The minimal excluded minors for surfaces

Sarah Houdaigoui

Doctor of Philosophy



Informatics Program,
Graduate Institute for Advanced Studies, SOKENDAI

September 2026

Abstract in English

As part of the graph minor project, Robertson and Seymour showed in 1990 that, for every surface, the class of graphs embeddable in that surface can be characterized by a finite set of minimal excluded minors [RS90]. However, the proof is purely existential and therefore provides no explicit information about these excluded minors. In 1993, Seymour established the first upper bound on the order of such minimal excluded minors [Sey93], showing that it is at most double-exponential in the genus of the surface.

More than thirty years later, this thesis provides the first improvement of this bound. As a first step, we reduced the upper bound to a quasi-polynomial function of the genus of the surface [HK26b]. Despite this progress, a significant gap remained between this bound and the known linear lower bound $\Omega(g)$ (where g denotes the genus). In particular, we conjectured in [HK26b] that the optimal upper bound should be polynomial. As a second step, we confirm this conjecture by proving that the order of the minimal excluded minors for a surface of genus g is $g^{2+o(1)}$ [HK26a]. This result considerably narrows the gap between the known lower and upper bounds, bringing the asymptotic behavior much closer to the conjectured optimum.

The key ingredient behind the quasi-polynomial bound is a new structural theorem showing that a minimal excluded minor for a surface of genus g contains at most $O(\log g)$ pairwise disjoint nested contractible cycles in some embedding of G. This improves a classical bound of Seymour [Sey93] of $O(g)$ and constitutes the main breakthrough enabling the transition from exponential to quasi-polynomial bound. To derive the almost quadratic bound, we further strengthen this result by establishing a separator-based variant: for every contractible subgraph $H \subseteq G$ with a separator of size s, the subgraph H contains $O(\log s)$ disjoint cycles that are contractible and nested in some embedding of G. This refinement replaces the dependence on the genus by a dependence on the separator size, which constitutes the new key ingredient in the proof of the almost quadratic bound.

As a further contribution, we improve Seymour's [Sey93] classical treewidth bound for minimal excluded minors on surfaces from $O(g^3)$ to $O(g \log g)$.

Abstract in Japanese

グラフマイナープロジェクトの一環として、Robertson と Seymour は 1990 年に、任意の曲面に対して、その曲面に埋め込み可能なグラフのクラスが有限個の極小禁止マイナーによって特徴付けられることを示した [RS90]。しかしながら、その証明は存在論的であり、これらの禁止マイナーに関する具体的な情報を与えるものではない。さらに Seymour は 1993 年に、このような極小禁止マイナーの位数に対する最初の上界を与え、それが曲面の種数に対して高々二重指数関数的であることを示した [Sey93]。

それから 30 年以上を経て、本論文はこの上界に対する初めての改良を与えるものである。まず第一段階として、極小禁止マイナーの位数に対する上界を種数の準多項式関数まで改善した [HK26b]。しかしながら、この結果と既知の線形下界 $\Omega(g)$（ここで g は曲面の種数を表す）との間には依然として大きな隔たりが残されていた。そこで我々は [HK26b] において、最適な上界は多項式であると予想した。本論文の第二段階では、この予想を肯定的に解決し、種数 g の曲面に対する極小禁止マイナーの位数が $g^{2+o(1)}$ で抑えられることを証明した [HK26a]。この結果により、既知の下界と上界の間の隔たりは大幅に縮小され、漸近的挙動は予想されていた最適値に大きく近づいた。

準多項式上界を導く鍵となるのは、新たに得られた構造定理である。この定理は、種数 g の曲面に対する極小禁止マイナーが、互いに頂点素で入れ子状の可縮閉路を高々 $O(\log g)$ 個しか含まないことを示している。これは Seymour [Sey93] による従来の $O(g)$ という上界を改善するものであり、指数関数的上界から準多項式上界への飛躍を可能にした主要なブレークスルーである。さらにほぼ二次の上界を導くために、この結果を分離集合に基づく形へと強化した。具体的には、分離集合の大きさが s である任意の可縮部分グラフ $H \subseteq G$ に対して、H に含まれる互いに頂点素で入れ子状可縮閉路の数は高々 $O(\log s)$ であることを示した。この改良により、種数への依存を分離集合の大きさへの依存へと置き換えることができ、ほぼ二次の上界の証明における新たな鍵となる構造的知見が得られた。

さらに、本論文では Seymour [Sey93] による極小禁止マイナーの木幅に関する古典的上界 $O(g^3)$ を $O(g \log g)$ に改善した。

Abstract in French

Dans le cadre du projet des graphes mineurs, Robertson et Seymour ont montré en 1990 que, pour toute surface, la classe des graphes qui s'y plongent peut être caractérisée par un ensemble fini de mineurs exclus minimaux [RS90]. Cependant, leur preuve, étant de nature existentielle, ne fournit aucune information explicite sur ces mineurs exclus. En 1993, Seymour a établi la première borne supérieure sur l'ordre de tels mineurs exclus minimaux [Sey93], montrant qu'il est au plus double exponentiel en le genre de la surface.

Plus de trente ans plus tard, cette thèse apporte la première amélioration de cette borne. Dans un premier temps, nous avons réduit cette borne supérieure à une fonction quasi-polynomiale en le genre de la surface [HK26b]. Malgré cette avancée, un écart important subsistait entre cette borne et la borne inférieure linéaire connue $\Omega(g)$, où g désigne le genre de la surface.Nous avons notamment conjecturé dans [HK26b] que la borne supérieure optimale devait être polynomiale. Dans un second temps, nous avons confirmé cette conjecture en montrant que l'ordre des mineurs exclus minimaux d'une surface de genre g est $g^{2+o(1)}$ [HK26a]. Ce résultat réduit considérablement l'écart entre les bornes inférieure et supérieure connues, rapprochant ainsi le comportement asymptotique de l'optimum conjecturé.

L'ingrédient principal derrière la borne quasi-polynomiale est un nouveau théorème structurel montrant qu'un mineur exclu minimal pour une surface de genre g contient au plus $O(\log g)$ cycles contractiles imbriqués deux à deux disjoints dans un certain plongement de G. Ce résultat améliore une borne classique de Seymour [Sey93], qui était de $O(g)$, et constitue la percée principale permettant de passer d'une borne exponentielle à une borne quasi-polynomiale. Pour obtenir la borne presque quadratique, nous renforçons ensuite ce résultat en établissant une variante fondée sur les séparateurs : pour tout sous-graphe contractile $H \subseteq G$ possédant un séparateur de taille s, le graphe H contient au plus $O(\log s)$ cycles contractiles deux à deux disjoints et imbriqués dans un certain plongement de G. Ce raffinement remplace la dépendance en le genre par une dépendance en la taille du séparateur, ce qui constitue le principal nouvel ingrédient dans la démonstration de la borne presque quadratique.

Enfin, nous améliorons également la borne classique de Seymour [Sey93] sur la largeur arborescente des mineurs exclus minimaux sur les surfaces, en la faisant passer de $O(g^3)$ à $O(g \log g)$.

Acknowledgements

First of all, I would like to express my deepest gratitude to my supervisor, Professor Ken-ichi Kawarabayashi, for giving me the opportunity to pursue a Ph.D. in Japan. During my internship at the National Institute of Informatics (NII) in 2019, I discovered not only a remarkable research environment but also a fascinating area of mathematics through Professor Kawarabayashi's introduction to graph coloring. It was also during this time that I developed a deep affection for Japan. After this internship, I had only one goal: to return to Japan and pursue my Ph.D. at the very same place. I am therefore profoundly grateful at Professor Kawarabayashi for placing his trust in me and accepting me as his Ph.D. student.

Throughout these three years, I have greatly benefited from his guidance, his trust, and the confidence he placed in me from the very beginning. He gave me the freedom to explore my own ideas, while always being available whenever I needed advice or wished to discuss my work. I am sincerely grateful for his kindness, his constant encouragement, and his unwavering support for all the activities and opportunities I chose to pursue during my Ph.D.

I would like to express my sincere gratitude to the members of my thesis committee, Professor Yoshida, Professor Uno, Professor Hirahara, and Professor Fujii. I am grateful to them for following the progress of my thesis from its very beginning and for their valuable advice and insightful comments.

I had the privilege of being part of Professor Kawarabayashi's Lab at NII. I would like to thank all its past and current members, as well as the group secretaries, for creating such a pleasant and welcoming working environment. My special thanks go to Lucas Picasarri-Arrieta, Marcelo Garlet Milani, Meike Hatzel, Arnab Char, Satoko Tsushima and Chiemi Kubo for the many delicious lunches at the restaurants you introduced me to, and for the enjoyable group gatherings we shared over the years.

I am particularly grateful to Lucas, not only for brightening the office with his jokes and his unmistakably French phlegm, but also for the guidance he provided throughout my Ph.D. Thank you for carefully proofreading my first paper and this thesis, for the many stimulating discussions we had, and for everything you taught me about academic life. I also greatly appreciated being included in your collaborations with visiting researchers, which gave me the opportunity to meet many wonderful people.

I would also like to thank Philip Wellnitz for making me feel part of the life of the institute and for his advice and support while I was looking for a postdoctoral position. Your arrival at NII truly strengthened the sense of community within the institute.

My thanks also go to all the visitors whom I had the pleasure of meeting and discussing research with, especially Dario Cavallaro, Max Dupré la Tour, Amadeus Reinald, Simon Döring and Ararat Harutyunyan.

Finally, I am sincerely grateful to Aki Sasaki for all the work she does in coordinating between NII and SOKENDAI, and for the kindness and patience with which she carries out this role.

I would like to thank my ENS tutor, Chien-Chung Huang, for introducing me to my supervisor and for kindly checking in on me throughout my Ph.D. I am equally grateful to my former internship supervisors, Pierre Aboulker and Nicolas Trotignon, for their kindness and for always taking the time to check on me whenever they had the opportunity.

I would especially like to thank Cléophée Robin for being my French mentor over the past two years. I learned a great deal about the French academic system thanks to you. You were always willing to answer my questions, even the silliest ones, and I am truly grateful for your constant encouragement and for checking in on me so regularly.

I am also grateful to my 2024 JGA group (Alexandre Vigny, Nacim Oijid and Matthieu Petiteau) for the wonderful collaboration that led to our paper. I would particularly like to thank Alex for organizing and leading our research sessions, for everything I learned throughout this project, and overall for being such a genuinely kind person.

During the summer of 2025, I had the opportunity to spend three weeks in France for research visits. I would like to thank Frédéric Havet and the COATI team for their warm welcome. I am especially grateful to Fred, Clément Rambaud and Caroline Silva for the many enjoyable and stimulating discussions we had. I would also like to thank Marthe Bonamy and Vincent Delecroix for welcoming me to LaBRI and for kindly hosting me at their home for one night. My thanks go to both Marthe and Fanny Hauser for the fascinating research we carried out that week, as well as for showing me around Bordeaux. Finally, I greatly enjoyed the SGT Summer School, and I would like to thank in particular Matthieu Liedloff for facilitating my introduction to other researchers, and Stephan Kratsch for the very pleasant discussions we had.

In January 2026, I went to Canada for the SODA conference and greatly enjoyed my time there. I would especially like to thank Oscar Fontaine and Luca Castelli Aleardi for our friendly conversations and the more in-depth discussions we had about graphs on surfaces. I would also like to thank Giannos Stamoulis, Sebastian Wiederrecht and Maximilian Gorsky for the very interesting discussions we had during the conference.

I also had the opportunity to visit Professor Bojan Mohar at Simon Fraser University after the conference, and I would like to express my sincere gratitude for his extremely kind hospitality in welcoming me into his group for a week, as well as for the stimulating discussions we had about my research.

I would like to express my gratitude to Akira Suzuki for bringing to my attention the job posting in his group at Tohoku University and for encouraging me to apply. I believe this is one of the best postdoctoral positions I could have hoped for, and I might not have had the courage to apply without his encouragement. I am very much looking forward to working with him and meeting the whole group.

Finally, I would like to say a few words to my family and friends, in French. Je tiens à remercier mes parents, ma sœur et mon frère pour leur soutien dans mon projet de faire une thèse au Japon. Merci particulièrement à mes parents pour tous leurs encouragements et leur sollicitude, malgré leur mécompréhension de mon sujet de thèse (non, je ne fais plus du coloriage, et le coloriage de graphes, c'est du sérieux, d'abord). Je remercie Malabar pour son écoute silencieuse et son soutien émotionnel qui valait tout l'or du monde.

Je remercie mes très bons amis Jules, Fanny et Adrien pour leur soutien malgré la distance. Je tiens en particulier à remercier Adrien pour sa disponibilité, son écoute et ses conseils. Tu es une personne formidable et flamboyante, et je suis chanceuse de t'avoir à mes côtés.

Un grand merci également à Véro et Chris qui m'ont félicitée pour toutes mes réussites tout au long de ma thèse, et dont le soutien m'a réchauffé le cœur.

Enfin, le meilleur pour la fin : merci de tout cœur à mon fiancé Thomas, sans qui ces années de thèse auraient été beaucoup moins plaisantes. Merci d'avoir été à mes côtés pour partager mes réussites et m'emmener les fêter dans de bons restaurants. Merci de m'avoir remonté le moral quand je doutais de moi. Merci de m'avoir toujours encouragée et d'avoir cru en moi. Tu as embelli mon quotidien et, en repensant à ces années de thèse, je pense avant tout à toutes ces aventures que nous avons vécues ensemble au Japon.

Contents

1 Summary **15**

2 Introduction **25**
2.1 Preliminaries . . . 25
2.1.1 Graph theory . . . 26
2.1.2 Graphs on surfaces . . . 31
2.2 An introduction to excluded minors for surfaces . . . 39
2.2.1 An introduction to minor-closed classes of graphs . . . 39
2.2.2 A historical retrospective on excluded minors for surfaces . . . 42
2.2.3 Algorithmic consequences . . . 45
2.3 Outline of this thesis . . . 46

3 Results and techniques **49**
3.1 A quasi-polynomial bound for the excluded minors for a surface . . . 49
3.1.1 Main results . . . 50
3.1.2 Technical overview . . . 51
3.1.3 Overview of the proof . . . 52
3.2 An almost quadratic bound for the excluded minors for a surface . . . 53
3.2.1 Main results . . . 53
3.2.2 Technical contributions over [HK26b] . . . 53
3.2.3 Overview of the proof . . . 54

4 Preliminary results **57**
4.1 Preliminary results for graphs on surfaces . . . 57
4.2 Basic results on excluded minors for surfaces . . . 61
4.2.1 Embedding of G . . . 61
4.2.2 Bounds for tree decompositions of G . . . 62
4.2.3 Separators in G . . . 63
4.2.4 Structure of $G - e$ . . . 65

5 Structural results **67**
5.1 Our inspiration: Seymour's structural result . . . 68
5.2 Isolated paths . . . 69
5.3 Nested cycles . . . 73
5.3.1 Contractible nested cycles . . . 73
5.3.2 Homotopic nested cycles . . . 85
5.4 Consequences of the structural results . . . 93

5.4.1 Treewidth and tree decomposition . . . 93
5.4.2 Excluded grid topological minor . . . 93
5.4.3 Bound on the number of almost disjoint Π-noncontractible cycles . . . 94
5.5 Links with facewidth . . . 95
5.6 Discussion on the optimality of the structural results . . . 98
5.6.1 Optimality of the bound for the nested/homotopic cycles . . . 98
5.6.2 The scope of the forbidden structure cannot be extended . . . 98

6 Maximum degree of a vertex and size of a face in (G, Π) 101
6.1 Preliminaries . . . 102
6.2 Maximum degree of a vertex and size of a face of (G, Π) . . . 106
6.3 Maximum degree of a vertex and size of a face of a Π-contractible subgraph of (G, Π) with a small separator . . . 115
6.4 Discussion . . . 120
6.4.1 Duality between vertices and faces in (G, Π) . . . 121
6.4.2 Toward tight bounds for the maximum degree and maximum size of a face in (G, Π) . . . 122

7 Bounds on the order of the excluded minors for surfaces 125
7.1 Balanced separator decompositions . . . 126
7.2 Finding a subgraph G_0 of (G, Π) with Π-contractible bridges . . . 129
7.3 Bounding the order of a 2-connected Π-contractible subgraph of (G, Π) . . . 130
7.4 A quasi-polynomial bound for the order of the excluded minors for surfaces . . . 132
7.5 An almost quadratic bound for the order of the excluded minors for surfaces . . . 135

8 Conclusion and open problems 141

A Lower bound for the excluded minors for a surface 143
A.1 Proof of the lower bound . . . 143
A.2 Heawood's bound . . . 144

Bibliography 144

Publication List 151

Index 153

Symbol index 157

Chapter 1

Summary

A famous mathematical puzzle, known as the *three utilities problem*, asks whether it is possible to connect three houses to each of three utilities (water, gas, and electricity) using pipes laid on the same level, in such a way that no two pipes intersect. A representation of the problem is shown in Figure 1.1.

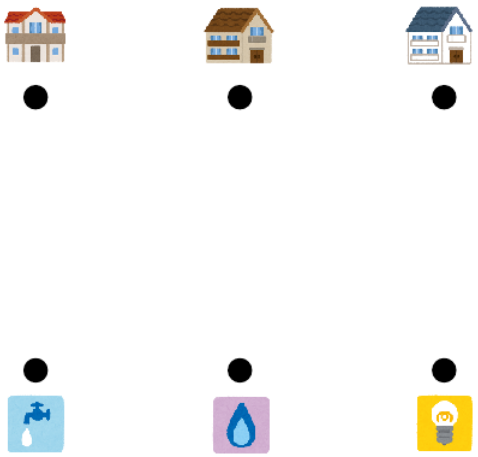


Figure 1.1: The three utilities problem

Although the exact origin of this mathematical puzzle is unknown, it dates back at least to 1900. While the formulation of the problem has varied over the years, its central idea has always remained the same: three elements on one side, denoted by A, and three elements on the other side, denoted by B, are placed on a plane (for instance on a sheet of paper), and the goal is to draw connections between every element of A and every element of B without the connections intersecting.

The three utilities problem admits no solution. Regardless of how the houses and utilities are arranged on the plane, and regardless of how the pipes are drawn, it is impossible to connect every house to every utility without some pipes crossing each other. In fact, the best possible arrangement still requires one crossing, as illustrated in Figure 1.2.

This mathematical puzzle can be reformulated in the language of graph theory. In graph theory, a *graph* is a mathematical structure consisting of *vertices*, which can be represented by

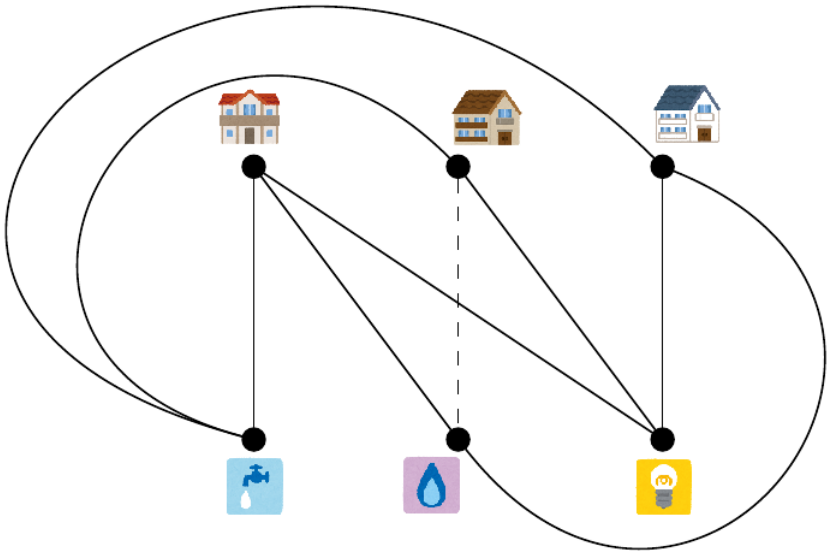


Figure 1.2: An attempt at solving the three utilities problem

points drawn on the plane, and *edges* connecting pairs of vertices, which can be represented by lines drawn between two points.

The three utilities problem can then be translated into the following graph-theoretic question. The graph $K_{3,3}$ consists of 6 vertices divided into two parts consisting of 3 vertices, where every vertex in one part is connected to every vertex in the other group. A representation of the graph $K_{3,3}$ is shown in Figure 1.3. The question is whether the graph $K_{3,3}$ can be drawn in the plane without any two edges crossing.

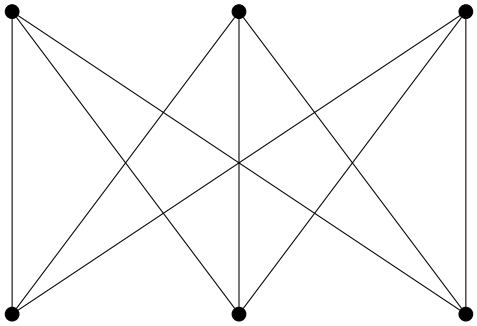


Figure 1.3: The graph $K_{3,3}$

The definition of a graph is very simple, which makes graphs remarkably versatile. Graphs can be used to represent social networks, road or railway networks, and many other real-world systems. Among them, some graphs can be drawn on the plane in such a way that none of their edges cross. These special graphs are called *planar graphs*.

Planar graphs naturally appear in many real-world applications of graph theory. For example, a road map can be represented by a graph in which vertices correspond to road intersections

and edges correspond to the roads connecting them. Such graphs are often planar, since roads typically only cross at intersections. Nevertheless, bridges, underground tunnels, and elevated highways may create crossings that prevent the graph from being planar.

Another example arises in computer engineering. A *CPU* (Central Processing Unit) is the main processor of a computer, often described as its "brain". A CPU is a flat electronic chip, and in early or rudimentary processors, the electrical connections all lie on the same surface. These connections should not cross, as this could create short circuits. Consequently, a CPU can be represented by a graph whose vertices correspond to electronic components and whose edges represent the electrical connections between them. For simple CPUs, these graphs are therefore planar, since all components and connections lie on the same flat surface.

It is important to note that a graph is simply a way of representing relationships between objects: the objects themselves are represented by vertices, and the relationships between them are represented by edges. As a consequence, a graph is not attached to a single drawing in the plane. The same graph can be drawn in many very different ways, as illustrated in Figure 1.4.

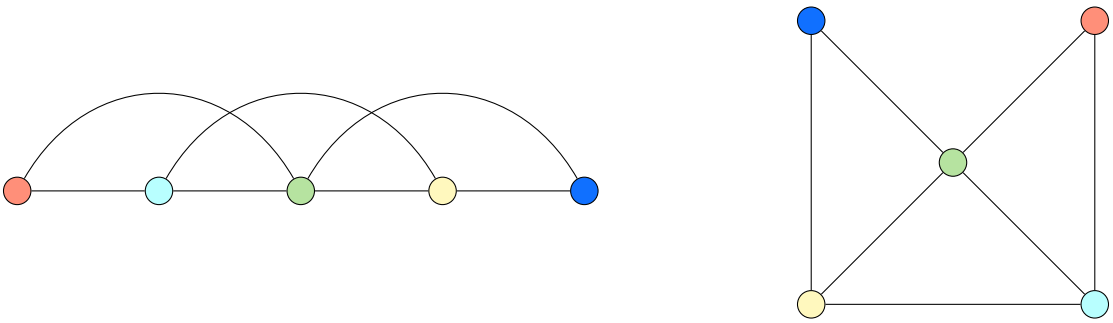


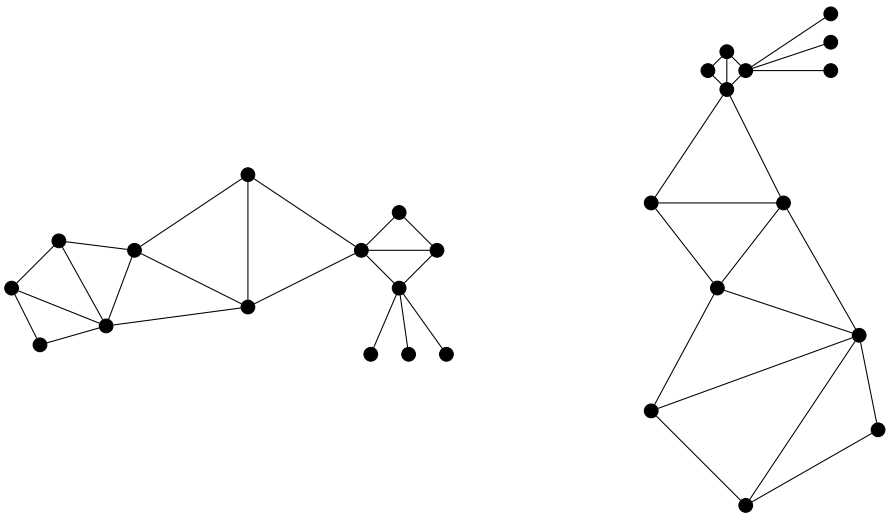
Figure 1.4: Two drawings of the same graph

In particular, even a planar graph can be drawn in a way that creates edge crossings. A graph is considered planar if there exists at least **one** way to draw it in the plane without any edges crossing. It is therefore not required that every drawing of the graph has this property; in fact, this is usually impossible.

Remark also that, in a representation of a planar graph in the plane without crossing edges, any deformation of the drawing, by translation, rotation, scaling (expanding or shrinking), or any combination of these transformations, still yields a drawing of the same planar graph in the plane without edge crossings. For an illustration of such a deformation, see Figure 1.5.

There is a natural equivalence between drawings of planar graphs in the plane without crossings and drawings on the sphere without crossings. To see this, first consider a graph G drawn in the plane without crossings. One can enclose the drawing in a bounded region (for instance, a curved square) and then smoothly deform the drawing so that it fits entirely inside this region on the sphere.

Conversely, suppose that G is drawn on the sphere without crossings. By slightly shrinking and repositioning the drawing, we can fit it inside a small curved square on the sphere. This region can then be "flattened", yielding a drawing of G in the plane without crossings, as illustrated in Figure 1.6.



(a) (b)

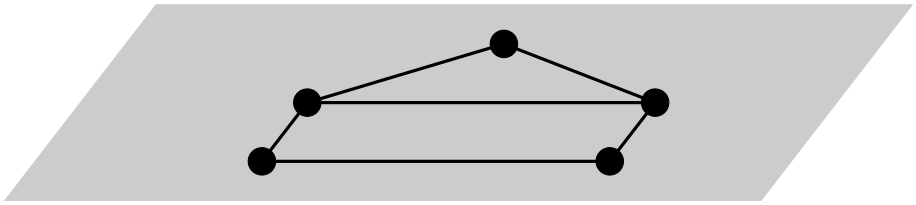
Figure 1.5: Deformation of a drawing of a planar graph in the plane without edge crossings. In both figures, the same graph is represented. In Figure (b), the drawing from Figure (a) has been deformed: it has been rotated by 90° counterclockwise, its left part has been stretched, and part of its right side has been compressed.

After these observations on graphs and their representations, a natural question is whether every graph can be drawn in the plane without edge crossings. In other words, are all graphs planar? The answer is clearly no. The graph $K_{3,3}$ introduced in the mathematical puzzle above provides a counterexample. Thus, not all graphs are planar.

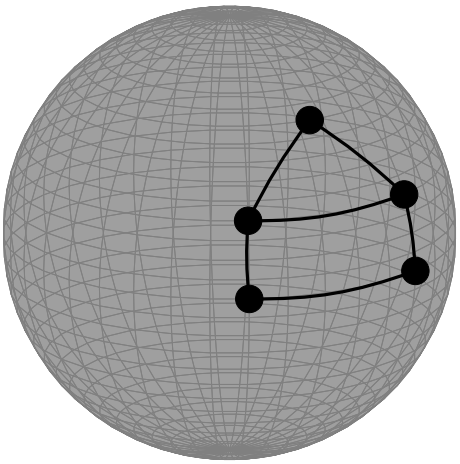
It is possible to draw non-planar graphs without edge crossings if we allow the plane to be augmented with additional structures such as bridges (or, equivalently, tunnels). Consider, for instance, the graph $K_{3,3}$. This graph is not planar and, as shown in Figure 1.2, the best possible drawing in the plane still contains a single crossing between two edges. However, if we introduce a bridge to "lift" one of the crossing edges, we obtain a way to draw $K_{3,3}$ without crossings in the plane plus a bridge (see Figure 1.7).

This extended notion of the plane, augmented with additional "bridges", is also useful for modeling real-world situations that go beyond purely planar structures. For instance, modern road networks include bridges and tunnels to avoid direct crossings. This is particularly common when local roads pass over or under highways. Another example is that of CPUs: modern processors are no longer built on a single flat layer but on several stacked levels. These layers communicate with each other, and can be interpreted as introducing "bridges" between different parts of the structure.

Remark that the exact position of the bridges in the plane does not matter. Indeed, since drawings without crossings can be continuously deformed without changing the underlying graph or creating crossings, we are free to modify the positions, sizes, and shapes of the bridges, as well as their placement in the plane. In other words, a bridge can be "stretched", "moved", or "reshaped" within the plane, and the resulting drawing still represents the same graph and does not have crossings.



(a) The graph H drawn in the plane without crossing edges



(b) The graph H drawn on the sphere without crossing edges

Figure 1.6: A graph H, commonly called the *house graph*, drawn in the plane and on the sphere without edge crossings. The two drawings illustrate the equivalence between representations of graphs in the plane and on the sphere.

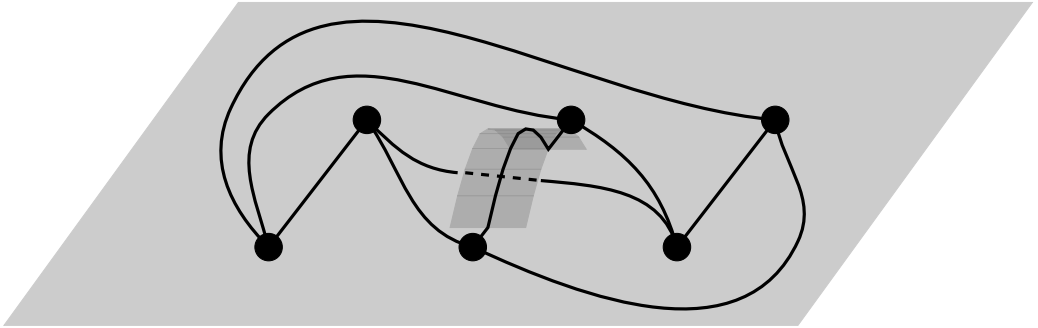


Figure 1.7: The graph $K_{3,3}$ drawn without crossings on the plane plus a bridge



Figure 1.8: A canonical plane with 3 bridges

In particular, this means that we can fix a specific model of the plane equipped with k bridges, by choosing once and for all the positions, sizes, and shapes of these bridges. Then, any graph that admits a drawing in the plane without crossings with k added bridges also admits a drawing without crossings in this fixed model of the plane with k bridges. In other words, the exact geometric placement of the bridges is irrelevant: any such drawing can be continuously deformed so as to fit into this chosen configuration. For an illustration, see Figure 1.8.

Remark then that if a graph can be drawn in a plane with k bridges, it can also be drawn in a plane with $k+m$ bridges for any $m \geq 0$. Indeed, it is enough to add additional bridges in regions of the drawing where they do not interfere with the existing edges.

Exactly as for the plane and the sphere, there is a correspondence between drawings without crossings in the plane with k bridges and drawings without crossings on the k-holed torus.

We call the torus, double torus and triple torus the 1-holed, 2-holed and 3-holed tori, respectively. We also identify the sphere with the 0-holed torus. For $k \geq 0$, we denote by $\mathbb{S}_k$ the k-holed torus .An illustration of $\mathbb{S}_1$, $\mathbb{S}_2$, and $\mathbb{S}_3$ is given in Figure 1.9. The family $(\mathbb{S}_k)_{k\geq 0}$ consists of the so-called *orientable surfaces* and the k-holed torus $\mathbb{S}_k$ is called the *orientable surface of genus k*.

Finally, we can easily show that any graph can be drawn on one of the surfaces $\mathbb{S}_k$ (for $k \geq 0$). To see this, take any graph G and start by drawing it in the plane, possibly with edge crossings. For each crossing, we introduce a bridge in order to eliminate the intersection; several bridges may overlap, which is not an issue. Suppose that this process introduces m bridges. We have seen that a drawing in the plane with m bridges is equivalent to a drawing on the m-holed torus $\mathbb{S}_m$. It follows that G admits a drawing on one of the surfaces in the family $(\mathbb{S}_k)_{k\geq 0}$.

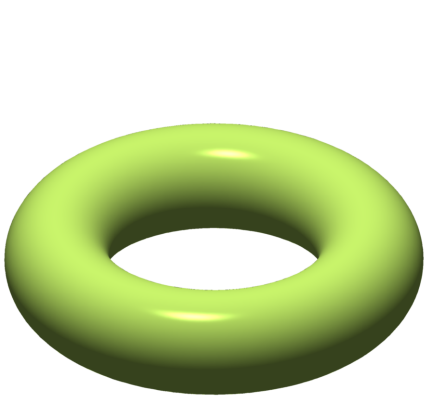
For a graph G, we define the *genus* of G as the smallest integer k such that G admits a drawing in the surface $\mathbb{S}_k$. As shown above, if a graph G has genus k, then it can also be drawn in any surface of genus $k+m$, for any $m \geq 0$.

A natural question is then: how can we determine the genus of a graph? More generally, how can we distinguish between graphs of genus at most g and graphs of genus at least $g+1$?

First, let G be a graph drawn in $\mathbb{S}_k$ without crossings. By deleting a vertex or an edge of G, we obtain a drawing of the resulting graph in $\mathbb{S}_k$ without crossings.

We say that H is a *subgraph* of G if it can be obtained from G by a sequence of vertex and edge deletions. The observation above shows that if G can be drawn in $\mathbb{S}_k$ without crossings, then so can all of its subgraphs. In particular, for any subgraph H of G, the genus of H is at most the genus of G.

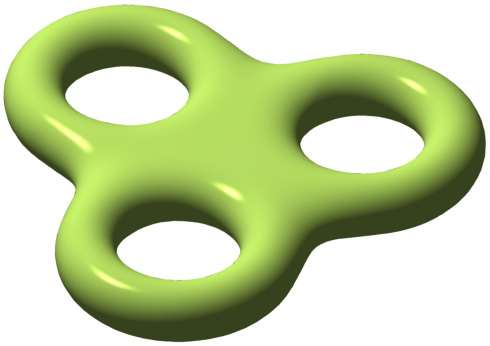
Another operation that preserves drawings is *vertex contraction*. Let G be a graph drawn in $\mathbb{S}_k$ without crossings, and let v be a vertex of G with exactly two neighbors. Contracting v



(a) The torus (or 1-holed torus)



(b) The double torus (or 2-holed torus)



(c) The triple torus (or 3-holed torus)

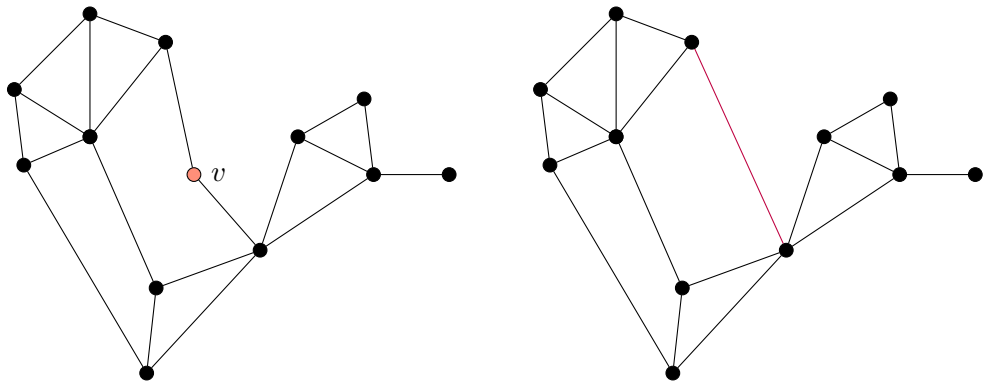
Figure 1.9: Illustrations of the torus, double torus and triple torus (respectively 1-holed, 2-holed and 3-holed tori)

consists of deleting the vertex v and adding an edge between its two neighbors (see Figure 1.10). Remark that the resulting graph obtained from G by contracting v also admits a drawing in $\mathbb{S}_k$ without crossings.

We say that H is a *topological minor* of G if H can be obtained from G by a sequence of vertex deletions, edge deletions, and contractions of vertices of degree two. The discussion above shows that if G can be drawn in $\mathbb{S}_k$ without crossings, then so can all of its topological minors. In particular, for any topological minor H of G, the genus of H is at most the genus of G. Note that every subgraph of G is also a topological minor of G.

A direct consequence of the result above is that if a graph G contains a graph H of genus k as a topological minor, then the genus of G is at least k. In other words, finding a topological minor of large genus (for instance genus k) inside a graph G certifies that G itself has large genus (at least k).

For example, if a graph G contains $K_{3,3}$ (see Figure 1.3 for a representation of $K_{3,3}$) as a topological minor, then, since $K_{3,3}$ has genus 1, the graph G has genus at least 1, and is therefore



(a) A graph G with a vertex v that has two neighbors highlighted in blue

(b) The graph obtained from G by contracting v

Figure 1.10: A vertex contraction

non-planar. Hence, one way to prove that a graph is non-planar is to show that it contains $K_{3,3}$ as a topological minor.

We define the graph K_5 as the graph with 5 vertices which are all pairwise adjacent. Exactly as $K_{3,3}$, K_5 is not planar. A classical theorem of Kuratowski [Kur30], dating back to 1930, states that planar graphs are exactly those graphs that do not contain $K_{3,3}$ or K_5 (see Figure 1.11 for a representation of K_5) as a topological minor.

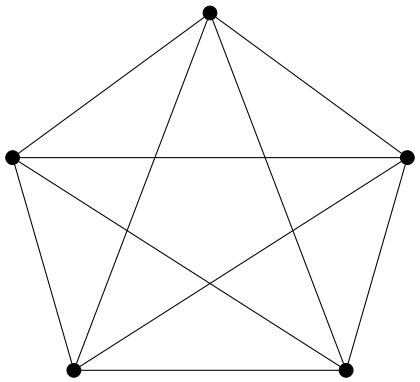


Figure 1.11: The graph K_5

Theorem (Kuratowski's theorem [Kur30])**.** *A graph is planar if and only if it does not contain $K_{3,3}$ or K_5 as a topological minor.*

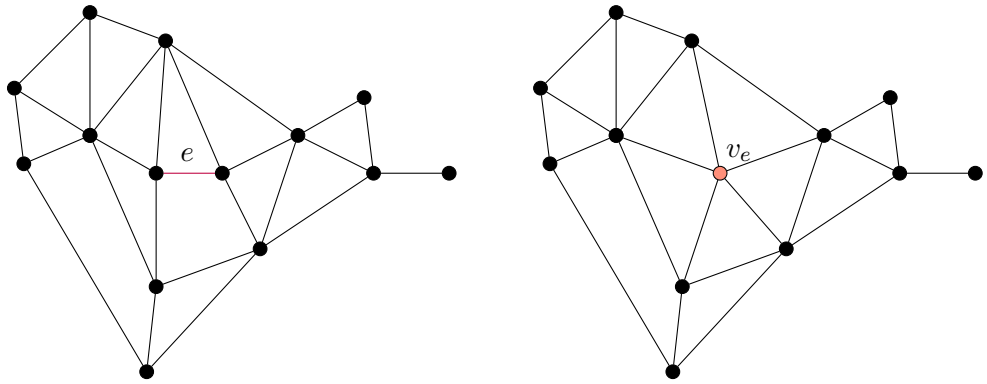
Does a similar characterization exist for graphs of genus at most g, beyond the case of planar graphs (which correspond to genus 0)?

The graph minor theorem [RS04], established by Robertson and Seymour in a long series of papers, implies that such a characterization indeed exists for every surface. More precisely, for

each $g \in \mathbb{N}$, there exists a finite family of graphs $\mathcal{G}_g$ such that a graph G has genus at most g if and only if it does not contain any graph in $\mathcal{G}_g$ as a topological minor.

The graph minor theorem does not rely on the notion of topological minors, but instead on a closely related notion called *minors*. The contraction operation can be extended to edges as follows. Let G be a graph and let e be an edge of G. *Contracting* e consists of identifying its two endpoints into a single vertex v_e, which is then connected to all former neighbors of these endpoints (see Figure 1.12). We denote by G/e the graph obtained from G by contracting the edge e.

We observe that a crossing-free drawing of G in $\mathbb{S}_k$ naturally induces a crossing-free drawing of G/e in $\mathbb{S}_k$: by continuously contracting the edge e on the surface until its endpoints merge, we obtain a crossing-free drawing of G/e in $\mathbb{S}_k$.



(a) A graph G with an edge e highlighted in blue

(b) The graph obtained from G by contracting e

Figure 1.12: An edge contraction

We say that H is a *minor* of G if H can be obtained from G by a sequence of vertex deletions, edge deletions, and edge contractions. The discussion above shows that if G can be drawn in $\mathbb{S}_k$ without crossings, then so can all its minors. In particular, for any minor H of G, the genus of H is at most the genus of G.

We also observe that if v is a vertex of G with exactly two neighbors, then contracting v is equivalent to contracting one of the two edges connected to v. It follows that every topological minor of G is also a minor of G.

Kuratowski's theorem has an analog, using minors instead of topological minors. This classical theorem was proved by Wagner [Wag37] in 1939:

Theorem (Wagner's theorem [Wag37])**.** *A graph is planar if and only if it does not contain $K_{3,3}$ or K_5 as a minor.*

The graph minor theorem, when applied to surfaces, can therefore be stated as follows:

Theorem (Graph minor theorem for surfaces [RS86b])**.** *For $g \in \mathbb{N}$, there exists a finite list of graphs $\mathcal{G}_g$ so that a graph has genus at most g if and only if it does not contain any graph from $\mathcal{G}_g$ as a minor.*

This is in fact equivalent to the corresponding statement formulated in terms of topological minors: it can be shown that there exists a finite set of forbidden minors if and only if there exists a finite set of forbidden topological minors.

Although the graph minor theorem answers the question of obtaining an exact characterization of graphs of genus at most g, the proof of Robertson and Seymour is purely existential and only guarantees the existence of a finite family $\mathcal{G}_g$ for each $g \in \mathbb{N}$. However, it does not provide any explicit information about the graphs in $\mathcal{G}_g$.

The new direction in this research area therefore became the determination of the exact lists of forbidden minors, at least for small values of g. For $g = 1$, extensive work has been carried out to identify the list of forbidden minors (see [DH72, Dec78, BW86, Neu93, Juv95, NM97, Hla97, Cha02, Woo06, GMC09, MŠ14, MW18]). However, this effort has not yet succeeded: more than 17,500 forbidden minors have been found, but there is no evidence that this constitutes the complete list of elements of $\mathcal{G}_1$.

The forbidden minor sets $\mathcal{G}_g$ ($g \in \mathbb{N}$) play an important role in several algorithms for graphs on surfaces. They are used in algorithms that decide whether a graph has genus at most g, compute the genus of a graph, or construct an embedding of a graph into a given surface.

Because of their algorithmic importance as well as their intrinsic theoretical interest, research on forbidden minors for surfaces did not stop after the failure to determine the complete list $\mathcal{G}_1$. Instead, the focus shifted from identifying the exact forbidden minors to studying their structural properties and deriving quantitative bounds. In particular, much effort has been devoted to obtaining upper bounds on the order (i.e., the number of vertices) and on the treewidth (see Section 2.1 for a definition) of the forbidden minors in $\mathcal{G}_g$ for $g \in \mathbb{N}$. Indeed, such bounds can often be used in place of the explicit list of forbidden minors in algorithmic applications like the one cited above.

Seymour [Sey93] was the first to establish such bounds in 1993. He showed that, for $g \in \mathbb{N}$, the forbidden minors in $\mathcal{G}_g$ have treewidth $O(g^3)$ and order $O(2^{2^k})$, where $k = (3g + 9)^9$.

In this thesis, we improve upon these bounds. It was widely presumed that his double-exponential upper bound on the order of the forbidden minors in $\mathcal{G}_g$ for $g \in \mathbb{N}$ is not asymptotically optimal. In particular, it was largely believed (private communication) that this bound should be polynomial in g, or even almost linear.

We improve the upper bound on the treewidth of the forbidden minors in $\mathcal{G}_g$ for $g \in \mathbb{N}$ to $O(g \log g)$, and the upper bound on their order to $g^{2+o(1)}$. This constitutes the first improvement on the bounds for both the treewidth and the order of the forbidden minors in $\mathcal{G}_g$ in more than thirty years. Moreover, the new bound on the order of these graphs comes close to the conjectured optimal bound $\tilde{O}(g)$.

Together with these two main results, we establish several additional results on the forbidden minors in $\mathcal{G}_g$ for $g \in \mathbb{N}$. In particular, we improve the bound on the size of the forbidden structure introduced by Seymour [Sey93] from $O(g)$ to $O(\log g)$ (see Chapter 5). We also refine a result of Thomassen [Tho97], which states that a forbidden minor in $\mathcal{G}_g$ does not contain a grid of size $O(g\sqrt{g})$ as a topological minor, by improving this bound to $O(\sqrt{g} \log g)$ (see Subsection 5.4.2).

Chapter 2

Introduction

2.1 Preliminaries 25
2.1.1 Graph theory 26
2.1.2 Graphs on surfaces 31
2.2 An introduction to excluded minors for surfaces 39
2.2.1 An introduction to minor-closed classes of graphs 39
2.2.2 A historical retrospective on excluded minors for surfaces 42
2.2.3 Algorithmic consequences 45
2.3 Outline of this thesis 46

In this chapter, we present the topic of this manuscript, namely excluded minors for surfaces. This chapter provides a broad introduction, starting from basic notions in graph theory and graphs on surfaces and culminating in the current state-of-the-art results on excluded minors for surfaces.

We begin by recalling fundamental notions from graph theory and graphs on surfaces in Section 2.1.

We then introduce the framework of minor-closed classes of graphs in Section 2.2. This section also contains a detailed historical overview of the problem of excluded minors for surfaces, as well as a discussion of the algorithmic consequences of the current theory developed in this context.

Finally, we outline the structure of the remainder of the manuscript.

2.1 Preliminaries

In this section, we introduce the main concepts and objects used throughout this manuscript. We present standard definitions and notation from graph theory and graphs on surfaces.

For a more complete introduction to these two areas, we refer the interested reader to classical textbooks commonly used as references for the definitions and notions presented below: [Die24, BM07] for graph theory and [MT01] for graphs on surfaces.

