i (note that every integer coprime with i is also coprime with p_i). Using computer assistance, we have verified for all $\ell \in \{1, \dots, 210\}$ (the equality is attained for $\ell = 210$) that

$$\sum_{i=1}^{\ell} \rho_{p_i} + \alpha_{p_i} \leq \frac{4946}{3675} \cdot \ell. \quad (8.9)$$

We now combine (8.8) with (8.7) and (8.9) to get that

$$\sum_{i=1}^h |Q_i| \leq \frac{3264\pi}{10255} \cdot k + \sum_{i=1}^{210} \left\lceil \frac{h-i+1}{210} \right\rceil \cdot (\rho_{p_i} + \alpha_{p_i}) \leq \frac{3264\pi}{10255} \cdot k + \frac{4946}{3675} \cdot h,$$

where the first inequality follows since $|\{j : 1 \leq j \leq h \text{ and } p_j = p_i\}| = \lceil \frac{h-i+1}{210} \rceil$. Since $(1, 0)$ is the only element of Q with a zero y -coordinate, the statement of the lemma follows. \square

8.4 Bounding the height of a nice set

In this section, we show that if k is sufficiently large, we may assume that the height of a maximum size k -nice set is at most three. We start with showing a sublinear upper bound.

Lemma 8.6. *For any k -nice set Q , there exists a k -nice set equivalent to Q that has height at most $\sqrt{2k}$.*

Proof. Consider a k -nice set Q and let Q_0 be a set equivalent to Q with the smallest possible height h ; by negating points if needed, we may assume that Q_0 is y -non-negative.

If $h \leq \sqrt{2k}$, we are done. Suppose for contradiction that $h > \sqrt{2k}$, and let (x_0, h) be the point in Q_0 with the maximum second coordinate. By considering the set $A^m Q_0$ for a suitable $m \in \mathbb{Z}$, where A is the matrix

$$A = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix},$$

we can assume that $|x_0| \leq h/2$. Observe that if $(x, y) \in Q_0$, then

$$\left| x - \frac{x_0}{h}y \right| \leq \frac{k}{h},$$

which yields

$$|x| \leq \frac{|x_0|}{h}y + \frac{k}{h} \leq |x_0| + \frac{k}{h} \leq \frac{h}{2} + \frac{k}{h}.$$

We conclude that the width of Q_0 is at most $\frac{h}{2} + \frac{k}{h}$. In particular, the height of the set $A'Q$, where A' is the matrix

$$A' = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

is at most

$$\frac{h}{2} + \frac{k}{h} < \frac{h}{2} + \frac{\sqrt{2k}}{2} < h,$$

which contradicts the choice of Q_0 as a set equivalent to Q that has the smallest possible height. \square

We next consider the following linear program (LP_ℓ) with 2ℓ variables $\sigma_1, \dots, \sigma_\ell$ and τ_1, \dots, τ_ℓ defined as follows:

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^{\ell} \rho_i(\tau_i - \sigma_i) \\ & \text{subject to} && \tau_i \geq \sigma_i \geq 0 && \text{for all } 1 \leq i \leq \ell, \text{ and} \\ & && -1 \leq i\tau_j - j\sigma_i \leq 1 && \text{for all } 1 \leq i, j \leq \ell. \end{aligned}$$

The objective value of the program (LP_ℓ) is denoted by γ_ℓ for $\ell \in \mathbb{N}$; again, the values of γ_ℓ for $\ell \in \{1, \dots, 20\}$ can be found in Table 8.2; we will show that $\gamma_\ell < 1$ for every $\ell \geq 4$ (see Lemma 8.9). Finally, define $\beta_0 = 1$ and $\beta_\ell = \beta_{\ell-1} + \alpha_\ell + \rho_\ell$; again, the values of β_ℓ for $\ell \in \{1, \dots, 20\}$ are in Table 8.2.

We next relate the linear program (LP_ℓ) to the sizes of k -nice sets.

Lemma 8.7. *Let $k \in \mathbb{N}$. For every $h \in \{1, \dots, k\}$, every k -nice set $Q \subseteq \mathbb{Z}^2$ with height exactly h has at most $\gamma_h k + \beta_h$ elements.*

Proof. Fix $k \in \mathbb{N}$ and $h \in \{1, \dots, k\}$. Let Q be a k -nice set with height exactly h . By negating some of the elements of Q , we may assume that Q is y -non-negative, and by considering the set $A^m Q$ for sufficiently large $m \in \mathbb{N}$ (if needed), where

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix},$$

ℓ	ρ_ℓ	α_ℓ	γ_ℓ	β_ℓ
1	1.0000	0.0000	1.0000	2.0000
2	0.5000	0.5000	1.0000	3.0000
3	0.6667	0.6667	1.0000	4.3333
4	0.5000	0.5000	0.9722	5.3333
5	0.8000	0.8000	0.9917	6.9333
6	0.3333	1.0000	0.9667	8.2667
7	0.8571	0.8571	0.9752	9.9810
8	0.5000	0.5000	0.9687	10.9810
9	0.6667	0.6667	0.9695	12.3143
10	0.4000	1.2000	0.9586	13.9143
11	0.9091	0.9091	0.9679	15.7325
12	0.3333	1.0000	0.9601	17.0658
13	0.9231	0.9231	0.9680	18.9120
14	0.4286	1.2857	0.9645	20.6262
15	0.5333	1.3333	0.9605	22.4929
16	0.5000	0.5000	0.9553	23.4929
17	0.9412	0.9412	0.9617	25.3753
18	0.3333	1.0000	0.9576	26.7086
19	0.9474	0.9474	0.9634	28.6033
20	0.4000	1.2000	0.9615	30.2033

Table 8.2: The numerical values of ρ_ℓ , α_ℓ , γ_ℓ and β_ℓ for $\ell \in \{1, \dots, 20\}$.

we can also assume that the set Q is x -non-negative. Finally, since the height of Q is $h \leq k$, this means that $Q \cup \{(1, 0)\}$ is k -nice and we can therefore assume $(1, 0) \in Q$.

Let s_i and t_i for $i = 1, \dots, h$ be the minimum and maximum reals such that (s_i, i) and (t_i, i) are in the convex hull of Q (note that the values are well-defined since $(1, 0) \in Q$ and the height of Q is exactly h). Since any two points (x, y) and (x', y') of Q satisfy $|xy' - x'y| \leq k$, it also holds that $|xy' - x'y| \leq k$ for any two points (x, y) and (x', y') of the convex hull of Q . Since $\sigma_i = s_i/k$ and $\tau_i = t_i/k$, $i = 1, \dots, h$, is a feasible solution of (LP_h) , we obtain that

$$\sum_{i=1}^h \rho_i(t_i - s_i) = k \sum_{i=1}^h \rho_i(\tau_i - \sigma_i) \leq \gamma_h k. \quad (8.10)$$

For $i \in \{0, \dots, h\}$, let Q_i be the set of points contained in Q with their second coordinate equal to i . Note that $|Q_0| = 1$ and

$$|Q_i| \leq \rho_i(t_i - s_i + 1) + \alpha_i = \rho_i(t_i - s_i) + (\alpha_i + \rho_i)$$

for every $i \in \{1, \dots, h\}$. It follows using (8.10) and the definition of β_h that

$$|Q| = \sum_{i=0}^h |Q_i| = \left(\sum_{i=1}^h \rho_i(t_i - s_i) \right) + 1 + \sum_{i=1}^h (\alpha_i + \rho_i) \leq \gamma_h k + \beta_h.$$

This concludes the proof of the lemma. \square

8.4.1 Analysis of the linear program

We now show that the optimal value of the linear program (LP_ℓ) is smaller than 1 for every $\ell \geq 4$. We use a relaxed version, which is denoted by (LP'_ℓ) , to analyze its optimum value. The linear program (LP'_ℓ) for $\ell \in \mathbb{N}$ also has 2ℓ variables $\sigma_1, \dots, \sigma_\ell$ and τ_1, \dots, τ_ℓ and is defined as follows (note that the objective function is the same):

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^{\ell} \rho_i(\tau_i - \sigma_i) \\ & \text{subject to} && \tau_i, \sigma_i \geq 0 && \text{for all } 1 \leq i \leq \ell, \\ & && \frac{\tau_j}{j} - \frac{\sigma_i}{i} \leq \frac{1}{ij} && \text{for all } 1 \leq i, j \leq \ell. \end{aligned}$$

Since any feasible solution of (LP_ℓ) is also a feasible solution of (LP'_ℓ) , we obtain the following.

Lemma 8.8. *For every $\ell \in \mathbb{N}$, the optimum value of (LP_ℓ) is at most the optimum value of (LP'_ℓ) .*

We are now ready to provide an upper bound on the optimum value of the linear program (LP_ℓ) , which is obtained using the linear program (LP'_ℓ) .

Lemma 8.9. *Let $\ell \in \mathbb{N}$. If $\ell \in \{1, 2, 3\}$, then the optimum value of (LP_ℓ) is equal to one, and if $\ell \geq 4$, then the optimum value of (LP_ℓ) is strictly smaller than one.*

Proof. We first show that the optimum value of (LP_ℓ) is at least one if $\ell \in \{1, 2, 3\}$. Indeed, the following feasible solutions of (LP_ℓ)

$$\begin{aligned} \ell = 1 : \quad & \sigma_1 = 0, \tau_1 = 1 \\ \ell = 2 : \quad & \sigma_1 = 0, \tau_1 = 3/4, \sigma_2 = 1/2, \tau_2 = 1 \\ \ell = 3 : \quad & \sigma_1 = 0, \tau_1 = 5/9, \sigma_2 = 1/3, \tau_2 = 7/9, \sigma_3 = 2/3, \tau_3 = 1 \end{aligned}$$

have objective value equal to 1. Hence, it remains to show that the optimum value of (LP_ℓ) is at most one if $\ell \in \{1, 2, 3\}$ and strictly smaller than one if $\ell \geq 4$.

Fix $\ell \in \mathbb{N}$ and consider the dual of the linear program (LP'_ℓ) , which will be denoted by (LD_ℓ) . Recall that any feasible solution of the dual program (LD_ℓ) provides an upper bound on the optimal value of the primal program (LP'_ℓ) and so the value of any feasible solution of (LD_ℓ) is an upper bound on the optimum value of (LP_ℓ) by Lemma 8.8.

We now present the linear program (LD_ℓ) in a form suitable for our analysis. Let $v_\ell \in \mathbb{R}^\ell$ be the vector whose i -th coordinate is equal to $1/i$. Also recall that Euler's totient function $\phi(k)$ is defined as follows: $\phi(1) = 1$ and $\phi(k)$ for $k \geq 2$ is the number of positive integers smaller than k that are coprime with k . The optimization problem of (LD_ℓ) asks for minimizing $v_\ell^T A v_\ell$ over all *non-negative* matrices $A \in \mathbb{R}^{\ell \times \ell}$ such that for each $i \in \{1, \dots, \ell\}$, the sum of the entries of the i -th row is bounded from above by Euler's totient function $\phi(i)$, and the sum of entries of the i -th column is bounded from below by $\phi(i)$, i.e.,

$$\sum_{j=1}^{\ell} A_{i,j} \leq \phi(i) \quad \text{and} \quad \sum_{j=1}^{\ell} A_{j,i} \geq \phi(i).$$

Observe that the entries of the matrix A are the variables of the dual of (LP'_ℓ) , where $A_{i,j}$ is the variable associated with the constraint for i and j in the definition of (LP'_ℓ) .

We illustrate the presentation of (LD_ℓ) by giving examples of two feasible solutions for $\ell = 4$:

$$A = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 2 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix}.$$

Observe that the definition of (LD_ℓ) yields that the sum of the entries of A in the i -th row is at most $\phi(i)$ and the sum of the entries of A in the i -th column is at least $\phi(i)$. This implies the sum of all entries of A is equal to $\phi(1) + \dots + \phi(\ell)$ and so all inequalities in the definition of (LD_ℓ) must hold with equality. It follows that we could assume without loss of generality that the matrix A in (LD_ℓ) is symmetric as if A is a feasible solution of (LD_ℓ) , then $(A + A^T)/2$ is also a feasible solution with the same objective value.

We now find a feasible solution of (LD_ℓ) with objective value equal to one, which will imply that the optimum value of (LP_ℓ) is at most one. We remark that the existence of this solution is implicitly established in the proof of [6, Proposition 4.4]. Let $A_\ell \in \mathbb{R}^{\ell \times \ell}$ be the zero-one matrix such that the entry in the i -th row and j -th column is equal to one iff $i + j \geq \ell + 1$ and

We conclude that the objective value of the feasible solution A_ℓ of (LD_ℓ) is at most one.