2.1.1 Graph theory

For graph theory notions, our definitions and notation follow for the most part the textbook *Graph Theory* [Die24].

Graph, vertex and edge

We define a *graph* G to be a couple (V, E) consisting of two sets: V is a finite set called the *vertex set* of G and $E \subseteq \binom{V}{2}$ (the elements of E are unordered subsets of V of size 2) is called the *edge set* of G. For a graph G, we denote by $V(G)$ and $E(G)$ respectively its vertex set and its edge set. The elements in $V(G)$ are the *vertices* of G and the elements in $E(G)$ are the *edges* of G.

We call the number of vertices of G its *order* and denote it by $|V(G)|$. We denote by $|E(G)|$ the number of edges of G.

Let G be a graph. Let $e \in E(G)$. Then, we say that a vertex $v \in V(G)$ is *incident* to e if $v \in e$. The two vertices incident with e are called the *endpoints* of e. Two vertices u, v of G are said to be *adjacent* or *neighbors* if uv is an edge of G.

We define a *multigraph* G to be a couple (V, E) consisting of two sets V and E, which are finite sets called respectively the *vertex set* and the *edge set* of G, together with a map $E \to V \cup V^2$ that assigns to each edge either one or two vertices that are its *endpoints*. Informally, multigraphs are special types of graphs that can have *loops* (edges whose endpoints are the same vertex) and *multiple edges* or *parallel edges* (several edges between the same pair of vertices).

In this manuscript, unless specified otherwise, we consider simple graphs: that is to say graphs without multiple edges and without loops. Moreover, the graphs we consider are undirected: for a graph G, if $u, v \in V(G)$ are the endpoints of $e \in E(G)$, then $uv \in E(G)$ and $vu \in E(G)$ are two notations of the same edge.

Degree

The *degree* of a vertex v of G is the number of edges incident to v; we denote it by $d_G(v)$ (or $d(v)$ if G is clear from the context). The *maximum degree* of G is the maximum degree of a vertex of G and is denoted $\Delta(G)$. We define the *average degree* of G to be the number

$$\frac{\sum_{v \in V(G)} d_G(v)}{|V(G)|} = \frac{2|E(G)|}{|V(G)}$$

Induced subgraph and subgraph

A *subgraph* G' of G is a graph such that $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$. We write $G' \subseteq G$ to signify that G' is a subgraph of G. Moreover, if $G' \subseteq G$ and G' contains all the edges $xy \in E(G)$ for $x, y \in V(G')$, then G' is said to be an *induced subgraph* of G.

We denote by $G - v$ (resp. $G - e$) the graph obtained from G after the deletion of a vertex v (resp. an edge e), and we denote $G - X$ the graph obtained from G after the deletion of a set X of vertices or edges.

Walk, path, circuit and cycle

A *walk* in G is a finite sequence of vertices of G in which each pair of consecutive vertices is adjacent. A *path* is a walk whose vertices are all distinct, except possibly its two endpoints. A *circuit* (or *closed walk*) is a walk whose first and last vertices coincide. A *cycle* is a circuit whose vertices are all distinct except for its two endpoints.

We identify walks, paths, circuits, and cycles with the subgraphs they induce. Moreover, we identify circuits and cycles that differ only by the choice of their starting vertex.

Let P be a path of G and let v_1 and v_2 be the two endpoints of P. The *interior* $\text{int}(P)$ of P is the set of vertices of P obtained after removing its two endpoints, that is to say it is the restriction $P - \{v_1, v_2\}$ of P.

Topological minor and minor

The *contraction* of an edge $e = uv \in E(G)$ is the operation of replacing the vertices u and v in G by a unique new vertex w that is adjacent to all the neighbors of u and v in G.

A *topological minor* H of G is a graph obtained from G by successive operations of deletion of a vertex, deletion of an edge and contraction of an edge which has at least one of its endpoints of degree 2.

A *subdivision* G of a graph H is a graph obtained from G by replacing its edges by paths of arbitrary length between the same endpoints.

Remark that the graph G is a subdivision of the graph H if and only if H is obtained from G by successive contraction of edges which has at least one of its endpoints of degree 2. Remark moreover that the definitions of topological minor and subdivision imply that H is a topological minor of G if and only G contains a subdivision of H as a subgraph.

A *minor* H of G is a graph obtained from G by successive operations of deletion of a vertex, deletion of an edge, and contraction of an edge. A *proper minor* of G is a minor of G different from G.

We denote by G/e the graph obtained from G after the contraction of an edge e. We denote by G/X the graph obtained from G after the contraction of a set X of edges (remark that the order in which the edges in X are contracted is not important).

Union and intersection of graphs

For two graphs G and H, we define the union $G \cup H$ of G and H as the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$. Similarly, we define the intersection $G \cap H$ of G and H as the graph with vertex set $V(G) \cap V(H)$ and edge set $E(G) \cap E(H)$.

Connectivity

A graph G is *connected* if, for every pair of vertices $u, v \in V(G)$, there is a path in G between u and v. Let $k \in \mathbb{N}^*$, G is *k-connected* if, for every set $V \subseteq V(G)$ of size at most $k-1$, $G - V$ is still connected. A *separator* of G is a set $V \subseteq V(G)$ so that $G - V$ is not connected. Let $k \geq 1$, if the set V has size k, then we call it a *k-separator* of G. If G contains a 1-separator $\{v\}$, v is called a *cutvertex* of G. Remark that, for $k \geq 1$, a graph contains no $(k-1)$-separator if and only if it is k-connected.

A *separation* of G is a pair (A, B) of subgraphs of G such that $A \cup B = G$ and $A \cap B$ contains no edge. Remark that $V(A \cap B)$ is a separator of G. For $k \in \mathbb{N}$, we say that (A, B) is a *k-separation* of G if (A, B) is a separation of G and $|V(A \cap B)| = k$.

A graph G is *disconnected* if it is not connected. In that case, we call the *connected components* of G its maximal connected subgraphs.

Classical classes of graphs

A class of graphs is a collection of graphs (defined up to isomorphism), usually defined by some specific property.

A *complete graph* is a graph K in which each pair of distinct vertices is adjacent. We denote by K_n the unique (up to isomorphism) complete graph of order n.

Let $m, n \in \mathbb{N}^*$. A $m \times n$ *grid* is (up to isomorphism) the graph with vertex set $\{1, ..., m\} \times \{1, ..., n\}$ and, for $1 \leq i, i' \leq m$ and $1 \leq j, j' \leq n$, (i, j) is adjacent to (i', j') if and only if $i = i'$ and $|j - j'| = 1$ or $j = j'$ and $|i - i'| = 1$. See Figure 2.1 for an illustration.

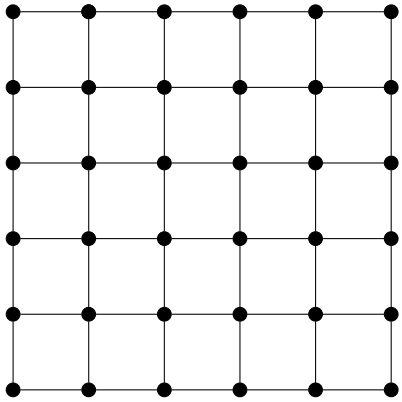


Figure 2.1: The 6×6 grid

A *bipartite graph* is a graph whose vertex set admits a partition into two sets A and B so that any pair of vertices in A (resp. B) are not adjacent. The bipartite graphs are exactly the graphs that do not contain odd cycles.

Lemma 2.1.1 ([Die24, Proposition 1.6.1])**.** *A graph is bipartite if and only if it contains no odd cycle.*

A *tree* T is a connected graph without any cycle. A vertex of T with degree at most 1 is a *leaf* of T, the vertices of T with degree at least 2 are the *internal vertices* of T. Every non-empty tree has a leaf (consider for instance the endpoint of a longest path in T).

We can distinguish a vertex v of T, called the *root*. In that case, T is said to be *rooted* in v. We define the *diameter* of a tree as the order of its longest path.

A *spanning tree* T of a graph G is a subgraph of G that is a tree and contains all the vertices of G. A *forest* F is a graph that is a union of trees on disjoint vertex sets. Remark that forests are therefore precisely the graphs that do not contain any cycle. Forests are in particular bipartite graphs.

The following two lemmas are classical results on forests.

Lemma 2.1.2. *Let F be a forest, then its average degree is strictly less than 2.*

Proof. Let us prove the result by induction on the order of F.

For $|V(F)| = 1$, it is clear.

Suppose that the result is true for k and let us prove it for $k+1$: let F be a forest of order $k+1$.

Let l be a leaf of F, then $F - l$ has order k and thus average degree strictly less than 2.

Now $\sum_{v \in V(F)} d(v) \leq 1 + \sum_{v \in V(F)-l} d(v) < 1 + 2k < 2(k+1)$. Finally, it implies that the average degree of F is strictly less than 2. □

Lemma 2.1.3. *Let F be a forest so that each internal vertex of F has degree at least 3, then there are more leaves in F than internal vertices.*

Proof. By Lemma 2.1.2, F has average degree strictly less than 2. Therefore

$$\sum_{v \in V(F)} d(v) < 2|V(F)|$$

Let l and i be respectively the number of leaves and internal vertices in F. Then,

$$f + 3i < 2(f + i)$$
$$i < f$$

Finally, this proves that there are more leaves in F than internal vertices. □

Tree decomposition

A *tree decomposition* of G is a pair $(T, (V_t)_{t \in V(T)})$ with T a tree and, for every $t \in V(T)$, $V_t \subseteq V(G)$ with the following properties:

- $\bigcup_{t \in V(T)} V_t = V(G)$,
- for every $e = uv \in E(G)$, there exists $t \in V(T)$ so that $u, v \in V_t$,
- for $t, t', t'' \in V(T)$ so that t' is on the path between t and t'' in T, $V_t \cap V_{t''} \subseteq V_{t'}$.

The *width* of a tree decomposition $(T, (V_t)_{t \in V(T)})$ of G is $\max_{t \in V(T)} |V_t| - 1$ and the *treewidth* $tw(G)$ of G is the minimal width of its tree decompositions.

A *path decomposition* of G is defined analogously to a tree decomposition of G with the difference that the tree in the tree decomposition is required to be a path in the path decomposition. Formally, a *path decomposition* of G is a tree decomposition $(P, (V_p)_{p \in P})$ of G with P being a path.

The *width* of a path decomposition $(P, (V_p)_{p \in P})$ of G is defined as its width as a tree decomposition and the *pathwidth* $pw(G)$ of G is the minimal width of its path decompositions.

We say that a tree decomposition $(T, (V_t)_{t \in V(T)})$ of G, possibly subject to additional constraints, is *minimal*, if $|V(T)|$ is chosen minimum among all tree decompositions of G satisfying these constraints.

The treewidth of forests and of the $k \times k$ grid ($k \in \mathbb{N}^*$) are classical results:

Proposition 2.1.4. *The treewidth of a forest is 1.*

Proposition 2.1.5. *For $k \in \mathbb{N}^*$, the $k \times k$ grid has treewidth k.*

Block-cutvertex graph

We define the following equivalence relation on $E(G)$: $e_1 \sim e_2$ if $e_1 = e_2$, or there is in G a cycle that contains e_1 and e_2. A *(2-connected) block* of G is the subgraph of G containing the edges of an equivalence class of the relation $\sim$ together with their endpoints. Remark that H is a block of G if and only if H is a maximal 2-connected subgraph of G.

Let C be the set of cutvertices of G and let B be the set of its 2-connected blocks. We can then construct a natural bipartite graph with vertex set $B \cup C$ and with edges bc for $b \in B$ and $c \in C$ such that c is in the block b. This bipartite graph is called the *block-cutvertex graph* of G and is a forest.

Lemma 2.1.6. *Let G be a graph, then the block-cutvertex graph of G is a forest.*

Proof. Let F be the block-cutvertex graph G. By definition, F is a bipartite graph. Let B, C be the partition of its vertex set into vertices that correspond respectively to blocks and to cutvertices.

By contradiction, suppose that F contains a cycle $b_1, c_1, ..., b_k, c_k$ ($k \geq 2$) with, for $1 \leq i \leq k$, $b_i \in B$ and $c_i \in C$.

Let e_1 (resp. e_2) be an edge in the block corresponding to b_1 (resp. b_2). Remark that there is a cycle in G that contains both e_1 and e_2: this cycle goes through the cutvertices $c_1, ..., c_k$ and through the blocks $b_1, ..., b_k$. However, by definition of a 2-connected block, e_1 and e_2 must then belong to the same block. A contradiction. □

A leaf in the block-cutvertex graph of G corresponds to a block of G, because every cutvertex is contained in at least two blocks.

Bridges

Let H be a graph and H_0 be a subgraph of H. A *bridge* B of H on H_0 is either an edge with both ends in H_0 (and its endpoints do not belong to the bridge) or a maximal connected subgraph of $H - V(H_0)$ together with all the edges which have one end in this subgraph and the other end in H_0 (the ends in H_0 does not belong to the bridge). Remark that the bridges of H on H_0 partition $E(H) - E(H_0)$. We say that $w \in H_0$ is an *attach* of a bridge B on H_0 if there exists $v \in V(B)$ such that $e = vw \in E(B)$.

Coloring and chromatic number

A *(proper) coloring* of a graph G is a map $c : V(G) \to S$ such that, for $vw \in E(G)$, we have $c(v) \neq c(w)$. The elements of the set S are called the available *colors*.

For $k \in \mathbb{N}^*$, we say that G is *k-colorable* or that G has a *k-coloring* if G has a coloring $c : V(G) \mapsto \{1, .., k\}$. The *chromatic number* of G is defined as the least integer k so that G is k-colorable. We denote the chromatic number of G by $\chi(G)$.

Let G be a graph, an *independent set* X of G is a set of vertices of G so that the graph $G[X]$ induced by X, that is to say the induced subgraph of G with vertex set X and edge set $E(G) \cap (X \times X)$, does not contain any edge.

Remark that a k-coloring of G is exactly a partition of its vertices into k independent sets (the vertices assigned to the same color form a part in this partition), called the *color classes*.

A graph G consisting of isolated vertices (vertices of degree 0) has chromatic number 1. Moreover, it is easy to show that a graph has chromatic number at most 2 if and only if it is bipartite:

Proposition 2.1.7. *Let G be a graph. Then, G has chromatic number at most 2 if and only if it is bipartite.*

Proof. $\Leftarrow$ Let (A, B) be the bipartite partition of the vertices of G. Then, we define $c : V(G) \mapsto \{1, 2\}$ by $c(v) = 1$ if $v \in A$ and $c(v) = 2$ if $v \in B$. By definition of bipartite graphs, two adjacent vertices v and w of G cannot both belong to A or B. Therefore, v and w have different colors with respect to c. Finally, c is a proper 2-coloring of G and $\chi(G) \leq 2$.

$\Rightarrow$ We suppose that $\chi(G) \leq 2$. Let $c : V(G) \mapsto \{1, 2\}$ be a proper 2-coloring of G. Then, we define $A = \{v \in V(G), c(v) = 1\}$ and $B = \{v \in V(G), c(v) = 2\}$. We will show that G is bipartite with vertex partition (A, B). As c is a proper coloring, no two vertices from A (resp. B) are adjacent. Therefore, G is indeed bipartite with vertex partition (A, B). □

The *degeneracy* is a notion with close links to the chromatic number.

Definition 2.1.8 (Degeneracy)**.** Let H be a graph and $k \geq 0$. We say that H is k-degenerate if every non empty subgraph of H has a vertex of degree at most k.

Indeed, a graph G that is k-degenerate can be showed to be $k + 1$-colorable.

Proposition 2.1.9. *Let G be a k-degenerate graph. Then, G is $k + 1$-colorable.*

Proof. We will proceed by induction on the order of G.

The result is trivial for $|V(G)| = 1$. Suppose that the result holds for graphs of order n. Suppose that G is of order $n+1$. Then, let v be a vertex of G of degree at most k. By hypothesis, the graph $G-v$ is $k+1$-colorable. Let $c' : V(G-v) \mapsto \{1, ..., k+1\}$ be a proper coloring of $G-v$. Remark that, as v has at most k neighbors in G, then there is a color i in $\{1, ..., k + 1\}$ that is not used by any of its neighbors with respect to c'. We finally define $c : V(G) \mapsto \{1, ..., k + 1\}$ by $c(v) = i$ and, for every $w \in V(G) - \{v\}$, $c(w) = c'(w)$. The function c is a proper coloring of G and G is thus $k + 1$-colorable. □

2.1.2 Graphs on surfaces

To describe graphs on surfaces, we follow the definitions and notation of Mohar and Thomassen's book *Graphs on Surfaces* [MT01], with the sole exception that we use the term genus to denote the Euler genus.

Surfaces and simple closed curves

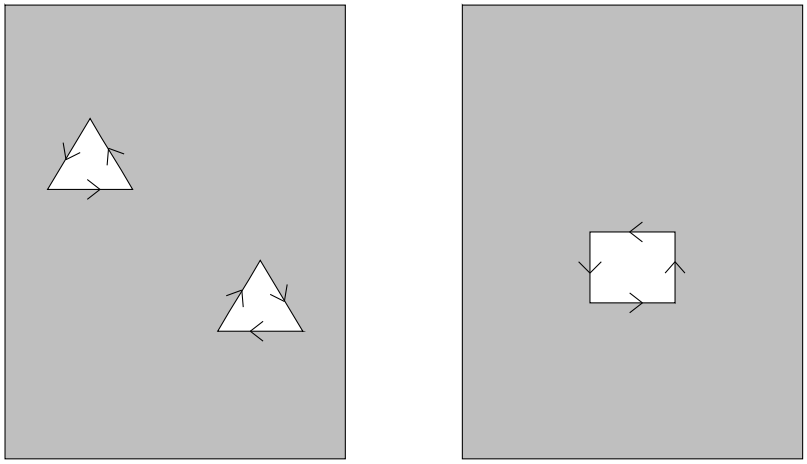
We define a *surface* as a connected compact Hausdorff topological space S so that each point of S has an open neighborhood homeomorphic to the open unit disk in $\mathbb{R}^2$. The surfaces considered here are also referred to as the closed connected surfaces, with "closed" here meaning compact and without boundary.

A *simple closed curve* in a surface S is a continuous 1-1 function $f : [0, 1] \mapsto S$ with $f(0) = f(1)$. We say that a simple closed curve on S *bounds a disk* if it can be continuously deformed into a single point.

After giving a very rigorous definition of surfaces, we will give to the reader a better intuition of what this definition encompasses.

The sphere is defined as the unit sphere in $\mathbb{R}^3$, it is a surface.

Let us now define two operations on surfaces: adding a handle and adding a crosscap. Let S be a surface, we consider two triangles T_1, T_2 that bound a disk on S. We form a new surface S' by deleting from S the interior of T_1 and T_2 and identifying T_1 with T_2 such that the clockwise orientations around T_1 and T_2 disagree, see Figure 2.2(a). We say that the surface S'' is obtained from S by *adding a handle*.



(a) Adding a handle (b) Adding a crosscap

Figure 2.2: Adding a handle or a crosscap to a surface

Now, let T be a quadrangle (with equilateral sides) that bounds a disk on S. We let S'' be the surface obtained by deleting the interior of T and identifying diametrically opposite points of the quadrangle, see Figure 2.2(b). We say that the surface S'' is obtained from S by *adding a crosscap*.

We now consider all surfaces obtained from the sphere $\mathbb{S}_0$ by adding handles and crosscaps. If we add h handles to $\mathbb{S}_0$, we obtain the surface $\mathbb{S}_h$ called the *orientable surface of genus* h ($h \geq 0$). If we add c crosscaps to $\mathbb{S}_0$, we obtain the surface $\mathbb{N}_c$ called the *nonorientable surface of genus* c ($c \geq 1$).

The surfaces $\mathbb{S}_1, \mathbb{S}_2, \mathbb{N}_1, \mathbb{N}_2$ are also called the *torus*, the *double torus*, the *projective plane* and the *Klein bottle*.

We have the following result about adding handles and crosscaps to the sphere:

Proposition 2.1.10 ([MT01, Proposition 3.1.2])**.** *Let S be the surface obtained from the sphere by adding h handles and c crosscaps. If $c = 0$, then $S = \mathbb{S}_h$. Otherwise, $S = \mathbb{N}_{2h+c}$.*

The *classification theorem of closed surfaces* states the following:

Theorem 2.1.11 ([MT01, Theorem 3.1.3])**.** *Every surface is homeomorphic to precisely one of the surfaces $\mathbb{S}_h$ ($h \geq 0$), or $\mathbb{N}_c$ ($c \geq 1$).*

For an orientable surface $\mathbb{S}_h$, the *orientable genus* is defined as h, while for a nonorientable surface $\mathbb{N}_c$, the *nonorientable genus* is defined as c.

The *Euler genus* of a surface is twice its orientable genus if the surface is orientable, and is equal to its nonorientable genus if the surface is nonorientable.

Throughout this manuscript, the *genus* of a surface will always refer to its Euler genus.

Embedding and facial walk

Embeddings of graphs on surfaces can be defined topologically or combinatorially. Both definitions lead to the exact same notion. Here, we give the combinatorial definition of embedding.

Let G be a graph and S a surface, we define an *embedding* Π of G in S to be a pair $\Pi = (\pi, \lambda)$ where $\pi = \{\pi_v | v \in V(G)\}$ associates to each vertex $v \in V(G)$ a cyclic permutation of the edges incident to v and $\lambda : E(G) \mapsto \{-1, 1\}$ associates to each edge $e \in E(G)$ a sign (either -1 or 1) called its *signature*.

A *local change* of an embedding $\Pi = (\pi, \lambda)$ of G changes the clockwise ordering to anticlockwise at some vertex $v \in V(G)$, i.e. π_v is replaced by its inverse π_v^{-1}, and $\lambda(e)$ is replaced by $-\lambda(e)$ for every edge e that is incident with v. Two embeddings of G are *equivalent* if one can be obtained from the other by a sequence of local changes.

The *face traversal procedure* is the following: We start with a vertex $v \in V(G)$ and an edge $e = vw \in E(G)$. Let us traverse the edge from v to w, if the signature of e is 1, then we continue the walk along the edge $e' = \pi_w(e)$, otherwise ($\lambda(e) = -1$), we change π_u by π_u^{-1} for every $u \in V(G)$ and we continue the walk along the edge $e' = \pi_w(e)$. And so forth. The walk is completed when the initial edge e is encountered in the same direction (from v to w) and with the same orientation (π is the same as before we began the walk).

A Π-*facial walk* (or, if clear in the context, *face*) is a closed walk in G determined by the face traversal procedure. We identify the faces with the subgraphs they induce. We denote by $F(G, \Pi)$ the set of the Π-facial walks of (G, Π). We define the *size* of a face to be the number of edges that the walk contains. We define the *maximum face degree* of (G, Π) to be $\Delta_F(G, \Pi) = \max\{|f|, f \text{ is a face in } (G, \Pi)\}$.

Let C be a cycle of G, we define the *signature* of C in Π to be $\lambda(C) = \prod_{e \in C} \lambda(e)$. We say that C is *two-sided* if $\lambda(C) = 1$ and *one-sided* otherwise. We say that an embedding Π of G is *orientable* if every cycle C of G is two-sided, otherwise we say that it is *nonorientable*.

For a disconnected graph G', an *embedding* Π' of G' is defined as the union of the embeddings of its connected components.

Let H be a subgraph of G. Then, the *induced embedding* $\Pi(H)$ of H with respect to Π is obtained from Π by ignoring all vertices in $V(G) - V(H)$ and all edges in $E(G) - E(H)$. More precisely, we restrict the signature λ of Π to $E(H)$, and if $e \in E(H)$ is an edge incident to $v \in V(H)$, then the successor of e in the clockwise ordering around v of $\Pi(H)$ is the first edge of H in the sequence $\pi_v(e), \pi_v^2(e), ...$

Genus and embeddability

Let G be a connected graph with embedding Π. We define the Euler characteristic of Π to be $\chi(\Pi) = |V(G)| - |E(G)| + |F(G,\Pi)|$.

The expression $g(\Pi) = 2 - \chi(\Pi)$ denotes the *genus* of Π, and $g(G)$ denotes the genus of an embedding of G with minimal genus. The formula

$$\chi(\Pi) = 2 - g(\Pi)$$

is called the *Euler formula.*

Let G' be a disconnected graph and let Π' be an embedding of G'. The *genus* of Π' is the sum of the genera of the embeddings corresponding to each of its connected components. Finally, we define $g(G')$ to be the sum of the genera of its connected component.

Proposition 2.1.12. *Let H be a graph and $H_1, ..., H_p$ ($p \geq 1$) be its connected components. We define the genus $g(H)$ of H by*

$$g(H) = g(H_1) + ... + g(H_p)$$

This definition of genus is consistent because the embeddings Π and Π' of the graphs G and G' correspond to topological embeddings of G and G' either in the surface $\mathbb{S}_{g(\Pi)/2}$ if Π is orientable or in the surface $\mathbb{N}_{g(\Pi)}$ if Π is nonorientable.

The genus of a connected graph can be expressed as the sum of the genus of its 2-connected blocks. A proof of this result can be found in [MT01].

Proposition 2.1.13 ([MT01, Theorem 4.4.2])**.** *Let H be a connected graph and $H_1, ..., H_p$ ($p \geq 1$) be its 2-connected blocks. Then,*

$$g(H) = g(H_1) + ... + g(H_p)$$

We say that G is *embeddable* in a surface S, if there exists an embedding Π of G so that $g(\Pi) = g(S)$ and Π is *orientable* if and only if S is orientable.

Separating cycles

Let $C = v_0e_1v_1e_2...v_{l-1}e_lv_0$ be a Π-two-sided cycle of a Π-embedded graph G. Suppose that the signature of each edge of C is positive in Π. We define the *left graph* and the *right graph* of C as follows: for $1 \leq i \leq l$, if $e_{i+1} = \pi_{v_i}^{k_i}(e_i)$, then all the edges $\pi_{v_i}(e_i), \pi_{v_i}^2(e_i), ..., \pi_{v_i}^{k_i-1}(e_i)$ are said to be on the left side of C and the left graph of C, denoted by $G_l(C,\Pi)$, is defined as the union of all bridges on C that attach to C by at least one edge on the left side of C. The right graph $G_r(C,\Pi)$ is defined analogously.

A cycle C of a Π-embedded graph G is Π-*separating* if C is two-sided and $G_l(C,\Pi)$ and $G_r(C,\Pi)$ have no edges in common. If $\Pi(G_l(C,\Pi))$ or $\Pi(G_r(C,\Pi))$ is an induced embedding of genus 0, then we say that C is Π-*contractible.* Suppose without loss of generality that $\Pi(G_l(C,\Pi))$ is an induced embedding of genus 0, we define $\text{int}(C,\Pi_H)$ (resp. $\text{ext}(C,\Pi_H)$) to be the bridges onto C in $G_l(C,\Pi)$ (resp. $G_r(C,\Pi)$). Moreover, we define $\text{Int}(C,\Pi_H) =$

$\text{int}(C, \Pi_H) \cup C$ and $\text{Ext}(C, \Pi_H) = \text{ext}(C, \Pi_H) \cup C$. If it is clear in the context, the mention of the embedding is removed, and we write $\text{Int}(C)$, $\text{int}(C)$, $\text{Ext}(C)$, $\text{ext}(C)$.

If two contractible cycles intersect in a path, then their symmetric difference is also contractible, which is a classical result for graphs on surfaces:

Lemma 2.1.14 ([MT01, Proposition 4.3.1])**.** *Let H be a graph, Π_H-embedded in a surface S_H. Let $x, y \in V(H)$ and P_1, P_2, P_3 be three internally disjoint paths from x to y. Let $C_{i,j} = P_i \cup P_j$ for $1 \leq i < j \leq 3$.*

If two of the three cycles $C_{i,j}$ $(1 \leq i < j \leq 3)$ are Π_H-contractible, then so is the third cycle.

For contractible cycles, we can define the notion of nested cycles:

Definition 2.1.15 (Nested cycles)**.** Let H be a graph with embedding Π_H. Let C and C' be two Π_H-contractible cycles of H such that $C' \subseteq \text{Int}(C, \Pi_H)$. We say that C' is Π_H-*nested* (or *nested* if it is clear in the context) in C.

Let $k \geq 1$ and let $C_0, ..., C_k$ be Π_H-contractible cycles of H. We say that $C_0, ..., C_k$ are Π_H-*nested* if for $0 \leq i \leq k-1$, C_i is Π_H-nested in C_{i+1}.

For instance, in Figure 2.3, the cycle C' is nested in C.

For two Π-contractible nested cycles C and C', we denote by $\text{Int}(C \cup C', \Pi)$ the subgraph of G inside the surface bounded by C and C' in Π, together with C and C'. Therefore, $\text{Int}(C \cup C', \Pi) = \text{Int}(C, \Pi) - \text{int}(C', \Pi)$. In Figure 2.3, $\text{Int}(C \cup C', \Pi)$ is the subgraph whose edges are indicated in blue and the surface bounded by C and C' in Π is indicated in light blue.

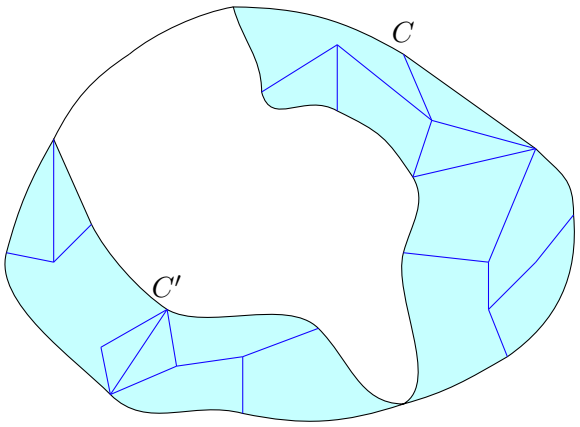


Figure 2.3: The cycle C' is nested in C. The subgraph $\text{Int}(C \cup C', \Pi)$ is represented in blue and the surface bounded by C and C' in Π is represented in light blue.

Planar graphs

A *planar graph* is a graph embeddable on the sphere. Let G be a planar graph. A *planar embedding* Π of G is an embedding of G on the sphere with a distinguished face called the *outer face*.

Remark that any embedding on the sphere induces an embedding in the plane: by distinguishing an arbitrary face to be the outer face and correspond to the unbounded face of G in the plane, we can transform an embedding on the sphere to an embedding in the plane.

Let C be a cycle of G, then C is two-sided and Π-separating. Let $G_l(C,\Pi)$ and $G_r(C,\Pi)$ be the left and right graph of G and suppose without loss of generality that the outer face of Π is in $G_r(C,\Pi)$. Then, we define $\mathrm{Int}(C,\Pi)$, $\mathrm{int}(C,\Pi)$, $\mathrm{Ext}(C,\Pi)$, $\mathrm{ext}(C,\Pi)$ as above, by making sure that the outer face is in the exterior of C.

Planar graphs is a well-studied class of graphs, with many interesting properties. Thereafter, we detailed a few classical results that will be used in this manuscript.

First, Whitney showed in 1933 [Whi92] that, for a given planar graph H, every embedding of H in the plane (or equivalently the sphere) can be obtained from a given embedding by a sequence of "flippings":

Definition 2.1.16 (Flipping)**.** Let H be a 2-connected planar graph with embedding Π_H in the plane. Let C be a cycle of H such that only two vertices v and w of C have incident edges in $\mathrm{ext}(C,\Pi_H)$. Then, we define a *flipping* of H with respect to C as a reembedding of H such that the embedding in $\mathrm{ext}(C,\Pi_H)$ is unchanged. The embedding of $H' = H \cap \mathrm{Int}(C,\Pi_H)$ is changed so that the new embedding of H' is equivalent to the original one, but the clockwise orientations of all the facial cycles are reversed.

Proposition 2.1.17 (Whitney's Theorem [MT01, Theorem 2.6.8], [Whi92])**.** *Let H be a 2-connected plane graph and Π_H be an embedding of H in the plane. Then, any embedding of H in the plane can be obtained from Π_H by a sequence of flippings.*

Then, we show two classical results on degeneracy for planar graphs.

Lemma 2.1.18. *Let H be a planar graph, then H is 5-degenerate.*

Proof. Let H' be a subgraph of H, then H' is also a planar graph. Let Π be an embedding of H' in the sphere. Let $v = |V(H')|$, $e = |E(H')|$ and $f = |F(H',\Pi)|$. By Euler formula,

$$v - e + f = 2$$

Moreover, as H' is a simple graph, every face of (H',Π) has size at least 3 and therefore

$$2e = \sum_{F \in F(H',\Pi)} |F| \geq 3f$$

Consequently

$$\begin{aligned} 3v - 3e + 3f &= 6 \\ 3v - 3e + 2e &\geq 6 \\ 3v - 6 &\geq e \end{aligned}$$

Finally, the average degree of H' is at most $\frac{2e}{v} < 6$ and there exists at least a vertex in H' with degree at most 5. □

Lemma 2.1.19. *Let H be a bipartite planar graph, then H is 3-degenerate.*

Proof. Let H' be a subgraph of H, then H' is also a bipartite planar graph. Let Π be an embedding of H' in the sphere. Let $v = |V(H')|$, $e = |E(H')|$ and $f = |F(H', \Pi)|$. By Euler formula,

$$v - e + f = 2$$

Moreover, as H' is a simple graph and is bipartite, every face of (H', Π) has size at least 4 and therefore

$$2e = \sum_{F \in F(H', \Pi)} |F| \geq 4f$$

Consequently

$$\begin{aligned} 4v - 4e + 4f &= 8 \\ 4v - 4e + 2e &\geq 8 \\ 4v - 8 &\geq 2e \end{aligned}$$

Finally, the average degree of H' is at most $\frac{2e}{v} < 4$ and there exists at least a vertex in H' with degree at most 3. □

Finally, we present the famous Four Color Theorem that states that every planar graph is 4-colorable.

Theorem 2.1.20 (Four Color Theorem [AH77])**.** *Every planar graph is 4-colorable.*

This result was conjectured by Guthrie in 1852 and was solved more than a century later, in 1977, by Appel and Haken that presented a computer-assisted proof [AH77]. As of now, there is no proof of the Four Color Theorem that does not require computer assistance. Another proof of the Four Color Theorem was then found by Robertson, Sanders, Seymour and Thomas [RSST96]. This proof was shorter, simpler and provided a quadratic time algorithm for the 4-coloring of planar graphs. Finally, in 2026, Inoue, Kawarabayashi, Miyashita, Mohar, Thomassen and Thorup [IKM$^+$26] gave a new proof of the Four Color Theorem that provides a near-linear time algorithm for the 4-coloring of planar graphs.

Cutting along a cycle

Let G be a Π-embedded graph. Let C be a Π-separating cycle, then cutting along C gives rise to two graphs $G_l(C, \Pi) \cup C$ and $G_r(C, \Pi) \cup C$ and their induced embedding $\Pi(G_l(C, \Pi))$ and $\Pi(G_r(C, \Pi))$.

Let C be a two-sided Π-nonseparating cycle. Let $\overline{G}$ be the graph obtained from G by replacing C with two copies of C such that all the edges of C on the left side of C are incident one copy of C and all the edges on the right side of C are incident with the other copy of C. We say that $\overline{G}$ is the graph obtained by *cutting along* C, and we call the induced embedding $\overline{\Pi}$.

Let $C = v_0 e_1 v_1 e_2 ... v_{l-1} e_l v_0$ be a one-sided Π-nonseparating cycle. Let us first define edges on the left side of G and edges on the right side of G at each vertex of C in a similar way as in the case of two-sided cycles: suppose first that the signature λ of Π satisfies $\lambda(e_i) = 1$ for $1 \leq i < l$ and $\lambda(e_l) = -1$. Then, we use the pairs of consecutive edges e_i, e_{i+1} on C $(1 \leq i < l)$ to define edges on the left side of C incident with the vertex v_i. Then we construct $\overline{G}$ by replacing C in G by the cycle $\overline{C} = v_0 e_1 ... e_l \overline{v}_0 \overline{e}_1 ... \overline{e}_l v_0$. The edges on the left side of C are adjacent to $v_0, ..., v_{l-1}$ and the edges on the right side of C are incident to $\overline{v}_0, ..., \overline{v}_{l-1}$. We extend λ by

putting $\lambda(\overline{e}_0) = ... = \lambda(\overline{e}_{l-1}) = 1$ and $\lambda(\overline{e}_l) = -1$ and hence obtain an embedding $\overline{\Pi}$ of $\overline{G}$. We say that $\overline{G}$ is the graph obtained by *cutting along* C.

Cutting along a cycle C induces a new embedding of $\overline{G}$ with the following properties:

Proposition 2.1.21 ([MT01, Lemma 4.2.4])**.** *Suppose that C is a Π_H-nonseparating cycle of a Π_H-embedded graph H in a surface. Let $\overline{H}$ be the graph obtained by cutting along C and let $\overline{\Pi}_H$ be its resulting embedding. Then all Π_H-facial walks are $\overline{\Pi}_H$-facial walks in $\overline{H}$, where edges of C are replaced by their copies in $\overline{H}$. If C is Π_H-two-sided, then $g(\overline{\Pi}_H) = g(\Pi_H) - 2$ and the two copies of C are the new Π_H-facial cycles. If C is Π_H-one-sided, then $g(\overline{\Pi}_H) = g(\Pi_H) - 1$ and $\overline{C}$ is a facial cycle in $\overline{H}$.*

Homotopic cycles

Let C and C' be two-sided cycles of a Π-embedded graph G. Suppose that C and C' are either disjoint or share a path (that might be reduced to a vertex). We say that C and C' are Π-*homotopic* if cutting along C and C' results in a graph which has a component D which contains precisely one copy of C and one copy of C' minus $\text{int}(C \cap C')$ and the induced genus of the embedding $\Pi(D)$ is 0. We then write $D = \text{Int}(C \cup C', \Pi)$.

Topologically, two cycles C and C' of a Π-embedded graph G in a surface S are Π-homotopic if the closed curve corresponding to C can be continuously deformed on S into the closed curve corresponding to C' (see Figure 2.4).

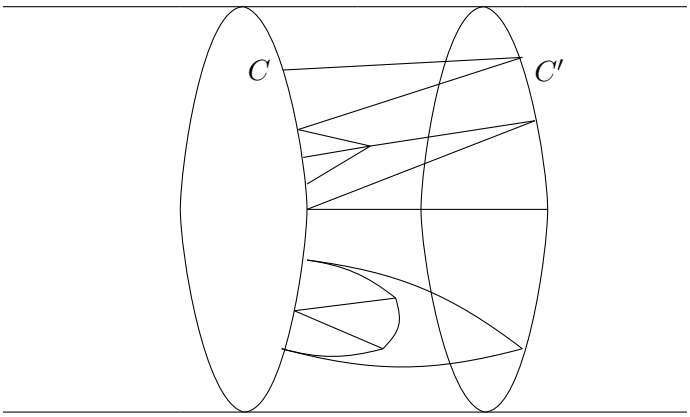


Figure 2.4: Two homotopic cycles C and C'

Remark that the definition of Π-homotopic cycles includes cycles that are Π-contractible. Two Π-contractible cycles are Π-homotopic if and only if one of them is Π-nested into the other one.

A classical result bounds the number of pairwise Π-nonhomotopic cycles in any graph embedded on a surface. The proof of this result can be found in [MT01].

Proposition 2.1.22 ([MT01, Proposition 4.2.7])**.** *Let H be a Π_H-embedded graph and a, b vertices of H (possibly $a = b$). If $P_0, ..., P_k$ are pairwise internally disjoint paths (or cycles) from a to b such that no two of them are Π-homotopic, then*

$$k \leq \begin{cases} g(\Pi_H) & \text{if } g(\Pi_H) \leq 1 \\ 3g(\Pi_H) - 3 & \text{if } g(\Pi_H) \geq 2 \end{cases}$$

Cycles bounding a disk or cylinder

Let G be a graph Π-embedded in a surface S. Let C be a Π-contractible cycle in G. Then, we say that C is the *boundary* of $\text{Int}(C, \Pi)$. Moreover, we say that C *bounds an empty disk in* Π if $\text{int}(C, \Pi)$ is empty. Let C and C' be two cycles that are disjoint or share a path (that might be reduced to a vertex) and Π-homotopic. We say that C and C' are the *boundaries* of $\text{Int}(C \cup C', \Pi)$. Moreover, we say that C and C' *bound an empty cylinder in* Π if $\text{int}(C \cup C', \Pi)$ is empty.

2.2 An introduction to excluded minors for surfaces

In this section, we provide a detailed introduction to the topic of excluded minors for surfaces. We begin with a general presentation of minor-closed classes of graphs, which form the broader framework encompassing excluded minors for surfaces. We then give a comprehensive historical overview of the study of excluded minors for surfaces. Finally, we discuss the algorithmic consequences of the theory of excluded minors for surfaces, emphasizing how improved bounds on the size of excluded minors directly lead to faster algorithms.

2.2.1 An introduction to minor-closed classes of graphs

A class of graphs $\mathcal{C}$ is said to be *closed under taking* respectively induced subgraphs, subgraphs, topological minors or minors if, for every graph $G \in \mathcal{C}$ and H respectively an induced subgraph, subgraph, topological minor or minor of G, H belongs to $\mathcal{C}$.

Naturally, we have the following property:

- A class of graphs closed under taking subgraphs is closed under taking induced subgraphs.
- A class of graphs closed under taking topological minors is closed under taking subgraphs.
- A class of graphs closed under taking minors is closed under taking topological minors.

Classes of graphs closed under taking induced subgraphs, subgraphs, topological minors or minors naturally appear in graph theory, as many studied properties of graphs are preserved by taking induced subgraphs, subgraphs, topological minors or minors.

For example, the class of complete graphs is closed under taking induced subgraphs, but it is not closed under taking subgraphs. The class of bipartite graphs is closed under taking subgraphs, but it is not closed under taking topological minors. The class of trees, the class of planar graphs and the class of graphs that have genus at most g are closed under taking minors.

Proposition 2.2.1. *The class of graphs that have genus at most g is closed under taking minors. In particular, the class of planar graphs (which are graphs with genus at most 0) is closed under taking minors.*

Proof. Let $\mathcal{C}$ be the class of graphs of genus at most g and let $G \in \mathcal{C}$.

It is enough to show that removing a vertex, an edge or contracting an edge from G induces a graph of genus at most g. Let Π be an embedding of G of genus $g(G)$.

Let $v \in V(G)$ and $H = G - v$. Let H' be the connected component of G containing v. Remark that

$$\begin{aligned} |V(H)| &= |V(H')| - 1 \\ |E(H)| &= |E(H')| - d_G(v) \\ |F(H, \Pi(H))| &\geq |F(H', \Pi(H'))| - d_G(v) + 1 \end{aligned}$$

Therefore, using Euler formula, we can see that the embedding of H induced by Π has genus at most $g(G)$. Therefore, $g(H) \leq g(\Pi(H)) \leq g(G)$.

Let $e \in E(G)$ and $H = G - e$. Let H' be the connected component of G containing e. Remark that

$$\begin{aligned} |V(H)| &= |V(H')| \\ |E(H)| &= |E(H')| - 1 \\ |F(H, \Pi(H))| &\geq |F(H', \Pi(H'))| - 1 \end{aligned}$$

Therefore, using Euler formula, we can see that the embedding of H induced by Π has genus at most $g(G)$. Therefore, $g(H) \leq g(\Pi(H)) \leq g(G)$.

Let $e \in E(G)$ and $H = G/e$. Let H' be the connected component of G containing e. Remark that

$$\begin{aligned} |V(H)| &= |V(H')| - 1 \\ |E(H)| &= |E(H')| - 1 \\ |F(H, \Pi(H))| &\geq |F(H', \Pi(H'))| \end{aligned}$$

Therefore, using Euler formula, we can see that the embedding of H induced by Π has genus at most $g(G)$. Therefore, $g(H) \leq g(\Pi(H)) \leq g(G)$.

Finally, as any proper minor H is obtained from G by a series of vertex deletion, edge deletion and edge contraction, it follows that $g(H) \leq g(G) \leq g$. Consequently, H belongs to the class $\mathcal{C}$ of graphs of genus at most g and $\mathcal{C}$ is closed under taking minors. □

A *graph parameter* is a function from the class of all graphs to $\mathbb{R}$ (usually with value in $\mathbb{N}$). Examples of graph parameters include the order of a graph, its number of edges, its maximum degree, its treewidth and its genus.

We say that a parameter p defined on graphs is *minor-monotone* if, for every graph G and every minor H of G, $p(G) \geq p(H)$.

A simple consequence of Proposition 2.2.1 is that the genus is a minor-monotone graph parameter.

Proposition 2.2.2. *Let G be a graph and H be a minor of G, then*

$$g(H) \leq g(G)$$

Let $\mathcal{C}$ be a class of graphs closed under minors. Then, a graph G is an *excluded minor* (resp. *excluded topological minor*) for $\mathcal{C}$ if $G \notin \mathcal{C}$ and every proper minor (resp. topological minor) of G belongs to $\mathcal{C}$. This notion of excluded minors can be adapted for classes of graphs closed under taking induced subgraphs and subgraphs.

For example, the forests have a unique excluded minor K_3 and the famous Wagner theorem [Wag37] shows that the excluded minors for the planar graphs are exactly $K_{3,3}$ and K_5.

We define the *minor ordering* $\leq_m$ on graphs as a partial ordering so that $H \leq_m G$ if and only if H is a minor of G.

Remark that the definition of excluded minor implies that an excluded minor for a class $\mathcal{C}$ is minimal in $\overline{\mathcal{C}}$ with respect to the minor ordering.

The excluded minors for a given class of graphs $\mathcal{C}$ entirely characterizes the class $\mathcal{C}$:

Proposition 2.2.3. *Let $\mathcal{C}$ be a class of graphs and let $\mathcal{E}$ be the set of its excluded minors. Then, G belongs to $\mathcal{C}$ if and only if G contains no $H \in \mathcal{E}$ as a minor.*

Proof. $\Rightarrow$ Let G be a graph in $\mathcal{C}$. By contradiction, suppose that $H \in \mathcal{E}$ is a minor of G. As $\mathcal{C}$ is closed under minors, $H \in \mathcal{C}$. However, H is an excluded minor for $\mathcal{C}$ and thus does not belong to $\mathcal{C}$ by definition. A contradiction: G does not contains any $H \in \mathcal{E}$ as a minor.

$\Leftarrow$ Let us show that if a graph does not belong to $\mathcal{C}$ then it contains some $H \in \mathcal{E}$ as a minor. Let G be a graph that does not belong to $\mathcal{C}$.

Let $\mathcal{M} = \{G' \leq_m G, G' \notin \mathcal{C}\}$ be the set of minors of G that does not belong to $\mathcal{C}$. Let $H \in \mathcal{M}$ be minimal in $\mathcal{M}$ with respect to the minor ordering: no proper minor of H belongs to $\mathcal{M}$. Then, let us show that H is an excluded minor for $\mathcal{C}$. Indeed, as $H \in \mathcal{M}$, H does not belong to $\mathcal{C}$. Moreover, every proper minor of H is a minor of G but does not belong to $\mathcal{M}$, therefore it belongs to $\mathcal{C}$. Finally, H an excluded minor for $\mathcal{C}$ and G indeed contains $H \in \mathcal{E}$ as a minor. □

Clearly, for $\mathcal{C}$ a class of graphs closed under minors, every excluded minor for $\mathcal{C}$ is an excluded topological minor for $\mathcal{C}$.

The converse does not hold. Nevertheless, the list of excluded topological minors for a class $\mathcal{C}$ can be constructed from the list of excluded minors for $\mathcal{C}$, as shown in [MT01, Proposition 6.1.1].

Proposition 2.2.4 ([MT01, Proposition 6.1.1])**.** *For each graph G_0, there is a finite collection $G_1, ..., G_p$ of graphs such that an arbitrary graph G has G_0 as a minor if and only if G contains a subdivision of one of $G_1, ..., G_p$.*

In particular, remark that a class of graphs is characterized by a finite set of excluded minors if and only if it is characterized by a finite set of excluded topological minor.

Although the list of excluded minors for a minor-closed class of graphs could a priori be infinite, the graph minor theorem [RS04], established by Neil Robertson and Paul Seymour in 2004, shows that every minor-closed class of graphs has only finitely many excluded minors.

Proposition 2.2.5 ([RS04])**.** *Every class of graphs closed under minors can be characterized by a finite set of excluded minors.*

A similar result does not hold for excluded induced subgraphs, excluded subgraphs or excluded topological minors. For instance, the set of excluded subgraphs for the bipartite graphs is the set of odd cycles which is infinite.

However, the proof of the graph minor theorem is purely existential and consequently provides no effective method for determining the excluded minors of a given minor-closed class.

2.2.2 A historical retrospective on excluded minors for surfaces

The characterization of planar graphs by excluded subdivisions attracted a lot of interest in the late 1920s. In 1930, Kuratowski [Kur30] showed that the planar graphs can be described as graphs that do not contain K_5 nor $K_{3,3}$ as a topological minor. Researchers from both Russia, Pontryagin and Kuratowski, and the United States of America, Frink and Smith were working independently on proving this characterization theorem. Kuratowski was first to publish his result and therefore to claim credit for it [Hud04].

This line of research was most likely motivated by the then-open conjecture that every planar graph is 4-colorable. Indeed, discovering new structural properties and characterizations of planar graphs may have been viewed as a possible route toward resolving this longstanding and central conjecture in graph theory. This interpretation is further supported by the fact that Kuratowski initially believed that excluding only K_5 would suffice to characterize planarity [Hud04]. Since K_5 is both an obvious obstruction to 4-colorability and a simple example of a nonplanar graph, this perspective appears natural in hindsight.

Subsequently, in 1937, Wagner [Wag37] showed that the planar graphs can be described as graphs that do not contain K_5 nor $K_{3,3}$ as a minor. This result and Kuratowski's theorem can easily be shown to be equivalent.

Finding an analog to Kuratowski's or Wagner's theorems for higher genus surfaces was for a long time a central open problem.

In 1978, Glover and Huneke showed that there is a finite number of excluded topological minors for the projective plane [GH78]; then Glover, Huneke and Wang presented in 1979 a list of 103 excluded topological minors [GHW79]; finally, Archdeacon proved in 1980-81 that this list is complete [Arc80, Arc81].

For the torus, several characterizations of excluded minors and excluded topological minors satisfying additional structural conditions have been obtained, see [DH72, Dec78, BW86, Neu93, Juv95, NM97, Hla97, Cha02, Woo06, GMC09, MŠ14, MW18]. In particular, encouraged by the access to more powerful computers, recent works relied on computational power to compute excluded minors for the torus. However, to this day, an exhaustive list is not known. Among all of these results, there are at least 17,500 excluded minors for the torus [MW18].

Recently, Mohar and Škoda [MŠ24a, MŠ24b] found the complete list of excluded minors of connectivity 2 for the Klein bottle. Along similar lines, Robertson and Seymour recently published results, originally proved in the 1990s, that establish a link between excluding Kuratowski graphs and being almost embeddable in a surface [RS26a, RS26b].

In 1989, Archdeacon and Huneke extended the result from the projective plane to every non orientable surface by showing that every non orientable surface admits only a finite number of excluded topological minors [AH89].

Robertson and Seymour finally settled the question by showing that the number of excluded minors for a surface is always finite in a preliminary result of the celebrated graph minor theorem [RS90]. Simpler proofs of the graph minor theorem for surfaces were subsequently obtained by Mohar [Moh01] and Thomassen [Tho97].

Robertson and Seymour actually proved a much more general result, of which the finiteness of the set of excluded minors for a fixed surface is only an intermediate consequence. Indeed, the graph minor theorem [RS04], proved by Robertson and Seymour in 2004, states that every minor-closed class of graphs can be characterized by a finite set of excluded minors.

Theorem 2.2.6 (Graph minor theorem [RS04]). *For every class of graphs $\mathcal{C}$ closed under minors, there exists a finite list of graphs $\mathcal{G}(\mathcal{C})$ such that a graph is in $\mathcal{C}$ if and only if it does not contain any graph from $\mathcal{G}(\mathcal{C})$ as a minor.*

This remarkable result, notable for its broad scope, its monumental proof, and the many developments it subsequently inspired, sparked considerable interest in understanding the structure of excluded minors for various minor-closed classes of graphs.

Especially, exploring the link between graph parameters that are minor-monotone and the presence of specific graphs as minor has since become a very fruitful research area. The most iconic example of such work is the grid minor theorem [RS86b] which states that there exists a function f such that, for every grid H, every graph with treewidth at least $f(|V(H)|)$ contains H as a minor. There has been extensive work on lower and upper bounds for the function f, see [CC16, CT21] for the current state of the art. Beyond the function f itself, the structure of the excluded minors characterizing bounded treewidth has also been studied [APC90, Ram97, Lag98, Chl02]. Analogous results to the grid minor theorem are known for pathwidth [Kin92, TUK94] and for treedepth [Nd12, KR21] (see [Nd12] for a definition of the latter parameter). Recently, Rambaud [Ram25] strengthened both the grid minor theorem and the excluded-minor characterization of bounded treedepth, obtaining an excluded-minor characterization for infinitely many parameters interpolating between treedepth and treewidth. We refer the reader to [SST23] for an excellent survey of excluded minors for various parameters and graph classes, and of their interplay with parameterized complexity.

We now turn to another minor-monotone graph parameter that has long attracted considerable attention: the genus. The study of graph embeddings on surfaces predates the graph minor project by several decades. However, the graph minor project profoundly renewed this area by reframing the study of surfaces through the perspective of excluded minors.

Despite the strength of the graph minor theorem, its proof is purely existential and therefore provides no effective method for determining the excluded minors for a given surface. As a consequence, exact lists of excluded minors are known only for the plane and the projective plane. Considerable effort has been devoted to the torus, and several families of excluded minors have been identified, yet no complete characterization has been obtained. Consequently, it appears unrealistic, at least in the foreseeable future, to expect explicit descriptions of the excluded minors for surfaces of higher genus.

This difficulty progressively shifted the focus of the area from determining explicit lists of excluded minors to understanding their structural properties. In particular, much attention has been devoted to bounding parameters such as the order and the treewidth of excluded minors for a surface as functions of its genus.

In 1993, Seymour gave the first and, until a recent joint work with Kawarabayashi [HK26b], the only known upper bound on the order of an excluded minor for a surface [Sey93].

Theorem 2.2.7 ([Sey93]). *Let S be a given surface of genus g. Every excluded minor for S has at most 2^{2^k} vertices where $k = (3g+9)^9$.*

He also gave in the same paper the first and, until [HK26b], the only known upper bound on the treewidth of an excluded minor for a surface [Sey93].

Theorem 2.2.8 ([Sey93, (3.3)]). *The treewidth of G is bounded by a polynomial in g:*

$$tw(G) \leq T_S(g)$$

with $T_S(g) = 3(g+3)^2(3g+16) - 3 = O(g^3)$.

Thomassen gave in 1997 a bound on the size of an excluded grid topological minor in a 2-connected excluded minor for a surface.

Theorem 2.2.9 ([Tho97])**.** *Let H be a 2-connected graph such that $g(H) = g$ and $g(H - e) < g$ for every edge e in H. Then H contains no subdivision of the $4k \times 2k$ grid, with $k = \lceil 800g^{3/2} \rceil$.*

The upper bound on the order of excluded minors for a surface given by Seymour is double exponential in the genus g. On the other hand, the lower bound on the maximal order of excluded minors for a surface is easily seen to be $\Omega(g)$. Indeed, a graph consisting of the disjoint union of $g+1$ copies of $K_{3,3}$ (see Figure 2.5 for an illustration) has order $6(g+1)$ and is an excluded minor for every surface of genus g.

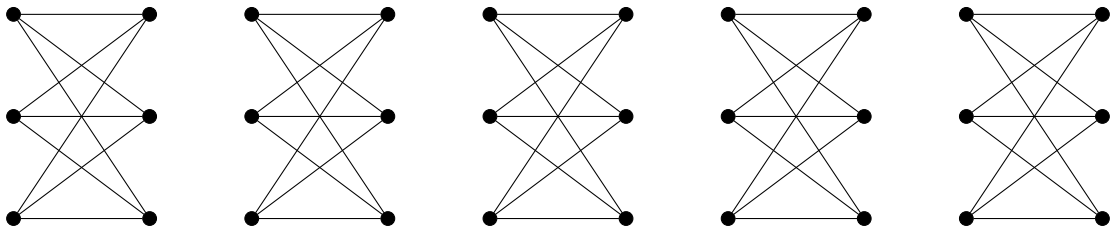


Figure 2.5: The graph consisting of the union of 5 disjoint copies of $K_{3,3}$

Proposition 2.2.10. *For every $g \geq 0$, the graph H_g consisting of the union of $g+1$ disjoint copies of $K_{3,3}$ is an excluded for every surface of genus g.*

Proof. Let S be a surface of genus g.

The graph H has genus $(g+1)g(K_{3,3}) = g+1$. Therefore, H cannot be embedded on S.

Moreover, let H' be a proper minor of H. As H' is a proper minor of H, there is at least a connected component C' of H' that is a proper minor of a connected component of H: C' is a proper minor of $K_{3,3}$ which is an excluded minor for the projective plane. Consequently, C' is planar.

Finally, H' is a minor of the graph H'' consisting of the union of g copies of $K_{3,3}$ and of C'. The graph H'' has genus at most g and therefore so does H' as it is its minor.

To conclude, each proper minor of H has genus g and therefore H is an excluded minor for the surface S. □

Therefore, the gap between the known lower and upper bounds was substantial, and Seymour's upper bound was widely believed to be far from optimal.

It was the prevailing view among leading researchers in the area that the upper bound should be polynomial in the genus of the surface, and possibly even quadratic or linear. This belief was also shared with me by both Mohar and Kawarabayashi.

Conjecture 1 (Private communication, Mohar and Kawarabayashi)**.** *Let S be a surface of genus g. The order of an excluded minor for S is bounded by a polynomial in g.*

The results presented in this manuscript are the most recent contributions to this line of research and settle the conjecture mentioned above (see Theorem 3.2.1 below), which has been a central question in the field.

2.2.3 Algorithmic consequences

From a structural and theoretical perspective, the characterization of graphs on surfaces via excluded minors is motivated by the search for a purely combinatorial description of surface-embeddable graphs, as opposed to a topological one. Beyond this conceptual motivation, a central driving force behind the study of excluded minors for surfaces is algorithmic. Indeed, either an explicit list or bounds on the order or treewidth of the excluded minors for a given surface have several important algorithmic consequences.

A classical example of the algorithmic relevance of such characterizations is planarity testing. By Kuratowski's theorem, a graph is planar if and only if it contains neither K_5 nor $K_{3,3}$ as a subdivision (or topological minor). This characterization underlies classical linear-time planarity-testing algorithms, including the path-addition algorithm of Hopcroft and Tarjan [HT74] and the edge-addition algorithm of Boyer and Myrvold [BM04], which produce a subdivision of K_5 or $K_{3,3}$ as a certificate of non-planarity.

The same approach extends naturally from planar graphs to graphs on surfaces. Indeed, a natural way to tackle various algorithmic problems on graphs embedded in a fixed surface, such as computing the genus of a graph or finding an embedding into a given surface, is to exploit the characterization of surface-embeddable graphs via a finite list of excluded minors. Consequently, the efficient computation of these excluded minors is a central problem in algorithms for graphs on surfaces.

As a consequence of the graph minor theorem, Robertson and Seymour [RS95] showed in 1995 that, for every minor-closed class of graphs $\mathcal{G}$, there exists a cubic-time algorithm deciding membership in $\mathcal{G}$, later improved to quadratic time by Kawarabayashi, Kobayashi, and Reed [KKR12] in 2012, and to quasi-linear time by Korhonen, Pilipczuk, and Stamoulis [KPS24] in 2024. However, this algorithm is only explicit in the sense that its existence and running time are guaranteed: actually instantiating it for a specific class $\mathcal{G}$ requires holding, in hand, the finite set of excluded minors that characterizes $\mathcal{G}$, a set which is, in general, unknown. Determining the excluded minors of a given graph class is therefore the missing ingredient that turns the abstract existence of a fast algorithm into an actual one.

This is made precise by a result of Fellows and Langston [FL89], who showed in 1989 that no algorithm can, given a Turing machine implementing a membership test for a minor-closed class $\mathcal{G}$, compute the excluded minors of $\mathcal{G}$: the task is undecidable in general. Fellows and Langston showed, however, that this obstacle disappears as soon as an explicit upper bound on the order of the excluded minors of $\mathcal{G}$ is known, in which case the excluded minors can be computed by testing minor-minimality among the finitely many candidate graphs of at most that order, see also Adler, Grohe, and Kreutzer [AGK08] for an algorithm along the same lines. Crucially, the running time of this procedure grows exponentially in the bound on the order of the excluded minors.

We also wish to emphasize the algorithmic significance of obtaining a small upper bound on the order of excluded minors for surfaces. Bounds of this type have repeatedly served as a fundamental ingredient in parameterized algorithms on graphs embedded in surfaces, where they enable algorithms whose dependence on the input size is linear or quadratic while isolating the complexity in the surface parameter.

For instance, Kawarabayashi, Mohar, and Reed [KMR08] presented at FOCS 2008 a linear-time FPT algorithm for embedding graphs into an arbitrary surface, parameterized by its genus,

whose correctness and parameter dependence rely on an upper bound on the size of the excluded minors for the surface. Likewise, Grohe, Kawarabayashi, and Reed [GKR13] used such a bound in their quadratic-time algorithm for computing graph minor decompositions.

Although replacing Seymour's double-exponential bound by an almost quadratic one does not immediately improve the asymptotic running times of these algorithms, it significantly strengthens one of their key structural ingredients. More broadly, recent progress in graph minor algorithms has repeatedly shown that sharper structural bounds eventually translate into substantially more efficient algorithms. We therefore expect our result to play a similar role in the development of algorithms for bounded-genus graphs.

Moreover, we hope that the techniques and results developed in this paper will prove useful to computational approaches, such as those of [Neu93, NM97, Cha02, MW18], aimed at finding the excluded minors of small-genus surfaces such as the torus and the Klein bottle. In particular, they may help reduce the search space involved in such approaches, thereby enabling more efficient algorithms and, ultimately, a more complete understanding of the excluded minors for these two surfaces.

Beyond graphs on surfaces, a major line of research in graph minor theory has been to turn the qualitative structural results of Robertson and Seymour into quantitative ones with polynomial bounds. Such improvements are driven by algorithmic applications, since polynomial dependencies are essential for obtaining efficient parameterized algorithms and practical structural decompositions.

This line of research has seen remarkable progress in recent years. Chekuri and Chuzhoy [CC16] established the first polynomial excluded grid theorem. More recently, building on [KTW20], Gorsky, Seweryn, and Wiederrecht [GSW25] proved that all quantitative bounds in the graph minor structure theorem can be made polynomial in the size of the excluded minor.

Our result places excluded minors for surfaces within this broader program. By replacing Seymour's double-exponential bound with an almost quadratic one, we provide the first polynomial bound for this fundamental parameter and thereby strengthen one of the remaining structural ingredients underlying algorithms for bounded-genus graphs.

2.3 Outline of this thesis

We chose to present the contributions of [HK26b] and [HK26a] in parallel in the structure of this manuscript, in order to highlight the similarities between the results and to better explain the improvements obtained between the first and second papers. Accordingly, the overall structure of this thesis closely follows the organization of both papers [HK26b, HK26a].

We also include more detailed discussions of our results, in particular regarding their optimality.

In Chapter 3, we provide an overview of the results and techniques developed during this thesis, which are then presented in details in the subsequent chapters. This part is divided into two sections, corresponding to the two papers produced during this work.

In Chapter 4, we first present basic results on graphs on surfaces, most of which are variations of classical results in this area. We then establish several elementary properties of excluded minors for a given surface. In particular, we show that an excluded minor of genus g has genus at most $g+2$. We also study embeddings of an excluded minor after the deletion of an edge, and

show that it is sufficient to restrict attention to 2-connected excluded minors in order to obtain bounds on their order. Finally, we recall known bounds on the treewidth of excluded minors for a given surface.

In Chapter 5, we present our main structural results on excluded minors for surfaces. We begin by recalling the structural theorem of Seymour [Sey93] and provide a simpler alternative proof of his result. We then extend this framework by establishing generalizations involving isolated paths, contractible nested cycles, and homotopic nested cycles.

We next derive several consequences of these improved structural results. In particular, we obtain improved bounds for the tree decomposition and for the characteristic size of an excluded grid topological minor of an excluded minor. We also prove a bound on the number of almost disjoint noncontractible cycles in an excluded minor with a fixed embedding.

We then discuss the relationship between our structural results and the embedding parameter known as facewidth. Finally, we address the optimality of these structural results.

In Chapter 6, we establish results on the maximum degree of a vertex and on the maximum face size in an excluded minor for a surface with a fixed embedding, both in the general case and in contractible subgraphs with a small separator. We then discuss the apparent duality between vertices and faces that emerges from the results of this chapter and those of Chapter 5. Finally, we address the optimality of our bounds on the maximum degree and maximum face size.

In Chapter 7, we present the proofs of both the quasi-polynomial and almost quadratic bounds on the order of excluded minors for surfaces. We begin by introducing the notion of balanced separator decompositions and by developing a tailored lemma based on the celebrated balanced separator theorem of Robertson and Seymour [RS86a].

We then identify, within an excluded minor for a given surface, a subgraph G_0 of G with separator A whose bridges on A are contractible. Next, we show how to bound the order of a 2-connected contractible subgraph of an excluded minor in terms of the size of its separator. Finally, we establish both the quasi-polynomial and almost quadratic bounds. The quasi-polynomial bound follows directly from the results of Chapter 6 combined with the structural results presented above, whereas the almost quadratic bound relies on finding, within G_0, a 2-connected contractible subgraph whose separator has logarithmic size.

In Appendix A, we determine the exact order of the smallest excluded minor for a surface of genus g. This order is closely related to Heawood's bound, a classical upper bound on the chromatic number of graphs embeddable in a given surface.

Chapter 3

Results and techniques

3.1 A quasi-polynomial bound for the excluded minors for a surface . . 49
3.1.1 Main results 50
3.1.2 Technical overview 51
3.1.3 Overview of the proof 52
3.2 An almost quadratic bound for the excluded minors for a surface . . 53
3.2.1 Main results 53
3.2.2 Technical contributions over [HK26b] 53
3.2.3 Overview of the proof 54

In this chapter, we outline the main results and techniques developed during this thesis, which are presented in detail in the subsequent chapters.

The contributions of this thesis are divided into two parts, corresponding to two results established sequentially during my Ph.D. In the first part, we prove that the order of the excluded minors for a surface is bounded by a quasi-polynomial function of the genus [HK26b]. In the second part, we show that this bound can in fact be improved to an almost quadratic bound in the genus [HK26a].

Although the second result strictly strengthens the first one, we present both contributions separately, since the second builds on techniques developed for the first. This organization highlights the progression of ideas and makes it easier to identify the key improvements between the two results.

In Section 3.1 and Section 3.2, we present the main results and techniques obtained respectively in [HK26b] and [HK26a].

3.1 A quasi-polynomial bound for the excluded minors for a surface

This section presents the results obtained in the paper *A quasi-polynomial bound the minimal excluded minors for a surface* [HK26b] which is joint work with Kawarabayashi. The paper was

published in the ACM-SIAM Symposium On Discrete Algorithms (SODA) 2026 conference, and its full version is available on ArXiv.

3.1.1 Main results

The main result of this paper is a quasi-polynomial bound on the order of the excluded minors for a surface:

Theorem 3.1.1 ([HK26b, Theorem 1.2])**.** *Let S be a given surface of genus g. Every excluded minor for S is of order at most $U(g) = g^{O(\log^3 g)}$.*

This result improves the first, and until our paper, the only known upper bound on the order of an excluded minor for a surface, proved by Seymour in 1993 [Sey93] (see Theorem 2.2.7).

It constituted the first improvement of this bound in thirty years, and we drastically lower it from double-exponential to quasi-polynomial in the genus g of the surface.

As a corollary of the structural results we obtain in this paper, we improve the bound on the treewidth given by Seymour in 1993 [Sey93] (see Theorem 2.2.8):

Theorem 3.1.2 ([HK26b, Theorem 1.4])**.** *Let S be a given surface of genus g. The treewidth of every excluded minor G for S is bounded by a function of g:*

$$tw(G) \leq T(g)$$

with $T(g) = 264(g+2)(m+1) - 1 = O(g \log g)$, where $m = 2(\lfloor \log_q(3g+4) \rfloor + 2)$ and $q = \frac{9073}{9072}$.

Moreover, as a subsidiary result, we obtain an improvement on an important result from Thomassen [Tho97], stated in Theorem 2.2.9. This result established a bound on the size of a excluded grid topological minor. We improved this bound from $O(g\sqrt{g})$ to $O(\sqrt{g} \log g)$:

Theorem 3.1.3 ([HK26b, Theorem 1.6])**.** *Let G be a 2-connected excluded minor for a surface of genus g. Then G contains no subdivision of the $4k \times 2k$ grid, with $k = O(\sqrt{g} \log g)$.*

Finally, we give a lower bound on the order of all excluded minors for a given surface that is tight in the sense that there is an excluded minor whose order is exactly the given lower bound (see Appendix A):

Theorem 3.1.4 ([HK26b, Theorem 1.7])**.** *Let S be a surface of genus g and let $L(g) = \left\lfloor \frac{7+\sqrt{1+24g}}{2} \right\rfloor + 1$. There exists an excluded minor for S of order*

$$|V(G)| = L(g)$$

and no excluded minor for S is of order less than $L(g)$.

3.1.2 Technical overview

Let us highlight the technical aspects of our result.

Let S, S' be surfaces with S' of genus g and S of genus $g+1$ or $g+2$. Let G be an excluded minor for the surface S' and suppose that G can be embedded in surface S with embedding Π.

The main catalyzer to reach a quasi-polynomial bound is a breakthrough on the characteristic size of a forbidden structure for G: it is not hard to show that G does not contain $O(g)$ disjoint cycles that are Π-contractible and nested in Π (the definition of *nested* can be found in Section 5.1). Seymour proved such a result in [Sey93, (2.2)] and we give a shorter and conceptually simpler proof of this result in Section 5.1.

We show that this bound can be lowered to $O(\log g)$ (Chapter 5, Corollary 5.3.15). A related result was obtained in [KS19]. Moreover, we extend this result to Π-noncontractible homotopic cycles of G.

These structural results are essential to obtaining the quasi-polynomial bound in this paper. Indeed, the techniques developed here, together with an $O(g)$ bound on the number of disjoint Π-contractible nested cycles, would only yield a single-exponential bound on the order of G. Such a bound is easily derived from a careful application of the structural tools introduced by Seymour in [Sey93]. The improvement from an $O(g)$ to an $O(\log g)$ bound on the size of this forbidden structure constitutes the main breakthrough of our work, enabling us to go beyond an exponential bound and achieve a quasi-polynomial one.

Another important ingredient of our proof is a $g^{O(\log^2 g)}$ upper bound on the maximum degree of G and on the maximum size of a face in (G, Π) (Chapter 6, Theorem 6.2.2). It was, to our understanding, the first such bound. The almost quadratic bound established later in this work only yields a trivial almost quadratic upper bound on the maximum degree of an excluded minor (see Section 6.4).

Then, we develop a new method to bound the diameter of the tree in a tree decomposition of G and we obtain a bound that is drastically lower than the one found by Seymour [Sey93, claim (6) in (4.1)]. Indeed, we lower it from a single exponential in g to a quasi-polynomial $g^{O(\log^3 g)}$ in g.

Finally, by switching from tree decomposition to path decomposition, and by taking advantage of the bound on the diameter of the tree of a tree decomposition of G (that still holds when applied to a path decomposition), we finally reach a $g^{O(\log^3 g)}$ bound on the order of G.

The two techniques on tree decompositions described above are not presented in this manuscript, as we develop in [HK26a] a new method for bounding the order of G that is simpler and yields stronger bounds. We therefore rely on this new technique to obtain the quasi-polynomial bound, rather than the original approach.

Our proof bears similarity with Seymour's on two main points:

- It uses tree decomposition and a polynomial bound on the treewidth of G and the maximum degree of the tree of a tree decomposition of G as central tools. Both proofs aim at bounding the order of G by bounding the treewidth, diameter and degree of a tree decomposition of G.
- The main structural results used in this paper are improvements and extensions of a structural result from Seymour [Sey93, 2.2] that is central to his proof.

The rest of the proof does not bear much similarities with the proof in [Sey93] and we develop our own method to reach a quasi-polynomial bound.

3.1.3 Overview of the proof

We define a *piece* as a vertex or a face of (G, Π).

To prove that G is of order bounded by a quasi-polynomial in g, we proceeded in [HK26b] in broad outline as follows:

- **We first prove that the degree of G and the size of faces in (G, Π) are bounded by a quasi-polynomial in g (Section 6.2).**

  The proof proceeds by contradiction. We make the assumption that there is a piece p that is either a vertex of degree $> \Delta(g)$ or a face of size $> \Delta(g)$ for some function Δ (defined in Section 6.2), and deduce that there is a Π-contractible fan on p (a fan F on a piece p is roughly a subgraph consisting of internally disjoint paths from p to a path, see Definition 6.1.2) whose size is smaller than $\Delta(g)$ by a polynomial factor in g. Then, we show that this fan extends to a Π-contractible subgraph of G that contains a forbidden structure consisting of $O(\log g)$ nested cycles that intersect only in p and that all are Π-contractible (see Section 5.3 for the proper statement).

- **We then show that the diameter of the tree of a tree decomposition of G of width w is bounded by a quasi-polynomial in g and w.**

  The proof proceeds by contradiction. We assume that there is a path of length $> P(g, w)$ (for some function P) in the tree of a tree decomposition of G of width w that induces a Π-contractible subgraph G' of G. On one hand, we show that it implies that $|V(G')| > P(g, w)$. On the other hand, we take advantage of the natural separations given by the tree decomposition to show that G' is separated from the rest of G by a small (polynomial in g) separation, and then to deduce that $|V(G')|$ must be bounded by $P(g, w)$. This leads to a contradiction. Finally, there is no such path in the tree of a tree decomposition of G and we deduce a bound on the diameter of the tree of a tree decomposition of G.

- **We finally put everything together to show our main result.**

  Seymour (Theorem 2.2.8, [Sey93, (3.3) and claim (5) in (4.1)]) shows that both the treewidth of G and the maximum degree of the tree in an optimal tree decomposition of G (with additional conditions) are bounded by $O(g^3)$. As subsidiary results of the structural results, we improve both of these results to $O(g \log g)$. These two bounds, together with the bound on the diameter of the tree, give a first bound, single exponential in g, on the order of G. Then, by looking at the pathwidth and a path decomposition of G, and by taking advantage of the bound on the diameter of the tree of a tree decomposition of G (that still holds when applied to a path decomposition), we finally reach a $g^{O(\log^3 g)}$ bound on the order of G.

We prove two main structural results that describe forbidden structures in Π-contractible subgraphs of (G, Π) and that will be used repetitively to find a contradiction in the main proof:

- There are at most $O(g)$ internally disjoint paths from one piece to another piece (Chapter 5, Proposition 5.2.3).

- There are at most $O(\log g)$ disjoint cycles that are either Π-contractible and nested in Π or Π-noncontractible homotopic (Chapter 5, Corollary 5.3.15 and Proposition 5.3.27). These two results are improvements on a result of Seymour (Proposition 5.1.1 in Section 5.1, [Sey93, (2.2)]) and are central to obtain the main result of this paper.

3.2 An almost quadratic bound for the excluded minors for a surface

This section presents the results obtained in the paper *An almost quadratic bound for the minimal excluded minors for a surface* [HK26a] which is joint work with Kawarabayashi. This paper is currently submitted to a conference, and its full version is available on ArXiv.

3.2.1 Main results

The main result is an almost quadratic bound on the order of the excluded minors for a surface:

Theorem 3.2.1 ([HK26a, Theorem 1.3])**.** *Let S be a given surface of genus g. Every excluded minor for S is of order at most $U(g) = g^{2+o(1)}$.*

3.2.2 Technical contributions over [HK26b]

Let us highlight the technical contributions. Let G be an excluded minor for a surface of genus g.

Our approach is based on a new structural theorem for a minimal excluded minor G for a surface of Euler genus g. In [HK26b], we proved that G contains $O(\log g)$ pairwise disjoint cycles that are contractible and nested in some embedding of G. We refine this result by establishing a separator-based variant: for any contractible subgraph H of G with a separator of size s, that is, any planar induced subgraph H that is separated from the rest of the graph G by a separator on its outer face boundary of size s, H contains $O(\log s)$ disjoint cycles that are contractible and nested in some embedding of G. This allows us to shift from a genus-dependent bound to a separator-dependent bound.

At a high level, the proof of this result builds on the argument developed in [HK26b], showing that G contains $O(\log g)$ pairwise disjoint cycles that are contractible and nested in some embedding of G. However, the conceptual shift underlying our approach is essential to establishing our main result: it allows us to overcome a barrier that has so far prevented previous approaches from obtaining an almost quadratic bound. Indeed, by identifying a contractible subgraph of G that admits a separator of size asymptotically smaller than g, we can focus on a region of the graph in which the number of disjoint contractible nested cycles is smaller than the $O(\log g)$ bound that applies to the entire graph and that is believed to be asymptotically tight, see Subsection 5.6.1. In the rest of our proof, we isolate a contractible subgraph of G with a separator of size $O(\log^2 g)$, which therefore contains $O(\log \log g)$ pairwise disjoint contractible nested cycles. Consequently, we can "forbid" $\Omega(\log \log g)$ nested cycles, rather than $\Omega(\log g)$ nested cycles, which allows us to obtain the almost quadratic bound in terms of the genus g (we can bound the number of vertices of such a contractible subgraph, which is one of the key ideas of the proof). This constitutes the paper's main new ingredient.

In the remainder of the proof, we build on the techniques developed in [HK26b] and use the celebrated balanced tree separator theorem from Robertson and Seymour [RS86a], to obtain our almost quadratic bound.

3.2.3 Overview of the proof

Let S, S' be surfaces with S' of genus g and S of genus $g+1$ or $g+2$. Let G be an excluded minor for the surface S' and suppose that G can be embedded in the surface S with embedding Π. Recall that a *piece* is a vertex or a face of (G, Π).

To prove that G is of order bounded by an almost quadratic function of g, we proceed in four steps. The main idea is to repeatedly use balanced separator decompositions to isolate a contractible region of G in which our separator-sensitive bound on nested contractible cycles can be applied.

- **We first isolate a large subgraph G_0 whose bridges with respect to a small separator are all Π-contractible (Section 7.2, Proposition 7.2.1).**

  We start with a tree decomposition of G of width $O(g \log g)$, as given by [HK26b]. We apply a recursive balanced separator decomposition to this tree decomposition. The decomposition produces $O(g \log^2 g)$ parts of comparable size, while the intersection of any one part with the other parts has size $O(g \log^2 g)$. Each part corresponds to a subgraph of G and we consider the components of these subgraphs after deleting their boundaries.

  The key point is that the components of one of these parts must all be Π-contractible. Indeed, suppose that every part contained a Π-noncontractible component. From each such component, we can extract a Π-noncontractible cycle meeting the boundary in at most one vertex. The resulting cycles are almost disjoint, and their number is $O(g \log^2 g)$. This contradicts the bound of Corollary 5.4.6 on the number of almost disjoint Π-noncontractible cycles. Hence, there is a subgraph G_0 separated from the rest of G by a set A of size $O(g \log^2 g)$, such that every bridge on A contained in G_0 is Π-contractible.

  This reduction is important for two reasons. First, since every bridge of G_0 is contractible, each such bridge can be treated as a planar graph inside a disk of the embedding Π. Second, the balanced decomposition guarantees that G_0 contains a large fraction of the vertices of G, up to a factor $Q(g) = O(g \log^2 g)$. Thus, it is enough to obtain a sufficiently strong bound on the size of G_0.

- **We then reduce each bridge on A to a 2-connected contractible subgraph with a logarithmic-size separator (Section 7.5, Proposition 7.5.1).**

  Consider a bridge B on A in G_0, and let $A_B \subseteq A$ be a separator of B of size between a and $2a-1$, for some $a \geq 3$. Since B is Π-contractible, the global bound of Corollary 5.3.15 on nested contractible cycles implies that B cannot contain a large grid minor: a sufficiently large grid minor would contain too many disjoint nested cycles. By the planar grid–treewidth relation (see [RST94, (6.2)]), this gives $tw(B) = O(\log g)$.

  We can therefore find a balanced separator decomposition of B with $O(g \log^2 g)$ parts. This decomposition, together with the fact that B is attached to the rest of G through A_B, allows us to isolate a 2-connected subgraph G_1 of B whose separator from the rest of B has size $O(\log g)$. More precisely, G_1 can be chosen to be one of the 2-connected components obtained after cutting the decomposition at an appropriate balanced separator. The remaining pieces of B can then be controlled in terms of g, a, and $|V(G_1)|$.

The crucial gain is that the separator of G_1 is logarithmic in g, rather than of order g. This is what allows us to apply the separator-sensitive nested-cycle result of Corollary 5.3.16 with $s = O(\log g)$.

- **We prove that such a subgraph G_1 has sub-polynomial order in g (Section 7.3, Proposition 7.3.1).**

  The separator-sensitive result of Corollary 5.3.16 implies that G_1 contains only $O(\log\log g)$ pairwise disjoint Π-contractible cycles that are nested in Π. This is the decisive improvement over the $O(\log g)$ bound that applies to the whole graph.

  We use this stronger restriction to control the local structure of G_1. First, if a vertex of G_1 had very large degree, or if a face of (G_1, Π) had very large size, we could find a large Π-contractible fan based at that piece. The fan can then be iteratively extended to produce many Π-contractible nested cycles, contradicting the $O(\log\log g)$ bound. Consequently, both the maximum degree of G_1 and the maximum size of a face of (G_1, Π) are bounded by a sub-polynomial function of g (see Section 6.3).

  We then use the planar structure of G_1 to propagate these local bounds through successive layers of nested cycles (see Proposition 7.3.1). Starting from an outer cycle, the vertices and bridges between consecutive layers grow by at most a sub-polynomial factor at each step. Since the number of nested layers that can occur in the relevant configuration is controlled by the $O(\log\log g)$ bound above, this yields a sub-polynomial bound $|V(G_1)| = g^{o(1)}$ (more precisely, $|V(G_1)|$ is bounded by the function $P(g)$ defined in Section 7.3).

  Thus, the large-scale structure of each bridge is reduced to a small 2-connected core G_1, while the dependence on the size a of its attachment to A is kept explicit.

- **We finally combine the bounds according to the size of the separators of the bridges (Section 7.5, Proposition 7.5.2).**

  We partition the bridges on A according to the size of their attachment separator. For $a_i = 3\cdot 2^i$, the i-th class consists of bridges whose separator has size between a_i and $2a_i - 1$. There are only $O(\log g)$ such classes, since the size of A is $O(g\log^2 g)$.

  For a fixed class, the number of bridges is bounded using the fact that each bridge has at least a_i edges incident with its separator, while the total number of edges in the contractible part is controlled by its number of vertices and the genus of the ambient embedding. This gives a bound on the number of bridges in the class of order $O\left(\frac{g\log^2 g}{a_i}\right)$.

  On the other hand, Proposition 7.5.1 bounds the size of each bridge by a function that is essentially linear in a_i and sub-polynomial in g. The factor a_i therefore cancels with the inverse dependence in the number of bridges in each class. Summing over the $O(\log g)$ classes gives a bound on $|V(G_0)|$ that is almost linear in g.

  Finally, Proposition 7.2.1 gives $|V(G)| \leq Q(g)|V(G_0)|$ where $Q(g)$ is nearly linear in g, and combining the two bounds yields

  $$|V(G)| \leq g^2(\log g)^{O((\log\log g)^3)} = g^{2+o(1)}$$

  The argument above is first carried out for 2-connected excluded minors. The general case then follows from the decomposition into 2-connected components given by Corollary 4.2.4.

We prove three main structural results that describe forbidden structures in (G, Π) and that will be used repetitively to find a contradiction in the main proof:

- There are $O(\log g)$ disjoint cycles that are Π-contractible and nested in Π (Subsection 5.3.1, Corollary 5.3.15).
- There are $O(\log g)$ disjoint cycles that are Π-noncontractible homotopic (Subsection 5.3.2, Proposition 5.3.27).
- In a subgraph H of G with a separator of size s, there are $O(\log s)$ disjoint cycles that are Π-contractible and nested in Π (Subsection 5.3.1, Corollary 5.3.16).

These three results improve on a result of Seymour [Sey93, (2.2)] and, more recently, on results in [HK26b, Section 5]. They are central to obtaining our main result.

Chapter 4

Preliminary results

4.1 Preliminary results for graphs on surfaces 57
4.2 Basic results on excluded minors for surfaces 61
4.2.1 Embedding of G 61
4.2.2 Bounds for tree decompositions of G 62
4.2.3 Separators in G 63
4.2.4 Structure of $G - e$ 65

In this chapter, we give preliminary results for general graphs on surfaces and results that are specific to excluded minors for surfaces.

The results for general graphs on surfaces are given in Section 4.1 and have rather simple proofs.

The results specific to excluded minors for surfaces are given in Section 4.2: they consist of preliminary remarks that are easily formulated when looking at excluded minors for surfaces, together with a presentation of the results on treewidth and tree decomposition for the excluded minors for surfaces.

4.1 Preliminary results for graphs on surfaces

We first present two variants of Proposition 2.1.22 that extend its scope of application in two different directions.

Definition 4.1.1 (Almost disjoint cycles)**.** Let $k \in \mathbb{N}$. Let $C_0, ..., C_k$ be cycles of a graph H. We say that $C_0, ..., C_k$ are *almost disjoint cycles* if each cycle C_i $(0 \leq i \leq k)$ has at most one vertex in common with the subgraph $\cup_{j \neq i} C_j$.

See Figure 4.1 for an illustration of almost disjoint cycles.

Proposition 4.1.2 ([HK26b, Proposition 4.8])**.** *Let H be a Π_H-embedded connected graph. If $C_1, ..., C_k$ are cycles of H that are almost disjoint, Π_H-noncontractible and such that no two of them are Π_H-homotopic, then*

$$k \leq \begin{cases} g(\Pi_H) & \text{if } g(\Pi_H) \leq 1 \\ 3g(\Pi_H) - 3 & \text{if } g(\Pi_H) \geq 2 \end{cases}$$

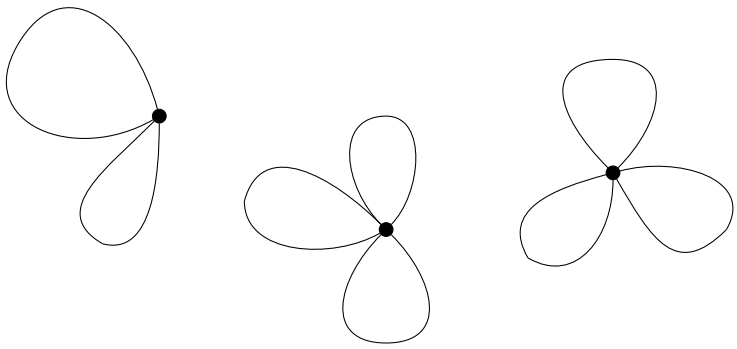


Figure 4.1: Almost disjoint cycles. The almost disjoint cycles are depicted in solid black lines. The black dots are vertices shared by several almost disjoint cycles.

Proof. This proof is adapted from the proof of Proposition 2.1.22 in [MT01].

If $g(\Pi_H) = 0$, then H is embedded in the sphere, and there are indeed no Π_H-noncontractible cycles.

If $g(\Pi_H) = 1$, then H is embedded in the projective plane. If there is a Π_H-noncontractible cycle C, let us cut along it, getting a new embedding Π'_H. Then, by Proposition 2.1.21, Π'_H is an embedding in the plane. Hence, every Π_H-noncontractible cycle almost disjoint from C is Π_H-homotopic to C and $k \leq 1$.

Let us suppose $g(\Pi_H) \geq 2$. For $1 \leq i \leq k$, let e_i be an edge of C_i. Let T be a spanning tree of H containing $C_i - e_i$ for $1 \leq i \leq k$ (such a spanning tree exists because $\cup_{1 \leq i \leq k}(C_i - e_i)$ is an acyclic subgraph of H). After contracting $E(T)$, the resulting multigraph $H/E(T)$ is embedded in the same surface with embedding $\tilde{\Pi}_H$. It contains only one vertex, the loops $e_1, ..., e_k$, and some additional loops. Let H', Π'_H be the graph and embedding obtained from $H/E(T)$ by removing every loop distinct from $e_1, ..., e_k$ that is contained in two $\tilde{\Pi}_H$-facial walks. Then, Π'_H is an embedding in the same surface.

Let W be a Π'_H-facial walk. If W contains an edge e distinct from $e_1, ..., e_k$, then e occurs twice in W (as otherwise it would have been removed previously), and hence the length of W is at least 3. If W consists of loops from $\{e_1, ..., e_k\}$ only, then W is not a single loop e_i since C_i would then be Π_H-contractible. Similarly, if $W = e_i e_j$, then C_i and C_j are Π_H-homotopic. Therefore, the length of W is at least 3. It follows that $2|E(H')| \geq 3f$ with f the number of Π'_H-facial walks. Now, by Euler's formula:

$$\begin{aligned} |V(H')| - |E(H')| + f &= 2 - g(\Pi_H) \\ |E(H')| - f &= g(\Pi_H) - 1 \\ 3|E(H')| - 3f &= 3g(\Pi_H) - 3 \end{aligned}$$

Finally, as $2|E(H')| \geq 3f$, it follows that

$$k \leq |E(H')| \leq 3|E(H')| - 3f = 3g(\Pi_H) - 3$$

This concludes the proof. □

Definition 4.1.3 (Cycles on a spanning tree)**.** Let H be a graph and a a vertex of H. Let $k \geq 1$ and let $C_0, ..., C_k$ be subgraphs of H. We say that $C_0, ..., C_k$ are *cycles on a spanning tree* rooted in a if:

- there exists a spanning tree T of $C_0 \cup ... \cup C_k$ rooted in a such that $(C_0 \cup ... \cup C_k) - T$ consists of $k+1$ edges $\{e_0, ..., e_k\}$ and, for $0 \leq i \leq k$, e_i belongs solely to C_i;
- let C_i' be the unique cycle induced by T and e_i, then C_i is the subgraph consisting of C_i' together with the unique path in T from a to C_i' (if C_i' contains a then $C_i = C_i'$).

Figure 4.2 shows an example of cycles on a spanning tree.

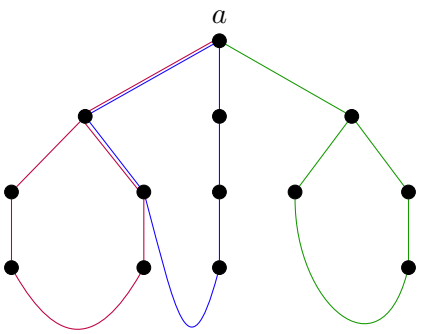


Figure 4.2: Cycles on a spanning tree rooted in a. The three colored subgraphs are cycles on a spanning tree rooted in a.

Proposition 4.1.4 ([HK26b, Proposition 4.10])**.** *Let H be a Π_H-embedded graph and a a vertex of H. Let $C_0, ..., C_k$ be cycles of H on a spanning tree rooted in a that are Π_H-noncontractible and pairwise Π_H-nonhomotopic. Then,*

$$k \leq \begin{cases} g(\Pi_H) & \text{if } g(\Pi_H) \leq 1 \\ 3g(\Pi_H) - 3 & \text{if } g(\Pi_H) \geq 2 \end{cases}$$

Proof. Let T be a spanning tree and and let a be its root. Let $C_0, ..., C_k$ be cycles on T with the same properties as in Definition 4.1.3.

Let H' be the multigraph obtained from H by first removing every vertex and edge in $H - (C_0 \cup ... \cup C_k)$ and then contracting T into the vertex a. Let Π_H' be the embedding of H' induced by Π_H. Then, $e_0, ..., e_k$ are loops on a in H'. Moreover, remark that, for every $0 \leq i \leq k$, e_i in Π_H' is in the same homotopy class as C_i in Π_H. Hence, by Proposition 2.1.22,

$$k \leq \begin{cases} g(\Pi_H) & \text{if } g(\Pi_H) \leq 1 \\ 3g(\Pi_H) - 3 & \text{if } g(\Pi_H) \geq 2 \end{cases}$$

This concludes the proof. □

Then, we prove two simple lemmas for graphs on surfaces that will be useful later.

Lemma 4.1.5 ([HK26b, Lemma 4.11])**.** *Let H be a 2-connected Π_H-embedded graph. Let C be a Π_H-contractible cycle. Then, every face in $Int(C, \Pi_H)$ is a cycle and, therefore, every edge in $int(C, \Pi_H)$ is contained in two distinct faces.*

Proof. If a face f of (H, Π_H) in $\mathrm{Int}(C, \Pi_H)$ is not a cycle, then H contains a cutvertex. However, this contradicts the fact that H is 2-connected. Hence, f must be a cycle.