For $\ell \geq 4$, we show that A_ℓ can be perturbed to a feasible solution of (LD_ℓ) with objective value strictly smaller than one. Observe that if $\ell \geq 4$, then the entries $(A_\ell)_{\ell-2,\ell-1}$, $(A_\ell)_{\ell-1,\ell-2}$, $(A_\ell)_{\ell-1,\ell}$ and $(A_\ell)_{\ell,\ell-1}$ are equal to one. Let A'_ℓ be the matrix obtained from A_ℓ by subtracting one from each of these four entries, adding one to the entries $(A_\ell)_{\ell,\ell-2}$, $(A_\ell)_{\ell-2,\ell}$, and adding two to the entry $(A_\ell)_{\ell-1,\ell-1}$. For example, if $\ell = 7$, we have the following:

$$A_7 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 \end{bmatrix} \quad \text{and} \quad A'_7 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 2 \\ 0 & 0 & 0 & 0 & 0 & 2 & 0 \\ 1 & 1 & 1 & 1 & 2 & 0 & 0 \end{bmatrix}.$$

Since the row and column sums of the matrices A_ℓ and A'_ℓ are the same, the matrix A'_ℓ is a feasible solution of (LD_ℓ) . The objective value of A'_ℓ is equal to

$$\begin{aligned} v_\ell^T A'_\ell v_\ell &= v_\ell^T A_\ell v_\ell - \frac{2}{(\ell-2)(\ell-1)} - \frac{2}{(\ell-1)\ell} + \frac{2}{(\ell-2)\ell} + \frac{2}{(\ell-1)^2} \\ &= 1 - 2 \left(\frac{1}{\ell-2} - \frac{1}{\ell-1} \right) \left(\frac{1}{\ell-1} - \frac{1}{\ell} \right) < 1. \end{aligned}$$

It follows that A'_ℓ is a feasible solution of (LD_ℓ) with objective value strictly smaller than one. \square

8.4.2 Constant height

Since $\gamma_\ell < 1$ for every $\ell \geq 4$ by Lemma 8.9, Lemma 8.7 asserts that k -nice sets with a fixed height of four or more has size less than k when k is sufficiently large. As we see next, this implies that there always exists a maximum k -nice set of height at most three if k exceeds some constant threshold k_0 . Note that the following lemma only gives a very crude upper bound on the threshold k_0 . In the Appendix A we employ computer assistance to decrease k_0 enough such that we can (again, by computer assistance) directly compute maximum sizes of k -nice sets for $k < k_0$.

Lemma 8.10. *There exists k_0 such that for every $k \geq k_0$ there exists a k -nice set of maximum size that has height at most three.*

Proof. Let $h_0 = 41020$, $C = \frac{3264\pi}{10255}$ and $B = \frac{4946}{3675}$. Note that Lemma 8.5 implies that every k -nice set with height $h \geq h_0$ has size at most $Ck + Bh + 1$. Further, let

$$\Gamma = \max\{\gamma_4, \dots, \gamma_{h_0}\}.$$

Note that Lemma 8.9 implies that $\Gamma \in (0, 1)$. Define k_0 to be

$$k_0 = \max \left\{ \frac{2B^2}{(1-C)^2}, \frac{\beta_{h_0} - 1}{1-\Gamma} \right\}.$$

Fix $k \geq k_0$ and consider a k -nice set $Q \subseteq \mathbb{Z}^2$ of maximum size. Note that $|Q| \geq k + 2$ as the set $\{(1, 0), (0, 1), (1, 1), \dots, (k, 1)\}$ is a k -nice set of size $k + 2$. By Lemma 8.6, we can assume that Q is y -non-negative with height $h \leq \sqrt{2k}$. If $h \geq h_0$, then Lemma 8.5 yields that the size of Q is at most

$$Ck + Bh + 1 \leq Ck + B\sqrt{2k} + 1 \leq Ck + \frac{B\sqrt{2}}{\sqrt{k_0}}k + 1 = Ck + (1 - C)k + 1 = k + 1.$$

If $4 \leq h \leq h_0$, then Lemma 8.7 yields that the size of Q is at most

$$\gamma_h k + \beta_h \leq \Gamma k + \beta_{h_0} \leq \Gamma k + (1 - \Gamma)k_0 + 1 \leq k + 1.$$

Since the size of the set Q is at least $k + 2$, its height h is at most three. □

8.5 Nice sets with height at most three

It is left to consider k -nice sets of height at most three. It turns out that the maximum size of a k -nice set of height at most three is given by the following pattern:

Lemma 8.11. *For every $k \geq 3$, the maximum size of a k -nice set of height at most 3 is*

- $k + 4$ if $k \bmod 6 = 2$,
- $k + 3$ if $k \bmod 6 \in \{1, 3, 5\}$, and
- $k + 2$ otherwise.

Proof. Fix $k \geq 3$ and let $N_{k,h}$ be the maximum size of a k -nice y -non-negative set with height $h \in \{1, 2, 3\}$. The statement of the lemma follows from the following four claims, which we prove next (note that every k -nice set is equivalent to a y -non-negative set with the same height).

Claim 8.12. $N_{k,1} = k + 2$.

Claim 8.13. $N_{k,2} \leq k + 3$, with equality if and only if $k \equiv 1 \pmod{2}$.

Claim 8.14. $N_{k,3} \leq k + 4$, with equality if and only if $k \equiv 2 \pmod{6}$.

Claim 8.15. $N_{k,3} \leq k + 2$ if $k \equiv 0 \pmod{6}$ or $k \equiv 4 \pmod{6}$.

In what follows, when Q is a k -nice y -non-negative set with height h and $i \in \{0, \dots, h\}$, we let $Q_i = \{x : (x, i) \in Q\}$ and, if Q_i is non-empty, we let s_i and t_i denote the smallest and largest elements of Q_i , respectively. Note that $|Q_0| \leq 1$ as Q_0 may only contain either 1 or -1 .

Proof of Claim 8.12. Consider a k -nice y -non-negative set Q with height one. Since the set Q is k -nice, $(s_1, 1) \in Q$ and $(t_1, 1) \in Q$, we obtain that $t_1 - s_1 \leq k$ and so $|Q_1| \leq k + 1$. It follows that $|Q| = |Q_0| + |Q_1| \leq k + 2$. This implies that $N_{k,1} \leq k + 2$. The bound is attained by the set $Q = \{(1, 0)\} \cup \{(x, 1) : 0 \leq x \leq k\}$. □

Proof of Claim 8.13. Consider a k -nice y -non-negative set Q with height two. If Q_1 is empty, we obtain $2t_2 - 2s_2 \leq k$ using that Q is k -nice, $(s_2, 2) \in Q$ and $(t_2, 2) \in Q$. So, the size of Q_2 is at most $k/2 + 1$ and so the size of Q is at most $|Q_0| + |Q_2| \leq 1 + k/2 + 1$, which yields that $|Q| \leq k + 2$. Hence, we can assume that Q_1 is non-empty and so Q contains the points $(s_1, 1)$ and $(t_1, 1)$.

Since the set Q is k -nice, we get that $t_2 - 2s_1 \leq k$ and $2t_1 - s_2 \leq k$, which yields that $s_1 \geq \left\lfloor \frac{t_2 - k}{2} \right\rfloor$ and $t_1 \leq \left\lfloor \frac{s_2 + k}{2} \right\rfloor$. It follows that

$$|Q_1| \leq t_1 - s_1 + 1 \leq \left\lfloor \frac{s_2 + k}{2} \right\rfloor - \left\lfloor \frac{t_2 - k}{2} \right\rfloor + 1.$$

Since the set Q_2 consists of odd numbers only, it follows that Q_2 has at most $\frac{t_2 - s_2}{2} + 1$ elements. We obtain using $|Q_0| \leq 1$ that

$$|Q| \leq \left\lfloor \frac{s_2 + k}{2} \right\rfloor - \left\lfloor \frac{t_2 - k}{2} \right\rfloor + \frac{t_2 - s_2}{2} + 3 \leq k + 3.$$

It follows that the size of Q is at most $k + 3$, with equality only if the above estimate is tight, which can only occur when $s_2 - k$ is divisible by two, i.e., when $k \equiv 1 \pmod{2}$ (recall that s_2 is odd). This establishes that $N_{k,2} \leq k + 3$ if $k \equiv 1 \pmod{2}$, and $N_{k,2} \leq k + 2$ otherwise. Moreover, if $k \equiv 1 \pmod{2}$, then the bound in Claim 8.13 is tight as the set $Q = \{(1, 0), (k, 2)\} \cup \{(x, 1) : 0 \leq x \leq k\}$ is k -nice. \square

Proof of Claim 8.14. Consider a k -nice y -non-negative inclusion-wise maximal set Q with height three. Observe that since Q_3 consists only of numbers that are not divisible by 3, we can compute the size of Q_3 as

$$|Q_3| = \frac{2}{3}(t_3 - s_3) + 1 + \rho, \tag{8.11}$$

where the constant ρ is (recall that neither s_3 nor t_3 is divisible by three)

$$\rho = \begin{cases} 0 & \text{if } t_3 \equiv s_3 \pmod{3}, \\ 1/3 & \text{if } s_3 \equiv 1 \pmod{3} \text{ and } t_3 \equiv 2 \pmod{3}, \text{ and} \\ -1/3 & \text{if } s_3 \equiv 2 \pmod{3} \text{ and } t_3 \equiv 1 \pmod{3}. \end{cases}$$

We first argue that $|Q| \leq k + 2$ unless Q_1 and Q_2 are both non-empty. Since the set Q is k -nice, it holds that $3t_3 - 3s_3 \leq k$ and $2t_2 - 2s_2 \leq k$ (assuming Q_2 is non-empty and so s_2 and t_2 are defined), which implies that $|Q_3| \leq \frac{2}{3}(t_3 - s_3) + 1 + \rho \leq \frac{2}{9}k + \frac{4}{3}$ and $|Q_2| \leq \frac{1}{2}(t_2 - s_2) + 1 \leq \frac{k}{4} + 1$ (note that this estimate also holds if Q_2 is empty). Therefore, if Q_1 is empty, we obtain (using $k \geq 3$) that the set Q has at most $|Q_0| + |Q_2| + |Q_3| \leq \frac{17}{36}k + \frac{10}{3} \leq k + 2$ elements. We now consider the case that Q_1 is non-empty and Q_2 is empty. If $|Q_1| + |Q_3| \leq 4$, then the size of Q is at most $|Q_0| + |Q_1| + |Q_3| \leq 5 \leq k + 2$. If $|Q_1| + |Q_3| \geq 5$, then the convex hull of Q contains a line segment of length at least 1.5 with endpoints of the form $((s_1 + s_3)/2, 2)$ and $((t_1 + t_3)/2, 2)$, so the convex hull contains two points of the form $(m, 2)$, $(m + 1, 2)$ for some integer m . In particular, since Q is inclusion-wise maximal, either $(m, 2)$ or $(m + 1, 2)$ should belong to Q , which contradicts Q_2 being empty.

In the rest of the proof of the claim, we assume that both Q_1 and Q_2 are non-empty. Since the set Q is k -nice, the following inequalities hold:

$$t_3 - 3s_1 \leq k, \quad 3t_1 - s_3 \leq k, \quad 2t_3 - 3s_2 \leq k \quad \text{and} \quad 3t_2 - 2s_3 \leq k.$$

We now define the following shorthand notation:

$$\begin{aligned} a_1 &= \left\lfloor \frac{t_3 - k}{3} \right\rfloor - \frac{t_3 - k}{3}, & b_1 &= \frac{s_3 + k}{3} - \left\lfloor \frac{s_3 + k}{3} \right\rfloor, \\ a_2 &= \left\lfloor \frac{2t_3 - k}{3} \right\rfloor - \frac{2t_3 - k}{3} \quad \text{and} & b_2 &= \frac{2s_3 + k}{3} - \left\lfloor \frac{2s_3 + k}{3} \right\rfloor. \end{aligned}$$

Note that all the four quantities a_1 , b_1 , a_2 and b_2 are non-negative. Moreover, since t_3 is not divisible by three, a_1 and a_2 cannot both be equal to zero, and since s_3 is not divisible by three, b_1 and b_2 cannot both be equal to zero. In other words, a_1 or a_2 is at least $1/3$ and b_1 or b_2 is at least $1/3$. We now obtain the following estimates on s_1 , t_1 , s_2 and t_2 :

$$\begin{aligned} s_1 &\geq \frac{t_3 - k}{3} + a_1, & t_1 &\leq \frac{s_3 + k}{3} - b_1, \\ s_2 &\geq \frac{2t_3 - k}{3} + a_2 \quad \text{and} & t_2 &\leq \frac{2s_3 + k}{3} - b_2. \end{aligned} \tag{8.12}$$

Using these four estimates, we obtain that

$$\begin{aligned} |Q_1| &= t_1 - s_1 + 1 \leq \frac{s_3 + k}{3} - \frac{t_3 - k}{3} + 1 - a_1 - b_1, \\ |Q_2| &= \frac{t_2 - s_2}{2} + 1 \leq \frac{1}{2} \left(\frac{2s_3 + k}{3} - \frac{2t_3 - k}{3} - a_2 - b_2 \right) + 1, \end{aligned}$$

These two estimates on the sizes of Q_1 and Q_2 and the estimate (8.11) now yields that

$$|Q| = |Q_0| + |Q_1| + |Q_2| + |Q_3| \leq k + 4 - \left(a_1 + b_1 + \frac{a_2}{2} + \frac{b_2}{2} \right) + \rho, \tag{8.13}$$

which implies that the size of Q is at most $k + 4$ (recall that $\rho \leq 1/3$).