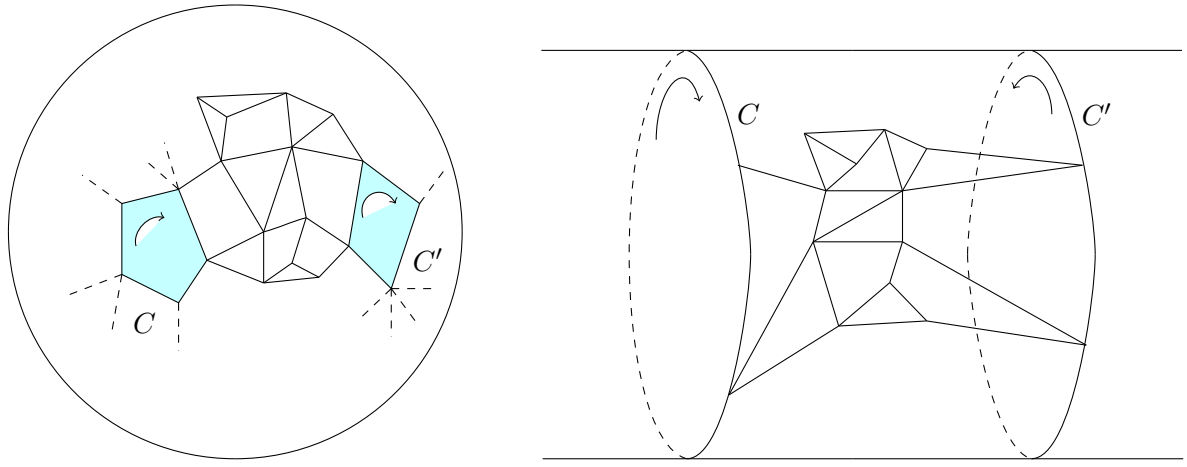
Moreover, as every face in $\mathrm{Int}(C, \Pi_H)$ is a cycle, the two sides of an edge e in $\mathrm{int}(C, \Pi_H)$ cannot be in the same face. Therefore, every edge in $\mathrm{int}(C, \Pi_H)$ is contained in two distinct faces. □

Definition 4.1.6 (Same relative orientation)**.** Let H be a graph that is Π_H-embedded in a surface S_H and Π'_H-embedded in a surface S'_H. Let C and C' be two almost disjoint cycles in H that are Π_H-contractible, both embedded in some disk or cylinder S_0 on S_H, and Π'_H-noncontractible homotopic.

Let $\overline{\Pi'_H}$ be the embedding in the sphere obtained from Π'_H by cutting along C and C'. We fix the clockwise orientation in both S_0 and $\overline{\Pi'_H}$ to be the orientation of some walk around C.

We say that C and C' *have the same relative orientation* in Π_H and Π'_H if the clockwise walk of C' in S_0 is the same as its clockwise walk in $\overline{\Pi'_H}$.

See Figure 4.3 for an illustration.



(a) Embedding of a subgraph of H in Π_H

(b) Embedding of a subgraph of H in Π'_H

Figure 4.3: Two cycles that have the same relative orientation. The cycles C and C' are Π_H-contractible (see Subfigure (a)) and Π'_H-noncontractible homotopic (see Subfigure (b)). They have the same relative orientation in Π_H and Π'_H.

Lemma 4.1.7 ([HK26b, Lemma 4.13])**.** *Let H be a graph that is Π_H-embedded in a surface S_H and $\Pi_{H'}$-embedded in a surface S'_H. Let C and C' be two almost disjoint cycles in H that are Π_H-contractible and Π'_H-noncontractible homotopic. Let (A, B) be the separation of H such that $V(C \cup C') = V(A \cap B)$ and $Int(C \cup C', \Pi_{H'}) = A$. Suppose moreover that $C \cup C' \cup A$ is a Π_H-contractible subgraph of H embedded in some disk or cylinder on S_H, and that C and C' have the same relative orientation in Π_H and in $\Pi_{H'}$.*

Then there exists an embedding Π'_H of H in S'_H so that $\Pi_H(A)$ and $\Pi'_H(A)$ are equivalent.

Proof. Remark that $\text{int}(C \cup C', \Pi'_H(H - A))$ is empty. Then, as C and C' have the same relative orientation in Π_H and in Π'_H, it is possible to embed A inside the cylinder $\text{int}(C \cup C', \Pi'_H(H - A))$ so that the embedding of A obtained is equivalent to $\Pi_H(A)$.

Then, by combining $\Pi'_H(H' - A)$ and this embedding of A inside $\text{int}(C \cup C', \Pi'_H(H' - A))$, we get an embedding of H in S'_H so that $\Pi_H(A)$ and $\Pi'_H(A)$ are equivalent. □

4.2 Basic results on excluded minors for surfaces

Let us fix the following notation for the rest of the thesis: Let S' be a surface of genus g and let G be an excluded minor for S'.

In this section, we first show that G can be embedded in a surface of genus at most $g + 2$ and therefore define an embedding of G in a surface S of genus at most $g + 2$. Then, we observe the consequences on G when removing an edge e from it. Next, we establish results on separators in G. In particular, we prove that it is enough to consider excluded minor G that are 2-connected. Finally, we present the results found by Seymour [Sey93] on the treewidth and tree decompositions of G and give an overview of our improvements on these results.

4.2.1 Embedding of G

We will show that G has genus at most $g + 2$. First, by definition, G is an excluded minor for the surface S' that has genus g. Therefore, G has no embedding in S', but every proper minor of G is embeddable in S'.

For $e \in E(G)$, we define $G_e = G - e$. As G_e is a proper minor of G, it can be embedded in S' and therefore has genus at most g. We define Π_e as an embedding of G_e of genus $g(G_e) \leq g$. We will extend the embedding Π_e of $G - e$ into an embedding of G and show that this extended embedding has genus at most $g + 2$.

Lemma 4.2.1. *Let Π'_e be an embedding of G that extends the embedding Π_e of $G - e$. Then,*

$$g(\Pi'_e) \leq g + 2$$

Proof. Let $C \subseteq G$ be the connected component of G containing e. There are two cases: either $C - e$ is connected or it is not.

- If $C - e$ is connected, then by Euler formula:

  $$|V(C - e)| - |E(C - e)| + |F(C - e, \Pi_e(C - e))| = 2 - g(\Pi_e)$$

  Moreover, remark that $|V(C)| = |V(C - e)|$, $|E(C)| = |E(C - e)| + 1$ and $|F(C, \Pi'_e(C))| \geq |F(C - e, \Pi_e(C - e))| - 1$.

  Therefore, we have

  $$\begin{aligned} 2 - g(\Pi'_e) &= |V(C)| - |E(C)| + |F(C, \Pi'_e(C))| \\ &\geq |V(C - e)| - |E(C - e)| + |F(C - e, \Pi_e(C - e))| - 2 \\ &\geq 2 - (g(\Pi_e) + 2) \end{aligned}$$

  hence implying $g(\Pi'_e) \leq g(\Pi_e) + 2$.

- If $C - e$ is not connected, let C_1 and C_2 be the two connected components of $C - e$. By Euler formula:

$$|V(C_1)| - |E(C_1)| + |F(C_1, \Pi_e(C_1))| = 2 - g(\Pi_e(C_1))$$
$$\text{and } |V(C_2)| - |E(C_2)| + |F(C_2, \Pi_e(C_2))| = 2 - g(\Pi_e(C_2))$$

  Moreover, remark that $|V(C)| = |V(C_1)| + |V(C_2)|$, $|E(C)| = |E(C_1)| + |E(C_2)| + 1$ and $|F(C, \Pi'_e(C))| = |F(C_1, \Pi_e(C_1))| + |F(C_2, \Pi_e(C_2))| - 1$.

  Therefore, we have

$$\begin{aligned} 2 - g(\Pi'_e) &= |V(C)| - |E(C)| + |F(C, \Pi'_e(C))| \\ &= |V(C_1)| - |E(C_1)| + |F(C_1, \Pi_e(C_1))| + |V(C_2)| - |E(C_2)| \\ &\qquad + |F(C_2, \Pi_e(C_2))| - 2 \\ &= 2 - g(\Pi_e(C_1)) + 2 - g(\Pi_e(C_2)) - 2 \\ &= 2 - (g(\Pi_e(C_1)) + g(\Pi_e(C_2))) \\ &= 2 - g(\Pi_e) \end{aligned}$$

  hence implying $g(\Pi'_e) = g(\Pi_e)$.

To conclude, in both cases, the embedding Π'_e has genus at most $g(\Pi_e) + 2 \leq g + 2$. □

Finally, the genus of G is at most the genus of its embeddings Π'_e and, by Lemma 4.2.1, $g(\Pi_e) \leq g + 2$. Consequently, we get $g(G) \leq g(\Pi'_e) \leq g + 2$.

From now on, we define S to be a surface of genus $g(G)$ in which G can be embedded and we define Π as an embedding of G in S.

4.2.2 Bounds for tree decompositions of G

Previous studies of the excluded minors for a surface lead to bounds regarding the tree decomposition of such excluded minors. Seymour [Sey93] proved a $O(g^3)$ bound on the treewidth of G and on the maximum degree of the tree T in a tree decomposition of G (for a tree decomposition with particular properties). We improve both of these bounds to $O(g \log g)$.

As stated in Subsection 2.2.2, the first and, until now, only known bound for the treewidth of G is the following:

Theorem 2.2.8 ([Sey93, (3.3)])**.** *The treewidth of G is bounded by a polynomial in g:*

$$tw(G) \leq T_S(g)$$

with $T_S(g) = 3(g+3)^2(3g+16) - 3 = O(g^3)$.

We improved this bound from $O(g^3)$ to $O(g \log g)$ in [HK26b]:

Theorem 3.1.2 ([HK26b, Theorem 1.4])**.** *The treewidth of G is bounded by the following function of g:*

$$tw(G) \leq T(g)$$

with $T(g) = 264(g+2)(m+1) - 1 = O(g \log g)$, where $m = 2(\lfloor \log_q(3g+4) \rfloor + 2)$ and $q = \frac{9073}{9072}$.

The proof of this result can be found at the end of Subsection 5.3.1.

Note that the constant q used in Theorem 3.1.2 and the one defined in [HK26b] are not the same: indeed, the constant q defined in [HK26b] was higher, equal to $\frac{1153}{1152}$.

Indeed, the generalized structural results established in [HK26a] that are detailed in Section 5.3 required to change the constant q to a smaller value.

Consequently, we state Theorem 3.1.2 with $q = \frac{9073}{9072}$ (which is still true because, with a smaller q, $T(g)$ is bigger) to avoid the multiplication of constants in this manuscript. Note furthermore that this does not change the asymptotic bound of $T(g)$ for the treewidth of G.

4.2.3 Separators in G

We analyze the connectivity of G. We show that there is no 2-separation between a Π-contractible subgraph B and the rest of G, unless B is reduced to an edge. Then we prove that it suffices to consider an excluded minor G that is 2-connected, since the other cases can be reduced to this case.

Lemma 4.2.2 ([HK26b, Lemma 4.14])**.** *Let (A, B) be a 2-separation of G and let $a, b \in V(G)$ so that $A \cap B = \{a, b\}$. Then, either B is an edge or is not contained in a disk in Π.*

Proof. If B is an edge then there is nothing to prove.

Suppose, by contradiction, that B is not an edge but is contained in a disk in Π. See Figure 4.4 for an illustration.

Then, let G' be the graph obtained by contracting $B - b$ into a single vertex a'. As G' is a proper minor of G, G' can be embedded on surface S' by embedding Π'.

Let $e = a'b \in E(G')$. In $\Pi'(G' - e)$, there exists an empty disk with a' and b on its boundary (e is embedded in Π' inside of it). Then, it is possible to embed B in this disk (because B can be embedded in a disk, as it was in Π). But, then, this embedding of B together with $\Pi'(G' - e)$ constitutes an embedding of G onto the surface S', a contradiction. □

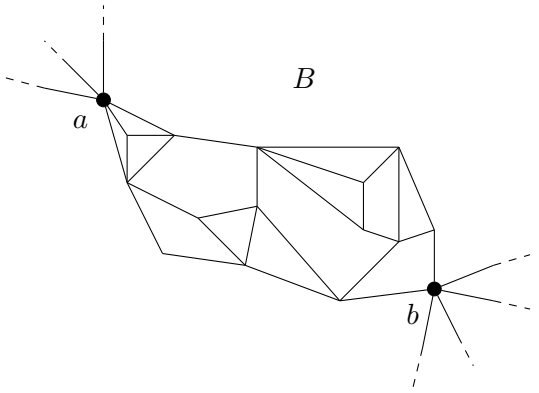


Figure 4.4: The subgraph B is 2-separated from the rest of the graph by $\{a, b\}$ and is contained in a disk in Π.

Lemma 4.2.3 ([HK26b, Lemma 4.15])**.** *Let $G_1, ..., G_p$ ($p \geq 1$) be the 2-connected blocks of G. Then, for $1 \leq i \leq p$, G_i is an excluded minor for some surface S_i.*

Proof. Let $1 \leq i \leq p$. It is sufficient to show that for any edge $e \in G_i$, $G_i - e$ and G_i/e has genus strictly smaller than $g(G_i)$.

Let $G' = G - e$ and $G'_i = G_i - e$ or $G' = G/e$ and $G'_i = G_i/e$. As G' is a proper minor of G, its genus is strictly smaller than $g(G)$. Moreover, $G_1, ..., G_{i-1}, G_{i+1}, ..., G_p$ are still blocks of G' and G'_i is either a block or contains several blocks of G'. By Proposition 2.1.13,

$$\begin{aligned} g(G) &= g(G_1) + ... + g(G_p) \\ &> g(G') = g(G_1) + ... + g(G_{i-1}) + g(G'_i) + g(G_{i+1}) + ... + g(G_p) \end{aligned}$$

Hence, $g(G'_i) < g(G_i)$. This concludes the proof. □

The following result is a refinement of [HK26b, Theorem 4.16]. By adding basic constraints on the bounding function N, we can extend the bound $N(g)$ on the order of the 2-connected excluded minors for a surface to a bound for all its excluded minors with the same function $N(g)$. This improves the result from [HK26b, Theorem 4.16], which had an added factor g.

Corollary 4.2.4 ([HK26a, Corollary 4.11])**.** *Suppose that, for any graph H that is an excluded minor for some surface S_H and that is 2-connected, $|V(H)| \leq N(g(S_H))$ with N having the two following properties:*

- *N is an increasing function, and*
- *$N(g_1 + g_2) \geq N(g_1) + N(g_2)$.*

Then, $|V(G)| \leq N(g)$.

Proof. If G is 2-connected, the result follows immediately.

Suppose not: let $G_1, ..., G_p$ ($p \geq 2$) be the 2-connected blocks of G. Then, by Lemma 4.2.3, for $1 \leq i \leq p$, G_i is an excluded minor for some surface S_i. Remark that, by Proposition 2.1.13, $g(G_1) + ... + g(G_p) = g(G)$.

Finally,

$$\begin{aligned} |V(G)| &\leq \sum_{i=1}^{p} |V(G_i)| && \text{by Proposition 2.1.13} \\ &\leq \sum_{i=1}^{p} N(g(S_i)) && \text{by hypothesis} \\ &\leq \sum_{i=1}^{p} N(g(G_i) - 1) && \text{because } g(G_i) > g(S_i) \\ &\leq N\left(\sum_{i=1}^{p} (g(G_i) - 1)\right) && \text{by using the second property of } N \\ &\leq N(g(G) - p) && \text{because } \sum_{i=1}^{p} g(G_i) = g(G) \\ &\leq N(g) && \text{because } g(G) \leq g + 2 \text{ and } p \geq 2 \end{aligned}$$

□

Corollary 4.2.4 implies that bounding the order of the 2-connected excluded minors for S' yields the same bound on the order of all the excluded minors for S'. Therefore, from this point on, we assume that G is 2-connected.

4.2.4 Structure of $G - e$

Let us prove the following result, which compares the embeddings (G, Π) and (G_e, Π_e).

Lemma 4.2.5 ([HK26b, Lemma 4.17])**.** *Let C be a Π-contractible cycle and let e be an edge with both endpoints in int(C, Π). Let f, f' be the faces in Π that contain e and let $C_e = f \cup f' - e$. Then, C_e is a contractible cycle in (G, Π) such that Int$(C_e, \Pi) = C_e + e$. Moreover, C_e is Π_e-nonseparating (and thus Π_e-noncontractible).*

Proof. As G is 2-connected and f and f' are faces in $\mathrm{Int}(C, \Pi)$, by Lemma 4.1.5, f and f' are distinct faces of (G, Π) and are cycles.

To show that C_e is a cycle, it is enough to show that f and f' intersect only in e. If $f \cap f'$ is disconnected, there are two vertices u and v so that $u, v \in f \cap f'$ and u and v are not consecutive in f nor f'. Then $\{u, v\}$ is a 2-separator in G that separates a Π-contractible subgraph H of G which contains two paths from u to v (respectively subpaths of the facial walks f and f') from the rest of the graph. However, as H is a Π-contractible subgraph and is not an edge, this contradicts Lemma 4.2.2.

We can, therefore, suppose that $f \cap f'$ is connected. Then suppose $f \cap f'$ is a path P of length at least 2. However, a 2-separator in G separates P from the rest of G. Lemma 4.2.2 leads to a contradiction. Finally, $f \cap f'$ is reduced to e, and C_e is indeed a cycle.

As f and f' are Π-contractible cycles and intersects on e, C_e is Π-contractible and $\mathrm{Int}(C_e, \Pi) = C_e + e$.

For the second part of the proof, suppose, by contradiction, that C_e is Π_e-separating.

Let $\mathcal{B}$ be the set of bridges on C_e in G_e. Let B_C be the bridge in $\mathcal{B}$ that contains C. See Figure 4.5 for an illustration of $\mathcal{B}$ in (G, Π). Remark that every bridge $B \in \mathcal{B} - B_C$ is so that $B \subseteq \mathrm{Int}(C, \Pi)$. Without loss of generality, B_C is on the left of C_e in Π_e. Then, let us reembed G_e so that $C_e \cup B_C$ is embedded as in Π_e and $\mathcal{B} - B_C$ is embedded as in Π on the left of C_e (which is possible because, for each bridge $B \in \mathcal{B} - B_C$, there is a disc in $G_l(C_e, \Pi_e)$ with the attaches of B on C_e on its boundary in their order on C_e). This new embedding Π'_e of G_e has genus at most g. Moreover, $G_r(C_e, \Pi'_e)$ is empty. Then, by embedding e in $G_r(C_e, \Pi'_e)$, we obtain an embedding of G in S', a contradiction. □

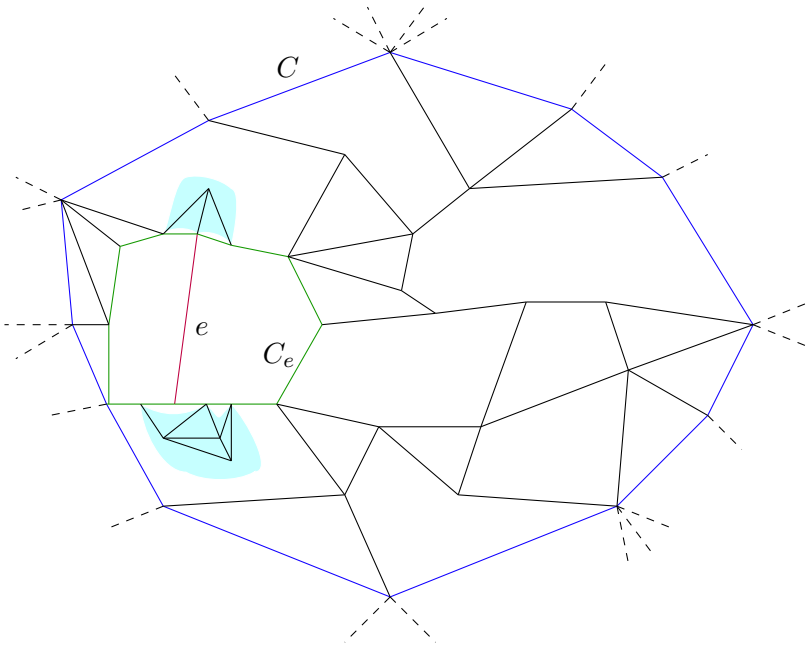


Figure 4.5: The subgraph $\mathrm{Int}(C, \Pi)$ of G is shown embedded in Π. The cycles C and C_e are represented in blue and green, respectively, and the edge e is represented in red. The bridges of $\mathcal{B} - B_C$ are highlighted in pale blue, while the bridge B_C consists of the remaining part of $G - (C_e \cup e)$.

Chapter 5

Structural results

5.1 Our inspiration: Seymour's structural result 68
5.2 Isolated paths 69
5.3 Nested cycles 73
5.3.1 Contractible nested cycles 73
5.3.2 Homotopic nested cycles 85
5.4 Consequences of the structural results 93
5.4.1 Treewidth and tree decomposition 93
5.4.2 Excluded grid topological minor 93
5.4.3 Bound on the number of almost disjoint Π-noncontractible cycles 94
5.5 Links with facewidth 95
5.6 Discussion on the optimality of the structural results 98
5.6.1 Optimality of the bound for the nested/homotopic cycles 98
5.6.2 The scope of the forbidden structure cannot be extended 98

In this chapter, we prove structural results that describe forbidden structures in (G, Π) and will be used repetitively to find contradictions in the main proof.

These results are all improvements over a structural result from Seymour [Sey93, (2.2)] that will be presented in Section 5.1.

In [HK26b, Section 5], we presented two different structural results called isolated paths and nested/homotopic cycles. The bound on the number of isolated paths was found to be $O(g)$ and the bound on the number of nested/homotopic cycles was found to be $O(\log g)$.

In [HK26a, Section 5], we presented results that both generalize and unify the two structural results from [HK26b, Section 5]. The two types of forbidden structures, isolated paths and nested/homotopic squares, from [HK26b] are in this paper unified into the same forbidden structure, called nested/homotopic cycles.

This unification allows us to find a better bound for the structure formerly called isolated paths: in [HK26b], it is proved that (G, Π) contains $O(g)$ isolated paths, whereas in [HK26a] we prove that (G, Π) contains $O(\log g)$ isolated paths. Apart from the extension of the scope of the structural results, the bound for nested/homotopic cycles remains asymptotically the same as in [HK26b].

While Corollary 5.3.15 and Proposition 5.3.27 are direct extensions of results proven in [HK26b], Corollary 5.3.16 is a variation of Corollary 5.3.15 which finds a bound on the number of nested cycles that depends on the size of a separator for the considered subgraph. This dependency on the size of a separator will be taken advantage of in the rest of the manuscript: by finding a contractible subgraph with a separator of logarithmic size in g, we can prove that it contains $O(\log \log g)$ nested cycles (see Proposition 6.3.1), which is a crucial ingredient to establish the almost quadratic bound on the order of an excluded minor for a surface of genus g.

In Section 5.1, we first present the structural result proved by Seymour [Sey93] and give a proof of his result that is simpler than his original proof and has the same asymptotic bound $O(g)$.

Then, in Section 5.2, we present the structure called isolated paths and prove that G contains $O(g)$ isolated paths. Even though the generalized structural result from [HK26a] encompasses this result on isolated paths by both generalizing it to a more general forbidden structure and by improving the asymptotic bound $O(\log g)$, we still decided to present this result as its proof is a lot simpler and shorter than the one of the generalized structural result.

Next, in Section 5.3, we present the generalized structures called nested and homotopic cycles and prove that G contains $O(\log g)$ nested and homotopic cycles. We also explain the difference between the nested/homotopic cycles defined in [HK26b] and the generalized definition of nested/homotopic cycles from [HK26a]. Moreover, we establish the separator variant of our structural result.

In Section 5.4, we give three consequences to our structural results: we improve the bound on the treewidth of G and the bound on the size of a grid subdivision in G. We also establish a bound on the number of almost disjoint Π-noncontractible cycles in (G, Π).

Finally, in Section 5.6, we discuss the optimality of our structural results and what room for improvement there is for these results.

5.1 Our inspiration: Seymour's structural result

In [Sey93], Seymour established the following structural result:

Proposition 5.1.1 ([Sey93, (2.2)])**.** *Let $e \in V(G)$ and suppose that there are $g+2$ disjoint cycles of $G_e = G - e$, all bounding disks in Π which include e. If $G - e$ can be drawn in S', then G can be drawn in S' as well.*

This result is one of the main tools allowing him to find the first bound on the excluded minors for a given surface.

In this section, we give an original proof for this result (that we reformulate) but with a bound of $3g - 2$ instead of $g + 1$ for the number of disjoint cycles. Even though our bound is worse than Seymour's, the two bounds are asymptotically equivalent at $O(g)$ and the proof we give is shorter and simpler than Seymour's proof.

Proposition 5.1.2. *Let $e \in V(G)$. Let $C_1, ..., C_{k+1}$ be disjoint Π-contractible nested cycles of G_e. Suppose moreover that $e \in Int(C_i, \Pi)$ for every $1 \leq i \leq k$. Then,*

$$k \leq \begin{cases} g & \text{if } g \leq 1 \\ 3g - 3 & \text{if } g \geq 2 \end{cases}$$

Proof. Let us show that C_i is Π_e-nonseparating for every $1 \leq i \leq k$. Suppose that C_i is Π_e-nonseparating.

First, let $\tilde{G} = G - \text{int}(C_i, \Pi)$. Let us modify the embedding $\Pi_e(\tilde{G})$ so that one side of the cycle C_i in $\Pi_e(\tilde{G})$ is empty. Let $\mathcal{B}$ be the set of bridges on C_i in $\tilde{G}$. Because C_i is $\Pi_e(\tilde{G})$-separating, then the set of bridges can be partitioned into two sets $\mathcal{B}_1$ and $\mathcal{B}_2$ so that the bridges in $\mathcal{B}_1$ and $\mathcal{B}_2$ are respectively on the left side and on the right side of C_i in $\Pi_e(\tilde{G})$. Let $B_0 \in \mathcal{B}$ be the bridge containing C_{k+1} and suppose that $B_0 \in \mathcal{B}_1$.

Let us show that every bridge in $\mathcal{B}_2$ can actually be embedded on the other side of C_i in $\Pi_e(\tilde{G})$ and let us change the embedding $\Pi_e(\tilde{G})$ accordingly. Indeed, $B \in \mathcal{B}_2$ is in $\text{int}(C_i \cup C_{k+1}, \Pi)$ and is therefore Π-contractible. Therefore, on the left side of C_i in $\Pi_e(\tilde{G})$, there is an empty disk with the attaches of B on C_i in the right order in $\Pi_e(\tilde{G})$. Finally, B can be embedded on the left side of C_i in $\Pi_e(\tilde{G})$ and we modify the embedding $\Pi_e(\tilde{G})$ so that it is the case.

Consequently, $\mathcal{B}_2$ is empty and one side of the cycle C_i in $\Pi_e(\tilde{G})$ is empty. Therefore, it is possible to find an empty cylinder with C_i on its boundary in $\Pi_e(\tilde{G})$. It is then possible to reembed $\text{int}(C_i, \Pi)$ in this empty cylinder as in Π.

Let Π'_e be this new embedding. However, remark that Π'_e is an embedding of G in the same surface as Π_e. Therefore, Π'_e is an embedding of G in a surface of genus g. A contradiction with the definition of G.

Finally, no cycle C_i for $1 \leq i \leq k$ is Π_e-separating.

To conclude, $C_1, ..., C_k$ are disjoint Π_e-nonseparating cycles. Consequently, they are also pairwise Π_e-nonhomotopic. Finally, by Proposition 4.1.2, $k \leq \max\{1, 3g(\Pi_e) - 3\} \leq \max\{1, 3g - 3\}$. □

This result from Seymour [Sey93, (2.2)] was a starting point of our research in order to improve the bound on the order of the excluded minors for surfaces. In what follows, we improve this bound from linear to logarithmic in g and extend this result in two ways:

- We prove that this result holds not only for disjoint cycles but also for cycles that touches in a specific way (on at most 2 pieces).
- We show that this result holds not only for Π-contractible nested cycles but also for Π-noncontractible homotopic ones.

5.2 Isolated paths

The main result of this subsection is Proposition 5.2.3, which shows that a subgraph of G, which is Π-contractible, contains at most $O(g)$ internally disjoint paths from a piece to another piece.

Recall that a *piece* of (G, Π) is either a vertex or a face of (G, Π).

Definition 5.2.1 (Degree of a piece)**.** Let p be a piece of (G, Π), we define $d_{(G,\Pi)}(p)$ (or $d(p)$ if it is clear in the context) to be $d_G(p)$ if p is a vertex and $|p|$ if p is a face.

Definition 5.2.2 (Isolated paths)**.** Let p and p' be two disjoint pieces of (G, Π). We say that G contains m *isolated paths from p to p' in* Π if:

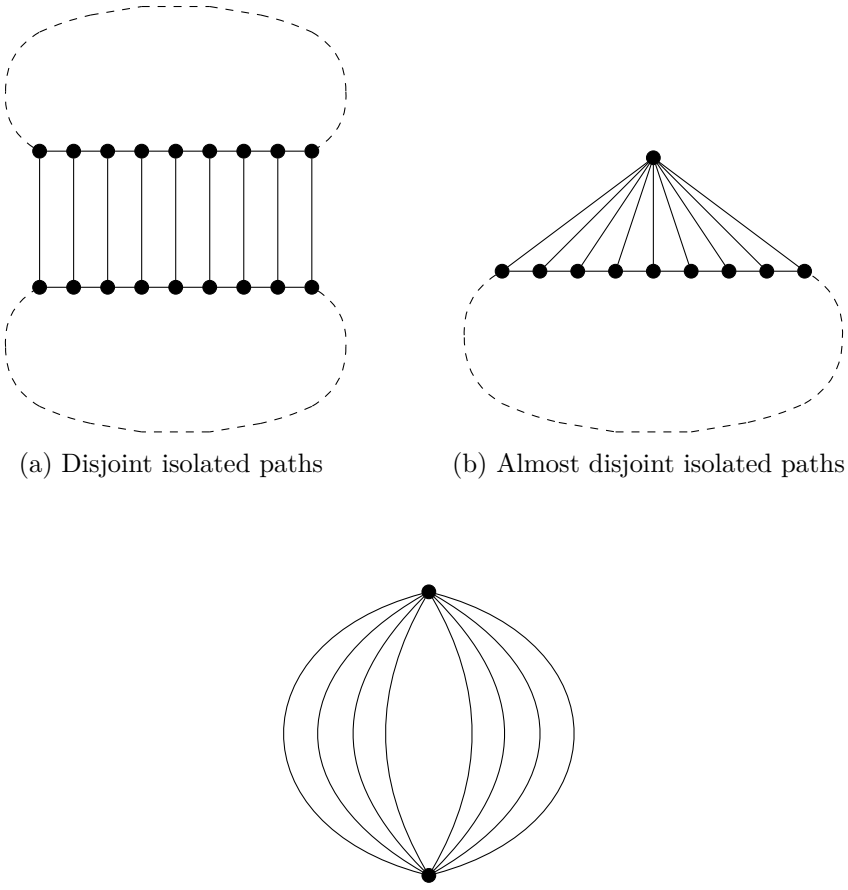
- G contains m internally disjoint paths $P_0, ..., P_{m-1}$ from p to p',

- For $0 \leq i \leq m-1$, let a_i and a'_i be respectively the endpoints of P_i on p and p'. If p (resp. p') is a face of (G, Π), then $a_0, ..., a_{m-1}$ (resp. $a'_0, ..., a'_{m-1}$) are pairwise disjoint and listed in this order on p (resp. p'), and
- For $0 \leq i \leq m-1$, the cycle induced by $P_i \cup P_{i+1} \cup p \cup p'$ is Π-contractible.

We say that the isolated paths are:

- *disjoint* if p and p' are both faces,
- *almost disjoint* if p (resp. p') is a vertex and p' (resp. p) is a face, and
- *joint* if p and p' are both vertices.

An example of the three types of isolated paths can be found in Figure 5.1.



(a) Disjoint isolated paths

(b) Almost disjoint isolated paths

(c) Joint isolated paths

Figure 5.1: Isolated paths. The solid lines indicate paths, whereas the dotted lines show the boundaries of the faces which the isolated paths use.

Proposition 5.2.3 ([HK26b, Proposition 5.3])**.** *G contains at most $4 \times (6g(\Pi) - 5) - 1 \leq 4 \times (6g + 7) - 1$ isolated paths in Π from a piece p to a piece p'.*

Proof. Suppose, by contradiction, that there exist at least $4 \times (6g+7)$ isolated paths from a piece p to a piece p' in Π. Let $P_0, ..., P_m$ with $m \geq 4 \times (6g+7) - 1$ be these paths.

For $0 \leq i \leq m-1$, let C_i be the unique cycle induced by $P_i \cup P_{i+1} \cup p \cup p'$ that does not contain the interior of p (resp. p') if p (resp. p') is a face. Let

$$\begin{aligned}
\mathcal{C}_1 &= \{C_0, C_2, ..., C_{2\times(6g+7)-2}\} \\
\overline{\mathcal{C}_1} &= \{C_0, C_1, ..., C_{2\times(6g+7)-2}\} \\
\mathcal{C}_2 &= \{C_{2\times(6g+7)}, C_{2\times(6g+7)+2}, ..., C_{4\times(6g+7)-2}\} \\
\overline{\mathcal{C}_2} &= \{C_{2\times(6g+7)}, C_{2\times(6g+7)+1}, ..., C_{4\times(6g+7)-2}\}
\end{aligned}$$

Let C_e be the unique cycle induced by $P_{2\times(6g+7)-2} \cup P_{2\times(6g+7)} \cup p \cup p'$ in Π that does not contain the interior of p (resp. p') if p (resp. p') is a face. Then, the interior of C_e is not empty as it contains $P_{2\times(6g+7)-1}$. Let e be the first edge of path $P_{2\times(6g+7)-1}$.

Now, let us modify the embedding Π_e of G_e into another embedding Π'_e of G_e on surface S'. Let us first show that both $\overline{\mathcal{C}_1}$ and $\overline{\mathcal{C}_2}$ contain a Π_e-contractible cycle:

- Suppose first that the isolated paths are disjoint or almost disjoint.

  Remark then that $\mathcal{C}_1 \cup \mathcal{C}_2$ contains almost disjoint cycles of G and $|\mathcal{C}_1| = |\mathcal{C}_2| = 6g+7$. By Proposition 4.1.2, as $|\mathcal{C}_1| \geq 6g+7$ and $g(\Pi) \leq g+2$, either $\mathcal{C}_1$ contains a Π_e-contractible cycle (in which case, we are done) or, by pigeonhole principle, there are three cycles $\tilde{C}_1, \tilde{C}_2, \tilde{C}_3 \in \mathcal{C}_1$ that are Π_e-noncontractible homotopic. Suppose that we are in the second case. There exist $0 \leq i < j < k \leq 6g+6$ such that $\tilde{C}_1 = C_{2i}$, $\tilde{C}_2 = C_{2j}$ and $\tilde{C}_3 = C_{2k}$. Remark that at least two of the cycles $\tilde{C}_1, \tilde{C}_2, \tilde{C}_3$ have the same relative orientation in Π and in Π_e, suppose without loss of generality that it is $\tilde{C}_1$ and $\tilde{C}_2$.

  By Lemma 4.1.7, it is then possible to reembed $\text{int}(C_{2i} \cup C_{2j}, \Pi_e)$ so that its new embedding is equivalent to $\Pi(\text{int}(C_{2i} \cup C_{2j}, \Pi))$. Moreover, it contains at least a cycle (for instance, C_{2i+1}) that becomes Π_e-contractible with the reembedding. The same reasoning applies to $\mathcal{C}_2$.

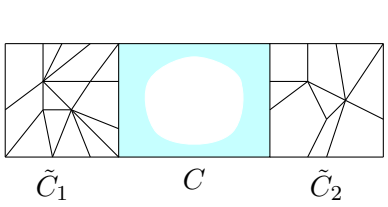
- Suppose now that the isolated paths are joint. Let

$$\begin{aligned}
\mathcal{P}_1 &= \{P_0, P_2, ..., P_{2\times(6g+7)-2}\} \\
\mathcal{P}_2 &= \{P_{2\times(6g+7)}, P_{2\times(6g+7)+2}, ..., P_{4\times(6g+7)-2}\}
\end{aligned}$$

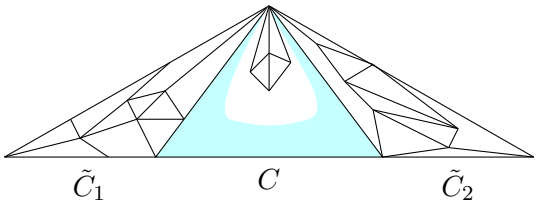
  Then, $\mathcal{P}_1 \cup \mathcal{P}_2$ contains internally disjoint paths from p to p' and $|\mathcal{P}_1| = |\mathcal{P}_2| = 6g+7$. By Proposition 2.1.22, as $|\mathcal{P}_1| \geq 6g+7$ and $g(\Pi) \leq g+2$, there exist two paths $\tilde{P}_1, \tilde{P}_2 \in \mathcal{P}_1$ that are Π_e-homotopic. Then there exist $0 \leq i < j \leq 3g+3$ such that $\tilde{P}_1 = P_{2i}$ and $\tilde{P}_2 = P_{2j}$ are Π_e-homotopic. However, it is then possible to reembed $\text{int}(P_{2i} \cup P_{2j}, \Pi_e)$ so that it is contained in a disk and it contains at least a cycle (for instance C_{2i}). The same reasoning applies to $\mathcal{P}_2$.

Finally, we know that there are two cycles $\tilde{C}_1 = C_i \in \overline{\mathcal{C}_1}$ and $\tilde{C}_2 = C_{2\times(6g+7)+j} \in \overline{\mathcal{C}_2}$ (for some $0 \leq i, j \leq 2 \times (6g+7) - 2$) that are Π_e-contractible. Let C be the cycle induced by $P_{2i+1} \cup P_{2\times(6g+7)+2j} \cup p \cup p'$ that does not contain the interior of p (resp. p') if p (resp. p') is a face.

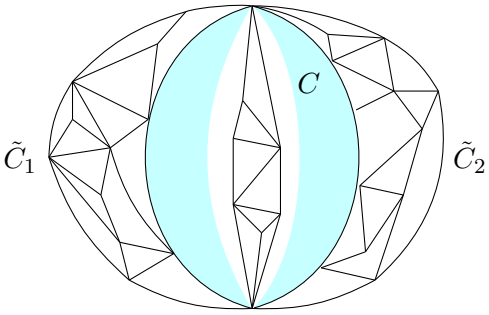
First, let us reembed the bridges in $\mathcal{B}_1 = \text{int}(\tilde{C}_1, \Pi)$ and $\mathcal{B}_2 = \text{int}(\tilde{C}_2, \Pi)$ on the interior of the disks delimited by $\tilde{C}_1$ and $\tilde{C}_2$ in Π_e: the bridges respectively in $\text{ext}(\tilde{C}_1, \Pi)$ and $\text{ext}(\tilde{C}_2, \Pi)$



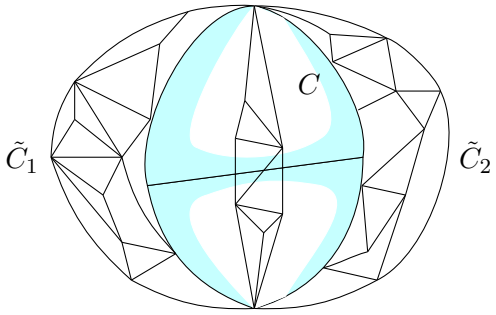
(a) When the isolated paths are disjoint.



(b) When the isolated paths are almost disjoint.



(c) When the isolated paths are joint and there is no path between P_{2i+1} and $P_{2\times(6g+7)+2j}$ in $\mathrm{int}(C,\Pi)$.



(d) When the isolated paths are joint and there is a path between P_{2i+1} and $P_{2\times(6g+7)+2j}$ in $\mathrm{int}(C,\Pi)$.

Figure 5.2: Illustration for the second part of the proof of Proposition 5.2.3. The disk(s) or cylinder in which $\mathrm{int}(C,\Pi)$ can be reembedded is depicted in blue in every one of the cases.

could not have been embedded in them, hence, in $\Pi_e(G_e - (\mathcal{B}_1 \cup \mathcal{B}_2))$, these disks are empty. It is then possible to embed $\mathcal{B}_1$ and $\mathcal{B}_2$ in their respective disks (as in Π, for example).

Let us now try to reembed $\mathrm{int}(C,\Pi)$ in a disk in $\Pi_e(G_e - \mathrm{int}(C,\Pi))$, the different cases are depicted in Figure 5.2.

- Suppose first that the isolated paths are disjoint or almost disjoint.

  Then, it is possible to find in embedding $\Pi_e(G_e - \mathrm{int}(C,\Pi))$ a cylinder with C on its boundary. As $\mathrm{int}(C,\Pi)$ is contained in a disk in S, then we can create a new embedding Π'_e on surface S' with $G_e - \mathrm{int}(C,\Pi)$ embedded as in Π_e and $\mathrm{int}(C,\Pi)$ embedded in this cylinder as in Π, a contradiction.

- Suppose now that the isolated paths are joint.

  First, if there is no path from the interior of P_{2i+1} to the interior of $P_{2\times(6g+7)+2j}$ in $\mathrm{int}(C,\Pi)$ then the bridges on C in $\mathrm{Int}(C,\Pi)$ can be partitioned into two sets $\mathcal{B}_1$ and $\mathcal{B}_2$ having attaches respectively only onto P_{2i+1} and only onto $P_{2\times(6g+7)+2j}$. Let $H_1 = \cup_{B\in\mathcal{B}_1} B$ and $H_2 = \cup_{B\in\mathcal{B}_2} B$. It is possible to find in embedding $\Pi_e(G_e - \mathrm{int}(C,\Pi))$ disjoint disks with respectively P_{2i+1} and $P_{2\times(6g+7)+2j}$ on their boundaries (because one side of each of the paths P_{2i+1} and $P_{2\times(6g+7)+2j}$ is empty). Then, we can create a new embedding Π'_e on surface S' with $G_e - \mathrm{int}(C,\Pi)$ embedded as in Π_e and H_1 and H_2 embedded in these disks as in Π, a contradiction.

Suppose that there is a path $\tilde{P}$ from the interior of P_{2i+1} to the interior of $P_{2\times(6g+7)+2j}$ in $\text{int}(C,\Pi)$. If $\tilde{P}$ has a negative signature, modify the embedding Π_e so that all the edges on $H = \tilde{P} \cup P_{2i+1} \cup (P_{2\times(6g+7)+2j} - \{p, p'\})$ have a positive signature, which is possible because H contains no cycle.

Then, the tree $T = P_{2i+1} \cup \tilde{P} \cup (P_{2\times(6g+7)+2j} - \{p, p'\})$ is contained in Π_e in a disk that has C on its boundary and that is empty in $\Pi_e(G_e - \text{int}(C,\Pi))$. Finally, we obtain a contradiction using the same argument as above. □

5.3 Nested cycles

In this section, we prove three structural results on G: Corollaries 5.3.15 and 5.3.16 and Proposition 5.3.27. These results describe similar forbidden structures in (G,Π).

The proofs in this section are from [HK26a, Section 5] which are in turn largely similar to the proofs in [HK26b, Section 5.2].

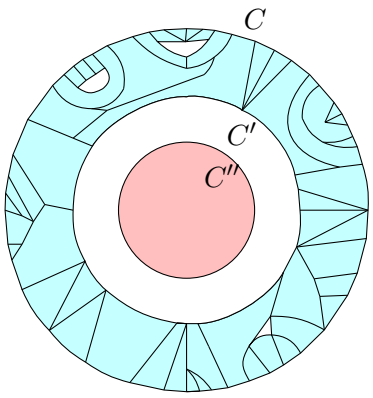
5.3.1 Contractible nested cycles

Let C be a Π-contractible cycle of G. Then, let $\mathcal{F}(C,\Pi)$ be the set of all faces inside of C in Π. Moreover, for a subgraph G_0 of G that is Π-contractible, let $\mathcal{F}(G_0,\Pi) = \bigcup_{C\subseteq G_0} \mathcal{F}(C,\Pi)$.

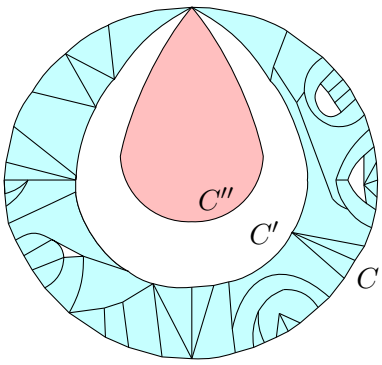
Definition 5.3.1 (Well nested)**.** Let H be a graph with embedding Π_H. Let C and C' be two Π_H-contractible cycles of H such that C' is Π_H-nested in C. We say that C' is Π_H*-well-nested* (or *well nested*, if it is clear in the context) in C if one of the following conditions holds:

- C and C' are disjoint (*(free,free)-well-nested* or *freely well nested*), or
- C and C' intersect at a vertex v (*(v,free)-well-nested* or *well nested pinched on the vertex* v), or
- C and C' both intersect the same face f of (G,Π) and if C intersects f in a subpath P (of length at least 3) then C' intersect f in a subpath that is in the interior of P (*(f,free)-well-nested* or *well nested pinched on the face* f), or
- C and C' intersect at two disjoint vertices v and v' (*(v,v′)-well-nested* or *well nested pinched on the vertices* v *and* v'), or
- C and C' both intersect the same two faces f and f' of (G,Π) and if C intersects f (resp. f') in a subpath P (resp. P'), of length at least 3, then C' intersect f (resp. f') in a subpath that is in the interior of P (resp. P') (*(f,f′)-well-nested* or *well nested pinched on the faces* f *and* f'), or
- C and C' both intersect at a vertex v and a face f of (G,Π) and if C intersects f in a subpath P (of length at least 3) then C' intersect f in a subpath that is in the interior of P (*(v,f)-well-nested* or *well nested pinched on the vertex* v *and the face* f).

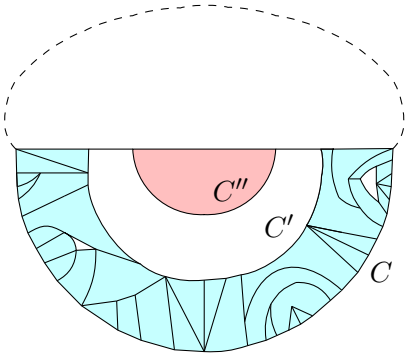
In the second and third cases, we say that C and C' are *well nested pinched on a piece* p. In the last three cases, we say that C and C' are *well nested pinched on two pieces* p *and* p'.



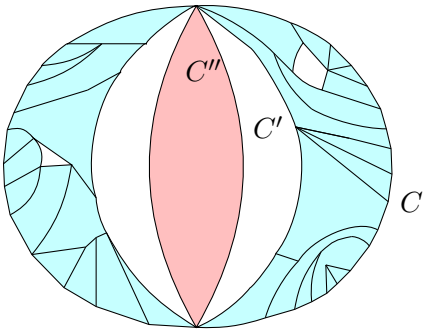
(a) Full contractible square



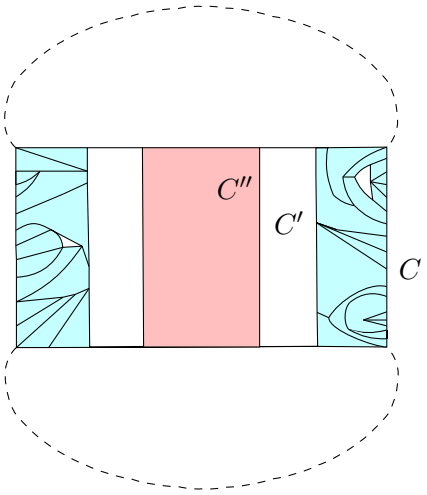
(b) Contractible square pinched on a vertex



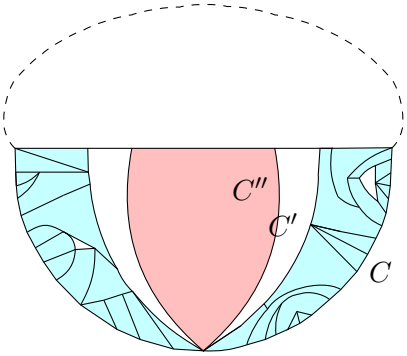
(c) Contractible square pinched on a face



(d) Contractible square pinched on two vertices



(e) Contractible square pinched on two faces



(f) Contractible square pinched on a vertex and a face

Figure 5.3: Contractible squares. The solid lines indicate paths, whereas the dotted line shows the face's boundary on which the contractible square in Figures (c), (e), and (f) is pinched. The zone represented in blue is $\mathcal{B}(C)$ (defined, e.g., in Definition 5.3.5), and the zone represented in pink is $\mathcal{I}(C)$ (defined in Definition 5.3.8).

Let $k \geq 1$ and let $C_0, ..., C_k$ be Π_H-contractible nested cycles of H. We say that $C_0, ..., C_k$ are Π_H-*well-nested* if either:

- for every $0 \leq i \leq k-1$, C_i is Π_H-freely-well-nested in C_{i+1}, or
- there is a piece p such that, for every $0 \leq i \leq k-1$, C_i is Π_H-well-nested in C_{i+1} pinched on p, or
- there are two pieces p and p' such that, for every $0 \leq i \leq k-1$, C_i is Π_H-well-nested in C_{i+1} pinched on p and p'.

The six figures of Figure 5.3 show cycles C, C', C'' that are respectively freely-well-nested, well-nested pinched on a vertex, well-nested pinched on a face, well-nested pinched on two vertices, well-nested pinched on two faces, and well-nested pinched on a vertex and a face in the reversed order.

Definition 5.3.2 (Closest)**.** Let H be a graph with embedding Π_H. Let C and C' be two Π_H-contractible cycles with C' nested in C of H. Let $\mathcal{P}$ be a property on the cycles of (H, Π_H).

Then, we say that C' is the cycle *cycle closest to a cycle* to C that has property $\mathcal{P}$ if, for every cycle C'' that is Π_H-contractible nested in C and has property $\mathcal{P}$, we have $C' \subseteq \text{Int}(C \cup C'', \Pi_H)$.

Definition 5.3.3 (Cycles in this order)**.** Let H be a graph with embedding Π_H.

Let $k \geq 1$ and let $C_0, ..., C_k$ be Π_H-noncontractible homotopic cycles of H. We say that $C_0, ..., C_k$ are *in this order* in Π_H if for $0 \leq i \leq k-1$ and $0 \leq j \leq k$, $C_j \cap \text{int}(C_i \cup C_{i+1}, \Pi_H) = \emptyset$.

For instance, in Figure 5.4(b), the cycles C_1, C_2, C_3 are homotopic in this order.

Definition 5.3.4. Let $C_1, ..., C_m$ be Π-contractible cycles in G that are Π-well-nested in this order. For $3 \leq i \leq m$, we define:

- $\mathcal{B}(C_i)$ as the set of faces in $\mathcal{F}(C_i \cup C_{i-1}, \Pi)$ that touch C_i (we call this set the *boundary* of C_i),
- $\mathcal{B}_{\text{Con}}(C_i, \Pi_e)$ as the set of the faces in $\mathcal{B}(C_i)$ that are Π_e-faces,
- $\mathcal{B}_{\text{N}}(C_i, \Pi_e)$ as the set of Π-faces in $\mathcal{B}(C_i)$ that are not Π_e-faces,
- $\mathcal{I}(C_i)$ as the set of Π-faces in $\text{Int}(C_{i-2}, \Pi)$, and
- $\mathcal{I}_{\text{N}}(C_i, \Pi_e)$ as a maximal size subset of Π-faces of $\mathcal{I}(C_i)$ that are almost-disjoint and that are not Π_e-faces.

Definition 5.3.5 (Contractible square)**.** Let C, C', C'' be Π-contractible cycles in G, well nested in the reverse order.

We say that (C, C', C'') is a *contractible square* with respect to C if C' is the cycle closest to C so that $\mathcal{B}(C) \subseteq \text{Int}(C \cup C', \Pi)$.

We say that a contractible square (C, C', C'') is respectively *full*, *pinched on* p and *pinched on* p *and* p' if C, C', C'' are respectively Π-fully-well-nested, Π-well-nested pinched on p and Π-well-nested pinched on p and p'.

The six figures of Figure 5.3 depict contractible squares (C, C', C'') that are respectively full, pinched on a vertex, pinched on a face, pinched on two vertices, pinched on two faces, and pinched on a vertex and a face.

Recall that, for $e \in E(G)$, we defined $G_e = G - e$ and Π_e as an embedding of G_e of genus $g(G_e) \leq g$. In this section, we ask that Π_e has the following property:

Definition 5.3.6 (Irreducible)**.** Let $C_1, ..., C_m = C$ be Π-contractible cycles that are Π-well-nested in this order so that $e \in \mathrm{Int}(C_1, \Pi)$.

For $2 \leq i \leq m-2$, we say that $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ is *irreducible* if there exists no other embedding Π'_e of G_e of genus g and $2 \leq i \leq m-2$ so that $\Pi_e(G - \mathcal{I}(C_{i+2}))$ and $\Pi'_e(G - \mathcal{I}(C_{i+2}))$ are identical embeddings and $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi'_e)$ is included into $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$.

We say that the embedding Π_e of G_e is *irreducible* with respect to $C_1, ..., C_m$ if, for $2 \leq i \leq m-2$, $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ is *irreducible*.

We show that, for any m Π-contractible Π-well-nested cycles, there indeed exists an embedding Π_e that is irreducible with respect to them.

Proposition 5.3.7. *Let $C_1, ..., C_m$ be Π-contractible cycles that are Π-well-nested in this order so that $e \in Int(C_1, \Pi)$.*

There exists an embedding Π_e of G_e that is irreducible with respect to $C_1, ..., C_m$.

Proof. Take Π_e to be an embedding of G_e of genus g so that $\mathcal{I}_{\mathrm{N}}(C_m, \Pi_e)$ is irreducible, which is always possible. We can easily see that $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ is irreducible for $2 \leq i \leq m-2$. Let $2 \leq i \leq m-2$ and suppose that $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ is not irreducible, then there exists an embedding Π'_e so that $\Pi_e(G - \mathcal{I}(C))$ and $\Pi'_e(G - \mathcal{I}(C))$ are identical embeddings and $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi'_e) \subsetneq \mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$. It implies that $\mathcal{I}_{\mathrm{N}}(C_m, \Pi'_e) \subsetneq \mathcal{I}_{\mathrm{N}}(C_m, \Pi_e)$ which contradicts the irreducibility of $\mathcal{I}_{\mathrm{N}}(C_m, \Pi_e)$. Therefore, $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ is irreducible for every $2 \leq i \leq m-2$ and Π_e is irreducible with respect to $C_1, ..., C_m$. □

For $C_1, ..., C_m$ Π-contractible cycles that are Π-well-nested in this order so that $e \in \mathrm{Int}(C_1, \Pi)$, we fix Π_e to be an embedding of G_e of genus $g(G_e) \leq g$ irreducible with respect to $C_1, ..., C_m$.

Now that Π_e is fixed, we omit its mention in $\mathcal{B}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ and $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ and only write $\mathcal{B}_{\mathrm{N}}(C_{i+2})$ and $\mathcal{I}_{\mathrm{N}}(C_{i+2})$ for every $2 \leq i \leq m-2$.

Definition 5.3.8 (Good/bad contractible square)**.** Let $e \in \mathrm{int}(C'', \Pi)$. Let (C, C', C'') be a contractible square with respect to C. We say that (C, C', C'') is a *bad square with respect to e and Π_e* if $\mathcal{I}_{\mathrm{N}}(C)$ contains more than $18 \times (42|\mathcal{B}_{\mathrm{N}}(C)| - 3)$ Π-faces.

Otherwise, we say that this square is *good with respect to e and Π_e*.

Definition 5.3.9 (Edge-sharing components)**.** Let H be a nonplanar graph and let Π_H be an embedding of H in a surface. Let $\mathcal{F}$ be a set of faces of (H, Π_H). We say that two faces of (H, Π_H) are edge-sharing if they share at least one edge. Let the *edge-sharing components* of $\mathcal{F}$ be its edge-sharing equivalence classes.

Let $\mathcal{C}$ be an edge-sharing component of $\mathcal{F}$. Suppose it induces a Π_H-contractible subgraph H' of H. Then, remark that, by Proposition 2.1.17, the Π_H-embedding of H' is the only planar embedding of H' up to equivalence. In particular, if H' is Π'_H-contractible for some other embedding Π'_H of H in a surface, the boundary of H' in Π'_H is the same as in Π_H.

Remark that the subgraphs induced by the edge-sharing components of $\mathcal{F}$ are 2-connected, but they do not always correspond to the 2-connected blocks of the subgraph induced by $\mathcal{F}$.

Lemma 5.3.10 ([HK26a, Lemma 5.1])**.** *Let (C, C', C'') be a contractible square. Let $e \in int(C'', \Pi)$.*

Let $F_1, ..., F_m \in \mathcal{I}_{\mathrm{N}}(C)$ be a maximal size set of almost disjoint Π-faces that are Π_e-noncontractible homotopic, then $m \leq 10$.

Proof. See Figure 5.4 for illustrations for the proof.

Suppose that there are at least 11 almost disjoint Π-faces in $\mathcal{I}_{\mathrm{N}}(C)$ that are Π_e-noncontractible homotopic. Let $\mathcal{C}$ be the set of these faces (that are cycles by Lemma 4.1.5).

If the square is pinched on two vertices v and v', let us select cycles from $\mathcal{C}$ so that one of v, v' does not intersect any of the selected cycles. First, remark that there is at most one cycle in $\mathcal{C}$ that intersects both v and v': indeed, the faces are chosen to be almost-disjoint. Finally, there are at most $\frac{11-1}{2} = 5$ cycles in $\mathcal{C}$ so that one of v, v' does not intersect any of the selected cycles.

If the square is pinched on two vertices v and v', let C_1, C_2, C_3, C_4, C_5 be the cycles from $\mathcal{C}$ so that one of v, v' does not intersect any of these cycles. If the square is not pinched on two vertices, let C_1, C_2, C_3, C_4, C_5 be distinct cycles from $\mathcal{C}$ (arbitrarily chosen). Without loss of generality, we have that C_1, C_2, C_3, C_4, C_5 are Π_e-homotopic in this order.

Remark that there are at least two of the cycles C_1, C_3, C_5 that have the same relative orientation in Π and Π_e, suppose without loss of generality that it is C_1 and C_3. Let us show that it is possible to modify Π_e so that C_2 becomes contractible.

First, note that:

- If the square is freely well nested: no vertex of C is in $\mathrm{Int}(C_1 \cup C_3, \Pi_e)$, otherwise the whole subgraph $\mathrm{Ext}(C, \Pi)$ would be in $\mathrm{Int}(C_1 \cup C_3, \Pi_e)$ which is impossible because $\mathrm{Ext}(C, \Pi)$ cannot not be Π_e-contractible. Therefore, $\mathrm{Int}(C_1 \cup C_3, \Pi_e) \subseteq \mathrm{Int}(C, \Pi)$.

- If the square is pinched on a piece p: the same reasoning for $C - p$ shows that no vertex of $C - p$ is in $\mathrm{Int}(C_1 \cup C_3, \Pi_e)$. If p is a face, it implies that $\mathrm{Int}(C_1 \cup C_3, \Pi_e) \subseteq \mathrm{Int}(C, \Pi)$ as there is no edge in $\mathrm{Ext}(C, \Pi)$ with an endpoint on p. If p is a vertex, then a bridge from $\mathrm{ext}(C, \Pi)$ in $\mathrm{Int}(C_1 \cup C_3, \Pi_e)$, if it exists, would be attached solely to v, which contradicts the fact that G is 2-connected and therefore has no cutvertex.

- If the square is pinched on two pieces p and p': the same reasoning for $C - (p \cup p')$ shows that no vertex of $C - (p \cup p')$ is in $\mathrm{Int}(C_1 \cup C_3, \Pi_e)$. If both p and p' are faces, it implies that $\mathrm{Int}(C_1 \cup C_3, \Pi_e) \subseteq \mathrm{Int}(C, \Pi)$ as there is no edge in $\mathrm{Ext}(C, \Pi)$ with an endpoint on p or p'. If p is a vertex and p' is a face, then a bridge from $\mathrm{ext}(C, \Pi)$ in $\mathrm{Int}(C_1 \cup C_3, \Pi_e)$, if it exists, would be attached solely to p, which contradicts the fact that G is 2-connected and therefore has no cutvertex. If both p and p' are vertices, then, as we showed before, $C_1 \cup C_3$ does not intersect both p and p', and we can conclude with the same reasoning as when the square is pinched on one vertex.

Therefore, $\mathrm{Int}(C_1 \cup C_3, \Pi_e) \subseteq \mathrm{Int}(C, \Pi)$ and $\Pi(\mathrm{Int}(C_1 \cup C_3, \Pi_e))$ is an embedding in a disk in S in which C_1, C_2 and C_3 are Π-facial walks.

Finally, by Lemma 4.1.7, we can modify the embedding Π_e so that C_2 is now a Π_e-contractible cycle. This contradicts the minimality of $\mathcal{I}_N(C)$. □

Lemma 5.3.11 ([HK26a, Lemma 5.2])**.** *Let $e \in G$ and let H_1, H_2 be subgraphs of G_e such that either:*

- *H_1 and H_2 are disjoint, or*

- *H_1 and H_2 intersect in a path P (possibly reduced to a vertex), and the internal vertices of P have degree 2 in H_1, or*

- *H_1 and H_2 intersect in two paths P and P' (possibly reduced vertices) and the internal vertices of P and P' have degree 2 in H_1.*

Then,

$$g(\Pi_e(H_1)) \leq g - g(\Pi_e(H_2)) + 6$$

Proof. Let $v = |V(H_1) \cap V(H_2)|$ and $e = |E(H_1) \cap E(H_2)|$. Let f_{H_1} and f_{H_2} be respectively the number of faces in an embedding of H_1 of genus $g(\Pi_e(H_1))$ and in an embedding of H_2 of genus $g(\Pi_e(H_2))$.

- If H_1 and H_2 are disjoint: Then, $g(\Pi_e(G_e)) \geq g(\Pi_e(H_1)) + g(\Pi_e(H_2))$.
- If H_1 and H_2 intersect in a path: Remark that $e = v - 1$ and that, when trying to merge embeddings of H_1 and H_2 into an embedding of $H_1 \cup H_2$, at least 1 face from each embedding is modified and at most 4 new faces appear.

  We have

  $$\chi(H_1) = V(H_1) - E(H_1) + f_{H_1} \quad \text{and} \quad \chi(H_2) = V(H_2) - E(H_2) + f_{H_2}$$

  This implies

  $$\begin{aligned}\chi(G_e) &\leq (V(H_1) + V(H_2) - v) - (E(H_1) + E(H_2) - e) + (f_{H_1} + f_{H_2} - 2 + 4) \\ &\leq \chi(H_1) + \chi(H_2) - v + v - 1 - 2 + 4 \\ &= \chi(H_1) + \chi(H_2) + 1\end{aligned}$$

  Therefore,

  $$\begin{aligned}g(G_e) &= 2 - \chi(G_e) \\ &\geq 2 - \chi(H_1) + 2 - \chi(H_2) - 2 - 1 \\ &= g(H_1) + g(H_2) - 3\end{aligned}$$

- If H_1 and H_2 intersect in two paths: Remark that $e = v - 2$ and that, when trying to merge embeddings of H_1 and H_2 into an embedding of $H_1 \cup H_2$, at least 1 face from each embedding is modified and at most 8 new faces appear.

  We have

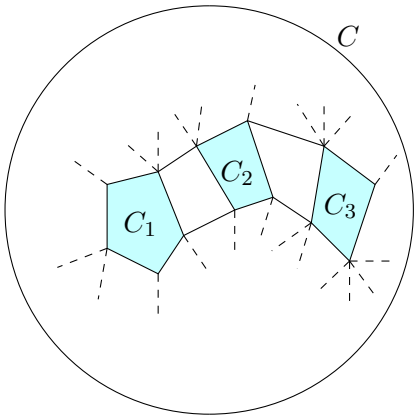
  $$\chi(H_1) = V(H_1) - E(H_1) + f_{H_1} \quad \text{and} \quad \chi(H_2) = V(H_2) - E(H_2) + f_{H_2}$$

  This implies

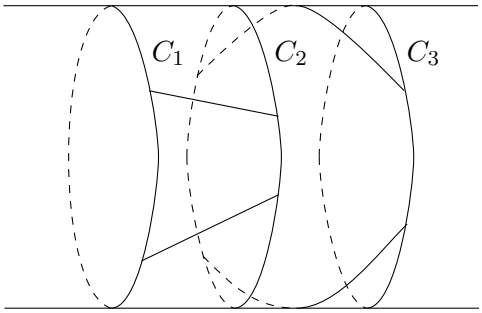
  $$\begin{aligned}\chi(G_e) &\leq (V(H_1) + V(H_2) - v) - (E(H_1) + E(H_2) - e) + (f_{H_1} + f_{H_2} - 2 + 8) \\ &\leq \chi(H_1) + \chi(H_2) - v + v - 2 - 2 + 8 \\ &= \chi(H_1) + \chi(H_2) + 4\end{aligned}$$

  Therefore,

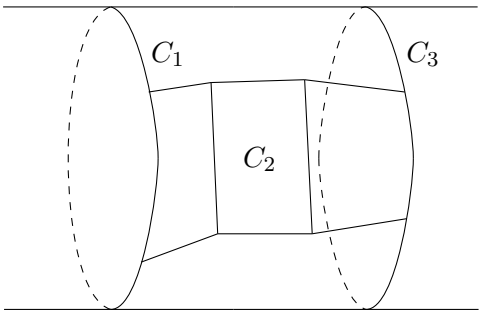
  $$\begin{aligned}g(G_e) &= 2 - \chi(G_e) \\ &\geq 2 - \chi(H_1) + 2 - \chi(H_2) - 2 - 4 \\ &= g(H_1) + g(H_2) - 6\end{aligned}$$



(a) The cycles C, C_1, C_2 and C_3 in Π



(b) The cycles C_1, C_2 and C_3 in Π_e



(c) The embedding constructed in Lemma 5.3.10

Figure 5.4: Illustration for the proof of Lemma 5.3.10

In every case, we get $g(\Pi_e(H_1)) \leq g(\Pi_e(G_e)) - (g(\Pi_e(H_2)) - 6)$. Finally,

$$g(\Pi_e(H_1)) \leq g - g(\Pi_e(H_2)) + 6$$

□

Lemma 5.3.12 ([HK26a, Lemma 5.3])**.** *Let $e \in int(C'', \Pi)$. Then, G contains no bad contractible square with respect to e and Π_e.*

Proof. Suppose by contradiction that G contains a bad contractible square (C, C', C'') with respect to e and Π_e. Therefore $|\mathcal{I}_{\mathrm{N}}(C)| \geq 18 \times (42|\mathcal{B}_{\mathrm{N}}(C)| - 3)$. First, let us show that there exists no bad contractible square with respect to e and Π_e with $|\mathcal{B}_{\mathrm{N}}(C)| = 0$.

Claim 5.3.12.1. $|\mathcal{B}_{\mathrm{N}}(C)| > 0$

Proof of the claim. Suppose that $|\mathcal{B}_{\mathrm{N}}(C)| = 0$.

Then, $\mathcal{B}(C)$ contains at most 2 edge-sharing components and remark that the embedding of each of these is the same in Π and Π_e.

Let $\tilde{G}$ be the graph obtained from G_e by removing $H = \text{int}(C', \Pi)$. $\tilde{G}$ has genus at most g (because it is a subgraph of G_e).

If (C, C', C'') are freely well nested or well nested pinched on a piece p, then $\mathcal{B}(C)$ contains a unique edge-sharing component, and remark that its embeddings in Π_e and Π are identical. It is then quite easy to find an empty disk in $\Pi_e(\tilde{G})$ with the vertices of C' on its boundary in this order (simply by following the walk on C' in $\Pi_e(\tilde{G})$).

Now suppose that (C, C', C'') are well nested pinched on two pieces. In that case, the two edge-sharing components of $\mathcal{B}(C)$ have identical embeddings in Π_e and Π. We only need to show that they are connected with the expected orientation.

Let B_1 and B_2 be the two edge-sharing components of $\mathcal{B}(C)$. If the square is pinched on two vertices v and v', let us define $r = v$ and $r' = v'$. If the square is pinched on a vertex v and a face f, let $r = v$ and $r' = C' \cap f$. If the square is pinched on two faces f and f', let $r = C' \cap f$ and $r' = C' \cap f'$.

We define D_1 and D_2 to be respectively $(B_1 \cup r \cup r') \cap C'$ and $(B_2 - (r \cup r')) \cap C'$. If there is no path in $\text{Int}(C', \Pi)$ from D_1 to D_2, then it is enough to find two empty disks in $\Pi_e(\tilde{G})$ with respectively on their boundaries the vertices of D_1 and D_2 in this order, which we can easily do. Otherwise, let P be a path from D_1 to D_2.

If P connects B_1 and B_2 with the expected orientation, then we can find an empty disk in $\Pi_e(\tilde{G})$ with the vertices of C' on its boundary in this order. Otherwise, suppose that no path P connects B_1 and B_2 with the expected orientation. In particular, if the square is pinched on a face f (maybe together with another piece), the path P_f between B_1 and B_2 that is a subwalk of f does not connect them with the expected orientation. We deduce that P and P_f have the same signature.

Now, let us look at the two faces f_1 and f_2 on the path P in $\Pi_e(\tilde{G} \cup P)$. We can see that $f_1 = f_2$ because it is the case in $\Pi_e(\tilde{G})$. Therefore, changing the signature of P does not increase the genus of the embedding; thus, by this manipulation, we return to the previous case.

In each case, we find an embedding of G in the surface S' of genus g.

This is a contradiction. ■

We can now suppose that $|\mathcal{B}_{\text{N}}(C)| > 0$.

Remove $H = \text{int}(C', \Pi)$ from G_e to obtain a graph $\tilde{G}$.

Claim 5.3.12.2. *The embedding of $\tilde{G}$ induced by Π_e has genus at most $g - 8|\mathcal{B}_{\text{N}}(C)|$.*

Proof of the claim. Remark that $\mathcal{I}_{\text{N}}(C) \subseteq \mathcal{I}(C)$ and therefore $\Pi_e(\mathcal{I}(C))$ contains $18 \times (42|\mathcal{B}_{\text{N}}(C)| - 3)$ almost disjoint Π_e-noncontractible cycles. Moreover, by Lemma 5.3.10, there are at least $\frac{18}{10} \times (42|\mathcal{B}_{\text{N}}(C)| - 3) \geq 42|\mathcal{B}_{\text{N}}(C)| - 3$ of them that are pairwise Π_e-nonhomotopic. Finally, by Proposition 4.1.2, $\Pi_e(\mathcal{I}(C))$ has genus at least $14|\mathcal{B}_{\text{N}}(C)|$.

By Lemma 5.3.11, we have $g(\Pi_e(\tilde{G})) \leq g - g(\Pi_e(\mathcal{I}(C)) + 6$.

As $\Pi_e(\mathcal{I}(C))$ has genus at least $14|\mathcal{B}_{\text{N}}(C)|$, we finally conclude that $\Pi_e(\tilde{G})$ has genus at most $g - 14|\mathcal{B}_{\text{N}}(C)| + 6 \leq g - 8|\mathcal{B}_{\text{N}}(C)|$. ■

Now, we will construct a new embedding of G from the embedding $\Pi_e(\tilde{G})$. See Figure 5.5 for an illustration.

Let $\tilde{\mathcal{B}_{\text{N}}}$ be the subset of faces of $\mathcal{B}_{\text{N}}(C)$ that touch C'.

The edge-sharing components of $\mathcal{B}_{\text{Con}}(C)$ that touch C' and the faces in $\tilde{\mathcal{B}_{\text{N}}}$ can be ordered with respect to C'. Then, we add a handle or crosscap between each consecutive edge-sharing component of $\mathcal{B}_{\text{Con}}(C)$/face of $\tilde{\mathcal{B}_{\text{N}}}$ x and y so that, for an edge-sharing component, the part

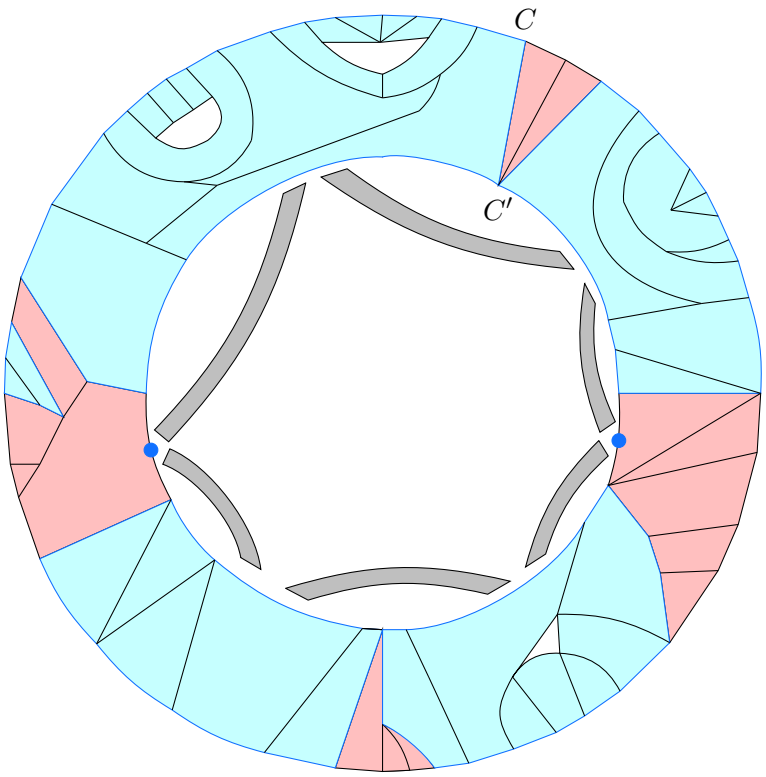


Figure 5.5: Illustration for the end of the proof of Lemma 5.3.12. The vertices in $\tilde{V}(C')$ are depicted in blue. The edge-sharing components of $\mathcal{B}_{\text{Con}}(C)$ are in light blue and have a blue outline. The bridges attached to only one edge-sharing component are in the white zones in the blue edge-sharing components. The faces from $\mathcal{B}(C)$ that are Π_e-noncontractible are depicted in pink. Finally, the added handles are depicted in grey. They are between consecutive edge-sharing components/vertices and can be handles or crosscaps.

of its boundary that contains endpoints of edges from H is reachable from the handle and has the expected orientation. Moreover, we add at most one handle or crosscap when the square is pinched into two pieces to make sure that the two edge-sharing components of $\mathcal{B}(C)$ are connected with the expected orientation.

By Claim 5.3.12.2, the embedding $\Pi_e(\tilde{G})$ has genus at most $g - 6|\mathcal{B}_{\text{N}}(C)|$. Remark that we add at most $2|\mathcal{B}_{\text{N}}(C)| + 1 \leq 3|\mathcal{B}_{\text{N}}(C)|$ handles and crosscaps, and therefore creates a new embedding of genus at most g. Moreover, this new embedding contains an empty cylinder with C' on one of its boundaries. Then, H can be embedded in a planar way (for example, as in Π) in this cylinder. The graph G can then be embedded in a surface of genus at most g, a contradiction. □

Lemma 5.3.13 ([HK26a, Lemma 5.4])**.** *Let C and C' be two Π-contractible nested cycles. Let $\mathcal{F}$ be a subset of the faces in $Int(C \cup C', \Pi)$ that intersect both C and C'.*

There exists a subset of faces from $\mathcal{F}$ of size at least $\frac{|\mathcal{F}|}{6}$ that are almost disjoint.

Proof. Let $\mathcal{B}(\mathcal{F})$ be the partition of $\mathcal{F}$ into its edge-sharing components.

We will then partition each edge-sharing component $B \in \mathcal{B}(\mathcal{F})$ into units $\mathcal{U}(B)$. Let $B \in \mathcal{B}(\mathcal{F})$. Remark that the faces in B are naturally (circularly) ordered inside B so that two faces

are consecutive in this order if they are edge-sharing. If the faces in B are linearly ordered, take the first face f_0 in this order, or any face f_0 of B if the faces in B are circularly ordered. Let P_0 be the path from C to C' in f_0 that is the first we go through when following the order of the faces of B from f_0 (if the order is circular, then P_0 is chosen so that we go through P_0 and then the other path from C to C' in f_0 consecutively in our traversal). We will define the units of B from f_0 using the following method: let f_1 be the first face vertex-disjoint from P_0 in order. Then, the first unit of B contains f_0 and all the faces between f_0 (included) and f_1 (excluded) in order. Iterate this method with f_1 to find the next unit and so on until every face of B belongs to a unit of B. See Figure 5.6 for an illustration.

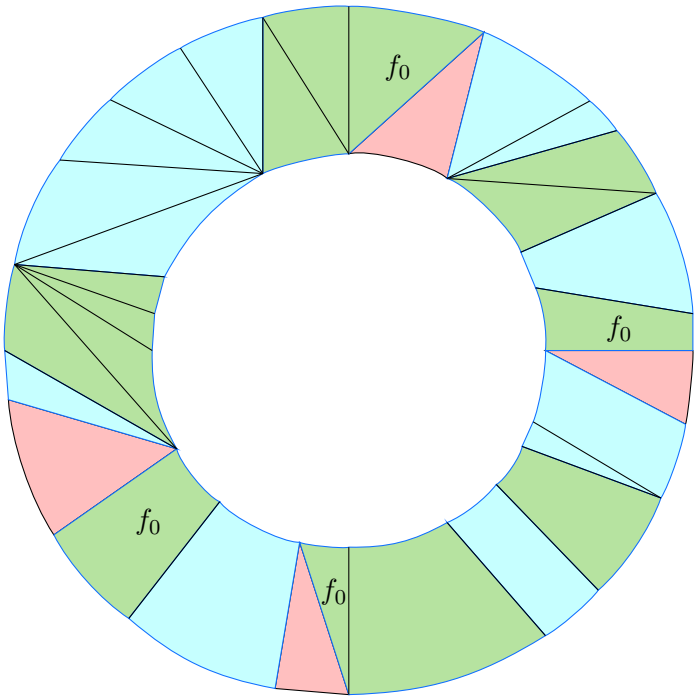


Figure 5.6: Illustration for the edge-sharing components and units in Lemma 5.3.13. The subgraphs in pink correspond to faces that are not in $\mathcal{F}$. The edge-sharing components in $\mathcal{B}(\mathcal{F})$ have a dark blue outline. The units of each edge-sharing component are in alternating blue and green. We also indicate the face f_0 selected for each edge-sharing component.

If there exists a vertex $v \in C$ (resp. $v \in C'$) such that units from different edge-sharing components intersect C (resp. C') only in v, then we merge these units.

Remark then that the units of an edge-sharing component are naturally ordered with respect to C and C' and that the faces from two units that are not (circularly) consecutive are vertex-disjoint.

Let $\mathcal{U}(\mathcal{F}) = \bigcup_{B \in \mathcal{B}(\mathcal{F})} \mathcal{U}(B)$. Let Γ_G be an auxiliary graph whose vertices are the units in $\mathcal{U}(\mathcal{F})$, and there is an edge between two units $U, U' \in \mathcal{U}(\mathcal{F})$ if they are not vertex-disjoint. By construction, Γ_G is a subgraph of a Hamiltonian cycle. Therefore, this graph is 3-colorable (because it is a subgraph of a cycle).

Therefore, by taking a color class of Γ_G which corresponds to a subset of $\mathcal{F}$ that has size at least $\frac{|\mathcal{F}|}{3}$ (this is possible because the color classes partition $\mathcal{F}$ into three sets), we get a subset of units that are vertex-disjoint and contain at least $\frac{|\mathcal{F}|}{3}$ of the faces in $\mathcal{F}$.

Moreover, remark that the faces in a unit $U \in \mathcal{U}(\mathcal{F})$ are naturally ordered with respect to C and C' and that by taking every other face in U, we get a subset of almost disjoint faces of U of size at least $\frac{|U|}{2}$.

By combining the two observations, we conclude that there exists a subset of $\mathcal{F}$ of almost disjoint faces of size at least $\frac{|\mathcal{F}|}{6}$. □

The following proposition follows the proof of [HK26b, Proposition 5.14]; however, the constant q differs significantly between the two results. This discrepancy arises from the broader notion of well-nested cycles adopted here (see Definition 5.3.1) compared to [HK26b].

Proposition 5.3.14 ([HK26a, Proposition 5.5])**.** *Let $q = \frac{9073}{9072}$ and $m' \in \mathbb{N}$. Let $C_1, ..., C_{2m'} = C$ be Π-contractible cycles of G that are Π-well-nested in this order. Let $e \in E(C_1 - C_2)$. Then,*

$$|\mathcal{I}_{\mathrm{N}}(C)| \geq q^{m'-2}$$

Proof. First, we define $\tilde{C}_{i+1}$ to be the cycle closest to C_{i+2} so that $\mathcal{B}(C_{i+2}) \subseteq \mathrm{Int}(C_{i+2} \cup \tilde{C}_{i+1}, \Pi)$. Then, by Lemma 5.3.12, $(C_{i+2}, \tilde{C}_{i+1}, C_i)$ is a good contractible square with respect to e and Π_e.

By Lemma 4.2.5, as $\mathrm{int}(C_2, \Pi)$ contains e, it is Π_e-noncontractible. Hence, $\mathcal{I}_{\mathrm{N}}(C_4)$ is not empty.

Let us first show that, for every $4 \leq i \leq 2m'$ there exists a subset of almost disjoint faces of $\mathcal{B}_{\mathrm{N}}(C_i)$ of size at least $\frac{|\mathcal{B}_{\mathrm{N}}(C_i)|}{12}$.

Claim 5.3.14.1. *For $4 \leq i \leq 2m'$, there exists a subset of almost disjoint faces of $\mathcal{B}_{\mathrm{N}}(C_i)$ of size at least $\frac{|\mathcal{B}_{\mathrm{N}}(C_i)|}{12}$.*

Proof of the claim. First, let us partition the faces in $\mathcal{B}(C_i)$ into sets $(\mathcal{F}_j)_{j \in \mathbb{N}^*}$. We define the sets $(\mathcal{F}_j)_{j \in \mathbb{N}^*}$ inductively as follows:

For $j = 1$, $\mathcal{F}_1$ is the set of faces from $\mathcal{B}(C_i)$ that intersect $\tilde{C}_{i-1}$. Moreover, we define C'_1 to be the cycle closest to C_i so that $\mathcal{B}(C_i) - \mathcal{F}_1 \subseteq \mathrm{Int}(C_i \cup C'_1, \Pi)$.

Suppose that for $j \geq 1$, $\mathcal{F}_j$ and C'_j have been defined. Then, we define $\mathcal{F}_{j+1}$ as the set of faces from $\mathcal{B}(C_i) - \cup_{k=1}^{j} \mathcal{F}_k$ that intersect C'_j. Moreover, we define C'_{j+1} to be the cycle closest to C_i so that $\mathcal{B}(C_i) - \bigcup_{k=1}^{j+1} \mathcal{F}_k \subseteq \mathrm{Int}(C_i \cup C'_{j+1}, \Pi)$.

Finally, $(\mathcal{F}_i)_{i \in \mathbb{N}^*}$ is defined. See Figure 5.7 for an illustration. As $\mathcal{B}(C_i)$ is finite, there exists $M \in \mathbb{N}^*$ such that $\mathcal{F}_j = \emptyset$ for all $j > M$.

Remark that by construction, for $j \geq 1$, $\mathcal{F}_j$ and $\mathcal{F}_{j+2}$ contain vertex-disjoint faces. Moreover, remark that, for $j \geq 1$, the faces in $\mathcal{F}_j \cap \mathcal{B}_{\mathrm{N}}(C_i)$ touch the two cycles C_i and C'_{j-1} (or $\tilde{C}_{i-1}$ if $j = 1$). Hence, by Lemma 5.3.13, there is a subset of faces from $\mathcal{F}_j \cap \mathcal{B}_{\mathrm{N}}(C_i)$ that are almost disjoint and this subset has size at least $\frac{|\mathcal{F}_j \cap \mathcal{B}_{\mathrm{N}}(C_i)|}{6}$.

To conclude, by taking the biggest subset in the partition of $\mathcal{B}_{\mathrm{N}}(C_i)$ into $\bigcup_{j \geq 1} \mathcal{F}_{2j} \cap \mathcal{B}_{\mathrm{N}}(C_i)$ and $\bigcup_{j \geq 0} \mathcal{F}_{2j+1} \cap \mathcal{B}_{\mathrm{N}}(C_i)$ and then applying Lemma 5.3.13, we get a subset of almost disjoint faces from $\mathcal{B}_{\mathrm{N}}(C_i)$ of size at least $\frac{|\mathcal{B}_{\mathrm{N}}(C_i)|}{12}$. ■

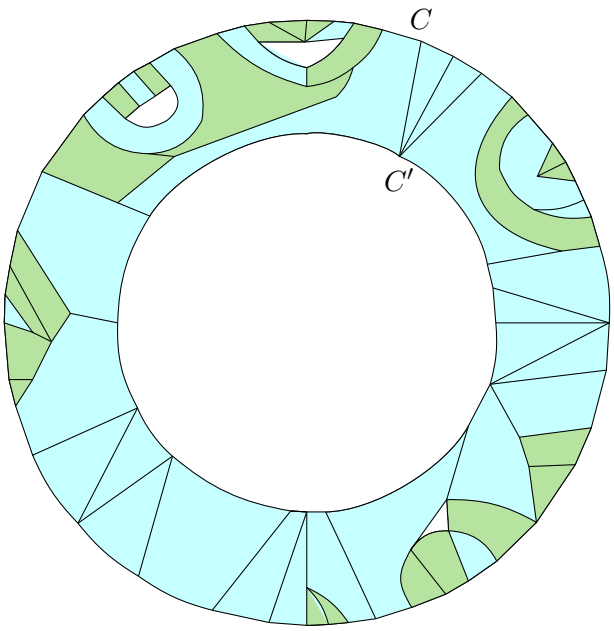


Figure 5.7: Illustration for the sets $(\mathcal{F}_i)_{i\in\mathbb{N}^*}$ in Claim 5.3.14.1. The subgraphs in white correspond to faces that are not in $\mathcal{B}(C)$ because they don't touch C. The sets $\bigcup_{j\geq 1}\mathcal{F}_{2j}$ and $\bigcup_{j\geq 0}\mathcal{F}_{2j+1}$ are depicted respectively in blue and green. Remark that, in this example, for $i > 5$, $\mathcal{F}_i$ is empty.

As $(C_{2i}, \tilde{C}_{2i-1}, C_{2i-2})$ is a good contractible square, for $2 \leq i \leq m' - 1$, $\mathcal{B}_{\mathrm{N}}(C_{2i})$ has size at least $\frac{|\mathcal{I}_{\mathrm{N}}(C_{2i})|+54}{756}$. Hence, by Claim 5.3.14.1, $\mathcal{I}_{\mathrm{N}}(C_{2i+2})$ has size at least

$$\begin{aligned}|\mathcal{I}_{\mathrm{N}}(C_{2i})| + \frac{|\mathcal{B}_{\mathrm{N}}(C_{2i})|}{12} &\geq |\mathcal{I}_{\mathrm{N}}(C_{2i})| + \frac{|\mathcal{I}_{\mathrm{N}}(C_{2i})| + 54}{9072} \\ &= \frac{9073|\mathcal{I}_{\mathrm{N}}(C_{2i})| + 54}{9072} \\ &\geq \frac{9073}{9072}|\mathcal{I}_{\mathrm{N}}(C_{2i})| = q|\mathcal{I}_{\mathrm{N}}(C_{2i})|\end{aligned}$$

By recurrence,

$$|\mathcal{I}_{\mathrm{N}}(C_{2m'})| \geq q^{m'-2}|\mathcal{I}_{\mathrm{N}}(C_4)| \geq q^{m'-2}$$

□

Corollary 5.3.15 ([HK26a, Corollary 5.6])**.** *Let $q = \frac{9073}{9072}$ and $m = 2(\lfloor \log_q(3g+4)\rfloor + 2)$. The graph G contains at most m cycles that are Π-well-nested.*

Proof. Let $m' \in \mathbb{N}$. Let $C_1, ..., C_{2m'}$ be Π-contractible cycles of G that are Π-well-nested in this order. Let $e \in E(C_1 - C_2)$.

By Proposition 5.3.14, $|\mathcal{I}_{\mathrm{N}}(C_{2m'})| \geq q^{m'-2}$.

If $m' - 2 \geq \log_q(3g+4)$, then $q^{m'-2} \geq 3g+4$. However, by Proposition 4.1.2, as $g(\Pi) \leq g+2$, $|\mathcal{I}_{\mathrm{N}}(C_{2m'})| \leq 3g+3$. Then, this implies that $m' \leq \lfloor \log_q(3g+4)\rfloor + 2$. Finally, there are at most $m = 2m' = 2(\lfloor \log_q(3g+4)\rfloor + 2)$ cycles that are Π-well-nested. □

Corollary 5.3.16 ([HK26a, Corollary 5.7])**.** *Let $q = \frac{9073}{9072}$. Let C be a Π-contractible cycle of G and suppose that ext(C, Π) has at most k attaches on C. Then, Int(C, Π) contains at most $2(\lfloor \log_q(60k + 180) \rfloor + 2)$ cycles that are Π-well-nested.*

Proof. Let $m' \in \mathbb{N}$. Let $C_1, ..., C_{2m'} = C$ be Π-contractible cycles of G that are Π-well-nested in this order. Let $e \in E(C_1 - C_2)$.

Let $\tilde{G} = \Pi_e(\text{Ext}(C, \Pi))$. First, let us show that $\tilde{G}$ has genus at most $g - \frac{|\mathcal{I}_N(C)|}{30} + 5$.

Claim 5.3.16.1. *The graph $\tilde{G}$ has genus at most $g - \frac{|\mathcal{I}_N(C)|}{30} + 7$.*

Proof of the claim. By Lemma 5.3.10, there are at most 10 Π_e-homotopic cycles in $\mathcal{I}_N(C)$. Therefore, as $\mathcal{I}_N(C) \subseteq \mathcal{I}(C)$, $\mathcal{I}(C)$ contains at least $\frac{|\mathcal{I}_N(C)|}{10}$ almost disjoint Π_e-noncontractible nonhomotopic cycles. Finally, by Proposition 4.1.2, $\Pi_e(\mathcal{I}(C))$ has genus at least $\frac{|\mathcal{I}_N(C)|}{30} - 1$.

Moreover, by Lemma 5.3.11, $g(\Pi_e(\tilde{G})) \leq g - g(\Pi_e(\mathcal{I}(C))) + 6$.

Therefore, $g(\Pi_e(\tilde{G})) \leq g - \frac{|\mathcal{I}_N(C)|}{30} + 7$. ■

By Claim 5.3.16.1, we know that $\Pi_e(\tilde{G})$ has genus at most $g - \frac{|\mathcal{I}_N(C)|}{30} + 7$. Moreover, as there are at most k attachments of ext(C, Π) on C, then by adding at most k handles or crosscaps, we create an empty cylinder with C on one of its boundaries. Then, int(C, Π) can be embedded (for example, as in Π) in this cylinder. Finally, as we add at most k handles and crosscaps, we manage to embed the graph G in a surface of genus at most $g - \frac{|\mathcal{I}_N(C)|}{30} + 7 + 2k$.

Because G has genus at least $g + 1$, we get that $g - \frac{|\mathcal{I}_N(C)|}{30} + 7 + 2k \geq g + 1$. Therefore, $\frac{|\mathcal{I}_N(C)|}{30} \leq 2k + 6$.

By Proposition 5.3.14, $|\mathcal{I}_N(C)| \geq q^{m'-2}$. Therefore,

$$\begin{aligned} & \frac{q^{m'-2}}{30} \leq 2k + 6 \\ \text{hence} \quad & q^{m'-2} \leq 60k + 180 \\ \text{and} \quad & m' \leq \lfloor \log_q(60k + 180) \rfloor + 2 \end{aligned}$$

Finally, there are at most $m = 2m' \leq 2(\lfloor \log_q(60k + 180) \rfloor + 2)$. □

5.3.2 Homotopic nested cycles

Definition 5.3.17 (Well-homotopic)**.** Let H be a graph with embedding Π_H. Let C and C' be two Π_H-noncontractible homotopic cycles of H. We say that C is Π_H*-well-homotopic* (or *well homotopic*, if it is clear in the context) in C' if one of the following condition holds:

- C and C' are disjoint (*(free,free)-well-homotopic* or *freely well homotopic*), or
- C and C' intersect in a vertex v (*(v,free)-well-homotopic* or *well homotopic pinched on the vertex v*), or
- C and C' both intersect the same face f of (G, Π) and if C intersect f in a path P (of length at least three) then C' intersect f in a subpath that is in the interior of P (*(f,free)-well-homotopic* or *well homotopic pinched on the face f*), or
- C and C' intersect at two disjoint vertices v and v' (*(v,v')-well-homotopic* or *well homotopic pinched on the vertices v and v'*), or

- C and C' both intersect the same two faces f and f' of (G, Π) and if C intersects f (resp. f') in a subpath P (resp. P'), of length at least 3, then C' intersect f (resp. f') in a subpath that is in the interior of P (resp. P') (*(f,f')-well-homotopic* or *well homotopic pinched on the faces f and f'*), or

- C and C' both intersect at a vertex v and a face f of (G, Π) and if C intersects f in a subpath P (of length at least 3) then C' intersect f in a subpath that is in the interior of P (*(v,f)-well-homotopic* or *well homotopic pinched on the vertex v and the face f*).