Suppose that $|Q| = k + 4$. Recall that a_1 or a_2 is at least $1/3$ and b_1 or b_2 is at least $1/3$. Therefore, the right side of (8.13) is smaller than $k + 4$ unless $a_1 = b_1 = 0$, $a_2 = b_2 = 1/3$, $\rho = 1/3$ and all four inequalities in (8.12) are tight. It follows from the definition of ρ that $s_3 \equiv 1 \pmod{3}$ and $t_3 \equiv 2 \pmod{3}$. We next derive that $k \equiv t_3 \equiv 2 \pmod{3}$ (as $a_1 = 0$ and so $k = t_3 + 3s_1$) and $k \equiv s_2 + 1 \equiv 0 \pmod{2}$ (as $a_2 = 1/3$ and so $k = 2t_3 - 3s_2 + 1$; recall that s_2 is odd). We conclude that if $|Q| = k + 4$, then $k \equiv 2 \pmod{6}$.

On the other hand, if $k \equiv 2 \pmod{6}$, we can construct a k -nice set Q with height three and $k + 4$ elements as follows. Let s be the smallest integer larger than $\frac{2}{3}k$ satisfying $s \equiv 1 \pmod{3}$, and consider the k -nice set

$$\begin{aligned} Q &= \{(1, 0)\} \cup \left\{ (0, 1), \dots, \left(\frac{k+s}{3}, 1 \right) \right\} \cup \\ &\quad \left\{ \left(\frac{k+1}{3}, 2 \right), \dots, \left(\frac{k+2s-1}{3}, 2 \right) \right\} \cup \{(s, 3), \dots, (k, 3)\}, \end{aligned}$$

which has $1 + \frac{k+s}{3} + 1 + \frac{k+2s-1-(k+1)}{6} + 1 + \frac{2(k-s+2)}{3} = k + 4$ elements. \square

$k \bmod 6$	$s_3 \bmod 6$	$t_3 \bmod 3$	a_1	b_1	a_2	b_2	c_2	S	ρ
0	1 or 4	1	2/3	1/3	1/3	2/3	1	2	0
0	1 or 4	2	1/3	1/3	2/3	2/3	1	11/6	+1/3
0	2 or 5	1	2/3	2/3	1/3	1/3	0	5/3	-1/3
0	2 or 5	2	1/3	2/3	2/3	1/3	0	3/2	0
4	1 or 4	1	0	2/3	2/3	0	1	3/2	0
4	1 or 4	2	2/3	2/3	0	0	1	11/6	+1/3
4	2 or 5	1	0	0	2/3	2/3	1	7/6	-1/3
4	2 or 5	2	2/3	0	0	2/3	1	3/2	0

Table 8.3: The values of the quantities $a_1, b_1, a_2, b_2, c_2, S = a_1 + b_1 + \frac{a_2}{2} + \frac{b_2}{2} + \frac{c_2}{2}$ and ρ in the cases considered in the proof of Claim 8.15.

Proof of Claim 8.15. Consider a k -nice y -non-negative inclusion-wise maximal set Q with height three. If Q_1 or Q_2 is empty, then Q has at most $k + 2$ elements as argued in the proof of Claim 8.14. We will need a refined version of (8.13). Define c_2 as

$$c_2 = \begin{cases} 1 & \text{if } \left\lfloor \frac{2s_3+k}{3} \right\rfloor \text{ is even, and} \\ 0 & \text{otherwise.} \end{cases}$$

Since t_2 is odd, we can strengthen the estimate (8.12) on t_2 to

$$t_2 \leq \frac{2s_3+k}{3} - b_2 - c_2,$$

which leads to the following stronger version of (8.13):

$$|Q| \leq k + 4 - \left(a_1 + b_1 + \frac{a_2}{2} + \frac{b_2}{2} + \frac{c_2}{2} \right) + \rho. \quad (8.14)$$

The values of the quantities a_1, b_1, a_2, b_2, c_2 and ρ for $k \equiv 0 \pmod{6}$ and $k \equiv 4 \pmod{6}$ and all possible values $s_3 \not\equiv 0 \pmod{3}$ and $t_3 \not\equiv 0 \pmod{3}$ can be found in Table 8.3, where $S = a_1 + b_1 + \frac{a_2}{2} + \frac{b_2}{2} + \frac{c_2}{2}$. Since the value of $S - \rho = a_1 + b_1 + \frac{a_2}{2} + \frac{b_2}{2} + \frac{c_2}{2} - \rho$ in each of the cases is larger than one, it follows that the size of Q is less than $k + 3$, i.e., it is most $k + 2$. \square

We now combine the claims to complete the proof of Lemma 8.11. Since a k -nice y -non-negative set with height 0 can only have at most 1 element, we get $N_k = \max\{N_{k,h} : h \in \{1, 2, 3\}\}$ and the desired result then follows from Claims 8.12, 8.13, 8.14, and 8.15. \square

As a direct consequence we get the main result of this chapter.

Theorem 8.16. *Let N_k for $k \in \mathbb{N}$ be the maximum size of a k -nice set. There exists k_0 such that it holds for every $k \geq k_0$ that*

$$N_k = \begin{cases} k + 4 & \text{if } k \bmod 6 = 2, \\ k + 3 & \text{if } k \bmod 6 \in \{1, 3, 5\}, \text{ and} \\ k + 2 & \text{otherwise.} \end{cases}$$

Proof. Follows directly from Lemma 8.10 and Lemma 8.11. \square

The (computer-aided) proof of the stronger Theorem 8.1 can be found in Appendix A.

Chapter 9

Open Problems

We conclude this thesis with some related questions that are left open. There are many obvious open problems: For only few entries in Tables 1.1 and 1.2 matching lower and upper bounds are known. A natural question asks what the correct integrality gaps and Erdős–Pósa ratios are in the various settings. Note that although the uncrossing property seems to be one key property to derive results both for the CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM, even in the vertex-disjoint case we know that e.g. for the family \mathcal{C}_{all} the special structure of the family has to be exploited to prove the correct Erdős–Pósa ratio.

Chapter 7 gives one of the relatively few examples where the currently best bound on the Erdős–Pósa ratio does not come from multiplying the bounds on the integrality gaps of the cycle packing and cycle transversal LP. Examining relations between integral cycle packings and cycle transversals without the “detour” through the fractional LP solutions, e.g. by directly combining approaches for the two problems, seems to be a promising approach to push the bounds further. On the other hand, however, the current state of the art does not allow us to rule out the possibility that the Erdős–Pósa ratio equals the product of the two integrality gaps for the cycle families in Tables 1.1 and 1.2.

Also in terms of approximability the numbers in the tables are not tight: We are not aware of any APX-hardness results in planar graphs for any of the cycle families, neither for the CYCLE PACKING PROBLEM nor the CYCLE TRANSVERSAL PROBLEM. The PTAS from Section 3.4 might be a first step towards deriving a PTAS for the general CYCLE PACKING PROBLEM for some cycle families.

Our results from Section 4.3 and 6 answer the natural question about the Erdős–Pósa property for many uncrossable cycle families on graphs that are embedded in more complex surfaces than a sphere. However, they also raise new questions: First, the bound of $O(g^2)$ in Theorem 4.11 does not seem optimal. The best known lower bound on the integrality gap of the LPs (1.1) and (1.3) is $\Omega(g)$ [48]. Note that for the separating cycles we already get a bound of $O(g)$, so only the case of the non-separating cycles remains.

For the CYCLE TRANSVERSAL PROBLEM in bounded-genus graphs our bound of $O(g)$ on the integrality gap seems natural, but we are not aware of any matching lower bound example. Also, our bounds here only hold for the sub-class of subgraph-defined uncrossable cycle families; and Sun [83] gave a similar bound for the family $\overrightarrow{\mathcal{C}}_{\text{all}}$. For arbitrary uncrossable cycle families not even the Erdős–Pósa property has been established. The most notable uncrossable family where this is open is the family $\mathcal{C}[l]$ of all shortest cycles from some uncrossable family \mathcal{C} w.r.t.

edge lengths l . Note that Theorem 6.25 reduces the problem of bounding the integrality gap by a constant to the case without any facial cycles.

Finally, there are some open questions that are raised directly from the results, proofs and methods in this thesis. First, the oracles that we review in Section 2.4 seem to be natural ways to implement access to the cycle family \mathcal{C} ; they are even sufficient to implement the algorithms for the CYCLE TRANSVERSAL PROBLEM by [41] and [16]. We do not know whether one of them implies the other; it could even hold that any uncrossable cycle family has these oracles.

In Section 2.6 we show how to “uncross” integral and even fractional sets of cycles among some uncrossable cycle family. To this end, Lemma 2.31 shows how to compute an almost optimum uncrossed LP solution, but it is not known how to compute an optimum uncrossed LP solution in polynomial time.

The Efficient Cycle Lemma in Chapter 4 is tight, but a slightly less general version might be improved to give a bound of 4 (cf. Section 4.2.3). This would not yield better bounds on the laminar cycle packing integrality gap than our results from Chapter 5, but it would immediately carry over to weighted cycle packing (cf. Section 4.4). Alternatively, it is still open to apply any of the results from Chapter 5 to weighted cycle packing. Note that this is not trivial because making the support of the LP solution in Algorithm 3 structured might lose LP value in the weighted setting.

In Section 6.2 we extend Goemans and Williamson’s [41] primal-dual approach for the planar CYCLE TRANSVERSAL PROBLEM to the facial cycles of an uncrossable cycle family in a bounded-genus graph. This is particularly interesting as the family of facial cycles of an uncrossable cycle family is in general not uncrossable any more. This raises the question whether the framework from [41] works in a more general setting than only for uncrossable cycle families. Currently, only very little is known about the CYCLE PACKING PROBLEM and the CYCLE TRANSVERSAL PROBLEM for cycle families that are not uncrossable.

In Chapter 8 we consider k -systems on the torus of maximum size. The case of the torus is completely resolved by our result (together with the results from Appendix A). For surfaces of higher genus $g > 1$ however, there is a gap between the lower bound of $\Omega(g^{k+1})$ and the upper bound of $O(g^{k+1} \log g)$ (both due to Greene [44]). This gap was very recently closed (up to a constant factor) by Aougab and Gaster [7] for the case $k = 1$, the remaining case $g > 1$ and $k > 1$ is still open. Note that for our application (in Theorem 4.11) not only a bound on the size of a maximum 1-system would be helpful: Given a 1-system on some surface, also the existence of a large subset of pairwise non-intersecting curves would be interesting (see [48] for an application of such a result to the DISJOINT PATHS PROBLEM).

Appendix A

Computer-assisted improvement over Theorem 8.16

In this section we prove Theorem 8.1 from Chapter 8. We use computer assistance to verify several inequalities and to directly compute the maximum size of a k -nice set for small k . The code of the programs we used can be found in Appendix B.

A.1 Decreasing the threshold for maximum height-3 sets

Lemma 8.10 shows the existence of a threshold k_0 such that all $k \geq k_0$ allow for a maximum-size k -nice set of height at most three, for which we can compute the exact size due to Lemma 8.11. It is left to compute maximum k -nice sets for $k < k_0$. To be able to do so we first need to decrease the threshold k_0 as much as possible, which can again be done using computer assistance.

Theorem A.1. *For every $k \geq 3225$, there exists a k -nice set of maximum size that has height at most three.*

Proof. Let Q be a k -nice set of maximum size. Let h_0 be the height of Q ; by Lemma 8.6, we may assume that $h_0 \leq \sqrt{2k}$. If Q has $k + 2$ elements, then there is nothing to prove since the set $\{(1, 0), (0, 1), (1, 1), \dots, (k, 1)\}$ is a k -nice set with $k + 2$ elements with height one. Hence, we will assume that Q has at least $k + 3$ elements.

We now show that if $h_0 \geq 4$, then $|Q| < k + 3$, a contradiction. To do so, we need to distinguish three cases based on the size of h_0 .

- Case $h_0 \geq 41020$. It is straightforward to verify that

$$\frac{3264\pi}{10255} \cdot \frac{h_0^2}{2} + \frac{4946}{3675} \cdot h_0 + 1 < \frac{h_0^2}{2} + 3.$$

In particular, since $k \geq h_0^2/2$ and $\frac{3264\pi}{10255} < 1$, we obtain using Lemma 8.5 that

$$|Q| \leq \frac{3264\pi}{10255} \cdot k + \frac{4946}{3675} \cdot h_0 + 1 < k + 3.$$

- Case $h_0 \in \{81, \dots, 41019\}$. We have verified with computer assistance that

$$\gamma_h \cdot \frac{h^2}{2} + \beta_h < \frac{h^2}{2} + 3$$

for every $h \in \{81, \dots, 41019\}$. Since $k \geq h_0^2/2$ and $\gamma_{h_0} < 1$ (by Lemma 8.9), we obtain using Lemma 8.7 that

$$|Q| \leq \gamma_{h_0}k + \beta_{h_0} \leq \gamma_{h_0} \frac{h_0^2}{2} + \left(k - \frac{h_0^2}{2}\right) + \beta_{h_0} < k + 3.$$

- Case $h_0 \in \{4, \dots, 80\}$. We have verified with computer assistance that $3225\gamma_h + \beta_h < 3225 + 3$ for every $h \in \{4, \dots, 80\}$. Since $\gamma_{h_0} < 1$ (by Lemma 8.9), we obtain using Lemma 8.7 that

$$|Q| \leq \gamma_{h_0}k + \beta_{h_0} \leq \gamma_{h_0}3225 + (k - 3225) + \beta_{h_0} < k + 3.$$

It follows that the height h_0 of the set Q is at most three. □

We remark that the bound of 3225 in Theorem A.1 is the best possible in the setting of the proof: consider $k = 3224$, $h = 80 \leq \sqrt{2k}$ and note that $\gamma_h k + \beta_h \approx 3227.039$, so Lemma 8.7 does not rule out the existence of a k -nice set of size $k + 3 = 3227$ and height $h = 80$.

Our next aim is to improve the bound in Lemma 8.6 when $k \leq 3224$, which will eventually lead to an improved version of Theorem A.1. We do so with computer assistance employing Algorithm 4, which is analyzed in the next lemma.

Lemma A.2. *If Algorithm 4 for input $k \in \mathbb{N}$ and $h \in \mathbb{N}$, $2 \leq h \leq k$, returns *verified*, then the following statement is true: every k -nice set with height h is equivalent to a k -nice set with height less than h .*

Proof. Fix $k \in \mathbb{N}$ and $h \in \mathbb{N}$, $h \leq k$, such that Algorithm 4 returned *verified*. Let Q be a k -nice set with height h . By negating a subset of the elements of Q , we may assume that Q is y -non-negative. Since the height of Q is h , there exists a point $(x_0, h) \in Q$. By considering the set $A^m Q$ for some $m \in \mathbb{Z}$ instead of Q , where A is the matrix

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix},$$

we may assume that $|x_0| \leq h/2$; by applying symmetry along the axis $x = 0$ to Q if needed, we can impose $0 < x_0 \leq h/2$.

First, observe that any point $(x, y) \in Q$ with $x < 0$ and $y \geq 1$ satisfies

$$|x| \leq \frac{k}{h} - \frac{x_0}{h} \cdot y \leq \frac{k}{h} - \frac{x_0}{h} = \frac{k - x_0}{h} < h,$$

where the last inequality holds since the main x_0 -loop in Algorithm 4 did not return *not verified* when testing $h \leq \frac{k-x_0}{h}$. If the set Q contains no point $(x, y) \in Q$ with $x \geq h$, then the width of Q is less than h and so the set $A'Q$, where A' is the matrix

$$A' = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

Algorithm 4: Algorithm that verifies the following: every k -nice set with height h is equivalent to a k -nice set with height less than h .

Input: positive integers $k \in \mathbb{N}$ and $h \in \mathbb{N}$ such that $2 \leq h \leq k$

Output: *verified* or *not verified*

```

 $x_0$ -loop for  $x_0 = 1$  to  $\lfloor h/2 \rfloor$  do
  if  $\gcd(x_0, h) \neq 1$  then continue
  if  $h \leq \frac{k-x_0}{h}$  then return not verified
  for  $y = 1$  to  $h$  do
     $x$ -loop for  $x = h$  to  $\lfloor \frac{x_0 y + k}{h} \rfloor$  do
      Continue 0 if  $\gcd(x, y) \neq 1$  then continue
       $z = \min\{y, x - y\}$ 
      Continue 1 if  $z + \frac{k}{x} < h$  then continue
       $w = 1$ 
      for  $y' = 1$  to  $h$  do
        for  $x' = \lfloor \frac{y'(x_0-h)-k}{h} \rfloor$  to  $\lfloor \frac{y'(x_0-h)+k}{h} \rfloor$  do
          if  $\gcd(x', y') \neq 1$  then continue
          if  $|x'y - (x - y)y'| > k$  then continue
          if  $|x'| > w$  then  $w = |x'|$ 
        end
      end
      Continue 2 if  $w < h$  then continue
      return not verified
    end
  end
end
return verified

```

is a set equivalent to Q with height at most $h - 1$. Hence, the set Q contains a point $(x, y) \in Q$ with $x \geq h$, and choose such a point so that x is as large as possible. Note that the point (x, y) is also a point with the largest first coordinate in the absolute value (as the first coordinate of all points with a negative first coordinate is greater than $-h$).

Since Algorithm 4 did not return *not verified*, the x -loop for (x, y) executed either **Continue 0**, **Continue 1** or **Continue 2**. As $(x, y) \in Q$, we must have $\gcd(x, y) = 1$ and so either **Continue 1** or **Continue 2** was executed in the x -loop.

We first analyze the case where **Continue 1** was executed in the x -loop. Consider the set Q' obtained from $A'Q$ by negating all points with negative second coordinate. Note that the set Q' contains the point $(-y, x)$ and $y \leq h \leq x$. The choice of the point (x, y) implies that $(-y, x)$ is a point with the largest second coordinate in Q' . We distinguish two cases: $y \leq x - y$ and $y > x - y$.

- If $y \leq x - y$, then it holds that $y + \frac{k}{x} < h$ (because **Continue 1** was executed). Since the set Q' is k -nice, the first coordinate of any point (x', y') contained in Q' has absolute value at most

$$|x'| \leq \frac{y}{x}y' + \frac{k}{x} \leq y + \frac{k}{x} < h.$$

It follows that the height of the set $A'Q'$, which is equivalent to the original set Q , is less than h .

- If $y > x - y$, we consider the set AQ' , which contains the point $(x - y, x)$. As **Continue 1** was executed and $x - y < y$, it holds that $(x - y) + \frac{k}{x} < h$. Since the set AQ' is k -nice, the first coordinate of any point (x', y') contained in AQ' has absolute value at most

$$|x'| \leq \frac{x - y}{x}y' + \frac{k}{x} \leq x - y + \frac{k}{x} < h,$$

and we conclude that the height of the set $A'AQ'$, which is equivalent to the original set Q , is less than h .

The analysis of the case where **Continue 1** was executed is now finished.

We next analyze the case where **Continue 2** was executed in the x -loop. Since

$$A^{-1} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix},$$

observe that the set $A^{-1}Q$ contains the points $(x_0 - h, h)$ and $(x - y, y)$, so that any point $(x', y') \in A^{-1}Q$ satisfies $\gcd(x', y') = 1$,

$$\frac{y'(x_0 - h) - k}{h} \leq x' \leq \frac{y'(x_0 - h) + k}{h} \quad \text{and} \quad |x'y - (x - y)y'| \leq k.$$

Since **Continue 2** was executed in the x -loop, it follows that $|x'| < h$ for all points $(x', y') \in A^{-1}Q$, i.e., the width of $A^{-1}Q$ is at most $h - 1$ and so the height of the set $A'A^{-1}Q$ is at most $h - 1$. \square

We are now ready to improve Lemma 8.6 for $k \in \{2, \dots, 3224\}$. We remark that the multiplicative constant $\sqrt{4/3}$ obtained in the following lemma cannot be improved in general as there exists a 3-nice set with 6 elements and height 2, but every 3-nice set with height one has at most 5 elements.

Lemma A.3. *Let $k \in \{2, \dots, 3224\}$. For every k -nice set Q , there exists a k -nice set equivalent to Q that has height at most $\sqrt{4k/3}$.*

Proof. We executed Algorithm 4 for all $k \in \{2, \dots, 3224\}$ and all $h \in \mathbb{N}$ such that $\sqrt{4k/3} < h \leq \sqrt{2k}$, and Algorithm 4 always returned *verified*. The implementation of Algorithm 4 is given in Section B.2 of Appendix B.

Consider a k -nice set Q and let Q_0 be a set equivalent to Q with the smallest possible height h . By Lemma 8.6, $h \leq \sqrt{2k}$. If $h > \sqrt{4k/3}$, we obtain a contradiction by Lemma A.2 applied with k and h . It follows that $h \leq \sqrt{4k/3}$. \square

We are now ready to prove the main theorem of this section.

Theorem A.4. *For every $k \geq 1892$, there exists a k -nice set of maximum size that has height at most three.*

Proof. Let Q be a k -nice set of maximum size. We can assume that $|Q| \geq k + 3$; if $|Q| = k + 2$, then the set $\{(1, 0), (0, 1), (1, 1), \dots, (k, 1)\}$ is an example of a k -nice set with height one. If $k \geq 3225$, the statement of the theorem follows from Theorem A.1, so we henceforth assume that $k \in \{1892, \dots, 3224\}$. By Lemma A.3, we may assume that the height h_0 of Q is at most $\sqrt{4k/3}$, so that $h_0 \leq 66$.

As in the proof of Theorem A.1, we now distinguish two cases based on the size of h_0 in order to show that if $h_0 \geq 4$, then $|Q| < k + 3$, a contradiction.

- Case $h_0 \in \{51, \dots, 66\}$. We have verified with computer assistance that

$$\gamma_h \cdot \frac{3h^2}{4} + \beta_h < \frac{3h^2}{4} + 3$$

for every $h \in \{51, \dots, 66\}$. Since $k \geq 3h_0^2/4$ and $\gamma_{h_0} < 1$ (by Lemma 8.9), we obtain using Lemma 8.7 that the size of Q is at most $\gamma_{h_0}k + \beta_{h_0} < k + 3$.

- Case $h_0 \in \{4, \dots, 50\}$. We have verified with computer assistance that $1892\gamma_h + \beta_h < 1892 + 3$ for every $h \in \{4, \dots, 50\}$. Since $\gamma_{h_0} < 1$ (by Lemma 8.9), we obtain using Lemma 8.7 that the size of Q is at most $\gamma_{h_0}k + \beta_{h_0} < k + 3$.

It follows that $h_0 \leq 3$, which completes the proof of the theorem. \square

Again, the bound of 1892 in Theorem A.4 is best possible in the setting of the proof, since for $k = 1891$ and $h = 50 \leq \sqrt{4k/3}$, we obtain that $\gamma_h k + \beta_h \approx 1894.036$.

A.2 Our algorithm for small k

In this section, we present our algorithm for computing the maximum size of a k -nice set, which we use for computing the maximum size of a k -nice set for $k \in \{3, \dots, 1891\}$. Before we do so, we need to establish the following lemma, which concerns the structure of a k -nice set that we may assume in the algorithm.

Lemma A.5. *Let $Q \subseteq \mathbb{Z}^2$ be an inclusion-wise maximal k -nice set with height h , $1 \leq h \leq k$. There exists an equivalent k -nice set Q' with height h such that*

- $(1, 0) \in Q'$, $(0, 1) \in Q'$ and $(1, 1) \in Q'$, and
- $Q' \subseteq \{0, \dots, k\} \times \{0, \dots, h\}$.