In the second and third cases, we say that C and C' are *well homotopic pinched on a piece p*. In the last three cases, we say that C and C' are *well homotopic pinched on two pieces p and p'*.

Let $k \geq 1$ and let $C_0, ..., C_k$ be Π_H-noncontractible homotopic cycles of H in this order in Π_H. We say that $C_0, ..., C_k$ are *Π_H-well-homotopic* if:

- for every $0 \leq i \leq k-1$, C_i is Π_H-freely-well-homotopic in C_{i+1}, or

- there is a piece p such that, for every $0 \leq i \leq k-1$, C_i is Π_H-well-homotopic in C_{i+1} pinched on p, or

- there are two pieces p and p' such that, for every $0 \leq i \leq k-1$, C_i is Π_H-well-homotopic in C_{i+1} pinched on p and p'.

The Figure 5.8 show cycles $C^1, C^{1'}, C^{1''}, C^{2''}, C^{2'}, C^2$ that are freely-well-homotopic in this order.

We extend the definition of closest cycles to non-contractible homotopic cycles:

Definition 5.3.18 (Closest)**.** Let H be a graph with embedding Π_H. Let C and C' be two Π_H-noncontractible cycles of H, Π_H-homotopic in this order. Let $\mathcal{P}$ be a property on the cycles of (H, Π_H).

Then, we say that C' is the cycle *closest* to C that has property $\mathcal{P}$ if, for every cycle C'' that is Π_H-homotopic to C in the order C, C'' and has property $\mathcal{P}$, we have $C' \subseteq \text{Int}(C \cup C'', \Pi_H)$.

Let C^1 and C^2 be two Π-well-homotopic cycles in G. Then, let $\mathcal{F}(C^1 \cup C^2, \Pi)$ be the set of all faces inside of $\text{Int}(C^1 \cup C^2, \Pi)$ in Π.

Definition 5.3.19. Let $C_1, ..., C_{2m}$ be Π-well-homotopic cycles of G. Suppose that $C_1, ..., C_{2m}$ are in this order in Π. For $4 \leq i \leq m$, we define $D_i = (C_{m-i+1}, C_{m+i})$. For $4 \leq i \leq m$, we define:

- $\mathcal{B}(D_i)$ as the union of the sets of faces in $\text{Int}((C_{2m'-i+1} \cup C_{2m'-(i-1)+1}, \Pi)$ that touch $C_{2m'-(i-1)+1}$ and of faces in $\text{Int}((C_{2m'+i} \cup C_{2m'+(i-1)}, \Pi)$ that touch $C_{2m'+i}$ (we call this set the *boundary* of D_i),

- $\mathcal{B}_{\text{Con}}(D_i, \Pi_e)$ as the set of the faces in $\mathcal{B}(D_i)$ that are Π_e-faces,

- $\mathcal{B}_{\text{N}}(D_i, \Pi_e)$ as the set of the faces in $\mathcal{B}(D_i)$ that are not Π_e-faces,

- $\mathcal{I}(D_i)$ as the set of Π-faces in $\text{Int}(C_{2m'-(i-2)+1} \cup C_{2m'+(i-2)}, \Pi)$, and

- $\mathcal{I}_{\text{N}}(D_i, \Pi_e)$ as a maximal size subset of Π-faces of $\mathcal{I}(D_i)$ that are almost-disjoint and that are not Π_e-faces.

Definition 5.3.20 (Homotopic square)**.** Let $C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''}$ be Π-noncontractible homotopic cycles in G, well homotopic in the following order: $C^1, C^{1'}, C^{1''}, C^{2''}, C^{2'}, C^2$. Let $C^{12} = C^1 \cup C^2$.

We say that $(C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''})$ is a *homotopic square* with respect to C^1 and C^2 if $C^{1'}$ and $C^{2'}$ are so that, for every other cycles $\tilde{C}^1$ and $\tilde{C}^2$ with $\mathcal{B}(C^{12}) \subseteq \text{Int}(C^1 \cup \tilde{C}^1, \Pi) \cup \text{Int}(C^2 \cup \tilde{C}^2, \Pi)$, we have $C^{1'} \subseteq \text{Int}(C^1 \cup \tilde{C}^1, \Pi)$ and $C^{2'} \subseteq \text{Int}(C^2 \cup \tilde{C}^2, \Pi)$.

We say that a homotopic square $(C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''})$ is respectively *full*, *pinched on* p and *pinched on* p *and* p' if $C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''}$ are respectively Π-freely-well-homotopic, Π-well-homotopic pinched on p, and Π-well-homotopic pinched on p and p' in the order $C^1, C^{1'}, C^{1''}$, $C^{2''}, C^{2'}, C^2$.

Figure 5.8 depicts the full homotopic square $(C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''})$.

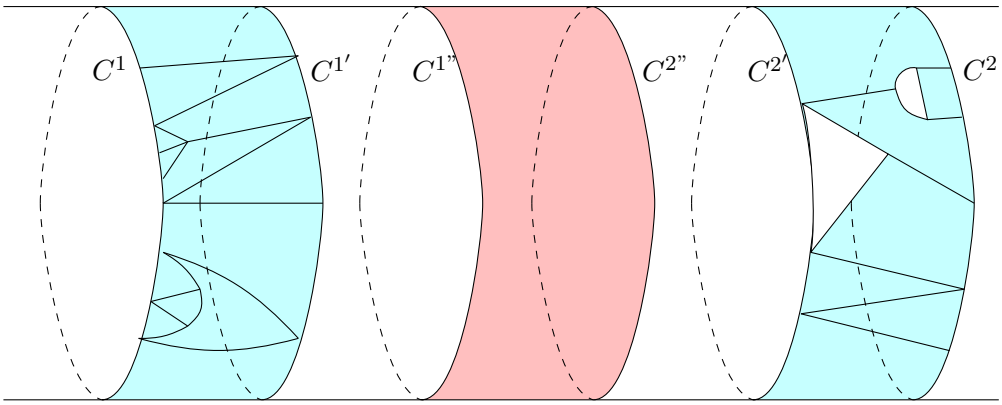


Figure 5.8: Homotopic square. The zone represented in blue is $\mathcal{B}(C^{12})$ (defined e.g. in Definition 5.3.20) and the zone represented in pink is $\mathcal{I}(C^{12})$ (defined in Definition 5.3.23) where $C^{12} = C^1 \cup C^2$.

Recall that, for $e \in E(G)$, we defined $G_e = G - e$ and Π_e as an embedding of G_e of genus $g(G_e) \leq g$. In this section, we ask that Π_e has the following property:

Definition 5.3.21 (Irreducible)**.** Let $C_1, ..., C_{2m}$ be Π-well-homotopic cycles of G. Suppose that $C_1, ..., C_{2m}$ are in this order in Π. For $4 \leq i \leq m$, we define $D_i = (C_{m-i+1}, C_{m+i})$. Let $e \in \text{Int}(C_m \cup C_{m+1}, \Pi)$.

For $2 \leq i \leq m-2$, we say that $\mathcal{I}_{\text{N}}(D_{i+2}, \Pi_e)$ is *irreducible* if there exists no other embedding Π'_e of G_e of genus g and $2 \leq i \leq m-2$ so that $\Pi_e(G - \mathcal{I}(D_{i+2}))$ and $\Pi'_e(G - \mathcal{I}(D_{i+2}))$ are identical embeddings and $\mathcal{I}_{\text{N}}(D_{i+2}, \Pi'_e)$ is included into $\mathcal{I}_{\text{N}}(D_{i+2}, \Pi_e)$.

We say that the embedding Π_e of G_e is *irreducible* with respect to $C_1, ..., C_{2m}$ if, for $2 \leq i \leq m-2$, $\mathcal{I}_{\text{N}}(D_{i+2}, \Pi_e)$ is *irreducible*.

We show that, for any $2m$ Π-well-homotopic cycles of G, there indeed exists an embedding Π_e that is irreducible with respect to them.

Proposition 5.3.22. *Let $C_1, ..., C_{2m}$ be Π-well-homotopic cycles of G. Suppose that $C_1, ..., C_{2m}$ are in this order in Π. For $4 \leq i \leq m$, we define $D_i = (C_{m-i+1}, C_{m+i})$. Let $e \in$ Int$(C_m \cup C_{m+1}, \Pi)$.*

There exists an embedding Π_e of G_e that is irreducible with respect to $C_1, ..., C_{2m}$.

Proof. Take Π_e to be an embedding of G_e of genus g so that $\mathcal{I}_{\mathrm{N}}(D_m, \Pi_e)$ is irreducible, which is always possible. We can easily see that $\mathcal{I}_{\mathrm{N}}(D_{i+2}, \Pi_e)$ is irreducible for $2 \leq i \leq m-2$. Let $2 \leq i \leq m-2$ and suppose that $\mathcal{I}_{\mathrm{N}}(D_{i+2}, \Pi_e)$ is not irreducible, then there exists an embedding Π'_e so that $\Pi_e(G - \mathcal{I}(D_{i+2}))$ and $\Pi'_e(G - \mathcal{I}(D_{i+2}))$ are identical embeddings and $\mathcal{I}_{\mathrm{N}}(D_{i+2}, \Pi'_e) \subsetneq \mathcal{I}_{\mathrm{N}}(D_{i+2}, \Pi_e)$. It implies that $\mathcal{I}_{\mathrm{N}}(D_m, \Pi'_e) \subsetneq \mathcal{I}_{\mathrm{N}}(D_m, \Pi_e)$ which contradicts the irreducibility of $\mathcal{I}_{\mathrm{N}}(D_m, \Pi_e)$. Therefore, $\mathcal{I}_{\mathrm{N}}(D_{i+2}, \Pi_e)$ is irreducible for every $2 \leq i \leq m-2$ and Π_e is irreducible with respect to $C_1, ..., C_{2m}$. □

Let $C_1, ..., C_{2m}$ be Π-well-homotopic cycles that are are in this order in Π. For $4 \leq i \leq m$, we define $D_i = (C_{m-i+1}, C_{m+i})$. Let $e \in \mathrm{Int}(C_m \cup C_{m+1}, \Pi)$. We fix Π_e to be an embedding of G_e of genus $g(G_e) \leq g$ irreducible with respect to $C_1, ..., C_{2m}$.

Now that Π_e is fixed, we omit its mention in $\mathcal{B}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ and $\mathcal{I}_{\mathrm{N}}(C_{i+2}, \Pi_e)$ and only write $\mathcal{B}_{\mathrm{N}}(C_{i+2})$ and $\mathcal{I}_{\mathrm{N}}(C_{i+2})$ for every $2 \leq i \leq 2m-2$.

Definition 5.3.23 (Good/bad homotopic square)**.** Let $e \in \mathrm{int}(C^{1''} \cup C^{2''}, \Pi)$. Let $(C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''})$ be a homotopic square with respect to C^1 and C^2. Let $C^{12} = C^1 \cup C^2$. We say that $(C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''})$ is a *bad square with respect to e and Π_e* if $\mathcal{I}_{\mathrm{N}}(C^{12})$ contains more than $18 \times (42|\mathcal{B}_{\mathrm{N}}(C^{12})| - 3)$ Π-faces.

Otherwise, we say that this square is *good with respect to e and Π_e*.

Lemma 5.3.24 ([HK26a, Lemma 5.8])**.** *Let $(C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''})$ be a homotopic square with respect to C^1 and C^2. Let $F_1, ..., F_m \in \mathcal{I}_{\mathrm{N}}(C^{12})$ be a maximal size set of almost disjoint Π-faces that are Π_e-noncontractible homotopic, then $m \leq 18$.*

Proof. Suppose that there are at least 19 almost disjoint Π-faces that are Π_e-noncontractible homotopic. Let $\mathcal{F}$ be the Π-faces in $\mathcal{I}_{\mathrm{N}}(C^{12})$. Contrary to Lemma 5.3.10, these faces may not be cycles (but either two cycles that intersect in a vertex or two disjoint cycles together with an edge that connects them). Let C_f be a cycle contained in the face $f \in \mathcal{F}$ (there is at least one) and let $\mathcal{C} = \{C_f, f \in \mathcal{F}\}$.

If the square is pinched on two vertices v and v', let us select cycles from $\mathcal{C}$ so that one of v, v' does not intersect any of the selected cycles. First, remark that there is at most one cycle in $\mathcal{C}$ that intersects both v and v': indeed, the faces are chosen to be almost-disjoint. Finally, there are at most $\frac{19-1}{2} = 9$ cycles in $\mathcal{C}$ so that one of v, v' does not intersect any of the selected cycles.

If the square is pinched on two vertices v and v', let $C_1, ..., C_9$ be the cycles from $\mathcal{C}$ so that one of v, v' does not intersect any of these cycles. If the square is not pinched on two vertices, let $C_1, ..., C_9$ be distinct cycles from $\mathcal{C}$ (randomly chosen). Without loss of generality, we have that $C_1, ..., C_9$ are Π_e-homotopic in this order.

There are at least two of the cycles C_1, C_5, C_9 that have the same relative orientation in Π and Π_e, suppose without loss of generality that it is C_1 and C_5. Let us show that it is possible to modify Π_e so that one of F_2, F_3, or F_4 becomes a Π_e-face.

First, note that:

- If the square is freely well homotopic: no vertex of C^1 (resp. C^2) is in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$, as otherwise, the whole subgraph $\mathrm{Ext}(C^{12}, \Pi)$ would be in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$ which is impossible because $\mathrm{Ext}(C^{12}, \Pi)$ cannot be Π_e-contractible. Therefore, $\mathrm{Int}(C_1 \cup C_5, \Pi_e) \subseteq \mathrm{Int}(C^{12}, \Pi)$.
- If the square is pinched on a piece p: the same reasoning for $C^1 - p$ (resp. $C^2 - p$) shows that no vertex of $C^1 - p$ (resp. $C^2 - p$) is in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$. If p is a face, it implies that

$\mathrm{Int}(C_1 \cup C_5, \Pi_e) \subseteq \mathrm{Int}(C^{12}, \Pi)$ as there is no edge in $\mathrm{Ext}(C^{12}, \Pi)$ with an endpoint on p. If p is a vertex, then a bridge from $\mathrm{ext}(C^{12}, \Pi)$ in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$, if it exists, would be attached solely to p, which contradicts the fact that G is 2-connected and therefore has no cutvertex.

- If the square is pinched on two pieces p and p': the same reasoning for $C^1 - (p \cup p')$ (resp. $C^2 - (p \cup p')$) shows that no vertex of $C^1 - (p \cup p')$ (resp. $C^2 - (p \cup p')$) is in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$. If both p and p' are faces, it implies that $\mathrm{Int}(C_1 \cup C_5, \Pi_e) \subseteq \mathrm{Int}(C^{12}, \Pi)$ as there is no edge in $\mathrm{Ext}(C^{12}, \Pi)$ with an endpoint on p or p'. If p is a vertex and p' is a face, then a bridge from $\mathrm{ext}(C^{12}, \Pi)$ in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$, if it exists, would be attached solely to p, which contradicts the fact that G is 2-connected and therefore has no cutvertex. If both p and p' are vertices, then as we showed before that $C_1 \cup C_5$ does not intersect both p and p', we can conclude with the same reasoning as when the square is pinched on one vertex.

Therefore, $\mathrm{Int}(C_1 \cup C_5, \Pi_e) \subseteq \mathrm{Int}(C^{12}, \Pi)$ and $\Pi(\mathrm{Int}(C_1 \cup C_5, \Pi_e))$ is an embedding in a disk in Π in which C_1 and C_5 are Π-facial walks.

Let us show that one face F of F_2, F_3, F_4 is contained in $\Pi(\mathrm{Int}(C_1 \cup C_5, \Pi_e))$ and therefore F is a Π-facial walk in this embedding. First, if one face F of F_2, F_3, F_4 is a cycle, then this face F is contained in $\Pi(\mathrm{Int}(C_1 \cup C_5, \Pi_e))$. Now, suppose that all of F_2, F_3, F_4 are not cycles. Let us show that one of these faces is in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$. Suppose, by contradiction, that none of them is entirely in $\mathrm{Int}(C_1 \cup C_5, \Pi_e)$. Then, without loss of generality, two of them (say F_2 and F_3) intersect C_1, and because the faces were chosen almost disjoint, they intersect C_1 in the same vertex v. Remark that v is either the vertex at which the two cycles in F_2 (resp. F_3) intersect or an endpoint of the edge in F_2 (resp. F_3) that is not part of its two cycles. However, note that v is then a cutvertex. Indeed, then there exists C'_2 and C'_3 cycles respectively in F_2 and F_3 such that v separates $\mathrm{Int}(C'_2 \cup C'_3, \Pi)$ from the rest of the graph. This contradicts the fact that G is 2-connected and therefore has no cutvertex.

Therefore, we indeed find that one of F_2, F_3, F_4 is contained in $\Pi(\mathrm{Int}(C_1 \cup C_5, \Pi_e))$ and therefore is a Π-facial walk in this embedding.

Finally, by Lemma 4.1.7, we can modify the embedding Π_e so that one of F_2, F_3, F_4 is now a Π_e-face.

This concludes the proof. □

Lemma 5.3.25 ([HK26a, Lemma 5.9])**.** *Let $e \in int(C^{1''} \cup C^{2''}, \Pi)$, G contains no bad homotopic square with respect to e and Π_e.*

Proof. Suppose by contradiction that G contains a bad homotopic square $(C^1, C^2, C^{1'}, C^{2'}, C^{1''}, C^{2''})$. Let $C^{12} = C^1 \cup C^2$. Therefore $|\mathcal{I}_{\mathrm{N}}(C^{12})| \geq 18 \times (42|\mathcal{B}_{\mathrm{N}}(C^{12})| - 3)$. First, let us show that there exists no bad homotopic square with respect to e and Π_e with $|\mathcal{B}_{\mathrm{N}}(C^{12})| = 0$.

Claim 5.3.25.1. $|\mathcal{B}_{\mathrm{N}}(C^{12})| > 0$

Proof of the claim. Suppose that $|\mathcal{B}_{\mathrm{N}}(C^{12})| = 0$.

Then, remark that $\mathcal{B}(C^{12})$ contains at most 4 edge-sharing components and remark that the embedding of each of these is the same in Π and Π_e.

Let $\tilde{G}$ be the graph obtained from G_e by removing $H = \mathrm{int}(C^{1'} \cup C^{2'}, \Pi)$. $\tilde{G}$ has genus at most g (because it is a subgraph of G_e).

- If the square is free or pinched on a piece: Let B_1 and B_2 be the two edge-sharing components of $\mathcal{B}(C^{12})$: $B_1 = \mathrm{Int}(C^1 \cup C^{1'}, \Pi)$ and $B_2 = \mathrm{Int}(C^2 \cup C^{2'}, \Pi)$. If the square is free,

let us define $r = \emptyset$. If the square is pinched on a vertex v, let $r = v$ and, if the square is pinched on a face f, let $r = (C^{1'} \Delta C^{2'}) \cap f$.

We define D_1 and D_2 to be respectively $(B_1 \cup r) \cap C^{1'}$ and $(B_2 - r) \cap C^{2'}$. If there is no path in $\mathrm{Int}(C^{1'} \cup C^{2'}, \Pi)$ from D_1 to D_2, then it is enough to find two empty disks in $\Pi_e(\tilde{G})$ with respectively on their boundaries the vertices of D_1 and D_2 in this order, which we can easily do. Otherwise, let P be a path from D_1 to D_2.

If P connects B_1 and B_2 with the expected orientation, then we can find an empty disk in $\Pi_e(\tilde{G})$ with the vertices of $C^{1'} \cup C^{2'}$ on its boundary in the right order. Otherwise, suppose that no path P connects B_1 and B_2 with the expected orientation. In particular, if the square is pinched on a face f, the path P_f between B_1 and B_2 that is a subwalk of f does not connect them with the expected orientation. We deduce that P and P_f have the same signature.

Now, let us look at the two faces f_1 and f_2 on the path P in $\Pi_e(\tilde{G} \cup P)$. We can see that $f_1 = f_2$ because it is the case in $\Pi_e(\tilde{G})$. Therefore, changing the signature of P does not increase the genus of the embedding; thus, by this manipulation, we return to the previous case.

- If the square is pinched on two pieces: Let B_1, B_1' and B_2, B_2' be the four edge-sharing components of $\mathcal{B}(C^{12})$ such that $B_1 \cup B_1' = \mathrm{Int}(C^1 \cup C^{1'}, \Pi)$ and $B_2 \cup B_2' = \mathrm{Int}(C^2 \cup C^{2'}, \Pi)$. If the square is pinched on two vertices v and v', let us define $r = v$ and $r' = v'$. If the square is pinched on a vertex v and a face f, let $r = v$ and $r' = (C^{1'} \Delta C^{2'}) \cap f$. If the square is pinched on two faces f and f', let $r = (C^{1'} \Delta C^{2'}) \cap f$ and $r' = (C^{1'} \Delta C^{2'}) \cap f'$.

  We define D_1, D_1' and D_2, D_2' to be respectively the connected component of B_1 in $(B_1 \cup r \cup r') \cap C^{1'}$, the connected component of B_1' in $(B_1' \cup r \cup r') \cap C^{1'}$ and $(B_2 - (r \cup r')) \cap C^{2'}$, $(B_2' - (r \cup r')) \cap C^{2'}$.

  Now, remark that bridges from $\mathrm{Int}(C^{1'} \cup C^{2'}, \Pi)$ are attached either to $D_1 \cup D_2$ or to $D_1' \cup D_2'$. Therefore, we can apply the same strategy as above for D_1 and D_2 (resp. D_1' and D_2').

In each case, we manage to find an embedding of G in the surface S' of genus g.
This is a contradiction. ■

We can now suppose that $|\mathcal{B}_{\mathrm{N}}(C^{12})| > 0$.

Remove $H = \mathrm{int}(C^{1'} \cup C^{2'}, \Pi)$ from G_e to obtain the graph $\tilde{G}$.

Claim 5.3.25.2. *The embedding of $\tilde{G}$ induced by Π_e has genus at most $g - 8|\mathcal{B}_{\mathrm{N}}(C)|$.*

Proof of the claim. Remark that $\mathcal{I}_{\mathrm{N}}(C) \subseteq \mathcal{I}(C)$ and therefore $\Pi_e(\mathcal{I}(C))$ contains $18 \times (42|\mathcal{B}_{\mathrm{N}}(C)| - 3)$ almost disjoint Π_e-noncontractible cycles. Moreover, by Lemma 5.3.24, there are at least $42|\mathcal{B}_{\mathrm{N}}(C)| - 3$ of them that are pairwise Π_e-nonhomotopic. Finally, by Proposition 4.1.2, $\Pi_e(\mathcal{I}(C))$ has genus at least $14|\mathcal{B}_{\mathrm{N}}(C)|$.

By Lemma 5.3.11, we have $g(\Pi_e(\tilde{G})) \leq g - g(\Pi_e(\mathcal{I}(C)) + 6$.

As $\Pi_e(\mathcal{I}(C^{12}))$ has genus at least $14|\mathcal{B}_{\mathrm{N}}(C^{12})|$, we finally conclude that $\Pi_e(\tilde{G})$ has genus at most $g - 14|\mathcal{B}_{\mathrm{N}}(C^{12})| + 6 \leq g - 8|\mathcal{B}_{\mathrm{N}}(C^{12})|$. ■

Now, we will construct a new embedding of G from the embedding $\Pi_e(\tilde{G})$.
Let $C^{12'} = C^{1'} \cup C^{2'}$ and let $\tilde{\mathcal{B}}_{\mathrm{N}}(C^{12})$ be the subset of faces of $\mathcal{B}_{\mathrm{N}}(C^{12})$ that touch $C^{12'}$.

Remark that there are at most 4 edge-sharing components of $\mathcal{B}(C^{12})$. Let $\mathcal{D}$ be the set of these edge-sharing components.

Let $D \in \mathcal{D}$. The edge-sharing components of $\mathcal{B}_{\text{Con}}(C^{12}) \cap D$ that touch $C^{12'}$ and the faces in $\tilde{\mathcal{B}}_{\text{N}}(C^{12}) \cap D$ can be ordered with respect to $C^{12'}$. Then, we add a handle or crosscap between each consecutive edge-sharing component of $\mathcal{B}_{\text{Con}}(C^{12}) \cap D$/face of $\tilde{\mathcal{B}}_{\text{N}}(C^{12}) \cap D$ x and y so that, for an edge-sharing component, the part of its boundary that contains endpoints of edges from H is reachable from the handle and has the expected orientation. Moreover, we add at most two handles or crosscaps to make sure that the edge-sharing components from $\mathcal{D}$ are connected with the expected orientation.

By Claim 5.3.25.2, the embedding $\Pi_e(\tilde{G})$ has genus at most $g - 8|\mathcal{B}_{\text{N}}(C)|$. Remark that we add at most $2|\mathcal{B}_{\text{N}}(C)| + 2 \leq 4|\mathcal{B}_{\text{N}}(C)|$ handles and crosscaps, and therefore creates a new embedding of genus at most g. Moreover, this new embedding contains one or two empty cylinders with $C^{12'}$ on their boundaries. Then, H can be embedded (for example, as in Π) in this cylinder or these cylinders. The graph G can then be embedded in a surface of genus at most g, a contradiction. □

Lemma 5.3.26 ([HK26a, Lemma 5.10])**.** *Let C_1, C_1', C_2', C_2 be Π-well-homotopic cycles of G in this order. Let $\mathcal{F}_1$ (resp. $\mathcal{F}_2$) be a subset of the faces in Int$(C_1 \cup C_1', \Pi)$ (resp. Int$(C_2 \cup C_2', \Pi)$) that intersect both C_1 and C_1' (resp. C_2 and C_2').*

There exists a subset of faces from $\mathcal{F}_1 \cup \mathcal{F}_2$ of size at least $\frac{|\mathcal{F}_1|+|\mathcal{F}_2|}{6}$ that are almost disjoint.

Proof.

- If $C_1' \cap C_2' = \emptyset$, or C_1' and C_2' intersect in a face or in two faces: Remark that $\mathcal{F}_1$ and $\mathcal{F}_2$ are disjoint, then, using Lemma 5.3.13 for both sets, we find that there exists a subset of faces from $\mathcal{F}_1 \cup \mathcal{F}_2$ of size at least $\frac{|\mathcal{F}_1|+|\mathcal{F}_2|}{6}$ that are almost disjoint.

- If C_1' and C_2' intersect in a vertex v or in a vertex v and a face f: Then faces from $\mathcal{F}_1$ and $\mathcal{F}_2$ can only intersect in v. Let $\mathcal{U}_1, \Gamma_1$ and $\mathcal{U}_2, \Gamma_2$ be the units and the auxiliary graphs obtained respectively from $\mathcal{F}_1$ and $\mathcal{F}_2$ by following the process in Lemma 5.3.13. By Lemma 5.3.13, Γ_1 and Γ_2 are cycles and therefore 3-colorable. Let Γ be the auxiliary graph obtained from $\mathcal{F}_1 \cup \mathcal{F}_2$ following the process in Lemma 5.3.13, then, as $\mathcal{U}_1$ and $\mathcal{U}_2$ intersect only in v, it is easy to see that Γ is 3-colorable. Finally, we conclude, by the same reasoning as Lemma 5.3.13, that there exists a subset of faces from $\mathcal{F}_1 \cup \mathcal{F}_2$ of size at least $\frac{|\mathcal{F}_1 \cup \mathcal{F}_2|}{6} = \frac{|\mathcal{F}_1|+|\mathcal{F}_2|}{6}$ that are almost disjoint.

- If C_1' and C_2' intersect in two vertices v and v': Then faces from $\mathcal{F}_1$ and $\mathcal{F}_2$ can only intersect in v and v'. Let $\mathcal{U}_1, \Gamma_1$ and $\mathcal{U}_2, \Gamma_2$ be the units and the auxiliary graphs obtained respectively from $\mathcal{F}_1$ and $\mathcal{F}_2$ by following the process in Lemma 5.3.13. By Lemma 5.3.13, Γ_1 and Γ_2 are cycles and therefore 3-colorable. Let Γ be the auxiliary graph obtained from $\mathcal{F}_1 \cup \mathcal{F}_2$ following the process in Lemma 5.3.13, then, as $\mathcal{U}_1$ and $\mathcal{U}_2$ intersect only in v and v', and it is easy to see that Γ is 3-colorable. Finally, we conclude, by the same reasoning as Lemma 5.3.13, that there exists a subset of faces from $\mathcal{F}_1 \cup \mathcal{F}_2$ of size at least $\frac{|\mathcal{F}_1 \cup \mathcal{F}_2|}{6} = \frac{|\mathcal{F}_1|+|\mathcal{F}_2|}{6}$ that are almost disjoint. □

Proposition 5.3.27 ([HK26a, Proposition 5.11])**.** *Let $q = \frac{9073}{9072}$ and $m = 2(\lfloor \log_q(3g+4) \rfloor + 2)$. G contains at most $2m$ Π-well-homotopic cycles.*

Proof. Let $m' \in \mathbb{N}$. Let $C_1, ..., C_{4m'}$ be Π-well-homotopic cycles of G. Suppose that $C_1, ..., C_{4m'}$ are in this order in Π.

Let $e \in E(C_{2m'} - C_{2m'-1})$. By Lemma 5.3.25, each pair of cycles $D_i = (C_{2m'-i+1}, C_{2m'+i})$ $(4 \leq i \leq 4m')$ induces a good homotopic square with respect to e and Π_e. We define $\tilde{C}_{2m'-(i+1)+1}$ and $\tilde{C}_{2m'+(i+1)}$ to be the cycles closest respectively to $C_{2m'-(i+2)+1}$ and $C_{2m'+(i+2)}$ so that $\mathcal{B}(D_{i+2}) \subseteq \text{Int}(C_{2m'-(i+2)+1} \cup C_{2m'-(i+1)+1}, \Pi) \cup \text{Int}(C_{2m'+(i+2)} \cup C_{2m'+(i+1)}, \Pi)$. By Lemma 5.3.25, for $2 \leq i \leq 2m' - 2$,

$$(C_{2m'-(i+2)+1}, \tilde{C}_{2m'-(i+1)+1}, C_{2m'-i+1}, C_{2m'+i}, \tilde{C}_{2m'+(i+1)}, C_{2m'+(i+2)})$$

is a good homotopic square with respect to e and Π_e.

By Lemma 4.2.5, as $\text{int}(D_2, \Pi)$ contains e, it is Π_e-noncontractible. Hence, $\mathcal{I}_{\text{N}}(D_4)$ is not empty.

Let us first show that, for every $4 \leq i \leq 2m'$ there exists a subset of almost disjoint faces of $\mathcal{B}_{\text{N}}(D_i)$ of size at least $\frac{|\mathcal{B}_{\text{N}}(D_i)|}{12}$.

Claim 5.3.27.1. *For $4 \leq i \leq 2m'$, there exists a subset of almost disjoint faces of $\mathcal{B}_{\text{N}}(D_i)$ of size at least $\frac{|\mathcal{B}_{\text{N}}(D_i)|}{12}$.*

Proof of the claim. First, let us partition the faces in $\mathcal{B}(D_i)$ into sets $(\mathcal{F}_j)_{j \in \mathbb{N}^*}$. We define the sets $(\mathcal{F}_j)_{j \in \mathbb{N}^*}$ inductively as follows:

For $j = 1$, $\mathcal{F}_1$ is the set of faces from $\mathcal{B}(D_i)$ that intersect $\tilde{C}_{2m'-(i-1)+1} \cup \tilde{C}_{2m'+(i-1)}$. Moreover, we define C'_1 and C''_1 to be the cycles closest to respectively $C_{2m'-i+1}$ and $C_{2m'+i}$ so that $\mathcal{B}(D_i) - \mathcal{F}_1 \subseteq \text{Int}(C_{2m'-i+1} \cup C'_1, \Pi) \cup \text{Int}(C_{2m'+i} \cup C''_1, \Pi)$.

Suppose that for $j \geq 1$, $\mathcal{F}_j$, C'_j and C''_j have been defined. Then, we define $\mathcal{F}_{j+1}$ as the set of faces from $\mathcal{B}(D_i) - \bigcup_{k=1}^{j} \mathcal{F}_k$ that intersect $C'_j \cup C''_j$. Moreover, we define C'_{j+1} and C''_{j+1} to be the cycles closest to respectively $C_{2m'-i+1}$ and $C_{2m'+i}$ so that $\mathcal{B}(D_i) - \bigcup_{k=1}^{j+1} \mathcal{F}_k \subseteq \text{Int}(C_{2m'-i+1} \cup C'_{j+1}, \Pi) \cup \text{Int}(C_{2m'+i} \cup C''_{j+1}, \Pi)$.

Finally, $(\mathcal{F}_i)_{i \in \mathbb{N}^*}$ is defined.

Remark that by construction, for $j \geq 1$, $\mathcal{F}_j$ and $\mathcal{F}_{j+2}$ contain vertex-disjoint faces. Moreover, remark that, for $j \geq 1$, the faces in $\mathcal{F}_j \cap \mathcal{B}_{\text{N}}(D_i)$ touch either the two cycles $C_{2m'-i+1}$ and C'_{j-1} (or $\tilde{C}_{2m'-(i-1)+1}$ if $j = 1$) or the two cycles $C_{2m'+i}$ and C''_{j-1} (or $\tilde{C}_{2m'+(i-1)}$ if $j = 1$). Hence, by Lemma 5.3.26, there is a subset of faces from $\mathcal{F}_j \cap \mathcal{B}_{\text{N}}(D_i)$ that are almost disjoint and this subset has size at least $\frac{|\mathcal{F}_j \cap \mathcal{B}_{\text{N}}(D_i)|}{6}$.

To conclude, by taking the biggest subset in the partition of $\mathcal{B}_{\text{N}}(D_i)$ into $\bigcup_{j \geq 1} \mathcal{F}_{2j} \cap \mathcal{B}_{\text{N}}(D_i)$ and $\bigcup_{j \geq 0} \mathcal{F}_{2j+1} \cap \mathcal{B}_{\text{N}}(D_i)$ and then applying Lemma 5.3.26, we get a subset of faces from $\mathcal{B}_{\text{N}}(D_i)$ of size at least $\frac{|\mathcal{B}_{\text{N}}(D_i)|}{12}$. ■

As the homotopic square with outer cycles D_{2i} is good with respect to e and Π_e for $2 \leq i \leq m' - 1$, $\mathcal{B}_{\text{N}}(D_{2i})$ has size at least $\frac{|\mathcal{I}_{\text{N}}(D_{2i})|+54}{756}$. Hence, by Claim 5.3.27.1, $\mathcal{I}_{\text{N}}(D_{2i+2})$ has size at least

$$\begin{aligned} |\mathcal{I}_{\text{N}}(D_{2i})| + \frac{|\mathcal{B}_{\text{N}}(D_{2i})|}{12} &\geq |\mathcal{I}_{\text{N}}(D_{2i})| + \frac{|\mathcal{I}_{\text{N}}(D_{2i})| + 54}{9072} \\ &= \frac{9073}{9072}|\mathcal{I}_{\text{N}}(D_{2i})| = q|\mathcal{I}_{\text{N}}(D_{2i})| \end{aligned}$$

By recurrence,

$$|\mathcal{I}_{\mathrm{N}}(D_{2m'})| \geq q^{m'-2}|\mathcal{I}_{\mathrm{N}}(D_4)| \geq q^{m'-2}$$

If $m' - 2 \geq \log_q(3g+4)$, $q^{m'-2} \geq 3g+4$. However, by Proposition 4.1.2, as $g(\Pi) \leq g+2$, $|\mathcal{I}_{\mathrm{N}}(D_{2m'})| \leq 3g+3$. Then, this implies that $m' \leq \lfloor \log_q(3g+4) \rfloor + 2$. Finally, there are at most $2m = 4m' = 2 \times (2\lfloor(\log_q(3g+4)\rfloor + 2))$ disjoint cycles that are Π-well-homotopic. □

5.4 Consequences of the structural results

5.4.1 Treewidth and tree decomposition

Proposition 5.3.14 makes it possible to find a better bound on the treewidth of G than the previously best known one from Seymour [Sey93], stated in Theorem 2.2.8, by combining it with the following result of Kawarabayashi and Sidiropoulos [KS19]:

Proposition 5.4.1 ([KS19])**.** *Let H be a graph embedded with embedding Π_H in a surface of genus g_H, and $tw(H) \geq 264 g_H \ell$. Then H contains ℓ Π_H-well-nested cycles.*

Theorem 3.1.2 ([HK26b, Theorem 1.4])**.** *The treewidth of G is bounded by the following function of g:*

$$tw(G) \leq T(g)$$

with $T(g) = 264(g+2)(m+1) - 1 = O(g \log g)$, where $m = 2(\lfloor \log_q(3g+4) \rfloor + 2)$ and $q = \frac{1153}{1152}$.

Proof. This result is obtained by combining Propositions 5.3.14 and 5.4.1. □

5.4.2 Excluded grid topological minor

Moreover, we can remark that the following results by Thomassen [Tho97] can be improved by adapting the original proof and then by using Proposition 5.3.14, in the case that G is 2-connected.

Definition 5.4.2 (Hexagonal grid)**.** We define J_k and C_{J_k} by induction on $k \geq 1$. Let J_1 be the cycle of length 6 C_6, and let C_{J_1} be this cycle. For each $k \geq 2$, let J_k be the union of J_{k-1} and all those 6-cycles in the hexagonal grid tiling which intersect J_{k-1} and let C_{J_k} be the cycle that bounds J_k in the hexagonal tiling.

Figure 5.9 represents the hexagonal grid J_3.

Proposition 5.4.3 ([Tho97], [MT01, Theorem 7.3.3])**.** *Let H be a 2-edge-connected graph such that $g(H) = g$ and $g(H_e) < g$ for every edge e in H. Then H contains no subdivision of the grid J_k, $k = \lceil 800 g^{3/2} \rceil$.*

Corollary 2.2.9. *Let H be a 2-edge-connected graph such that $g(H) = g$ and $g(H_e) < g$ for every edge e in H. Then H contains no subdivision of the $4k \times 2k$ grid, $k = \lceil 800 g^{3/2} \rceil$.*

Proof. The result follows from the easily verifiable fact that the $4k \times 2k$ grid contains J_k for all $k \geq 1$. □

Indeed, the proof of Proposition 5.4.3 relies on the two following lemmas:

Let H be a subgraph of G such that H is a subdivision of a hexagonal grid. Let C be the outer cycle of H. We say that H is *good* in G if the union of H and all the H-bridges which have an attach in $H - C$ is planar.

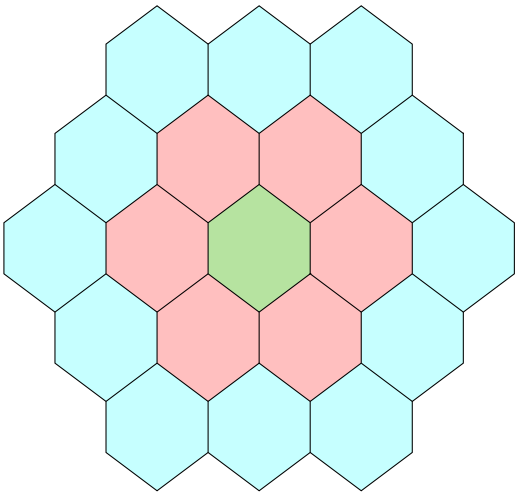


Figure 5.9: The hexagonal grid J_3. The colors show the construction of the grid by induction: J_1 is represented in green; then, by adding to J_1 the pink part, we obtain J_2; finally, by adding to J_2 the blue part, we get J_3.

Lemma 5.4.4 ([Tho97], [MT01, Proposition 7.3.1])**.** *Let H' be a subgraph of G such that H' is a subdivision of the hexagonal grid J_n. Let $k \geq 1$, if $n > 100k\sqrt{g}$, then H' contains a subdivision H of the hexagonal grid J_k such that H is good in G.*

Lemma 5.4.5 ([Tho97], [MT01, Proposition 7.3.2])**.** *Let H be a subgraph of G such that H is a subdivision of the hexagonal grid J_k and such that H is good in G. If $k \geq 4g + 14$ then the Π-genus of the subdivision of $J_{k-4g-12}$ in J_k is zero. In other words, Π induces a planar embedding of $J_{k-4g-12}$.*

By using Lemma 5.4.4 and Proposition 5.3.14 (instead of Lemma 5.4.5), we get a significant improvement on Proposition 5.4.3:

Theorem 3.1.3 ([HK26b, Theorem 1.6])**.** *G contains no subdivision of the grid J_k and no subdivision of the $4k \times 2k$ grid, $k = 100\sqrt{g} \times m = O(\sqrt{g} \log g)$.*

5.4.3 Bound on the number of almost disjoint Π-noncontractible cycles

We then prove the following simple consequence of Proposition 5.3.27:

Corollary 5.4.6 ([HK26a, Corollary 5.12])**.** *Let $\mathcal{C}$ be a set of almost disjoint Π-noncontractible cycles of G. Then, $\mathcal{C}$ contains at most $4m^2 \times (3g + 3)$ cycles.*

Proof. By Proposition 4.1.2, as $g(\Pi) \leq g + 2$, there are at most $3g + 3$ almost disjoint Π-noncontractible cycles that are not pairwise Π-homotopic in G.

Let us now show that $\mathcal{C}$ contains at most $4m^2$ Π-homotopic cycles. If there are more than $4m^2$ of them, then either more than $2m$ of them are disjoint or more than $2m$ of them all intersect in the same vertex. In both cases, $\mathcal{C}$ induces Π-well-homotopic cycles, either free or pinched at one vertex, which contradicts Proposition 5.3.27. We conclude that there are at most $4m^2$ Π-homotopic cycles in $\mathcal{C}$.

Finally, we get that $\mathcal{C}$ contains at most $4m^2 \times (3g + 3)$ cycles. □

This proof can easily be adapted for cycles that pairwise share at most 2 pieces. Here, we prove it for almost disjoint cycles because, in that case, the proof is short and because it is the version we need to prove the bound on the excluded minors for surfaces.

5.5 Links with facewidth

In this section, we define the notion of *facewidth* and explain the links between our structural results and the facewidth of (G, Π). Moreover, we present results that are quite similar to our structural results and whose proofs also bear resemblance with our proofs.

First, let us define the *facewidth* of an embedding. We will use the definition from [MT01].

Let Π_H be an embedding of a graph H on a surface S_H. We define the *facewidth* $fw(H, \Pi_H)$ of (H, Π_H) as the smallest number k such that there exist Π_H-facial walks $W_1, ..., W_k$ whose union contains a Π_H-noncontractible cycle. Equivalently, the facewith $fw(H, \Pi_H)$ is the biggest number k such that every noncontractible simple closed curve on S_H intersects the graph H in at least k points.

If Π_H is an embedding of genus 0, then we put $fw(H, \Pi_H) = \infty$. If no confusion is possible, we also write for short $fw(\Pi_H)$. The facewidth can also be called *representativity*, even though it is a more ancient terminology.

The two following results from Seymour and Thomas and from Mohar were published roughly at the same time (in 1995 and 1996) and prove that embeddings with large facewidth (at least logarithmic in the genus of the embedding) are unique.

Proposition 5.5.1 ([ST96])**.** *Let H be a 3-connected graph and let Π_H be an embedding of H of facewidth k.*

1. *If $k \geq 100 \log g(\Pi_H)/\log\log g(\Pi_H)$ (or $k \geq 100$ if $g(\Pi_H) \leq 2$), then any embedding Π'_H of H satisfies $g(\Pi'_H) \geq g(\Pi_H)$, and if $g(\Pi'_H) = g(\Pi_H)$, then Π_H and Π'_H are equivalent (in particular, they are embeddings in the same surface).*

2. *Let Π'_H be an embedding of H. If $fw(G, \Pi'_H) \geq 3$ and $k \geq \min\{400, \log g(\Pi_H)\}$, then either Π_H and Π'_H are equivalent, or $g(\Pi'_H) \geq g(\Pi_H) + 10^{-6}k^2$.*

Proposition 5.5.2 ([Moh95])**.** *Let H be a 2-connected graph and let Π_H and Π'_H be embeddings of H of facewidth k and k', respectively. Let $\kappa = \min\{k, k'\}$. Suppose that $k \geq 4$, $k' \geq 7$ and*

$$\left\lfloor \frac{k'-3}{2} \right\rfloor^{\lfloor \kappa/2 \rfloor - 2} > g(\Pi'_H)$$

Then, Π_H and Π'_H are embeddings in the same surface, and Π'_H is Whitney equivalent to Π_H.

The equation presented above in Proposition 5.5.2 implies that embeddings Π'_H with

$$fw(\Pi'_H) \geq \frac{3 \log g(\Pi'_H)}{\log\log g(\Pi'_H)} + 5$$

are unique embeddings with so large facewidth.

The proofs of these two results bear a lot of similarity with the proofs of Lemmas 5.3.12 and 5.3.25 and Propositions 5.3.14 and 5.3.27. Indeed, both sets of proofs require to look at two different embeddings Π_1 and Π_2 of the same graph and to monitor the Π_1-faces that become Π_2-noncontractible cycles.

Moreover, these results give us information on the facewidth of (G, Π). Let C be a Π-contractible cycle of G and suppose that $e \in \text{int}(C, \Pi)$. Let Π' be an embedding of G obtained from Π_e by adding e in the embedding in an arbitrary way. Then, remark that the embeddings Π and Π_e cannot be equivalent because, by Lemma 4.2.5, C is Π'-noncontractible.

Therefore, as Π and Π' are not equivalent embedding, according to the first result of Proposition 5.5.1, we get that $fw(G, \Pi) < 100\frac{\log g}{\log \log g}$.

These two results give us an upper bound on the facewidth of (G, Π) with respect to g. However, this is also a consequence of our structural results and that is also pointed out in [Moh95].

Indeed, suppose that (G, Π) has facewidth at least $2m + 1$. Then, let f be a face of (G, Π), by hypothesis, there is no noncontractible simple closed curve on S that intersect at most $2m$ Π-faces including f.

We define the notion of *radius*. It will also be used in Section 7.3.

Definition 5.5.3 (Radius)**.** Let H be a graph with embedding Π_H in a surface. Let C be a Π_H-contractible cycle of H and let $G_C = \text{Int}(C, \Pi_H)$. We define the *radius* $rad(f)$ of a face f in $\mathcal{F}(C, \Pi_H)$ inductively as follow: The faces that share a vertex with $V(C)$ have radius 1. For $i \geq 2$, suppose we have already determined the faces of radius 1 to i. Then, the faces of radius $i+1$ are those that are not of radius at most i but share a vertex with the faces of radius i. We say that G_C has *radius* k if $\max\{rad(f) | f \in \mathcal{F}(C, \Pi_H)\} = k$.

An illustration can be found in Figure 5.10.

We can easily show that there exists a cycle C_f of G that is Π-contractible and so that the radius of f with respect to C_f is $m+1$. Let f be the cycles C_1. Then, we define inductively C_i, for $2 \leq i \leq m+1$, as follows: for $2 \leq i \leq m+1$, the cycle C_i is the cycle bounding the faces that are touching C_{i-1}.

By hypothesis on the facewidth of Π, each of the cycles C_i is Π-contractible and remark that the radius of f with respect with C_i is i. Finally, we define $C_f = C_{m+1}$ and have the property we wanted.

By Corollary 5.3.15, G contains at most m disjoint Π-contractible nested cycles. However, $C_1, ..., C_{m+1}$ are $m+1$ disjoint Π-contractible nested cycles, a contradiction.

Finally, (G, Π) has facewidth at most $2m = O(\log g)$. Therefore, the result on the facewidth of (G, Π) that can be shown using Proposition 5.5.1 or Proposition 5.5.2 is also a simple consequence of our structural results.

Actually, our structural results are a lot stronger than this $O(\log g)$ bound on the facewidth of (G, Π).

To illustrate this, let us define a new notion called *maximum facewidth* that we invented. For a graph H with an embedding Π_H, we first define the *facewidth* $fw(f, \Pi_H)$ *of a face* f of (H, Π) as the smallest number k such that there exists Π_H-facial walks $W_1 = f, ..., W_k$ whose union contains a Π_H-noncontractible cycle.

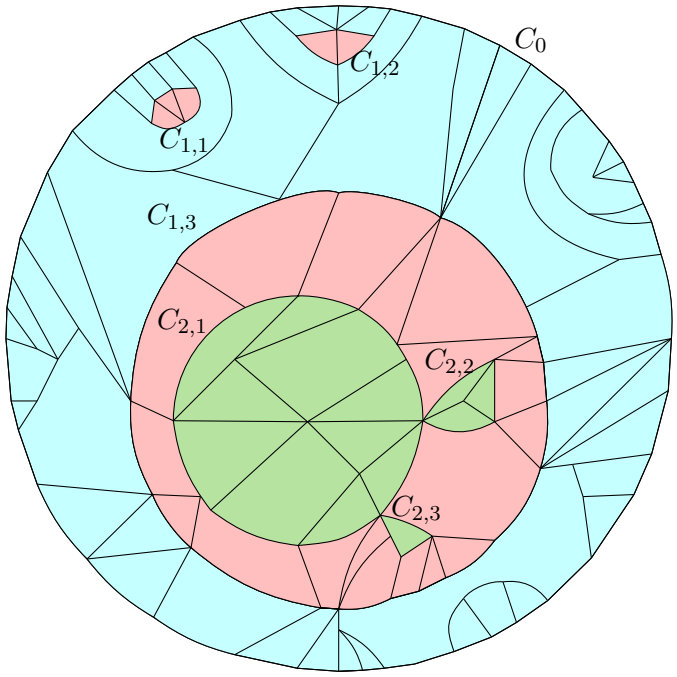


Figure 5.10: Radius of a contractible graph. The three colors indicate the radius of the faces with respect to the outer cycle C_0: faces in blue, pink, and green are respectively of radius 1, 2, and 3.

Remark that the facewidth of (H, Π_H) is then

$$fw(G, \Pi_H) = \min_{f \in F(H,\Pi_H)} fw(f, \Pi_H)$$

We define the *maximum facewidth* max-$fw(H, \Pi_H)$ of (H, Π_H) as the biggest facewidth of a face in (H, Π_H).

$$\text{max-}fw(H, \Pi_H) = \max_{f \in F(H,\Pi_H)} fw(f, \Pi_H)$$

It is clear from the definition that

$$fw(H, \Pi_H) \leq \text{max-}fw(H, \Pi_H)$$

Our structural results actually implies that the maximum facewidth of G is logarithmic in g:

$$\text{max-}fw(G, \Pi) \leq 2m$$

This can be seen by using the exact same proof as above.

To conclude, our structural results are a lot stronger than a bound on the facewidth of (G, Π). Indeed, the facewidth bounds the facewidth of the face of (G, Π) with the smallest facewidth whereas the maximum facewidth bounds the facewidth of all the faces of (G, Π). The $O(\log g)$ bound on the facewidth would not be enough to bound the order of G because we need a general bound that can be used all throughout the graph G, whereas the bound on the facewidth witnesses only a noncontractible simple closed curve on S that G intersects in $O(\log g)$ points.

5.6 Discussion on the optimality of the structural results

In this section, we discuss the optimality of the structural results that we obtained in this chapter. First, we discuss the optimality of the current bound for the nested/homotopic cycles. Then, we establish that there is no benefit in extending the definition of the forbidden structure even further.

5.6.1 Optimality of the bound for the nested/homotopic cycles

In [HK26b], we conjectured that the bounds we obtain for the number of nested and homotopic cycles are essentially tight.

There are two reasons supporting this conjecture. First, Mohar [Moh95] observed that examples constructed by Archdeacon [Arc92] show that the bound established in Proposition 5.5.2 is close to best possible. Similarly, Seymour and Thomas [ST96] described an example due to Thomassen [Tho90] showing that their corresponding bounds are also close to best possible.

Our results are closely related to those of Seymour and Thomas [ST96] and Mohar [Moh95], both in terms of the proof techniques and the resulting bounds. This similarity suggests that our bounds may likewise be close to optimal.

More precisely, we conjecture that the asymptotically optimal bound on the number of Π-well-nested and Π-well-homotopic cycles is

$$O\left(\frac{\log g}{\log\log g}\right)$$

Under this conjecture, our bounds are tight up to a factor of $\log\log g$.

Conjecture 2. *The graph G has at most $O\left(\frac{\log g}{\log\log g}\right)$ Π-well-nested cycles and Π-well-homotopic cycles.*

5.6.2 The scope of the forbidden structure cannot be extended

There is no benefit in extending the definition of nested/homotopic cycles further.

First, it is not possible to extend the definition to cycles that do not belong to the same homotopy class, that is, to cycles that are not all contractible or all noncontractible and homotopic. Indeed, the arguments in Claims 5.3.12.1 and 5.3.25.1 do not apply to cycles belonging to different homotopy classes.

Moreover, considering different homotopy classes is not necessary for our purposes, since we already have results such as Proposition 2.1.22 and its variants Propositions 4.1.2 and 4.1.4, which bound the number of pairwise nonhomotopic noncontractible cycles in an embedding of genus g. These bounds are easily seen to be optimal.

One might then consider extending the definition of nested/homotopic cycles by allowing them to be pinched on more than two pieces (vertices or faces), as illustrated in Figure 5.11.

However, this would not yield any stronger result, since this case can be reduced to nested/homotopic cycles pinched on at most two pieces. Indeed, as illustrated in Figure 5.11, $2k$ nested/homotopic cycles pinched on more than three pieces induce k nested/homotopic cycles on each pair of consecutive pieces. Applying the bounds for nested/homotopic cycles pinched on at most two pieces established in Corollary 5.3.15 and Proposition 5.3.27, we therefore obtain that there are at most $2m$ nested cycles and at most $4m$ homotopic cycles pinched on $k \geq 3$ pieces.

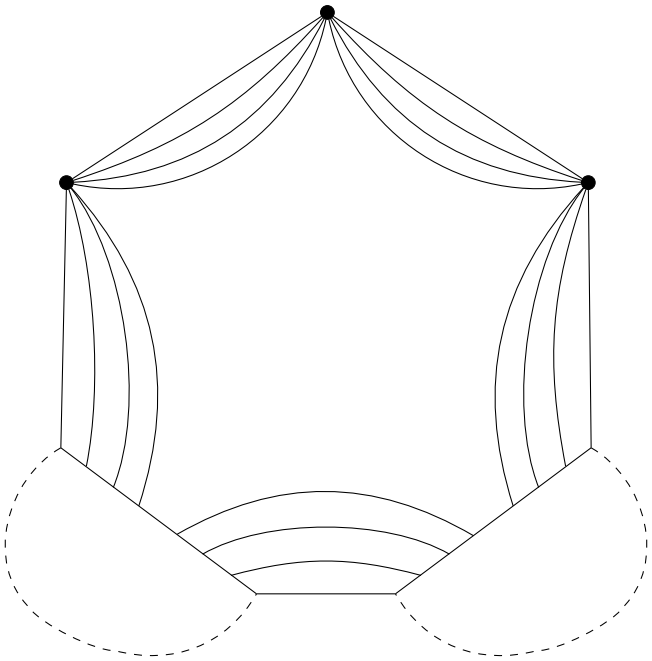


Figure 5.11: An illustration of four Π-well-nested cycles that are pinched on five pieces (three vertices and two faces). The solid lines indicate paths, whereas the dotted line shows the face's boundary on which the nested cycles are pinched.

Chapter 6

Maximum degree of a vertex and size of a face in (G, Π)

6.1 Preliminaries **102**

6.2 Maximum degree of a vertex and size of a face of (G, Π) **106**

6.3 Maximum degree of a vertex and size of a face of a Π-contractible subgraph of (G, Π) with a small separator **115**

6.4 Discussion **120**

6.4.1 Duality between vertices and faces in (G, Π) 121

6.4.2 Toward tight bounds for the maximum degree and maximum size of a face in (G, Π) 122

Recall that G is an excluded minor for a surface S' of genus g such that G has an embedding Π of genus at most $g+2$. In this chapter, we prove two bounds on the maximum degree of a vertex and the maximum size of a face in (G, Π) and in a subgraph G_1 of G that is Π-contractible and has a separator of size $O(\log g)$. These two results are used in Chapter 7 (precisely in Proposition 7.3.1) to show a bound on the order of a 2-connected Π-contractible of G that is the essential ingredient to find either a quasi-polynomial or an almost quadratic bound for the order of the excluded minors for surfaces.

Recall that q and m where defined in Section 5.3: $q = \frac{9073}{9072}$ and $m = 2(\lfloor \log_q(3g+4) \rfloor + 2)$. Let $m' = \lfloor \log_{\frac{4}{3}} 3(4m^2(3g+3)+1) \rfloor$.

Moreover, we define G_1 to be a Π-contractible 2-connected subgraph of G such that there exists a cycle C_1 of G so that $G_1 = \text{Int}(C_1, \Pi)$ and G_1 has a separator A_1 that separates it from the rest of G and has size at most $A(g) = 6\lfloor \log_{\frac{4}{3}} 3((T(g)+1) \times m' + 1) \rfloor \times (12m+8)$ with $T(g) = 264(g+2)(m+1) - 1$ as defined in Section 5.4.

Remark that, because G_1 has a separator of size logarithmic in g, Corollary 5.3.16 applies for G_1 and its separator. This motivates establishing a specific bound for the maximum degree and maximum size of faces in (G_1, Π) which ends up to be asymptotically a lot smaller than in the general case.

We thus show that the maximum degree of a vertex of G and the maximum size of a face in (G, Π) are bounded by $\Delta(g) = g^{O(\log^2 g)}$. Whereas the same bound in (G_1, Π) is found to be $\Delta'(g) = (\log g)^{O((\log\log g)^2)}$. Notice, in particular, that the bound for (G_1, Π) is sub-polynomial in g (i.e. asymptotically smaller than any function Cg^α for any two positive constants C and α) while the bound for (G, Π) is already quasi-polynomial.

The bound on the maximum degree of a vertex and the maximum size of a face in (G, Π) was proved in [HK26b]. The bound on the maximum degree of a vertex and the maximum size of a face in (G_1, Π) is adapted from the proof of Theorem 6.2.2 to apply in the specific context of G_1 and was proved in [HK26a].

These bounds are, to our understanding, the first existing bounds on the maximum degree and the maximum size of faces in (G, Π) or one of its subgraph. Indeed, the result of Seymour [Sey93] did not necessitate such a bound.

In this chapter, we first present in Section 6.1 definitions that are used for both Sections 6.2 and 6.3.

Then, in Section 6.2 and in Section 6.3, we establish the bounds on the maximum degree of a vertex of G and the maximum size of a face respectively in (G, Π) and (G_1, Π).

Finally, in Section 6.4, we discuss the results we obtain in this chapter and what room for improvement there is for these results. In particular, we discuss the limitations of the duality between vertices and faces in (G, Π) and propose open questions relative to the maximum degree and the maximum size of faces in (G, Π).

6.1 Preliminaries

First, we define a notion of faces of G_1 with respect to Π that restricts the faces of (G, Π) to G_1. This definition of faces makes it possible to still use properties that G has as an excluded minor, which would not have been possible if we had defined the faces of (G_1, Π) to be the faces of $(G_1, \Pi(G_1))$. Therefore, it is a more natural definition to use here.

Definition 6.1.1 (Face of G_1 with respect to Π)**.** Let $f \in \mathcal{F}(G, \Pi)$. We define the *faces of G_1 with respect to* Π to be the intersection between a face in $\mathcal{F}(G, \Pi)$ and G_1, that is to say if $f \in \mathcal{F}(G, \Pi)$ then each connected component of $f \cap G_1$ is a face of G_1 with respect to Π. We denote by $\mathcal{F}(G_1, \Pi)$ the set of these faces.

Moreover, we define the *size of a face* $f \in \mathcal{F}(G_1, \Pi)$ to be the number of vertices of f.

Finally, a *piece of* (G_1, Π) is either a vertex of G_1 or a face of G_1 with respect to Π.

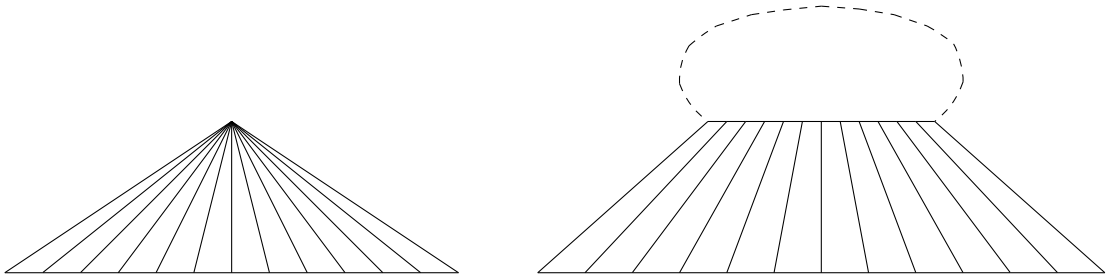
We then define a structure called a *fan (with an arch)* that will be used extensively in the proof of Theorem 6.2.2 and Theorem 6.3.2.

Definition 6.1.2 (Fan (with an arch))**.** Let p be a piece of (G, Π) (resp. (G_1, Π)). A *fan* H from p in G (resp. G_1) is a subgraph of G (resp. G_1) such that:

- H contain a path P (*horizontal path* of H), disjoint from p;
- there are M paths $Q_1, ..., Q_M$ from p to P for some $M \geq 1$, such that, for every $1 \leq j \neq k \leq M$, if p is a vertex, $Q_j \cap Q_k = p$ and, if p is a face, $Q_j \cap Q_k = \emptyset$, $Q_1, ..., Q_M$ intersect p in this order (*vertical paths* of H);

- for every $1 \leq k \leq M$, $Q_k \cap p$ and $Q_k \cap P$ is a vertex;
- $\Pi(H)$ induces a planar embedding.

The *size* of H is M.
An illustration of a fan from a vertex and a face can be found in Figure 6.1.



(a) Fan on a vertex (b) Fan on a face

Figure 6.1: Fans from a vertex and a face. The solid lines indicate paths, whereas the dotted line shows the boundary of the face from which the fan in Figure (b) is.

Let p be a piece of (G, Π) (resp. (G_1, Π)). A *fan with an arch* H from p in G (resp. G_1) is a subgraph of G (resp. G_1) such that:

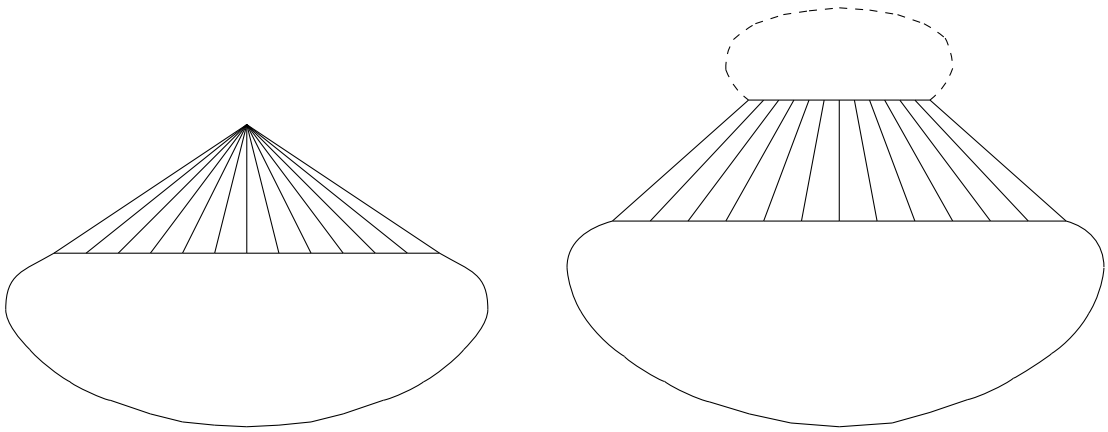
- H contains a fan H' from p;
- Let $Q_1, ..., Q_M$ be the vertical paths of H' and P be its horizontal path, H contains a path A (*arch path* of H) whose endpoints are on $Q_1 \cap P$ and $Q_M \cap P$, which is disjoint from H' except on its endpoints. Moreover, $H' \cup A$ is Π-contractible;
- There is at most one vertex in the interior of A (*intersection point* of A) such that there is no edge in $\mathrm{Int}(H, \Pi)$ which has an endpoint in $\mathrm{int}(A) - v$ with $v = \emptyset$ if A has no intersection point and v the intersection point of A if A has one.

An illustration of a fan with an arch from a vertex and a face can be found in Figure 6.2.

Now, we define several objects that will play an essential role throughout the proof.

Let H be a fan (with or without an arch) from p in G (resp. G_1). Let P be its horizontal path and $Q_1, ..., Q_M$ be its vertical paths with $M \geq 2$. We define a *column* of H to be the subpath of P between the two endpoints of two consecutive vertical paths from p: let $1 \leq k \leq M - 1$, then the k-th column of H is the subpath of P between the endpoints $Q_k \cap P$ (included) and $Q_{k+1} \cap P$ (excluded).

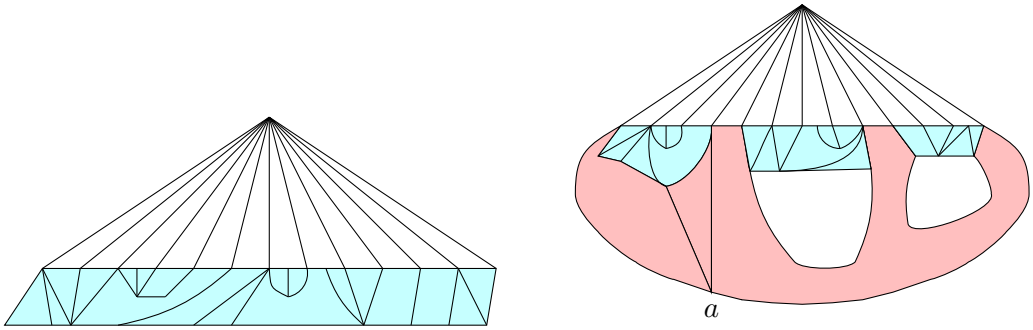
Then, we define $\mathcal{F}(H)$ to be the set of faces that touch the interior of the horizontal path P of H in $\mathrm{Ext}(H, \Pi)$. Assume now that H is a fan with an arch, let A be its arch path. Let C be the cycle induced by $A \cup p$ that does not contain p in Π and let C' be the cycle bounding $P \cup \bigcup_{1 \leq i \leq M} Q_i$ in Π, we define C_A to be the circuit that bounds $\mathrm{Int}(C \cup C', \Pi)$. Then, we define $\overline{\mathcal{F}}(H)$ to be the set of faces in $\mathrm{Int}(C_A, \Pi)$ that touch the horizontal path P of H and do



(a) Fan with an arch from a vertex

(b) Fan with an arch from a face

Figure 6.2: Fans with an arch from a vertex and a face. The solid lines indicate paths, whereas the dotted line shows the boundary of the face from which the fan in Figure (b) is.



(a) The faces in $\mathcal{F}(H)$ of the fan H are represented in blue.

(b) The intersection point of the fan with an arch H is a. The faces in $\mathcal{F}(H)$ of H are represented in blue, and the faces in $\mathcal{F}_{arch}(H)$ are represented in pink.

Figure 6.3: The set $\mathcal{F}(H)$ for a fan, and $\mathcal{F}(H)$ and $\mathcal{F}_{arch}(H)$ for a fan with an arch

not intersect A. We also define $\mathcal{F}_{arch}(H)$ to be the set of faces in $\text{Int}(C_A, \Pi)$ that touches the horizontal path P of H and intersect A. See Figure 6.3 for illustrations.

Let H be a fan (with or without an arch) from p with the horizontal path P. Suppose the signature on the edges of H in Π is positive. If H is a fan of (G_1, Π), we can modify the embedding Π so that the signature on the edges of H in Π is positive, because H is Π-contractible. Therefore, it is always possible to suppose that the signature on the edges of H in Π is positive.

Let C_H be the cycle of H so that $H \subseteq \text{Int}(C_H, \Pi)$. Then, a face $f \in \mathcal{F}(H)$ is an *arched face* if f does not intersect C_H in a single segment (vertex or path). Let $f \in \mathcal{F}(H)$ be a face; we define the *arches* of f to be the maximal sub walks of f whose interior vertices are not in C_H. Remark that the arches of any face $f \in \mathcal{F}(H)$ are paths with their ends on C_H. For an arch A

of a face $f \in \mathcal{F}(H)$, we define C_A to be the Π-contractible cycle induced by $A \cup C_H$ that does not contain H.

Moreover, remark that if an arch A of a face $f \in \mathcal{F}(H)$ has positive signature (which is always the case if H is a fan of (G_1, Π)), then either $\text{int}(C_A, \Pi)$ or $\text{ext}(C_A, \Pi)$ contain no edge with an endpoint on the interior of A (because A is a subwalk of a face in Π). We say that an arch A of a face $f \in \mathcal{F}(H)$ is Π-*noncontractible* if $A \cup H$ induces a Π-noncontractible cycle. See Figure 6.4 for an illustration.

Let $f, f' \in \mathcal{F}(H)$ be two faces intersecting in $G_1 - H$. Remark that, using Lemma 4.2.2, we can show that f and f' must intersect in a vertex or an edge. Let v_1 and v_2 be the first and last intersection points of f and f' in G_1 (resp. G_1), with potentially $v_1 = v_2$, and let $P_{1,f}, P_{2,f}, P_{1,f'}, P_{2,f'}$ be the subpaths of f and f' from H to v_1 and v_2. Then, we define the *arches* of f and f' to be $A_1 = P_{1,f} \cup P_{1,f'}$ and $A_2 = P_{2,f} \cup P_{2,f'}$. Remark that A_1 and A_2 are paths with both ends on H.

Moreover, remark that if an arch A of two faces $f, f' \in \mathcal{F}(H)$ has a positive signature (which is always the case if H is a fan of (G_1, Π)), then either $\text{int}(C_A, \Pi)$ or $\text{ext}(C_A, \Pi)$ contain no edge with an endpoint on the interior of A except on one sole vertex that we call the *intersection point* of A. We say that an arch A of two faces $f, f' \in \mathcal{F}(H)$ is Π-*noncontractible* if $A \cup H$ induces a Π-noncontractible cycle. See Figure 6.4 for an illustration.

We define $\mathcal{A}(H)$ to be the set of all arches induced by either one face from $\mathcal{F}(H)$ or by two faces from $\mathcal{F}(H)$ that intersect in $G - H$ (resp. $G_1 - H$).

Suppose that every arch $A \in \mathcal{A}(H)$ induces together with H a Π-contractible subgraph of G. As $G' = H \cup \mathcal{A}(H)$ is Π-contractible, we can suppose that the signature on every edge of G' is positive. This is always the case if H is a fan of (G_1, Π). We say that $A \in \mathcal{A}(H)$ is *empty above* (resp. *empty below*) if $\text{ext}(C_A, \Pi)$ (resp. $\text{int}(C_A, \Pi)$) contains no edge with an endpoint on the interior of A (except maybe on one sole vertex). Remark that, in that case, every arch $A \in \mathcal{A}(H)$ is either empty above or empty below. See Figure 6.4 for an illustration.

Let $A \in \mathcal{A}(H)$ and let $P_A = C_A \cap P$. We define the *size* of A to be the number of columns that P_A intersects.

Suppose that every arch $A \in \mathcal{A}(H)$ induces together with H a Π-contractible subgraph of G. This is always the case if H is a fan of (G_1, Π).

We define $\mathcal{A}_{max}(H)$ to be the set of arches in $\mathcal{A}(H)$ so that $A \in \mathcal{A}_{max}(H)$ if, for every $A' \in \mathcal{A}(H)$ with $A \neq A'$, $A \not\subseteq \text{Int}(C_{A'}, \Pi)$. We call the arches in $\mathcal{A}_{max}(H)$ the *maximal arches* of H. Remark that, by construction, every arch in $\mathcal{A}_{max}(H)$ is empty below. Remark that they are naturally ordered along P.

Let us first modify $\mathcal{A}_{max}(H)$ to ensure that the Π-contractible cycles induced by the arches in $\mathcal{A}_{max}(H)$ do not intersect. To do so, we proceed by order of the arches in $\mathcal{A}_{max}(H)$ along P: for the arch A in $\mathcal{A}_{max}(H)$ that is leftmost on P and has not yet been treated, we remove from $\mathcal{A}_{max}(H)$ every arch A' so that $\text{int}(C_A, \Pi) \cap \text{int}(C_{A'}, \Pi) \neq \emptyset$. Then, remark that two maximal arches are consecutive in this order if and only if they intersect.

Let us now modify $\mathcal{A}_{max}(H)$ so that no maximal arches of size at most 2 are adjacent by merging consecutive arches (this can also be done by order along P). See Figure 6.5 for an illustration.

For the purpose of initialization, we define all the notions defined for a fan (with or without an arch) for a piece p: We define the *horizontal path* and the *vertical paths* to be the piece p and therefore $\mathcal{F}(H)$ to be the set of faces that touch p. We define the notions of *arched faces*,

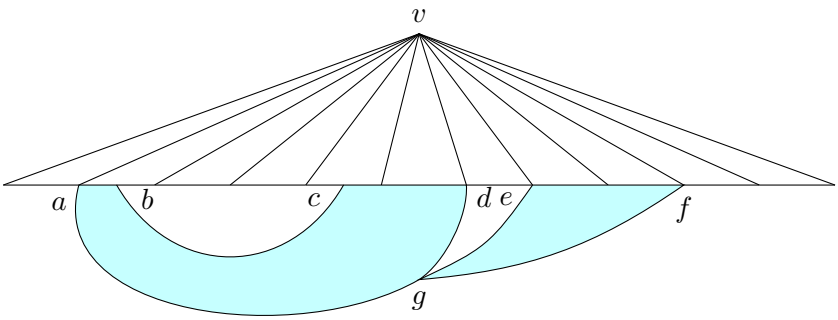


Figure 6.4: Illustration for arched faces and arches. The fan H is represented, and two faces from $\mathcal{F}(H)$ are represented in blue. The face F on the left is arched, but the face F' on the right is not. The paths from a to d and from b to c are the two arches of F, and the path from e to f is the only arch of F'. The paths from d to e through g and from a to f through g are both arches of F and F', and both have g as their intersection point. In this representation, the arches from a to d and from a to f are empty below, and the arches from b to c and from d to e are empty above.

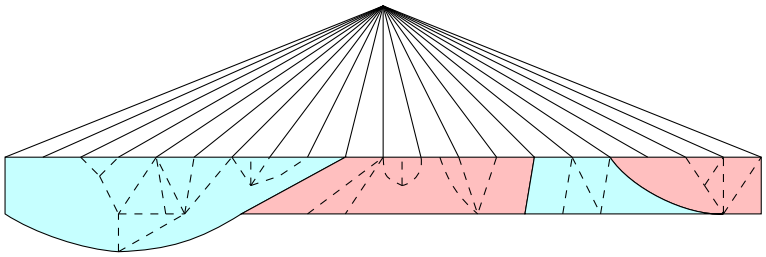


Figure 6.5: Illustration for maximal arches. The edges inside a maximal arch are represented in dotted. The fan H represented has 3 maximal arches, represented in alternating blue and pink.

arch on a face, *arch on two faces*, etc., as for the general case. We define a *column* of H to be an edge adjacent to p if p is a vertex and a vertex of p if p is a face.

6.2 Maximum degree of a vertex and size of a face of (G, Π)

Let $\tilde{g} = 4(6g + 7) - 1$.

We prove that the maximum degree of a vertex of G and the maximum size of a face in (G, Π) are bounded by

$$\Delta(g) = 2m(\tilde{g}+1)^4 \left(4m(\tilde{g}+1)^2\right)^{m^2}$$

To prove this, we proceed by contradiction and reach a contradiction by proving that a high degree of a vertex (resp. big size of a face) leads to a large number of Π-contractible nested cycles, which contradicts Corollary 5.3.15. The proof makes repetitively use of Propositions 4.1.4, 5.2.3 and 5.3.27 and Corollary 5.3.15.

We first establish the following definition that will be use in the proof of Theorem 6.2.2:

Definition 6.2.1 (Independent subgraphs)**.** Let H be a graph, let $k \geq 1$ and let $\mathcal{H} = \{H_0, ..., H_k\}$ be a set of subgraphs of H. We say that $\mathcal{H}$ is a set of *independent subgraphs* of H if it is possible to associate to each subgraph H_i an edge e_i so that $e_i \notin \cup_{j \neq i} H_j$.

We now prove the main result of this section:

Theorem 6.2.2 ([HK26b, Theorem 6.3])**.**

$$\Delta(G) \leq \Delta(g) \quad \textit{and} \quad \Delta_F(G, \Pi) \leq \Delta(g)$$

Proof. We prove several lemmas before proving the result.

Lemma 6.2.2.1. *Let p be a piece of (G, Π) and let H be a fan (with an arch) from p in (G, Π) or $H = p$. Let P be the horizontal path of H. Let $f \in \mathcal{F}(H) \cup \mathcal{F}_{arch}(H)$ ($\mathcal{F}_{arch}(H) = \emptyset$ if H is a fan without an arch or $H = p$) so that $f \cup H$ is contained either in a disk or in a cylinder in Π. Then, f touches at most $\tilde{g}$ columns of H.*

Proof. If $H = p$, then a face f so that $f \cup H$ is contained either in a disk or in a cylinder in Π cannot intersect several columns of p, as otherwise, p would either be a cutvertex of G which is in contradiction with the fact that G is 2-connected or have a 2-separator that separates a Π-contractible subgraph from the rest of the graph which is in contradiction with Lemma 4.2.2.

Suppose that H is a fan (with an arch). Suppose that there is a face $f \in \mathcal{F}(H) \cup \mathcal{F}_{arch}(H)$ that touches at least $\tilde{g} + 1$ columns of H. Then, f together with H forms $\tilde{g} + 1$ isolated paths on p. However, this is a contradiction according to Proposition 5.2.3. □

Lemma 6.2.2.2. *Let H be either a piece p or a fan from a piece p in (G, Π) of size at least M. Let P be the horizontal path of H. Suppose every arch in $\mathcal{A}(H)$ intersects H in $P \cup p$. Then, there exists at least $\frac{M}{4m(\tilde{g}+1)^4}$ consecutive columns that induce a fan H' from v so that the arches in $\mathcal{A}(H')$ are all Π-contractible.*

Proof. Let us show that there exist at most $(3g + 3) \times 2 \times 4m(\tilde{g}^2 + \tilde{g} + 1)$ columns so that the Π-noncontractible arches of H all touch one of these columns.

First, remark that it is enough to only consider independent Π-noncontractible arches of H: indeed if A is a Π-noncontractible arch of H that is not independent of a set $\mathcal{A}$ of Π-noncontractible arches of H, then A shares its first and last edges with arches in $\mathcal{A}$. Hence, $\mathcal{A}$ and $\mathcal{A} \cup \{A\}$ intersect the same columns.

Let $Q_1, ..., Q_M$ be the vertical paths of H with $M \geq 2$.

Let $\mathcal{A}$ be a maximal set of independent Π-noncontractible arches of H. For $A \in \mathcal{A}$, let $\mathcal{I}_A \subseteq [\![1, M]\!]$ be the indexes of the columns of H that A intersects ($|\mathcal{I}_A| \leq 2$) and let C_A be the Π-noncontractible cycle on p that contains A and does not contain the interior of p if p is a face induced by $A \cup P \cup p \cup \bigcup_{k \in \mathcal{I}} Q_k$. We then define the set $\mathcal{C} = \{C_A | A \in \mathcal{A}\}$ of cycles on p that correspond to the arches in $\mathcal{A}$.

By Proposition 4.1.4, as $g(\Pi) \leq g + 2$, there are at most $3g + 3$ pairwise Π-noncontractible nonhomotopic cycles in $\mathcal{C}$. Thus, it remains to prove that, for a given class of homotopy in Π, there exists a set of at most $8m(\tilde{g}^2 + \tilde{g} + 1)$ columns so that all the arches on p whose associated cycles all belong to this class of homotopy touch one of these columns.

Let $\tilde{\mathcal{A}}$ be a set of arches from $\mathcal{A}$ that are Π-noncontractible and whose associated cycles are Π-homotopic and let $\tilde{\mathcal{C}}$ be the set of associated cycles. Let us show that there exists a set of at most $8m(\tilde{g}^2 + \tilde{g} + 1)$ columns so that every arch in $\tilde{\mathcal{A}}$ touches one of these columns.

First, we modify the subgraph $H \cup \bigcup_{A \in \tilde{\mathcal{A}}} A$ so that it has a positive signature in Π. This is possible because, for each cycle $C \in \tilde{\mathcal{C}}$, C is two-sided and hence has a positive signature.

Let A_1 and A_1' be the two arches from $\tilde{\mathcal{A}}$ whose associated cycles C_1 and C_1' in $\tilde{\mathcal{C}}$ are so that, for every cycle $C \in \tilde{\mathcal{C}}$, $C \subseteq \mathrm{Int}(C_1 \cup C_1', \Pi)$. Let $\tilde{G} = \mathrm{Int}(C_1 \cup C_1', \Pi)$. The embedding $\Pi(\tilde{G})$ is an embedding in a cylinder bounded by C_1 and C_1'; hence it is a planar embedding. For $A \in \tilde{\mathcal{A}}$, we say that A is empty above (resp. empty below) if A is empty above (resp. empty below) in the planar embedding $\Pi(\tilde{G})$ with outer face C_1. Remark that every arch $A \in \tilde{\mathcal{A}}$ is either empty above or empty below.

We then partition $\tilde{\mathcal{A}}$ into two sets $\overline{\mathcal{A}}$ and $\underline{\mathcal{A}}$ which correspond to the set of arches of $\tilde{\mathcal{A}}$ which are empty respectively above and below. Let $\overline{\mathcal{C}}$ and $\underline{\mathcal{C}}$ be the sets of the associated cycles.

Let us show that there exists a set of at most $4m(\tilde{g}^2 + \tilde{g} + 1)$ columns so that every arch in $\overline{A}$ touches one of these columns. The same result for $\underline{A}$ can be proved using the same method.

Let us first find a maximal subset of internally disjoint arches from $\overline{\mathcal{A}}$. First, take one of the two arches from $\overline{\mathcal{A}}$ whose associated cycles C_1 and C_1' in $\overline{\mathcal{C}}$ are so that, for every cycle $C \in \overline{\mathcal{C}}$, $C \subseteq \mathrm{Int}(C_1 \cup C_1', \Pi)$. Let A_1 be such an arch. Then, select the arch A_2 from $\overline{\mathcal{A}}$ that is internally disjoint from A_1 and so the cycle associated with A_2 is the closest to the cycle associated with A_1. Iterate this process to get a maximal set $\overline{\mathcal{A}}_0 = \{A_1, ..., A_p\}$ of pairwise internally disjoint arches (for some $p \geq 1$). Let $\overline{\mathcal{C}}_0 = \{C_1, ..., C_p\}$ be the set of associated cycles. Let C_{p+1} be the cycle so that, for every cycle $C \in \overline{\mathcal{C}}$, $C \subseteq \mathrm{Int}(C_1 \cup C_{p+1}, \Pi)$ and A_{p+1} the associated arch. See Figure 6.6 for an illustration.

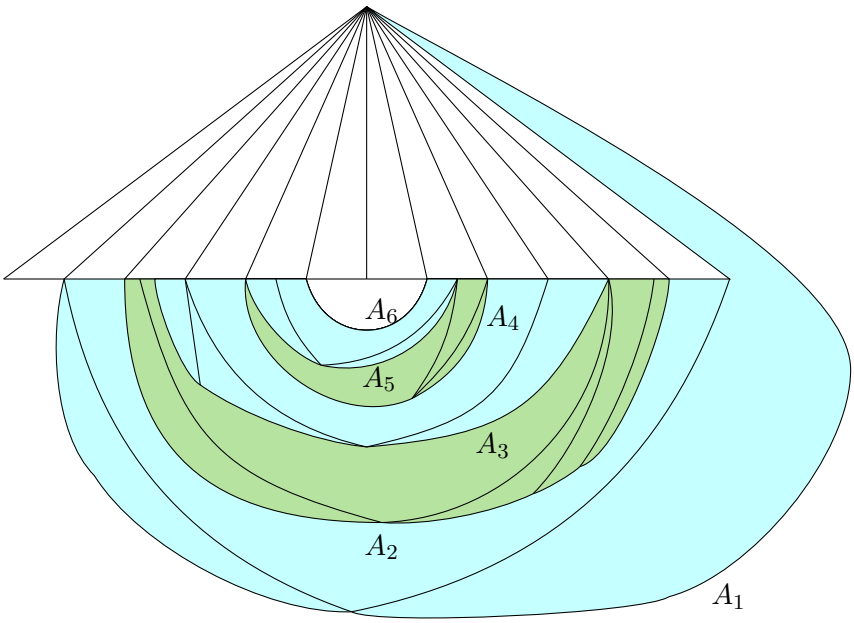


Figure 6.6: Illustration for the sets $\overline{\mathcal{A}}_0$ and $(\overline{\mathcal{A}}_i)_{1 \leq i \leq p}$. The internally disjoint arches in $\overline{\mathcal{A}}_0$ are in this example $A_1, A_2, A_3, A_4, A_5, A_6$. We have $p = 6$ and $A_7 = A_6$. Moreover, the disks containing the arches in $\overline{\mathcal{A}}_i$ are depicted in alternating blue and green.

Remark that if an arch of $\overline{\mathcal{A}}$ intersects $\mathrm{int}(C_i \cup C_{i+1}, \Pi)$ $(1 \leq i \leq p)$ then it is contained in $\mathrm{Int}(C_i \cup C_{i+1}, \Pi)$. Let $\overline{\mathcal{A}}_i$ be the set of arches in $\overline{\mathcal{A}} - \overline{\mathcal{A}}_0$ that lie in $\mathrm{Int}(C_i \cup C_{i+1}, \Pi)$ $(1 \leq i \leq p)$. Then, the arches in $\overline{\mathcal{A}}$ can be partitioned into the sets $\overline{\mathcal{A}}_i$ for $0 \leq i \leq p$. See Figure 6.6 for an illustration.

Claim 6.2.2.2.1. *There exists a set of at most $4m$ columns of H so that all arches in $\overline{\mathcal{A}}_0$ touch at least one of these columns.*

Proof of the claim. Let us create an auxiliary graph Γ whose vertices are the columns of H intersected by the arches in $\overline{\mathcal{A}}_0$. There is an edge between two columns if an arch in $\overline{\mathcal{A}}_0$ touches both columns. Remark that, by Proposition 5.3.27, there are at most $2m$ cycles from $\overline{\mathcal{C}}$ that correspond to a matching in Γ. Hence, a maximal matching of Γ has a size at most $2m$. Moreover, the set of the vertices in a maximal matching of Γ is a dominating set of Γ. Therefore, a set of at most $4m$ columns of H exists so that all arches in $\overline{\mathcal{A}}_0$ touch at least one of these columns. ■

Claim 6.2.2.2.2. *Suppose that $p > 4m$. Let $\mathcal{B}$ be the set of at most $4m$ columns of H obtained in Claim 6.2.2.2.1. Then, there exists a set $\mathcal{I} \subseteq [\![1, p]\!]$ of size at most $4m$ so that all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}} \overline{\mathcal{A}}_i$ touch at least one of these columns.*

Proof of the claim.

We will prove the property by induction on $p' = p - 4m$.

For $p' = 1$, we have $p = 4m + 1$. By the pigeonhole principle, there are 2 arches A_i and A_j of $\overline{\mathcal{A}}_0$ ($A \leq i \leq j \leq p$, potentially $i = j$) that touch the same column of $\mathcal{B}$.

Assume $i \neq j$. Remark then that every arch in $\overline{\mathcal{A}}_k$ for $i \leq k < j$ touches this same column. Finally, by taking $\mathcal{I} = [\![1, p]\!] - [\![i, j[\![$ which has size at most $4m$, we indeed have that all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}} \overline{\mathcal{A}}_i$ touch at least one of the $4m$ columns $\mathcal{B}$.

Assume $i = j$. Remark that every arch in $\overline{\mathcal{A}}_k$ for either $k \geq i$ or $k \leq i$ touches this same column. Finally, by taking $\mathcal{I} = [\![1, p]\!] - [\![i, p]\!]$ or $\mathcal{I} = [\![1, p]\!] - [\![1, i]\!]$ which has size at most $4m$, we indeed have that all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}} \overline{\mathcal{A}}_i$ touch at least one of the $4m$ columns $\mathcal{B}$.

Now, let $p' = p - 4m \geq 2$ and suppose that the property is true for every $p'' < p'$.

By the pigeonhole principle, there are arches A_i and A_j of $\overline{\mathcal{A}}_0$ ($1 \leq i \leq j \leq p$, potentially $i = j$) that touch the same column of $\mathcal{B}$.

Assume $i \neq j$. Remark then that every arch in $\overline{\mathcal{A}}_k$ for $i \leq k < j$ touches this same column. Let $\mathcal{I} = [\![1, p]\!] - [\![i, j[\![$ which has size at most $p - 1$. We have that that all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}} \overline{\mathcal{A}}_i$ touch at least one of the $4m$ columns $\mathcal{B}$.

Assume $i = j$. Remark that every arch in $\overline{\mathcal{A}}_k$ for either $k \geq i$ or $k \leq i$ touches this same column. Let $\mathcal{I} = [\![1, p]\!] - [\![i, p]\!]$ or $\mathcal{I} = [\![1, p]\!] - [\![1, i]\!]$ which has size at most $p - 1$. We have that all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}} \overline{\mathcal{A}}_i$ touch at least one of the $4m$ columns $\mathcal{B}$.

In both cases, we proved that there is $\mathcal{I} \subseteq [\![1, p]\!]$ of size at most $p - 1$ so that all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}} \overline{\mathcal{A}}_i$ touch at least one of the $4m$ columns $\mathcal{B}$. Let us apply induction on $\cup_{i \in \mathcal{I}} A_i$ and $\cup_{i \in \mathcal{I}} \overline{\mathcal{A}}_i$ with $p'' = |\mathcal{I}| - 4m < p'$. By induction hypothesis, there is a set $\mathcal{I}' \subseteq \mathcal{I}$ of size at most $4m$ so that all arches in $\cup_{i \in \mathcal{I}} A_i$ and $\cup_{i \in \mathcal{I} - \mathcal{I}'} \overline{\mathcal{A}}_i$ touch at least one of the $4m$ columns $\mathcal{B}$.

Finally, all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}'} \overline{\mathcal{A}}_i$ touch at least one of the $4m$ columns $\mathcal{B}$. ■

Claim 6.2.2.2.3. *For $1 \leq i \leq p$, suppose that A_i intersects the column c of H. Then, there exists a set of at most $\tilde{g}^2 + \tilde{g}$ columns of H so that every arch in $\overline{\mathcal{A}}_i$ touches either c or one of these columns.*

Proof of the claim.

- If A_i has no intersection point:

  Remark that, by construction, the arches in $\overline{\mathcal{A}}_i$ always intersect the interior of A_i either in a vertex or in a subpath of A_i. Remark moreover that, by construction, as the arches in $\overline{\mathcal{A}}_i$ are empty above, their paths in $\text{int}(C_i \cup C_{i+1}, \Pi)$ are pairwise internally disjoint.

For each arch $A \in \overline{\mathcal{A}}_i$, let v_A be its intersection point with A_i (if A shares a subpath with A_i, the intersection point is the vertex on which A and A_i diverge), we call the partition of A in the two subpaths of A with v_A as an endpoint its *half-arches*. Let $\mathcal{B}_1$ and $\mathcal{B}_2$ be the partition of the half-arches in $\overline{\mathcal{A}}_i$ so that, for two half-arches $B, B' \in \mathcal{B}_1$ (resp. $\mathcal{B}_2$), the paths from p to A_i induced by B and B' are Π-homotopic. Remark that every arch of $\overline{\mathcal{A}}_i$ has one of its half-arch in $\mathcal{B}_1$ and the other in $\mathcal{B}_2$. Remark, moreover, that either $\mathcal{B}_1$ or $\mathcal{B}_2$ (potentially none and not both) contains half-arches that touch p. Suppose, therefore, without loss of generality, that $\mathcal{B}_1$ does not contain any half-arch that touches p.