Proof. Consider an inclusion-wise maximal k -nice set Q with height h , $h \leq k$. We may assume that Q is y -non-negative by negating a subset of its elements. Since the height of Q is at most k and Q is an inclusion-wise maximal k -nice set, the set Q contains one of the points $(1, 0)$ or $(-1, 0)$; by negating the point if needed, we may assume that $(1, 0) \in Q$.

Since the set Q is y -non-negative, there exists $m \in \mathbb{Z}$ such that the set $A^m Q$ is x -non-negative, where A is the matrix

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix};$$

choose the smallest (possibly negative) $m \in \mathbb{Z}$ with this property. Note that m is well-defined since the set Q contains at least one element with positive second coordinate (otherwise, Q would not be an inclusion-wise maximal k -nice set).

Set $Q' = A^m Q$. Since Q is an inclusion-wise maximal k -nice set, the set Q' is also an inclusion-wise maximal k -nice set. The choice of m implies that there exists $(x_0, y_0) \in Q'$ such that $x_0 - y_0 < 0$ (otherwise, the set $A^{m-1} Q$ would also be x -non-negative). We claim that the first coordinate of any point in Q' is at most k . Suppose that Q' contains a point (x, y) such that $x > k$; note that y is at most k as the height of Q is at most k . It follows (note that $y \leq k < x$ and $x_0 < y_0$) that

$$|x_0 y - y_0 x| = x y_0 - x_0 y \geq x(x_0 + 1) - x_0 k = x + x_0(x - k) \geq x > k,$$

which is impossible since the set Q' is k -nice. We conclude that $Q' \subseteq \{0, \dots, k\} \times \{0, \dots, h\}$. Finally, since the set Q' is an inclusion-wise maximal k -nice set, it contains the point $(0, 1)$ (because $|x| \leq k$ for every $(x, y) \in Q'$) and the point $(1, 1)$ (because $|y - x| \leq k$ for every $(x, y) \in Q'$). \square

We are now ready to present and analyze the recursive algorithm for computing the maximum size of a k -nice set with given height.

Lemma A.6. *Let $k \in \mathbb{N}$ and $h \in \mathbb{N}$ such that $2 \leq h \leq k$. For every $N \in \mathbb{N}$, the Algorithm 5 returns*

- *the maximum size of a k -nice set with height h if there exists a k -nice set with height h with more than N elements, and*
- *the value N , otherwise.*

Proof. The Algorithm 5 calls the recursive procedure `backtrack` which is also presented in Algorithm 5. Note that the recursive calls within the procedure `backtrack` are made only on the lines marked as `Call 1` and `Call 2`. We observe that the parameters k and h during the recursive calls of the procedure `backtrack` never change, and the parameter ℓ always decreases by one. In addition, the main body of the procedure is executed only if $\ell > 0$. We fix $k \in \mathbb{N}$ and $h \in \mathbb{N}$, $2 \leq h \leq k$, for the rest of the proof.

We now prove by induction on the value of ℓ that the return value of the procedure `backtrack` is always at least the parameter N that it was called with; moreover, the value of N

Algorithm 5: Backtracking algorithm to compute the size of a maximum k -nice set.

Input: Values $k, h, N \in \mathbb{N}$ with $2 \leq h \leq k$

Output: The size of a maximum k -nice set if it is larger than N , otherwise N

$L[1] = 0, \quad U[1] = k$

for $i = 2$ **to** h **do**

 | $L[i] = 1, \quad U[i] = k$

end

return $\text{backtrack}(k, h, N, h, 0, L[1..h], U[1..h])$

procedure $\text{backtrack}(k, h, N, \ell, N_{>}, L[1..\ell], U[1..\ell])$

$M[i][a][b] = |\{z, a \leq z \leq b \text{ and } \gcd(z, i) = 1\}|$

$M_k[i][a][b] = \max_{a \leq a' \leq b' \leq b, |b' - a'| \leq k/i} M[i][a'][b']$

Update N

if $\ell = 0$ **then return** $N_{>} + 1$

if $\ell < h$ **then**

 | $N' = N_{>} + 1$

for $i = 1$ **to** $\ell - 1$ **do**

 | **if** $L[i] \leq U[i]$ **then** $N' = N' + M_k[i][L[i]][U[i]$

end

Call 1

if $N' > N$ **then** $N = \text{backtrack}(k, h, N, \ell - 1, N_{>}, L[1..\ell - 1], U[1..\ell - 1])$

end

if $L[\ell] > U[\ell]$ **then return** N

for $a = L[\ell]$ **to** $U[\ell]$ **do**

for $b = a$ **to** $U[\ell]$ **do**

 | **if** $\gcd(a, \ell) \neq 1$ **or** $\gcd(b, \ell) \neq 1$ **then continue**

 | **if** $(b - a)\ell > k$ **then continue**

 | $N' = N_{>} + M[\ell][a][b] + 1$

for $i = 1$ **to** $\ell - 1$ **do**

 | $L'[i] = \max \left\{ L[i], \left\lfloor \frac{ai - k}{\ell} \right\rfloor, \left\lfloor \frac{bi - k}{\ell} \right\rfloor \right\}$

 | $U'[i] = \min \left\{ U[i], \left\lfloor \frac{ai + k}{\ell} \right\rfloor, \left\lfloor \frac{bi + k}{\ell} \right\rfloor \right\}$

 | **if** $L'[i] \leq U'[i]$ **then** $N' = N' + M_k[i][L'[i]][U'[i]$

end

if $N' \leq N$ **then continue**

 | $N_{\geq} = N_{>} + M[\ell][a][b]$

Call 2

 | $N = \text{backtrack}(k, h, N, \ell - 1, N_{\geq}, L'[1..\ell - 1], U'[1..\ell - 1])$

end

end

return N

end

never decreases during the execution of the procedure `backtrack`. If $\ell = 0$, the procedure just returns $N_{>} + 1$ on the line marked `Update N`. The procedure `backtrack` was called from the instance for $\ell = 1$ either at the line marked `Call 1` or at the line marked `Call 2`. In the former case, it held that $N' = N_{>} + 1$ is larger than N (note the `if` condition just before `Call 1`) and so the return value of the instance of the procedure for $\ell = 0$ is larger than N . In the latter case, it held that $N' = N_{>} + M[\ell][a][b] + 1$ is larger than N (note the `if` condition before `Call 2`) and since the instance of the procedure for $\ell = 0$ was called with $N_{\geq} = N_{>} + M[\ell][a][b]$, its return value is larger than N . If $\ell > 0$, the value of N is only affected by recursive calls of the procedure `backtrack`, however, they never return a smaller value of N than the one that they were called with by induction. We conclude that the return value of the procedure `backtrack` is always at least the parameter N . Moreover, if the return value is larger, then the return command on the line marked `Update N` was executed during the recursion and the return value is actually equal to the largest value ever returned on the line marked `Update N`.

Consider an instance of the procedure `backtrack` for ℓ_0 ; note that the instance is at depth $h - \ell_0$ of the recursion. Let I_0 be the values of ℓ in the instances that made a recursive call on the line marked `Call 1` and let a_i and b_i be the values of the variables a and b when a recursive call was made on the line marked `Call 2` for the value of ℓ equal to $i \in \{\ell_0 + 1, \dots, h\} \setminus I_0$. In particular, if $\ell_0 = h$, then $I_0 = \emptyset$ and no a_i 's and b_i 's are defined. Note that $h \notin I_0$. Also, note that $(b_i - a_i)i \leq k$ for every $i \in \{\ell_0 + 1, \dots, h\} \setminus I_0$ and it holds that

$$N_{>} = \sum_{j \in \{\ell_0 + 1, \dots, h\} \setminus I_0} M[j][a_j][b_j].$$

Observe that the values of the array L and U satisfy the following for every $i \in \{1, \dots, \ell_0\}$:

$$\begin{aligned} L[i] &= \max \left\{ 1, \max_{j \in \{\ell_0 + 1, \dots, h\} \setminus I_0} \left\lceil \frac{ia_j - k}{j} \right\rceil, \max_{j \in \{\ell_0 + 1, \dots, h\} \setminus I_0} \left\lceil \frac{ib_j - k}{j} \right\rceil \right\} \\ U[i] &= \min \left\{ k, \min_{j \in \{\ell_0 + 1, \dots, h\} \setminus I_0} \left\lfloor \frac{ia_j + k}{j} \right\rfloor, \min_{j \in \{\ell_0 + 1, \dots, h\} \setminus I_0} \left\lfloor \frac{ib_j + k}{j} \right\rfloor \right\} \end{aligned}$$

with the exception of $L[1]$, which is equal to

$$L[1] = \max \left\{ 0, \max_{j \in \{\ell_0 + 1, \dots, h\} \setminus I_0} \left\lfloor \frac{a_j - k}{j} \right\rfloor, \max_{j \in \{\ell_0 + 1, \dots, h\} \setminus I_0} \left\lfloor \frac{b_j - k}{j} \right\rfloor \right\} = 0.$$

In particular, if the instance of the procedure `backtrack` for ℓ_0 makes a recursive call on the line marked `Call 2`, the inequality $L[\ell_0] \leq a \leq b \leq U[\ell_0]$ implies that then the values of a and b at that point satisfy

$$|aj - a_j\ell_0| \leq k, \quad |bj - a_j\ell_0| \leq k, \quad |aj - b_j\ell_0| \leq k \quad \text{and} \quad |bj - b_j\ell_0| \leq k$$

for every $j \in \{\ell_0 + 1, \dots, h\} \setminus I_0$. Since the same was true in the instances of the procedure `backtrack` for $\ell \in \{\ell_0 + 1, \dots, h\} \setminus I_0$, it follows that

$$|a_j j' - a_{j'} j| \leq k, \quad |b_j j' - b_{j'} j| \leq k \quad \text{and} \quad |a_j j' - b_{j'} j| \leq k \tag{A.1}$$

for all $j, j' \in \{\ell_0 + 1, \dots, h\} \setminus I_0$.

Suppose that $\ell_0 = 0$ and consider the following set $Q \subseteq \mathbb{Z}^2$:

$$Q = \{(1, 0)\} \cup \bigcup_{i \in \{1, \dots, h\} \setminus I_0} \{z : a_i \leq z \leq b_i \text{ and } \gcd(z, i) = 1\}.$$

Since $h \leq k$, $|b_i - a_i| \leq k/i$ for all $i \in \{1, \dots, h\} \setminus I_0$, and (A.1) holds for all $j, j' \in \{1, \dots, h\} \setminus I_0$, the set Q is a k -nice set. Note that $a_i \in Q$ and $b_i \in Q$ for every $i \in \{1, \dots, h\} \setminus I_0$ because $\gcd(a_i, i) = 1$, $\gcd(b_i, i) = 1$ and $a_i \leq b_i$. The size of Q is equal to

$$1 + \sum_{i \in \{1, \dots, h\} \setminus I_0} M[i][a_i][b_i] = 1 + N_{>},$$

and so the return value of the procedure `backtrack` is actually the size of the k -nice set Q . We conclude that if the procedure `backtrack` returns a value n larger than the parameter N that it was called with, then there exists a k -nice set with height h (note that $h \notin I_0$ and $(a_h, h) \in Q$) that has size n . It follows that if Algorithm 5 returns a value n that is larger than N , then there exists a k -nice set with height h that has size n .

To complete the proof of the lemma, we need to show that if there exists a k -nice set with height h that has $n > N$ elements, then the return value of Algorithm 5 is at least n . Fix a k -nice set Q with height h of maximum possible size n and assume that $n > N$. By Lemma A.5, we may assume that $(1, 0) \in Q$, $(0, 1) \in Q$ and $Q \subseteq \{0, \dots, k\} \times \{0, \dots, h\}$. Let I_0 be the set of those $y \in \{1, \dots, h\}$ that are not the second coordinate of any point in Q , and for every $i \in \{1, \dots, h\} \setminus I_0$, define a_i and b_i to be the minimum and maximum first coordinate of the points in Q with their second coordinate equal to i , respectively. We inspect the sequence of recursive calls made on the line marked `Call 1` for $\ell \in I_0$ and on the line marked `Call 2` in the loop for $a = a_\ell$ and $b = b_\ell$ for $\ell \in \{1, \dots, h\} \setminus I_0$. Consider the instance for ℓ_0 in this sequence of the recursive calls. Note that the invariants that we observed earlier imply that $L[i] \leq a_i$ and $b_i \leq U[i]$ for any $i \in \{1, \dots, \ell_0\} \setminus I_0$, and

$$N_{>} = \sum_{i \in \{\ell_0 + 1, \dots, h\} \setminus I_0} M[i][a_i][b_i] = |Q \cap \{0, \dots, k\} \times \{\ell_0 + 1, \dots, h\}|.$$

In particular, if $\ell_0 = 0$, then the procedure `backtrack` returns the size of the set $|Q| = n$ on the line marked `Update N`.

We now consider the case $\ell_0 \in I_0$ (note that $h \notin I_0$ as the height of Q is h). Note that the value of N' in the `if` condition just before the line marked `Call 1` is

$$\begin{aligned} N' &= N_{>} + 1 + \sum_{i \in \{1, \dots, \ell_0 - 1\}, L[i] \leq U[i]} M_k[i][L[i]][U[i]] \\ &\geq N_{>} + 1 + \sum_{i \in \{1, \dots, \ell_0\} \setminus I_0} M_k[i][L[i]][U[i]] \\ &\geq N_{>} + 1 + \sum_{i \in \{1, \dots, \ell_0\} \setminus I_0} M[i][a_i][b_i] = |Q|. \end{aligned}$$

Note that we are using $L[i] \leq a_i \leq b_i \leq U[i]$ and $|b_i - a_i| \leq k/i$ for $i \in \{1, \dots, \ell_0\} \setminus I_0$. In particular, if the recursive call on the line marked `Call 1` is not made, then the value of N is already at least $n = |Q|$ and so Algorithm 5 eventually returns a value that is at least n .

We next consider the case $\ell_0 \in \{1, \dots, h\} \setminus I_0$ and consider the iteration of the `for` cycles for $a = a_{\ell_0}$ and $b = b_{\ell_0}$. This iteration reaches at the least the `if` condition above the line marked `Call 2` because $L[\ell_0] \leq a_{\ell_0} \leq b_{\ell_0} \leq U[\ell_0]$, $\gcd(a_{\ell_0}, \ell_0) = 1$, $\gcd(b_{\ell_0}, \ell_0) = 1$ and $|b_{\ell_0} - a_{\ell_0}| \leq k/\ell_0$. Since $L'[i] \leq a_i \leq b_i \leq U'[i]$ for any $i \in \{1, \dots, \ell_0 - 1\} \setminus I_0$, we obtain that the value of N' in the `if` condition above the line marked `Call 2` is

$$\begin{aligned} N' &= N_{>} + 1 + M[\ell_0][a_{\ell_0}][b_{\ell_0}] + \sum_{i \in \{1, \dots, \ell_0 - 1\}, L'[i] \leq U'[i]} M_k[i][L'[i]][U'[i]] \\ &\geq N_{>} + 1 + M[\ell_0][a_{\ell_0}][b_{\ell_0}] + \sum_{i \in \{1, \dots, \ell_0 - 1\} \setminus I_0} M_k[i][L'[i]][U'[i]] \\ &\geq N_{>} + 1 + M[\ell_0][a_{\ell_0}][b_{\ell_0}] + \sum_{i \in \{1, \dots, \ell_0 - 1\} \setminus I_0} M[i][a_i][b_i] = |Q|. \end{aligned}$$

In particular, if the recursive call on the line marked `Call 2` is not made, then the value of N is already at least $n = |Q|$ and so Algorithm 5 eventually returns a value that is at least n .

We conclude that if the whole sequence of recursive calls consisting of those to be made on the line marked `Call 1` for $\ell \in I_0$ and on the line marked `Call 2` in the loop for $a = a_\ell$ and $b = b_\ell$ for $\ell \in \{1, \dots, h\} \setminus I_0$ is not made, then the value of N was already at least n . If the whole sequence of the recursive calls is made, it reaches the line marked `Update N` in the call with $\ell = 0$ and the procedure `backtrack` returns n . In either case, Algorithm 5 returns a value that is at least n . Since we have shown that if there exists a k -nice set with height h that has $n > N$ elements, then the return value of Algorithm 5 is at least n , the proof of the lemma is now complete. \square

We are now ready to determine the maximum size of a k -nice set for every $k \in \{3, \dots, 1891\}$.

Theorem A.7. *Let K_0 be the set containing the 59 integers listed in Table 8.1. For every $k \in \{3, \dots, 1891\} \setminus K_0$, the maximum size of k -nice set is*

- $k + 4$ if $k \bmod 6 = 2$,
- $k + 3$ if $k \bmod 6 \in \{1, 3, 5\}$, and
- $k + 2$, otherwise.

If $k \in K_0 \cap \{3, \dots, 1891\}$, then the maximum size of a k -nice set is the value $N(\mathbb{T}^2, k)$ given in Table 8.1.

Proof. Fix $k \in \{3, \dots, 1891\}$, and let N be the integer defined as

- $N = k + 4$ if $k \bmod 6 = 2$,
- $N = k + 3$ if $k \bmod 6 \in \{1, 3, 5\}$, and
- $N = k + 2$, otherwise.

Note that Lemma 8.11 yields that there exists a k -nice set of size N . For every $h \in \{2, \dots, \lfloor \sqrt{4k/3} \rfloor\}$ such that Algorithm 4 returned *not verified* we executed Algorithm 5. The implementation of Algorithm 5 can be found in Section B.2 of Appendix B. Lemma A.3 and A.6 yield that the maximum of the numbers returned by the procedures is the maximum size of a k -nice set. If $k \notin K_0$, all procedures returned N , and if $k \in K_0$, the maximum number returned by the procedures is the value of $N(\mathbb{T}^2, k)$ in Table 8.1. \square

A.3 Proof of Theorem 8.1

We are now ready to prove Theorem 8.1, which is restated here for convenience. Before we do so, for completeness, we present an argument on the maximum size of a 1-nice set and the maximum size of a 2-nice set.

Proposition A.8. *The maximum size of a 1-nice set is 3 and the maximum size of a 2-nice is 4.*

Proof. Fix $k \in \{1, 2\}$. Let Q be a k -nice set with maximum size and let h the height of Q . By Lemma 8.6, we may assume that $h \leq \sqrt{2k}$, i.e. $h \leq k$ (note that h is an integer). Finally, by Lemma A.5, we may also assume that $(1, 0) \in Q$, $(0, 1) \in Q$ and $(1, 1) \in Q$, and $Q \subseteq \{0, \dots, k\}^2$. Hence, if $k = 1$, the size of Q is at most 3 and the 1-nice set $\{(1, 0), (0, 1), (1, 1)\}$ witnesses that this bound is tight. If $k = 2$, observe that the set Q cannot contain both $(2, 1)$ and $(1, 2)$, which implies that the size of Q is at most 4. The 2-nice set $\{(1, 0), (0, 1), (1, 1), (2, 1)\}$ witnesses that this bound is tight. \square

Theorem 8.1. *Let K_0 be the set containing the 59 integers listed in Table 8.1. For every $k \in \mathbb{N} \setminus K_0$, it holds that*

$$N(\mathbb{T}^2, k) = \begin{cases} k + 4 & \text{if } k \bmod 6 = 2, \\ k + 3 & \text{if } k \bmod 6 \in \{1, 3, 5\}, \text{ and} \\ k + 2 & \text{otherwise.} \end{cases}$$

The values of $N(\mathbb{T}^2, k)$ for $k \in K_0$ are given in Table 8.1.

Proof. The values of $N(\mathbb{T}^2, 1)$ and $N(\mathbb{T}^2, 2)$ are determined by Proposition A.8 and the values of $N(\mathbb{T}^2, k)$ for $k \in \{3, \dots, 1891\}$ are determined in Theorem A.7. If $k \geq 1892$, Theorem A.4 implies that there exist a maximum size k -nice set that has height at most 3, and so the values of $N(\mathbb{T}^2, k)$ for $k \geq 1892$ are determined by Lemma 8.11. \square

Appendix B

Source code of our programs

B.1 Implementation of Algorithm 4

The following code implements Algorithm 4 from Section A.1 using the programming language C.

```
#include<stdio.h>
#include<stdlib.h>
#define MAXK 3224

void bug(char *s) {
    printf("There is a bug walking around.\n");
    if (s) printf("%s\n",s);
    exit(0);
}

int gcd(int p, int q) {
    if (p < 0) p = -p;
    if (q < 0) q = -q;
    while (1) {
        if (!p) return q;
        if (!q) return p;
        if (p > q) { p %= q; continue; }
        q %= p;
    }
}

// Algorithm 4 — returns 1 if for any given k-nice set Q of
// height h one can always find an equivalent set of height
// smaller than h
int eliminate(int k, int h) {
    int x0,x,y,w,z,xprime,yprime;
    if (h > k) bug("k_cannot_be_smaller_than_height_to_exclude");
    for (x0 = 0; x0 <= h/2; x0++) {
```

```

if (h <= (k-x0)/h) return 0;
for (y = 1; y <= h; y++)
  for (x = h; x <= (k + y*x0)/h; x++) // point (x,y) is
    assumed to have maximum first coordinate in the set Q,
    hence width of Q is equal to x
    if (gcd(x,y) == 1) {
      // computing z = min{y, x-y}
      z = y; if (z > x-y) z = x-y;
      if (z + k/h < h) continue;
      w = 1;
      for (yprime = 1; yprime <= h; yprime++) {
        for (xprime = (yprime*(x0-h)-k+h-1)/h; xprime <=
          (yprime*(x0-h)+k)/h; xprime++) {
          if (gcd(xprime, yprime) != 1) continue;
          if (abs(xprime*y - (x-y)*yprime) > k) continue;
          if (abs(xprime) > w) w = abs(xprime);
        }
      }
      if (w < h) continue;
      return 0;
    }
  }
return 1;
}

// for given k >= 2, finds as small h as possible satisfying the
// following: any k-nice set is equivalent to k-nice set with
// height at most h
int compute_height(int k) {
  if (k < 2) bug("k_must_be_at_least_2");
  int h;
  for (h = 1; h*h < 2*k; h++); // by Lemma 8.6, h*h is at most 2*k
  while (eliminate(k,h)) h--;
  return h;
}

int main() {
  int k, h, verified, maxk;
  verified = 1;
  for (k = 2; k <= MAXK; k++) {
    h = compute_height(k);
    if (3*h*h > 4*k) {
      verified = 0;
      printf("Algorithm_failed_for_k=%d:obtained_height_h=%d_
        is_too_large\n", k, h);
      break;
    }
  }
}

```

```

    }
  }
  if (verified) printf("Claim of Lemma A.3 verified for  $k \leq$ 
    %d\n", MAXK);
  return 0;
}

```

B.2 Implementation of Algorithm 5

The algorithm presented in the following implements Algorithm 5 together with the procedure `backtrack` from Section A.2. Note that some code from the implementation of Algorithm 4, presented in Section B.1 is reused. We even keep the implementation of the core of Algorithm 4 because for computing the maximum size of a k -nice set we only consider heights h where Algorithm 4 returns `not verified`.

```

#include<stdio.h>
#include<stdlib.h>
#define MAXH 110
#define MAXK 2000

int gcds[MAXK+1][MAXK+1]; // gcds[i][j] = 1 if gcd(i, j) = 1,
  otherwise, gcds[i][j] = 0
short M[MAXH+1][MAXK+1][MAXK+1]; // M[h][j1][j2] is the size of
  the set {z : j1 <= z <= j2, gcd(z, h) = 1}
short Mk[MAXH+1][MAXK+1][MAXK+1]; // given k, Mk[h][j1][j2] is the
  maximum of M[h][j1'][j2'] over j1 <= j1' <= j2' <= j2,
  |j1'-j2'| < k/h

void bug(char *s) {
  printf("There is a bug walking around.\n");
  if (s) printf("%s\n", s);
  exit(0);
}

int gcd(int p, int q) {
  if (p < 0) p = -p;
  if (q < 0) q = -q;
  while (1) {
    if (!p) return q;
    if (!q) return p;
    if (p > q) { p %= q; continue; }
    q %= p;
  }
}

// initialize the values of M and gcds

```

```

void init_M_and_gcds(void) {
  int i,j,k;
  for (i = 0; i <= MAXK; i++) for (j = 0; j <= MAXK; j++)
    gcds[i][j] = (gcd(i,j) == 1)?1:0;
  for (i = 1; i <= MAXH; i++) for (j = 0; j <= MAXK; j++) {
    for (k = 0; k < j; k++) M[i][j][k]=0;
    M[i][j][j] = gcds[i][j];
    for (k = j+1; k <= MAXK; k++) M[i][j][k] = M[i][j][k-1] +
      gcds[i][k];
  }
}

// Algorithm 4 — returns 1 if for any given k-nice set Q of
// height h one can always find an equivalent set of height
// smaller than h
int eliminate(int k, int h) {
  int x0,x,y,w,z,xprime,yprime;
  if (h > k) bug("k_cannot_be_smaller_than_height_to_exclude");
  for (x0 = 0; x0 <= h/2; x0++) {
    if (h <= (k-x0)/h) return 0;
    for (y = 1; y <= h; y++)
      for (x = h; x <= (k + y*x0)/h; x++) // point (x,y) is
        // assumed to have maximum first coordinate in the set Q,
        // hence width of Q is equal to x
        if (gcd(x,y) == 1) {
          // computing z = min{y, x-y}
          z = y; if (z > x-y) z = x-y;
          if (z + k/h < h) continue;
          w = 1;
          for (yprime = 1; yprime <= h; yprime++) {
            for (xprime = (yprime*(x0-h)-k+h-1)/h; xprime <=
              (yprime*(x0-h)+k)/h; xprime++) {
              if (gcd(xprime, yprime) != 1) continue;
              if (abs(xprime*y - (x-y)*yprime) > k) continue;
              if (abs(xprime) > w) w = abs(xprime);
            }
          }
          if (w < h) continue;
          return 0;
        }
  }
  return 1;
}

// Procedure backtrack

```

```

int backtrack(int k, int h, int N, int l, int Ngreater, int *L,
int *U) {
int Lprime[MAXH];
int Uprime[MAXH];
int a, b, i, Nprime, Ngeq;

if (l == 0) return Ngreater + 1;
if (l < h) {
  Nprime = Ngreater + 1;
  for (i = 1; i < l; i++) if (L[i] <= U[i]) Nprime +=
    Mk[i][L[i]][U[i]];
  if (Nprime > N) N = backtrack(k, h, N, l-1, Ngreater, L, U);
}
if (L[l] > U[l]) return N;
for (a = L[l]; a <= U[l]; a++) if (gcds[a][l]) //if
  ((!a) || (gcds[l][a]))
  for (b = a; b <= U[l]; b++) if (gcds[b][l]) { // if
    ((!b) || (gcds[l][b])) {
      if (((b - a)*l) > k) continue;
      Nprime = Ngreater + M[l][a][b] + 1;
      for (i = 1; i < l; i++) {
        // computing Lprime[i] = max{L[i], ceiling((a*i-k+)/l),
          ceiling((b*i-k)/l)}
        Lprime[i] = L[i];
        if (Lprime[i] < (a*i-k+1-1)/l) Lprime[i] = (a*i-k+1-1)/l;
        if (Lprime[i] < (b*i-k+1-1)/l) Lprime[i] = (b*i-k+1-1)/l;
        // computing Uprime[i] = min{U[i], floor((a*i+k)/l),
          floor((b*i+k)/l)}
        Uprime[i] = U[i];
        if (Uprime[i] > (a*i+k)/l) Uprime[i] = (a*i+k)/l;
        if (Uprime[i] > (b*i+k)/l) Uprime[i] = (b*i+k)/l;
        if (Lprime[i] <= Uprime[i]) Nprime +=
          Mk[i][Lprime[i]][Uprime[i]];
      }
      if (Nprime <= N) continue;
      Ngeq = Ngreater + M[l][a][b];
      N = backtrack(k, h, N, l-1, Ngeq, Lprime, Uprime);
    }
  }
return N;
}

```

// Algorithm 5

```

int compute(int k, int h, int N) {
int L[MAXH]; int U[MAXH]; // we never use values L[0] and U[0]
int i;
if (h > MAXH) bug("MAXH_is_too_small");

```

```

L[1] = 0; U[1] = k;
for (i = 2; i <= h; i++) {
    L[i] = 1;
    U[i] = k;
}
return backtrack(k,h,N,h,0,L,U);
}

// for a given k >= 2, returns the maximum size of a k-nice set
void compute_max(int k) {
    int i,j,h,maxh,N; // N - maximum size of a k-nice set; dif -
        // expected value of the difference N-k from Lemma 14
    if (k > MAXK) bug("MAXK_is_too_small");
    if (k < 2) bug("k_should_be_at_least_2");
    // initializing Mk[h][i][j]; Mk[h][j1][j2] is the maximum of
        // M[h][j1][j2] over j1 <= j1' <= j2' <= j2, |j1'-j2'| < k/h
    for (h = 1; h*h <= 2*k; h++) {
        for (i = 0; i <= k; i++) for (j = i; j <= k; j++)
            if ((j-i)*h <= k)
                Mk[h][i][j] = M[h][i][j];
            else
                Mk[h][i][j] = (Mk[h][i][j-1] > M[h][j-k/h][j]) ?
                    Mk[h][i][j-1] : M[h][j-k/h][j];
    }
    switch (k%6) {
        case 0: N = k + 2; break;
        case 2: N = k + 4; break;
        case 4: N = k + 2; break;
        default: N = k + 3;
    }
    for (maxh = 2; 3*maxh*maxh <= 4*k; maxh++); maxh--;
    for (h = 2; h <= maxh; h++)
        if (!eliminate(k,h)) // we check k-nice sets of height h only
            // for those pairs (k,h) for which Algorithm 4 returns 'not
            // verified'
            N = compute(k,h,N);
    printf("Maximum_for_k=%d_is_%d.\n",k,N);
}

int main() {
    int k;
    if (sizeof(short) != 2) bug("sizeof(short)!=2");
    init_M_and_gcds();
    for (k = 3; k <= 1891; k++) compute_max(k);
    return 0;
}

```

B.3 Python script to compute the values $\rho_\ell, \gamma_\ell, \alpha_\ell$ and β_ℓ

Here is a simple Python script that we used to compute the values ρ_ℓ, α_ℓ and β_ℓ as well as an upper bound on the value γ_ℓ for any given $\ell \in \mathbb{N}$.

```

from math import gcd, isqrt
from fractions import Fraction as Fc

MAX_L = 301 # compute values of rho_l, alpha_l, gamma_l, beta_l
            for l in range(1, MAX_L)

def prime_factors(n):
    # iterates over all prime factors of n, without repetition
    if n % 2 == 0:
        yield 2
        while n % 2 == 0:
            n //= 2
    i = 3
    while i < isqrt(n) + 1:
        if n % i == 0:
            yield i
            while n % i == 0:
                n //= i
        i += 2
    if n > 2:
        yield n

def totient(n):
    # returns the Euler's totient function phi(n)
    if n == 0:
        return 0
    result = n
    for p in prime_factors(n):
        result *= (1 - 1/p)
    return int(result)

def rho(l):
    # returns the value of rho_l
    result = 1
    for p in prime_factors(l):
        result *= (1 - Fc(1,p))
    return result

def recursive_alpha(l, a, b):
    # used by alpha(l)
    rho_l = rho(l)
    mid = (a + b) // 2

```

```

# compute the maximum over intervals containing the mid point
left_max = -rho_l
points = 0
for i in range(mid, a-1, -1):
    if gcd(i, l) == 1:
        points += 1
        left_max = max(left_max, points - rho_l*(mid-i+1))
right_max = -rho_l
points = 0
for i in range(mid, b+1):
    if gcd(i, l) == 1:
        points += 1
        right_max = max(right_max, points - rho_l*(i-mid+1))
mid_max = left_max + right_max + rho_l
if gcd(mid, l) == 1:
    mid_max -= 1
if b - a <= 1:
    return mid_max
# maximum may be also attained at some interval disjoint from
# the mid point
return max(mid_max, recursive_alpha(l, a, mid-1),
           recursive_alpha(l, mid+1, b))

def alpha(l):
    # returns the value of alpha_l
    return recursive_alpha(l, 1, 2*l)

def euler_dict(h):
    # computes a feasible solution to the linear program LD_h
    # in a form of non-zero entries of a matrix A_h

    E = dict() # E[i] = phi(i), Euler's totient function
    E[0] = 0
    for i in range(1, h+1):
        E[i] = totient(i)

    P = dict() # P[(i,j)] is the value of (i,j)-th entry of A_h
    i = 1
    j = h
    row_sum = 0
    row_max = E[h]
    while i < h or j > 0:
        if row_sum + E[i] < row_max:
            row_sum += E[i]
            P[(i, j)] = E[i]
            i += 1

```

```

    elif row_sum + E[i] == row_max:
        P[(i, j)] = E[i]
        row_sum = 0
        i += 1
        j -= 1
        row_max = E[j]
    else:
        P[(i, j)] = row_max - row_sum
        column_sum = row_max - row_sum
        column_max = E[i]
        j -= 1
        while column_sum < column_max:
            if column_sum + E[j] < column_max:
                P[(i, j)] = E[j]
                column_sum += E[j]
                j -= 1
            elif column_sum + E[j] == column_max:
                P[(i, j)] = E[j]
                column_sum += E[j]
                i += 1
                j -= 1
                row_sum = 0
                row_max = E[j]
            else:
                P[(i, j)] = column_max - column_sum
                row_sum = column_max - column_sum
                row_max = E[j]
                column_sum = column_max
                i += 1

    return P

def gamma_h(h):
    # returns an upper bound for the value of gamma_h
    A = euler_dict(h)
    result = 0.0
    for (i, j) in A:
        result += Fc(A[(i, j)], i*j)
    return result

ALPHAS = dict()
RHOS = dict()
print("computing alphas")
for k in range(1, MAX_L):
    ALPHAS[k] = alpha(k)
    RHOS[k] = rho(k)

```

```
print("computing_betas")
BETAS = dict()
BETAS[0] = 1
for k in range(1, MAX_L):
    BETAS[k] = BETAS[k-1] + ALPHAS[k] + RHOS[k]

print("computing_gammas")
GAMMAS = dict()
for k in range(1, MAX_L):
    GAMMAS[k] = gamma_h(k)

print("saving_to_file")
with open("torus_a_b_gamma_beta.txt", "w") as file:
    for n in range(1, MAX_L):
        a = float(RHOS[n])
        b = float(ALPHAS[n])
        c = float(GAMMAS[n])
        d = float(BETAS[n])
        file.write(f"{n}\t{a:.10f}\t{b:.10f}\t{c:.10f}\t{d:.10f}\n")
print("finished!")
```

Bibliography

- [1] Ian Agol. Bounds on exceptional Dehn filling. *Geometry & Topology*, 4(1):431–449, 2000.
- [2] Tarik Aougab. Constructing large k -systems on surfaces. *Topology and its Applications*, 176:1–9, 2014.
- [3] Tarik Aougab. Curves intersecting exactly once and their dual cube complexes. *Groups, Geometry and Dynamics*, 11:1061–1101, 2017.
- [4] Tarik Aougab. Local geometry of the k -curve graph. *Transactions of the American Mathematical Society*, 370(4):2657–2678, 2018.
- [5] Tarik Aougab, Ian Biringer, and Jonah Gaster. Packing curves on surfaces with few intersections. *International Mathematics Research Notices*, 2019(16):5205–5217, 2017.
- [6] Tarik Aougab and Jonah Gaster. Curves on the torus intersecting at most k times. *Mathematical Proceedings of the Cambridge Philosophical Society*, 174(3):569–584, 2023.
- [7] Tarik Aougab and Jonah Gaster. From arcs to curves: quadratic growth of 1-systems. *arXiv preprint, arXiv:2508.05555*, 2025.
- [8] Kenneth Appel and Wolfgang Haken. A proof of the four color theorem. *Discrete Mathematics*, 16:179–180, 1976.
- [9] Mark A. Armstrong. *Basic Topology*. Undergraduate Texts in Mathematics. Springer New York, 2013.
- [10] Stefan Arnborg, Derek G. Corneil, and Andrzej Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM Journal on Algebraic Discrete Methods*, 8(2):277–284, 1987.
- [11] Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994.
- [12] Kenneth L. Baker, Cameron Gordon, and John Luecke. Bridge number, heegaard genus and non-integral dehn surgery. *Transactions of the American Mathematical Society*, 367:5753–5830, 2015.
- [13] Roger C. Baker, Gözde Harman, and János Pintz. The difference between consecutive primes, II. *Proceedings of the London Mathematical Society*, 83(3):532–562, 11 2001.
- [14] Igor Balla, Marek Filakovský, Bartłomiej Kielak, Daniel Král', and Niklas Schlomberg. Curves on the torus with few intersections. *arXiv preprint, arXiv:2412.18002*, 2025.

-
- [15] Ann Becker and Dan Geiger. Optimization of pearl’s method of conditioning and greedy-like approximation algorithms for the vertex feedback set problem. *Artificial Intelligence*, 83(1):167–188, 1996.
- [16] Piotr Berman and Grigory Yaroslavtsev. Primal-dual approximation algorithms for node-weighted network design in planar graphs. In *Approximation, Randomization, and Combinatorial Optimization (Proceedings of APPROX 2012)*, pages 50–60, 2012.
- [17] Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996.
- [18] Hans L. Bodlaender. A partial k -arboretum of graphs with bounded treewidth. *Theoretical Computer Science*, 209(1-2):1–45, 1998.
- [19] Wouter Cames van Batenburg, Louis Esperet, and Tobias Müller. Coloring Jordan regions and curves. *SIAM Journal on Discrete Mathematics*, 31(3):1670–1684, 2017.
- [20] Alberto Caprara, Alessandro Panconesi, and Romeo Rizzi. Packing cuts in undirected graphs. *Networks*, 44(1):1–11, 2004.
- [21] Yuk Hei Chan and Lap Chi Lau. On linear and semidefinite programming relaxations for hypergraph matching. *Mathematical Programming*, 135(1):123–148, 2012.
- [22] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. An $O(\sqrt{n})$ approximation and integrality gap for disjoint paths and unsplittable flow. *Theory of Computing*, 2(1):137–146, 2006.
- [23] Chandra Chekuri, Marcelo Mydlarz, and F. Bruce Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Transactions on Algorithms*, 3(3):Article 27, 2007.
- [24] Hong-Bin Chen, Hung-Lin Fu, and Chih-Huai Shih. Feedback vertex set on planar graphs. *Taiwanese Journal of Mathematics*, 16:2077–2082, 2012.
- [25] Joseph Cheriyan, Howard Karloff, Rohit Khandekar, and Jochen Könemann. On the integrality ratio for tree augmentation. *Operations Research Letters*, 36(4):399–401, 2008.
- [26] Norishige Chiba, Takao Nishizeki, Shigenobu Abe, and Takao Ozawa. A linear algorithm for embedding planar graphs using pq-trees. *Journal of Computer and System Sciences*, 30(1):54–76, 1985.
- [27] Julia Chuzhoy, David H. K. Kim, and Rachit Nimavat. Almost polynomial hardness of node-disjoint paths in grids. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 1220–1233, New York, NY, USA, 2018. Association for Computing Machinery.
- [28] Michele Conforti, Samuel Fiorini, Tony Huynh, Gwenaël Joret, and Stefan Weltge. The stable set problem in graphs with bounded genus and bounded odd cycle packing number. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2896–2915, 2020.

-
- [29] Harald Cramér. Some theorems concerning prime numbers. *Arkiv för Matematik, Astronomi och Fysik*, 15(5):33, 1921.
- [30] Irit Dinur and Samuel Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162, 07 2004.
- [31] Jack Edmonds and Ellis L. Johnson. Matching, Euler tours and the Chinese postman. *Mathematical Programming*, 5(1):88–124, 1973.
- [32] Paul Erdős and Lajos Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics*, 17:347–352, 1965.
- [33] Samuel Fiorini, Nadia Hardy, Bruce Reed, and Adrian Vetta. Approximate min–max relations for odd cycles in planar graphs. *Mathematical Programming*, 110(1):71–91, 2007.
- [34] Zachary Friggstad and Mohammad R. Salavatipour. Approximability of packing disjoint cycles. *Algorithmica*, 60(2):395–400, 2011.
- [35] Michael R. Garey and David S. Johnson. The rectilinear steiner tree problem is np-complete. *SIAM Journal on Applied Mathematics*, 32(4):826–834, 1977.
- [36] Naveen Garg and Nikhil Kumar. Dual half-integrality for uncrossable cut cover and its application to maximum half-integral flow. In *28th Annual European Symposium on Algorithms (ESA 2020)*, pages 55:1–55:13, 2020.
- [37] Naveen Garg, Nikhil Kumar, and András Sebő. Integer plane multiflow maximisation: one-quarter-approximation and gaps. *Mathematical Programming*, 195:403–419, 2022.
- [38] Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996.
- [39] Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, May 1997.
- [40] Michel X. Goemans and David P. Williamson. The primal-dual method for approximation algorithms and its application to network design problems. In *Approximation Algorithms for NP-Hard Problems*, page 144–191. PWS Publishing Co., USA, 1996.
- [41] Michel X. Goemans and David P. Williamson. Primal-dual approximation algorithms for feedback problems in planar graphs. *Combinatorica*, 18(1):37–59, 1998.
- [42] Alexander Göke, Jochen Koenemann, Matthias Mnich, and Hao Sun. Hitting weighted even cycles in planar graphs. *SIAM Journal on Discrete Mathematics*, 36(4):2830–2862, 2022.
- [43] Andrew Granville. Harald Cramér and the distribution of prime numbers. *Scandinavian Actuarial Journal*, 1995(1):12–28, 1995.
- [44] Joshua E. Greene. On loops intersecting at most once. *Geometric and Functional Analysis*, 29:1828–1843, 2019.

-
- [45] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *49th Annual IEEE Symposium on Foundations of Computer Science*, pages 573–582, 2008.
- [46] Frank Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing*, 4(3):221–225, 1975.
- [47] Chien-Chung Huang, Mathieu Mari, Claire Mathieu, Kevin Schewior, and Jens Vygen. An approximation algorithm for fully planar edge-disjoint paths. *SIAM Journal on Discrete Mathematics*, 35:752–769, 2021.
- [48] Chien-Chung Huang, Mathieu Mari, Claire Mathieu, and Jens Vygen. Approximating maximum integral multiflows on bounded genus graphs. *Discrete & Computational Geometry*, 70:1266–1291, 2023.
- [49] Martin Juvan, Aleksander Malnič, and Bojan Mohar. Systems of curves on surfaces. *Journal of Combinatorial Theory, Series B*, 68:7–22, 1996.
- [50] Alexander V. Karzanov. How to tidy up a symmetric set-system by use of uncrossing operations. *Theoretical Computer Science*, 157(2):215–225, 1996.
- [51] Ken-Ichi Kawarabayashi and Atsuhiko Nakamoto. The Erdős–Pósa property for vertex- and edge-disjoint odd cycles in graphs on orientable surfaces. *Discrete Mathematics*, 307(6):764–768, 2007.
- [52] Philip N. Klein, Claire Mathieu, and Hang Zhou. Correlation clustering and two-edge-connected augmentation for planar graphs. *Algorithmica*, 85:3024–3057, 2023.
- [53] Jon Kleinberg and Amit Kumar. Wavelength conversion in optical networks. *Journal of Algorithms*, 38(1):25–50, 2001.
- [54] Daniel Král’, Jean-Sebastien Sereni, and Ladislav Stacho. Min-max relations for odd cycles in planar graphs. *SIAM Journal on Discrete Mathematics*, 26(3):884–895, 2012.
- [55] Daniel Král’ and Heinz-Jürgen Voss. Edge-disjoint odd cycles in planar graphs. *Journal of Combinatorial Theory, Series B*, 90(1):107–120, 2004.
- [56] Björn Kriepke and Matthias Schymura. On generic δ -modular integer matrices with two rows. *arXiv preprint, arXiv:2502.15394*, 2025.
- [57] Michael Krivelevich, Zeev Nutov, Mohammad R. Salavatipour, Jacques Verstraete, and Raphael Yuster. Approximation algorithms and hardness results for cycle packing problems. *ACM Transactions on Algorithms*, 3(4):Article 48, 2007.
- [58] Lap Chi Lau, R. Ravi, and Mohit Singh. *Iterative Methods in Combinatorial Optimization*. Cambridge University Press, 2011.
- [59] Claudio L. Lucchesi and Daniel H. Younger. A minimax theorem for directed graphs. *Journal of the London Mathematical Society II*, 17(3):369–374, 1978.
- [60] Jie Ma, Xingxing Yu, and Wenan Zang. Approximate min-max relations on plane graphs. *Journal of Combinatorial Optimization*, 26(1):127–134, 2013.

- [61] Justin Malestein, Igor Rivin, and Louis Theran. Topological designs. *Geometriae Dedicata*, 168(1):221–233, 2014.
- [62] Matthias Middendorf and Frank Pfeiffer. On the complexity of the disjoint paths problem. *Combinatorica*, 13(1):97–107, 1993.
- [63] Bojan Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM Journal on Discrete Mathematics*, 12(1):6–26, 1999.
- [64] Owen J. Murphy. Computing independent sets in graphs with large girth. *Discrete Applied Mathematics*, 35(2):167–170, 1992.
- [65] Guylain Naves, F. Bruce Shepherd, and Henry Xia. Maximum weight disjoint paths in outerplanar graphs via single-tree cut approximators. *Mathematical Programming*, 197:1049–1067, 2023.
- [66] Tomás Oliveira e Silva, Siegfried Herzog, and Silvio Pardi. Empirical verification of the even Goldbach conjecture and computation of prime gaps up to $4 \cdot 10^{18}$. *Mathematics of Computation*, 83(288):2033–2060, 2014.
- [67] János Pach, Gábor Tardos, and Géza Tóth. Crossings between non-homotopic edges. *Journal of Combinatorial Theory, Series B*, 156:389–404, 2022.
- [68] Piotr Przytycki. Arcs intersecting at most once. *Geometric and Functional Analysis*, 25:658–670, 2015.
- [69] Luise Puhlmann and Niklas Schlömborg. Improved Erdős–Pósa inequalities for odd cycles in planar graphs. *arXiv preprint, arXiv:2512.22865*, 2025.
- [70] Dieter Rautenbach and Bruce Reed. The Erdős–Pósa property for odd cycles in highly connected graphs. *Combinatorica*, 21(2):267–278, 2001.
- [71] Dieter Rautenbach and Friedrich Regen. On packing shortest cycles in graphs. *Information Processing Letters*, 109(14):816–821, 2009.
- [72] Bruce Reed. Mangoes and blueberries. *Combinatorica*, 19(2):267–296, 1999.
- [73] Bruce Reed, Neil Robertson, Paul Seymour, and Robin Thomas. Packing directed circuits. *Combinatorica*, 16(4):535–554, 1996.
- [74] Bruce A. Reed and F. Bruce Shepherd. The Gallai–Younger conjecture for planar graphs. *Combinatorica*, 16(4):555–566, 1996.
- [75] Neil Robertson, Daniel P. Sanders, Paul Seymour, and Robin Thomas. Efficiently four-coloring planar graphs. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 571–575. Association for Computing Machinery, 1996.
- [76] Neil Robertson and Paul D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986.

-
- [77] Paul Schmutz Schaller. Mapping class groups of hyperbolic surfaces and automorphism groups of graphs. *Compositio Mathematica*, 122:243–260, 2000.
- [78] Niklas Schlömlberg. An improved integrality gap for disjoint cycles in planar graphs. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297, pages 122:1–122:15, 2024.
- [79] Niklas Schlömlberg, Hanjo Thiele, and Jens Vygen. Packing cycles in planar and bounded-genus graphs. *SIAM Journal on Computing*, 54(2):469–502, 2025.
- [80] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.
- [81] Paul D. Seymour. On odd cuts and plane multicommodity flows. *Proceedings of the London Mathematical Society*, 3(1):178–192, 1981.
- [82] John Stillwell. *Classical topology and combinatorial group theory*. Graduate texts in mathematics 72. Springer, 1980.
- [83] Hao Sun. A constant factor approximation for directed feedback vertex set in graphs of bounded genus. In *Approximation, Randomization, and Combinatorial Optimization (Proceedings of APPROX/RANDOM 2024)*, volume 317, pages 18:1–18:20, 2024.
- [84] Éva Tardos and Vijay V. Vazirani. Improved bounds for the max-flow min-multicut ratio for planar and $K_{r,r}$ -free graphs. *Information Processing Letters*, 47(2):77–80, 1993.
- [85] Carsten Thomassen. Embeddings of graphs with no short noncontractible cycles. *Journal of Combinatorial Theory, Series B*, 48(2):155–177, 1990.
- [86] Carsten Thomassen. The Erdős–Pósa property for odd cycles in graphs of large connectivity. *Combinatorica*, 21(2):321–333, 2001.
- [87] Mihalis Yannakakis. Node-and edge-deletion np-complete problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing, STOC '78*, pages 253–264, New York, NY, USA, 1978. Association for Computing Machinery.
- [88] John W. T. Youngs. Minimal imbeddings and the genus of a graph. *Journal of Mathematics and Mechanics*, 12(2):303–315, 1963.