If a subset $\mathcal{B}$ of arches in $\mathcal{B}_1$ is so that every half-arch of $\mathcal{B}$ in $\text{int}(C_i \cup C_{i+1}, \Pi)$ intersects A_i on the same vertex v_i, then $\mathcal{B}$ intersects at most $\tilde{g}$ columns, as otherwise these arches would create $\tilde{g}+1$ isolated paths on p and v_i and this would be in contradiction with Proposition 5.2.3. Moreover, if a subset $\mathcal{B}$ of arches in $\mathcal{B}_1$ is so that the half-arches of $\mathcal{B}$ in $\text{int}(C_i \cup C_{i+1}, \Pi)$ intersect A_i on disjoint vertices, then $\mathcal{B}$ intersects at most $\tilde{g}$ columns, as otherwise these arches would create $\tilde{g}+1$ isolated paths on p and A_i and this would be in contradiction with Proposition 5.2.3.

Finally, in total, there exists a set of at most $\tilde{g}^2 \leq \tilde{g}^2 + \tilde{g}$ columns so that every arch in $\overline{\mathcal{A}}_i$ touches either c or one of these columns.

- If A_i has one intersection point:

  Let v_i be the intersection vertex of A_i and let A_i^1 and A_i^2 be the partition of A_i into two subpaths of A_i with v_i as an endpoint.

  Remark that, by construction, the arches in $\overline{\mathcal{A}}_i$ always intersect the interior of A_i either in a vertex or in a subpath of A_i^1 or A_i^2. Remark moreover that, by construction, as the arches in $\overline{\mathcal{A}}_i$ are empty above, their paths in $\text{int}(C_i \cup C_{i+1}, \Pi)$ are pairwise internally disjoint. Let $\overline{\mathcal{A}}_i^0$, $\overline{\mathcal{A}}_i^1$ and $\overline{\mathcal{A}}_i^2$ be the sets of arches in $\overline{\mathcal{A}}_i$ that intersect the interior of A_i in respectively a vertex, A_i^1 and A_i^2.

  Remark that two arches $A, A' \in \overline{\mathcal{A}}_i^0$ must intersect the interior of A_i in the same vertex (to avoid crossings). Let v_i' be the vertex of A_i that arches in $\overline{\mathcal{A}}_i^0$ all intersect (if $\overline{\mathcal{A}}_i^0$ is not empty). For each arch $A \in \overline{\mathcal{A}}_i^0$, let v_A be its intersection point with A_i (if A shares a subpath with A_i, the intersection point is the vertex on which A and A_i diverge), we call the partition of A in the two subpaths of A with v_A as an endpoint its *half-arches*. Let $\mathcal{B}_1$ and $\mathcal{B}_2$ be the partition of the half-arches in $\overline{\mathcal{A}}_i^0$ so that, for two half-arches $B, B' \in \mathcal{B}_1$ (resp. $\mathcal{B}_2$), the paths from p to A_i induced by B and B' are Π-homotopic. Remark that every arch of $\overline{\mathcal{A}}_i^0$ has one of its half-arch in $\mathcal{B}_1$ and the other in $\mathcal{B}_2$. Remark, moreover, that either $\mathcal{B}_1$ or $\mathcal{B}_2$ (potentially none and not both) contains half-arches that touch p. Suppose, therefore, without loss of generality, that $\mathcal{B}_1$ does not contain any half-arch that touches p. Remark that $\mathcal{B}_1$ intersects at most $\tilde{g}$ columns, as otherwise these arches would create $\tilde{g}+1$ isolated paths on p and v_i' and this would be in contradiction with Proposition 5.2.3. Therefore, there exists a set of at most $\tilde{g}$ columns so that every arch in $\overline{\mathcal{A}}_i^0$ touches one of these columns.

  Suppose, without loss of generality, that A_i^2 intersects c. The, the arches in $\overline{\mathcal{A}}_i^2$ all touch c. Therefore, it only remains to bound the number of columns that the arches in $\overline{\mathcal{A}}_i^1$ intersect.

  If a subset $\mathcal{A}^1$ of arches in $\overline{\mathcal{A}}_i^1$ is so that the arches of $\mathcal{A}^1$ in $\text{int}(C_i \cup C_{i+1}, \Pi)$ intersect A_i on the same vertex v_i^1, then $\mathcal{A}^1$ intersect at most $\tilde{g}$ columns of H, otherwise these arches

would create $\tilde{g}+1$ isolated paths on p and v_i^1 and this would be in contradiction with Proposition 5.2.3. Moreover, if a subset $\mathcal{A}^1$ of arches in $\overline{\mathcal{A}}_i^1$ is so that the arches of $\mathcal{A}^1$ in $\text{int}(C_i \cup C_{i+1}, \Pi)$ intersect A_i on disjoint vertices, then $\mathcal{A}^1$ intersect at most $\tilde{g}$ columns of H, as otherwise these arches would create $\tilde{g}+1$ isolated paths on p and A_i^1 and this would be in contradiction with Proposition 5.2.3.

Finally, there exists a set of at most $\tilde{g}^2$ columns so that every arch in $\overline{\mathcal{A}}_i^1$ touches one of these columns. And in total, there exists a set of at most $\tilde{g}^2+\tilde{g}$ columns so that every arch in $\overline{\mathcal{A}}_i$ touches either c or one of these columns.

■

By Claim 6.2.2.2.2, there exists a set $\mathcal{C}$ of at most $4m$ columns and a set $\mathcal{I} \subseteq [\![1, p]\!]$ of size at most $4m$ so that all arches in $\overline{\mathcal{A}}_0$ and in $\cup_{i \notin \mathcal{I}} \overline{\mathcal{A}}_i$ touch at least one of these columns. Moreover, let $i \in \mathcal{I}$ and c_i be a column of $\mathcal{C}$ that A_i touches. By Claim 6.2.2.2.3, there is a set of at most $\tilde{g}^2+\tilde{g}$ columns so that the arches in $\overline{\mathcal{A}}_i$ all touch either c_i or these columns. Finally, the arches in $\overline{\mathcal{A}}$ touch at most $4m(\tilde{g}^2+\tilde{g}+1)$ columns.

In total, there exists at most $(3g+3)\times 2\times 4m(\tilde{g}^2+\tilde{g}+1)$ columns so that the Π-noncontractible arches of H all touch one of these columns.

Now, let us show that there are at least $\frac{M}{4m(\tilde{g}+1)^4}$ consecutive columns of H so that, in the induced fan H', $\mathcal{A}(H')$ contains no Π-noncontractible arch.

Let $\mathcal{A}_p(H)$ and $\mathcal{A}_P(H)$ be respectively the Π-noncontractible arches of $\mathcal{A}(H)$ that touch p and the arches that touch solely P. Let us show that there are at most $(3g+3) \times 2 \times [4m(\tilde{g}^2+\tilde{g}+1) \times 2\tilde{g}+4 \times (3g+3) \times \tilde{g}] \leq 4m(\tilde{g}+1)^4$ columns so that:

1. the faces $\mathcal{F}_p(H)$ in $\mathcal{F}(H)$ that have an arch in $\mathcal{A}_p(H)$ touch only these columns, and

2. the faces $\mathcal{F}_P(H)$ in $\mathcal{F}(H)$ that have arches only in $\mathcal{A}_P(H)$ are so that each of their Π-noncontractible arches touch one of these columns.

This will guarantee that there are at least $\frac{M}{4m(\tilde{g}+1)^4}$ consecutive columns of H so that, in the induced fan H', $\mathcal{A}(H')$ contains no Π-noncontractible arch.

We start with the at most $(3g+3) \times 2 \times 4m(\tilde{g}^2+\tilde{g}+1)$ columns $\mathcal{C}_0$ found previously. Hence, the second property is already guaranteed.

Let $\mathcal{A}$ be a maximal subset of independent Π-noncontractible homotopic empty above arches from $\mathcal{A}_p(H)$. Let $\mathcal{C} \subseteq \mathcal{C}_0$ of size at most $4m(\tilde{g}^2+\tilde{g}+1)$ be the columns so that every arch in $\mathcal{A}$ intersect a column of $\mathcal{C}$. Let S_0 be the smallest cylinder in S that contains all the cycles associated with $\mathcal{A}$.

For $c \in \mathcal{C}$, let $\mathcal{A}_c \subseteq \mathcal{A}$ be the arches in $\mathcal{A}$ that intersect the column c, $(\mathcal{A}_c)_{c \in \mathcal{C}}$ is a partition of $\mathcal{A}$. Remark that for each column $c \in \mathcal{C}$, every face associated with the arches in $\mathcal{A}_c$ except two (that we will call f_c^1 and f_c^2) intersects only c. Hence, adding columns that intersect only these two faces is enough. Moreover, let $\mathcal{F}_{\mathcal{C}} = \{f_c^1, f_c^2 | c \in \mathcal{C}\}$, remark that every face in $\mathcal{F}_{\mathcal{C}}$ except two (f_1^* and f_2^*) is contained in S_0. By Lemma 6.2.2.1, the faces in $\mathcal{F}_{\mathcal{C}} - \{f_1^*, f_2^*\}$ intersect at most $\tilde{g}$ columns. Moreover, the faces f_1^* and f_2^* contain at most $2 \times (3g+3)$ Π-noncontractible arches on H whose associated cycles are pairwise Π-nonhomotopic and, by Lemma 6.2.2.1, their Π-homotopic arches intersect at most $\tilde{g}$ columns. Hence, in total, the faces f_1^* and f_2^* intersect at most $2 \times (3g+3) \times \tilde{g}$ columns.

Finally, there are at most $(3g+3) \times 2 \times [4m(\tilde{g}^2+\tilde{g}+1) \times 2\tilde{g}+4 \times (3g+3) \times \tilde{g}]$ columns $\mathcal{C}_1$ so that the faces in $\mathcal{F}_p(H)$ touch only these columns and the arches in $\mathcal{A}_P(H)$ touch at least one of these columns.

To conclude, remark then that every set of consecutive columns of H that does not contain $\mathcal{C}_1$ induces a fan H' so that $\mathcal{A}(H')$ does not contain a Π-noncontractible arch. Therefore, there is a set of at least $\frac{M}{4m(\tilde{g}+1)^4}$ consecutive columns that induce a fan H' from v so that the arches in $\mathcal{A}(H')$ are all Π-contractible. □

Lemma 6.2.2.3. *Let H be a fan with an arch from p in (G,Π) of size at least M. Then, there exists a subset of consecutive faces from $\mathcal{F}(H)$ that spans at least $\frac{M}{\tilde{g}^2}$ columns.*

Proof. Let P_{arch} be the arch path of H. By definition of $\mathcal{F}(H)$, no face in $\mathcal{F}(H)$ touches P_{arch}, but the faces in $\mathcal{F}(H)$ may not be consecutive.

First, by Lemma 6.2.2.1, each face of $\mathcal{F}_{arch}(H)$ intersect at most $\tilde{g}$ columns.

If P_{arch} has no intersection point, then $\mathcal{F}_{arch}(H)$ contains only one face and $\mathcal{F}_{arch}(H)$ intersect at most $\tilde{g}$ columns in total.

Suppose P_{arch} has 1 intersection point. Suppose that the faces in $\mathcal{F}_{arch}(H)$ touch at least $\tilde{g}^2+1$ columns, then it contains at least $\tilde{g}+1$ faces that touch pairwise disjoint columns. However, then there would be at least $\tilde{g}+1$ isolated paths on p and the intersection point of P_{arch} and we would find a contradiction with Proposition 5.2.3.

Hence, in any case, $\mathcal{F}_{arch}(H)$ touches at most $\tilde{g}^2$ disjoint columns.

Finally, it is possible to find a subset of at least $\frac{M}{\tilde{g}^2}$ consecutive columns that are touched by no face from $\mathcal{F}_{arch}(H)$. Therefore, consecutive faces from $\mathcal{F}(H)$ touch these consecutive columns. □

Lemma 6.2.2.4. *Let $1 \leq i \leq m$ and let H be a fan (with an arch) from p in (G,Π) of size at least $\frac{f(g,i)-2}{4m(\tilde{g}+1)^4}$. Suppose that every arch in $\mathcal{A}(H)$ induces, together with H, a Π-contractible subgraph of G. Then, there exists $1 \leq k \leq i$ so that there are at least $f(g,i-k)$ maximal arches from $\mathcal{A}_{max}(H)$ of size at least $f(g,k-1)+2$.*

Proof. Suppose, by contradiction, that, for every $1 \leq k \leq i$, there are less than $f(g,i-k)$ maximal arches from $\mathcal{A}_{max}(H)$ of size at least $f(g,k-1)+2$.

Remark that $\mathcal{A}_{max}$ can be partitioned into $i+1$ sets of arches $\mathcal{A}_0,...,\mathcal{A}_i$ so that, for $0 \leq k \leq i$, $A \in \mathcal{A}_{max}$ is in $\mathcal{A}_k$ if its size is in the range $f(g,k-1)+2$ to $f(g,k)+1$ (with $f(g,-1)=-1$). Remark, moreover, that, by hypothesis, $\mathcal{A}_i$ is empty.

If $i=1$, then $\mathcal{A}_{max}$ only contains maximal arches of size at most 2. However, this contradicts the definition of $\mathcal{A}_{max}$.

Hereafter, suppose $i \geq 2$. Therefore,

$$\begin{aligned}\sum_{A\in\mathcal{A}_{max}} |A| &= \sum_{k=0}^{i-1}\sum_{A\in\mathcal{A}_k} |A| \\ &= \sum_{A\in\mathcal{A}_0} |A| + \sum_{k=1}^{i-1}\sum_{A\in\mathcal{A}_k} |A| \\ &\leq 4\sum_{k=1}^{i-1} f(g,k) + \sum_{k=1}^{i-1}(f(g,i-k)-1)\times(f(g,k)+1) \\ &< 4\sum_{k=1}^{i-1} f(g,k) + \sum_{k=1}^{i-1} f(g,i-k)\times f(g,k)\end{aligned}$$

As, for $1 \leq k \leq i-1$,

$$\begin{aligned}f(g,i-k)\times f(g,k) &= \left(4\sqrt{2}m(\tilde{g}+1)^2\right)^{(i-k)^2+k^2} \\ &= \left(4\sqrt{2}m(\tilde{g}+1)^2\right)^{i^2+2k(k-i)} \\ &\leq \left(4\sqrt{2}m(\tilde{g}+1)^2\right)^{i^2-2} = \frac{f(g,i)}{32m^2(\tilde{g}+1)^4}\end{aligned}$$

and

$$\begin{aligned}4\sum_{k=1}^{i-1} f(g,k) &\leq 4\times(i-1)\times f(g,i-1) \\ &\leq (i-1)\times 4\times\left(4\sqrt{2}m(\tilde{g}+1)^2\right)^{i^2-2} = (i-1)\times\frac{4f(g,i)}{32m^2(\tilde{g}+1)^4}\end{aligned}$$

then

$$\begin{aligned}\sum_{A\in\mathcal{A}_{max}} |A| &< (i-1)\times\left(\frac{4f(g,i)}{32m^2(\tilde{g}+1)^4} + \frac{f(g,i)}{32m^2(\tilde{g}+1)^4}\right) \\ &\leq (i-1)\frac{f(g,i)-2}{4m^2(\tilde{g}+1)^4} \\ &< \frac{f(g,i)-2}{4m(\tilde{g}+1)^4}\end{aligned}$$

However, as the maximal arches of H cover all the horizontal paths of H, $\sum_{A\in\mathcal{A}_{max}} |A| \geq \frac{f(g,i)-2}{4m(\tilde{g}+1)^4}$, a contradiction. □

Lemma 6.2.2.5. *If there exists a piece p of (G, Π) such that*

$$d(p) \geq \Delta(g) = 4m(\tilde{g}+1)^4\left(4\sqrt{2}m(\tilde{g}+1)^2\right)^{m^2}$$

then G contains $m+1$ cycles that are Π-contractible well nested pinched on p.

Proof. Let p be a piece of (G, Π) of degree at least $\Delta(g)$.

Let $f(g,i) = \left(4\sqrt{2}m(\tilde{g}+1)^2\right)^{i^2}$ for $i \in \mathbb{N}$. Let us prove simultaneously the following properties by induction on i with $0 \leq i \leq m$:

1. If there is a fan H from p of size $f(g,i)$ in G, then G contains i Π-contractible well nested cycles with a column of H strictly contained in the innermost nested cycle.
2. If there is a fan with an arch H from p of size $f(g,i)$ in G, then there are i Π-contractible well nested cycles in $\text{int}(H, \Pi)$ with a column of H strictly contained in the innermost nested cycle.

First, for $i = 0$, both properties are trivial.

Then, let $1 \leq i \leq m$ and suppose that both properties are valid for every $k < i$. Let us prove each property for i.

1. Suppose that there is a fan H from p of size $f(g,i)$ in (G, Π).

   Let us prove Property (1) for i.

   Let P be the horizontal path of H.

   Remove the first and last columns of the fan H; therefore, no face in $\mathcal{F}(H)$ touches the first or last vertical path of H (except potentially in p).

   By Lemma 6.2.2.2, there exist at least $\frac{f(g,i)-2}{4m(\tilde{g}+1)^4}$ consecutive columns of H (let us call the associated fan H', P' its horizontal path) so that the arches in $\mathcal{A}(H')$ do not induce, together with H', Π-noncontractible cycles. Therefore, every arch in $\mathcal{A}(H')$ induces, together with H', a Π-contractible subgraph of G.

   By Lemma 6.2.2.4, there exists $1 \leq k \leq i$ so that there are at least $f(g, i-k)$ maximal arches from $\mathcal{A}_{max}(H')$ of size at least $f(g,k-1)+2$. Remark then that we can construct a fan H'' of size at least $f(g,i-k)$ so that each maximal arch of size at least $f(g,k-1)+2$ is contained in one of its columns. Remark also that, even though each maximal arch A has size at least $f(g,k-1)+2$, the column of a fan H'' containing A may not contain all the columns of H' that A touches. It, however, contains at least $f(g,k-1)$ of them.

   By induction hypothesis on Property (1) for H'', G contains $i-k$ Π-contractible well-nested cycles with a column of H'' strictly contained in the innermost nested cycle. Furthermore, by construction, this column contains an arch A from $\mathcal{A}_{max}(H')$ (hence empty below) of size at least $f(g,k-1)+2$. Let H_A be the fan with an arch induced by H' whose arch path is A (if A touches p, then we artificially add columns on the sides so that the fan structure touches the arch appropriately). The fan H_A has size at least $f(g,k-1)$. Therefore, by induction hypothesis on Property (2) for H_A, G contains $k-1$ Π-contractible well-nested cycles in $\text{int}(H_A, \Pi)$ with a column of H_A (which corresponds to a column of H) strictly contained in the innermost nested cycle.

   Finally, G contains i Π-contractible well-nested cycles with a column of H strictly contained in the innermost nested cycle: the $i-k$ nested cycles found in H'', the cycle associated to A and the $k-1$ nested cycles found in H_A.

2. Suppose that there is a fan with an arch H from p of size $f(g,i)$ in (G, Π).

   Let us prove Property (2) for i.

   Let P be the horizontal path of H and P_{Arch} be the arch path of H.

Remove the first and last columns of the fan H; therefore, no face in $\mathcal{F}(H)$ touches the first or last vertical path of H (except potentially in p).

By Lemma 6.2.2.3, we can restrict $\mathcal{F}(H)$ to a subset of consecutive faces from $\mathcal{F}(H)$ that spans at least $\frac{f(g,i)-2}{\tilde{g}^2} \geq \frac{f(g,i)-2}{4m(\tilde{g}+1)^4}$ columns so that this subset of faces does not touch P_{arch} (let us call the associated fan H', P' its horizontal path and P_{arch} is still its arch path).

By Lemma 6.2.2.4, there exists $1 \leq k \leq i$ so that there are at least $f(g, i-k)$ maximal arches from $\mathcal{A}_{max}(H')$ of size at least $f(g, k-1)+2$. Remark then that we can construct a fan with an arch H'', with the arch path that is a subpath of P_{Arch}, of size at least $f(g, i-k)$ so that each maximal arch of size at least $f(g, k-1)+2$ is contained in one of its columns. Remark that, even though each maximal arch A has size at least $f(g, k-1)+2$, the column of a fan H'' containing A may not contain all the columns of H' that A touches. It, however, contains $f(g, k-1)$ of them.

By induction hypothesis on Property (2) for H'', G contains $i-k$ Π-contractible well-nested cycles in $\text{int}(H'', \Pi)$ with a column of H'' strictly contained in the innermost nested cycle. Furthermore, by construction, this column contains an arch A from $\mathcal{A}_{max}(H')$ (hence empty below) of size at least $f(g, k-1)+2$. Let H_A be the fan with an arch induced by H' whose arch path is A. It has a size of at least $f(g, k-1)$. Therefore, by induction hypothesis on Property (2) for H_A, G contains $k-1$ Π-contractible well-nested cycles in $\text{int}(H_A, \Pi)$ with a column of H_A (which corresponds to a column of H) strictly contained in the innermost nested cycle.

Finally, G contains i Π-contractible well-nested cycles in $\text{int}(H, \Pi)$ with a column of H strictly contained in the innermost nested cycle: the $i-k$ nested cycles found in H'', the cycle associated to A and the $k-1$ nested cycles found in H_A.

This concludes the induction.

By Lemma 6.2.2.2, there are at least $\frac{\Delta(g)}{4m(\tilde{g}+1)^4} = f(g, m)$ columns of p that are consecutive in Π so that the faces that touch these columns induce no Π-noncontractible cycles.

Hence, these faces induce a 1-layer structure H from v of size $f(g, m)$.

Let us apply Property (1) to H with $i = m$. We obtain that G contains m Π-contractible well-nested cycles pinched on p with a column of H strictly contained in the innermost nested cycle in (G, Π). Finally, by taking the cycle on the boundary of this column as the last nested cycle, G contains $m+1$ Π-well-nested cycles pinched on p in (G, Π). □

Suppose either $\Delta(G) > \Delta(g)$ or $\Delta_F(G, \Pi) > \Delta(g)$. Then there exists a piece p of (G, Π) with $\Delta(p) > \Delta(g)$.

Then, by Lemma 6.2.2.5, (G, Π) contains $m+1$ Π-contractible well nested cycles pinched on p. However, this contradicts Corollary 5.3.15. Finally, $d(p) \leq \Delta(g)$ and therefore $\Delta(G) \leq \Delta(g)$ and $\Delta_F(G, \Pi) \leq \Delta(g)$. □

6.3 Maximum degree of a vertex and size of a face of a Π-contractible subgraph of (G, Π) with a small separator

Let $\tilde{m} = 2(\lfloor \log_q(60A(g) + 180) \rfloor + 2)$.

We prove that the maximum degree of a vertex of G_1 and the maximum size of a face in (G_1, Π) are bounded by

$$\Delta'(g) = \left(4\sqrt{2A(g)(2\tilde{m}+1)^4\tilde{m}^3}\right)^{\tilde{m}^2}$$

Similarly to the proof of Theorem 6.2.2, we proceed by contradiction and reach a contradiction by proving that a high degree of a vertex (resp. big size of a face) leads to a large number of Π-contractible nested cycles, which contradicts Corollary 5.3.16.

The proof uses mainly Corollary 5.3.16 to find contradiction, where Theorem 6.2.2 uses four different results. This highlights the success of our generalized results in Section 5.3, but is also due to restricting ourselves to the Π-contractible subgraph G_1 of (G, Π).

First, remark that, by using Corollary 5.3.16, we can show that there are at most $\tilde{m} = O(\log\log g)$ nested cycles in (G_1, Π):

Proposition 6.3.1 ([HK26a, Proposition 6.1])**.** *There are at most $\tilde{m}$ Π-well-nested cycles in (G_1, Π).*

Proof. By definition of G_1, there are at most $A(g)$ attachments on the outer cycle C_1 of G_1. Therefore, by Corollary 5.3.16, there are at most $\tilde{m} = 2(\lfloor \log_q(60A(g)+180)\rfloor + 2)$ nested cycles in (G_1, Π). □

We now prove the main result of this section:

Theorem 6.3.2 ([HK26a, Theorem 6.2])**.**

$$\Delta(G_1) \leq \Delta'(g) \quad \textit{and} \quad \Delta_F(G_1, \Pi) \leq \Delta'(g)$$

Proof. We prove several lemmas before proving the result.

Lemma 6.3.2.1. *Let p be a piece of (G_1, Π) and let H be a fan (with an arch) from p in (G_1, Π) or $H = p$. Let P be the horizontal path of H. Let $f \in \mathcal{F}(H) \cup \mathcal{F}_{arch}(H)$ ($\mathcal{F}_{arch}(H) = \emptyset$ if H is a fan without an arch or $H = p$). Then, f touches at most $2\tilde{m}+1$ columns of H.*

Proof. If $H = p$, then the face f cannot intersect three columns of p, as otherwise, p would either be a cutvertex of G or have a 2-separator that separates a Π-contractible subgraph from the rest of the graph which is in contradiction respectively with the fact that G is 2-connected and with Lemma 4.2.2.

Suppose that H is a fan (with an arch). Suppose that there is a face $f \in \mathcal{F}(H) \cup \mathcal{F}_{arch}(H)$ that touches at least $2\tilde{m}+2$ columns of H. Then, f together with H forms $\tilde{m}+1$ nested cycles ((vertex, face) if p is a vertex or (face, face) if p is a face). However, this contradicts Proposition 6.3.1. □

Lemma 6.3.2.2. *Let H be a fan p in (G_1, Π) of size at least M. Then, there exists a subset of consecutive faces from $\mathcal{F}(H)$ that does not touch $p \cup C_1$ and spans at least $\frac{M}{4A(g)(2\tilde{m}+1)^4\tilde{m}^2}$ columns.*

Proof. First, let $\mathcal{F}_{C_1}(H)$ be the set of faces in $\mathcal{F}(H)$ that touch C_1. We will show that it intersects at most $A(g) \times (2\tilde{m}+1)^3$ columns.

Suppose that $\mathcal{F}_{C_1}(H)$ intersects more than $A(g)\times(2\tilde{m}+1)^3$ columns. By Lemma 6.3.2.1, each face of $\mathcal{F}_{C_1}(H)$ intersects at most $2\tilde{m}+1$ columns. Therefore, there are more than $A(g)\times(2\tilde{m}+1)^2$ faces in $\mathcal{F}_{C_1}(H)$ that intersect pairwise disjoint columns.

Now, by the pigeonhole principle, there is a subpath P_0 of C_1 that corresponds to a facial subwalk of (G_1, Π) so that more than $(2\tilde{m}+1)^2$ faces of $\mathcal{F}_{C_1}(H)$ that intersect pairwise disjoint columns touch P_0.

However, by Proposition 6.3.1, there are at most $2\tilde{m}+1$ of them that touch the same vertex of P_0 and at most $2\tilde{m}+1$ of them that touch pairwise disjoint vertices of P_0, a contradiction.

We conclude that the faces in $\mathcal{F}_{C_1}(H)$ intersect at most $A(g) \times (2\tilde{m}+1)^3$ columns. We can therefore restrict to a fan H' that spans at least $\frac{M}{A(g)\times(2\tilde{m}+1)^3}$ columns and so that $\mathcal{F}(H')$ does not touch C_1.

Next, let $\mathcal{F}_p(H')$ be the set of faces in $\mathcal{F}(H')$ that touch p. We will show that it intersects at most $4(2\tilde{m}+1)\tilde{m}^2$ columns.

Suppose that $\mathcal{F}_p(H')$ intersects more than $4(2\tilde{m}+1)\tilde{m}^2$ columns. By Lemma 6.3.2.1, each face of $\mathcal{F}_p(H')$ intersects at most $2\tilde{m}+1$ columns. Therefore, there are more than $4\tilde{m}^2$ faces in $\mathcal{F}_p(H')$ that intersect pairwise disjoint columns.

Now, if there are more than $2\tilde{m}$ of them that touch p in the same vertex, then remark that this induces $\tilde{m}+1$ Π-well-nested cycles. Therefore, there are more than $2\tilde{m}$ of them that touch pairwise disjoint vertices of p (in that case p must be a face). However, if there are more than $2\tilde{m}$ of them that touch pairwise disjoint vertices of p, then this induces $\tilde{m}+1$ Π-well-nested cycles, a contradiction.

We conclude that the faces in $\mathcal{F}_p(H')$ intersect at most $4(2\tilde{m}+1)\tilde{m}^2$ columns. We can therefore restrict to a fan H'' that spans at least $\frac{M}{4A(g)(2\tilde{m}+1)^4\tilde{m}^2}$ columns and so that $\mathcal{F}(H'')$ does not touch C_1 nor p. □

Lemma 6.3.2.3. *Let H be a fan with an arch from p in (G_1, Π) of size at least M. Then, there exists a subset of consecutive faces from $\mathcal{F}(H)$ that spans at least $\frac{M}{(2\tilde{m}+1)^2}$ columns.*

Proof. Let P and P_{arch} be respectively the horizontal path and the arch path of H. By definition of $\mathcal{F}(H)$, no face in $\mathcal{F}(H)$ touches P_{arch}, but the faces in $\mathcal{F}(H)$ may not be consecutive on P.

First, by Lemma 6.3.2.1, each face of $\mathcal{F}_{arch}(H)$ intersects at most $2\tilde{m}+1$ columns.

If P_{arch} has no intersection point, then $\mathcal{F}_{arch}(H)$ contains only one face and $\mathcal{F}_{arch}(H)$ intersects at most $2\tilde{m}+1$ columns in total.

Suppose P_{arch} has one intersection point. Suppose that the faces in $\mathcal{F}_{arch}(H)$ touch at least $(2\tilde{m}+1)^2+1$ columns, then it contains at least $2\tilde{m}+2$ faces that touch pairwise disjoint columns. However, then there would be at least $\tilde{m}+1$ nested cycles (with p and the intersection point of P_{arch}), and we would find a contradiction with Proposition 6.3.1.

Hence, in any case, $\mathcal{F}_{arch}(H)$ touches at most $(2\tilde{m}+1)^2$ disjoint columns.

Finally, it is possible to find a subset of at least $\frac{M}{(2\tilde{m}+1)^2}$ consecutive columns that are touched by no face from $\mathcal{F}_{arch}(H)$. Therefore, consecutive faces from $\mathcal{F}(H)$ touch these consecutive columns. □

We define the following function for every $i \in \mathbb{N}$:

$$f(g,i) = \left(4\sqrt{2A(g)(2\tilde{m}+1)^4\tilde{m}^3}\right)^{i^2}$$

Lemma 6.3.2.4. *Let $1 \leq i \leq \tilde{m}$ and let H be a fan (with an arch) from p in (G_1, Π) of size at least $\frac{f(g,i)-2}{4A(g)(2\tilde{m}+1)^4\tilde{m}^2}$. Then, there exists $1 \leq k \leq i$ so that there are at least $f(g, i-k)$ maximal arches from $\mathcal{A}_{max}(H)$ of size at least $f(g,k-1)+2$.*

Proof. Suppose, by contradiction, that, for every $1 \leq k \leq i$, there are fewer than $f(g, i-k)$ maximal arches from $\mathcal{A}_{max}(H)$ of size at least $f(g, k-1)+2$.

Remark that $\mathcal{A}_{max}$ can be partitioned into $i+1$ sets of arches $\mathcal{A}_0, ..., \mathcal{A}_i$ so that, for $0 \leq k \leq i$, $A \in \mathcal{A}_{max}$ is in $\mathcal{A}_k$ if its size is in the range $f(g, k-1)+2$ to $f(g, k)+1$ (with $f(g, -1) = -1$). Remark, moreover, that, by hypothesis, $\mathcal{A}_i$ is empty.

If $i = 0$, then $\mathcal{A}_{max}$ only contains maximal arches of size at most 2. However, this contradicts the definition of $\mathcal{A}_{max}$. Therefore, we can suppose that $i \geq 1$.

Therefore,

$$\begin{aligned}
\sum_{A \in \mathcal{A}_{max}} |A| &= \sum_{k=0}^{i-1} \sum_{A \in \mathcal{A}_k} |A| \\
&= \sum_{A \in \mathcal{A}_0} |A| + \sum_{k=1}^{i-1} \sum_{A \in \mathcal{A}_k} |A| \\
&\leq 4 \sum_{k=1}^{i-1} f(g,k) + \sum_{k=1}^{i-1} (f(g, i-k) - 1) \times (f(g,k) + 1) \\
&< 4 \sum_{k=1}^{i-1} f(g,k) + \sum_{k=1}^{i-1} f(g, i-k) \times f(g,k)
\end{aligned}$$

As, for $1 \leq k \leq i-1$,

$$\begin{aligned}
f(g, i-k) \times f(g,k) &= \left(4\sqrt{2A(g)(2\tilde{m}+1)^4 \tilde{m}^3}\right)^{(i-k)^2+k^2} \\
&= \left(4\sqrt{2A(g)(2\tilde{m}+1)^4 \tilde{m}^3}\right)^{i^2+2k(k-i)} \\
&\leq \left(4\sqrt{2A(g)(2\tilde{m}+1)^4 \tilde{m}^3}\right)^{i^2-2} = \frac{f(g,i)}{32A(g)(2\tilde{m}+1)^4 \tilde{m}^3}
\end{aligned}$$

and

$$\begin{aligned}
4 \sum_{k=1}^{i-1} f(g,k) &\leq 4 \times (i-1) \times f(g, i-1) \\
&\leq (i-1) \times 4 \times \left(4\sqrt{2A(g)(2\tilde{m}+1)^4 \tilde{m}^3}\right)^{i^2-2} = (i-1) \times \frac{4f(g,i)}{32A(g)(2\tilde{m}+1)^4 \tilde{m}^3}
\end{aligned}$$

then

$$\begin{aligned}
\sum_{A \in \mathcal{A}_{max}} |A| &< (i-1) \times \left(\frac{4f(g,i)}{32A(g)(2\tilde{m}+1)^4 \tilde{m}^3} + \frac{f(g,i)}{32A(g)(2\tilde{m}+1)^4 \tilde{m}^3} \right) \\
&\leq (i-1) \frac{f(g,i)-2}{4A(g)(2\tilde{m}+1)^4 \tilde{m}^3} \\
&< \frac{f(g,i)-2}{4A(g)(2\tilde{m}+1)^4 \tilde{m}^2}
\end{aligned}$$

However, as the maximal arches of H cover all the horizontal paths of H, $\sum_{A \in \mathcal{A}_{max}} |A| \geq \frac{f(g,i)-2}{4A(g)(2\tilde{m}+1)^4 \tilde{m}^2}$, a contradiction. □

Lemma 6.3.2.5. *If there exists a piece p of (G_1, Π) such that*

$$d_{G_1}(p) \geq \Delta'(g) = \left(4\sqrt{2A(g)(2\tilde{m}+1)^4\tilde{m}^3}\right)^{\tilde{m}^2}$$

then G_1 contains $\tilde{m}+1$ cycles that are Π-contractible well nested pinched on p.

Proof. Let p be a piece of (G_1, Π) of degree or size at least $\Delta'(g)$.

Let us prove simultaneously the two following properties by induction on i with $0 \leq i \leq m$:

1. If there is a fan H from p of size $f(g,i)$ in G_1, then G_1 contains i Π-well-nested cycles with a column of H strictly contained in the innermost nested cycle.

2. If there is a fan with an arch H from p of size $f(g,i)$ in G_1, then there are i Π-well-nested cycles in $\text{int}(H, \Pi)$ with a column of H strictly contained in the innermost nested cycle.

First, for $i = 0$, both properties are trivial.

Then, let $1 \leq i \leq m$ and suppose that both properties are valid for every $k < i$. Let us prove each property for i.

1. Suppose that there is a fan H from p of size $f(g,i)$ in (G_1, Π).

   Let us prove Property (1) for i.

   Let P be the horizontal path of H.

   Remove the first and last column of the fan H; therefore, no face in $\mathcal{F}(H)$ touches the first or last vertical path of H (except potentially in p).

   By Lemma 6.3.2.2, there exist at least $\frac{f(g,i)-2}{4A(g)(2\tilde{m}+1)^4\tilde{m}^2}$ consecutive columns of H (let us call the associated fan H', and let P' be its horizontal path) so that the arches in $\mathcal{A}(H')$ do not touch either p or C_1.

   By Lemma 6.3.2.4, there exists $1 \leq k \leq i$ so that there are at least $f(g, i-k)$ maximal arches from $\mathcal{A}_{max}(H')$ of size at least $f(g,k-1)+2$. Remark then that we can construct a fan H'' of size at least $f(g,i-k)$ so that each maximal arch of size at least $f(g,k-1)+2$ is contained in one of its columns. Remark also that, even though a maximal arch A has size at least $f(g,k-1)+2$, the column of a fan H'' containing A may not contain all the columns of H' that A touches. It, however, contains at least $f(g,k-1)$ of them.

   By induction hypothesis on Property (1) for H'', G_1 contains $i-k$ Π-well-nested cycles with a column of H'' strictly contained in the innermost nested cycle. Furthermore, by construction, this column contains an arch A from $\mathcal{A}_{max}(H')$ (hence empty below) of size at least $f(g,k-1)+2$. Let H_A be the fan with an arch induced by H' whose arch path is A. The fan H_A has size at least $f(g,k-1)$. Therefore, by induction hypothesis on Property (2) for H_A, G_1 contains $k-1$ Π-well-nested cycles in $\text{int}(H_A, \Pi)$ with a column of H_A (which corresponds to a column of H) strictly contained in the innermost nested cycle.

   Finally, G_1 contains i Π-well-nested cycles with a column of H strictly contained in the innermost nested cycle: the $i-k$ nested cycles found in H'', the cycle associated to A, and the $k-1$ nested cycles found in H_A.

2. Suppose that there is a fan with an arch H from p of size $f(g,i)$ in (G_1,Π).
   Let us prove Property (2) for i.

   Let P be the horizontal path of H and P_{Arch} be the arch path of H.

   By Lemma 6.3.2.3, we can restrict $\mathcal{F}(H)$ to a subset of consecutive faces from $\mathcal{F}(H)$ that spans at least $\frac{f(g,i)}{(2\tilde{m}+1)^2} \geq \frac{f(g,i)-2}{4A(g)(2\tilde{m}+1)^4\tilde{m}^2}$ columns so that this subset of faces does not touch P_{arch} (let us call the associated fan H', let P' be its horizontal path).

   By Lemma 6.3.2.4, there exists $1 \leq k \leq i$ so that there are at least $f(g,i-k)$ maximal arches from $\mathcal{A}_{max}(H')$ of size at least $f(g,k-1)+2$. Remark then that we can construct a fan with an arch H'', with arch path P_{Arch}, of size at least $f(g,i-k)$ so that each maximal arch of size at least $f(g,k-1)+2$ is contained in one of its columns. Remark that, even though each maximal arch A has size at least $f(g,k-1)+2$, the column of a fan H'' containing A may not contain all the columns of H' that A touches. It, however, contains $f(g,k-1)$ of them.

   By induction hypothesis on Property (2) for H'', G_1 contains $i-k$ Π-well-nested cycles in $\text{int}(H,\Pi)$ with a column of H'' strictly contained in the innermost nested cycle. Furthermore, by construction, this column contains an arch A from $\mathcal{A}_{max}(H')$ (hence empty below) of size at least $f(g,k-1)+2$. Let H_A be the fan with an arch induced by H' whose arch path is A. It has a size of at least $f(g,k-1)$. Therefore, by induction hypothesis on Property (2) for H_A, G_1 contains $k-1$ Π-well-nested cycles in $\text{int}(H_A,\Pi)$ with a column of H_A (which corresponds to a column of H) strictly contained in the innermost nested cycle.

   Finally, G_1 contains i Π-well-nested cycles in $\text{int}(H,\Pi)$ with a column of H strictly contained in the innermost nested cycle: the $i-k$ nested cycles found in H'', the cycle associated to A, and the $k-1$ nested cycles found in H_A.

This concludes the induction.

By hypothesis, $d_{G_1}(p) > \Delta'(g)$. Therefore, $H = p$ is a fan of size $f(g,m)$ in G_1.

Let us apply Property (1) to H with $i = \tilde{m}$. We obtain that G_1 contains $\tilde{m}$ Π-well-nested cycles pinched on p with a column of H strictly contained in the innermost nested cycle in (G_1,Π). Finally, by taking the cycle on the boundary of this column as the last nested cycle, G_1 contains $\tilde{m}+1$ Π-well-nested cycles pinched on p in (G_1,Π). □

Suppose either $\Delta(G_1) > \Delta'(g)$ or $\Delta_F(G_1,\Pi) > \Delta'(g)$. Then there exists a piece p of (G_1,Π) with $d_{G_1}(p) > \Delta'(g)$.

Then, by Lemma 6.3.2.5, (G_1,Π) contains $\tilde{m}+1$ Π-well-nested cycles pinched on p. However, this contradicts Proposition 6.3.1. Finally, $d_{G_1}(p) \leq \Delta'(g)$ and therefore $\Delta(G_1) \leq \Delta'(g)$ and $\Delta_F(G_1,\Pi) \leq \Delta'(g)$. □

6.4 Discussion

In this section, we discuss the results presented above on bounds for the maximum degree and maximum size of a face of (G,Π). We first examine the apparent duality between vertices and faces of (G,Π) that emerges from the results of Chapter 5 and this chapter. We then discuss how tight bounds can be obtained for the maximum degree and maximum face size of (G,Π), and derive several open problems from this analysis.

6.4.1 Duality between vertices and faces in (G, Π)

As can be seen from the structural results proved in Chapter 5 (on isolated paths and nested/homotopic cycles) and the results of this chapter (on maximum degree and maximum size of a face), there is a striking symmetry between vertices and faces of (G, Π). Our results even use the notion of *piece* to encompass both vertices and faces in a single definition. In fact, the corresponding statements for vertices and faces have exactly the same form and the same bounds.

Nevertheless, this symmetry between vertices and faces of (G, Π) is not perfect. This becomes apparent when examining more closely the proofs of Proposition 5.2.3 and Lemmas 5.3.12 and 5.3.25 in Chapter 5 and those of Theorems 6.2.2 and 6.3.2 in this chapter. At several points in these proofs, it is necessary to distinguish between the case where a piece is a vertex and the case where it is a face, as the two cases do not behave in exactly the same way.

The apparent symmetry between vertices and faces of (G, Π) leads us to consider the geometric dual multigraph of (G, Π).

Let us first define the *geometric dual multigraph* of a graph with a given embedding. We use the definition from [MT01].

Let H be a connected graph with embedding $\Pi_H = (\pi, \lambda)$. We define the *geometric dual multigraph* H^* of H with respect to Π_H (or Π_H*-dual multigraph*) and its embedding $\Pi_H^* = (\pi^*, \lambda^*)$, called the *dual embedding* of Π_H^*, as follows. The vertices of H^* correspond to the Π_H-facial walks. The edges of H^* are in bijective correspondence $e \mapsto e^*$ with the edge of H, and the edge e^* joins the vertices corresponding to the Π_H-facial walks containing e (if both sides of e are in the same Π_H-facial walk, then e^* is a loop). If $W = e_1...e_k$ is a Π_H-facial walk and w its corresponding vertex in H^*, then $\pi_w^* = (e_1^*, ..., e_k^*)$. For $e^* = ww'$, we set $\lambda^*(e^*) = 1$ if and only if the Π_H-facial walks W and W' corresponding to w and w' traverse the edge e in opposite direction.

It is easy to see that H is the geometric dual multigraph of H^* with respect with Π_H. Moreover, by Euler formula, Π_H^* is an embedding in the same surface as Π_H.

The dual multigraphs of planar graphs with respect to their embeddings in the plane have been extensively studied and possess many interesting properties; see, for instance, [MT01, Section 2.6]. However, most of these properties do not carry over from graphs embedded in the plane to graphs embedded in arbitrary surfaces.

One of the few properties that does carry over is the following:

Proposition 6.4.1 ([MT01, Proposition 4.1.6])**.** *Let H be a connected Π_H-embedded graph and H^* its geometric dual multigraph with respect to Π_H. We have $(H/e)^* = H^* - e^*$ and, if e is contained in two facial walks, then $(H - e)^* = H^*/e^*$.*

At first glance, this property seems very promising for showing that the geometric dual multigraph (G^*, Π^*) of (G, Π) is also an excluded minor for the surface S'.

However, Proposition 6.4.1 does not imply that G^* is an excluded minor, because the geometric dual multigraph depends on a fixed embedding of the original graph. Indeed, for $e \in E(G)$, the graph $G - e$ is a proper minor of G and is therefore embeddable in the surface S of genus g. However, the graph $(G - e)^* = G^*/e^*$ is defined with respect to an embedding $\Pi(G - e)$, which may be an embedding of $G - e$ of genus greater than g and therefore not an embedding in S. Consequently, G^*/e^* may have the same genus as G^*, and G^* may not be an excluded minor for any surface.

To conclude, the apparent duality in our results between vertices and faces of (G, Π) is not a consequence of geometric duality for graphs on surfaces. We believe that this apparent duality is primarily due to the small separators induced by subgraphs whose only intersection with the rest of the graph consists of vertices or faces (namely, the vertices and faces on which the isolated paths or nested/homotopic cycles are pinched).

6.4.2 Toward tight bounds for the maximum degree and maximum size of a face in (G, Π)

For the quasi-polynomial bound, we were able to obtain bounds on the maximum degree and maximum face size of (G, Π) that were asymptotically better than the bound on the order of G. Indeed, we obtained a $g^{O(\log^2 g)}$ bound on the maximum degree and maximum size of a face, while the order of G was bounded by $g^{O(\log^3 g)}$.

However, for the almost quadratic bound, we did not obtain a new bound on the maximum degree and maximum face size of the whole graph (G, Π). Indeed, we only established such a bound for a 2-connected Π-contractible subgraph of G with a separator of size $O(\log^2 g)$.

Therefore, in the current state of our results, the only known bound on the maximum degree and maximum size of a face of (G, Π) is the trivial bound implied by the order of G:

$$\Delta(G) \leq |V(G)| \leq U'(g) \quad \text{and} \quad \Delta_F(G, \Pi) \leq |V(G)| \leq U'(g)$$

It might be the case that, for every surface S, there exists an excluded minor for S that has a *dominating vertex*. For a graph H, we say that $v \in V(H)$ is a *dominating vertex* of H if v is adjacent to every other vertex of H.

Problem 1. *Does there exist an excluded minor for each surface that contains a dominating vertex?*

In Appendix A, we prove that, for every surface of genus g, there exists an excluded minor H of order $L(g) = \left\lfloor \frac{7+\sqrt{1+24g}}{2} \right\rfloor + 1$. We can further show that this excluded minor is dense and, in particular, has maximum degree of the same asymptotic order as $L(g)$:

Proposition 6.4.2. *Let H be an excluded minor for a surface S_H of genus g so that $|V(H)| = L(g)$. Then,*

$$|E(H)| \geq L(g) + g - 2 \quad \textit{and} \quad \Delta(H) \geq \frac{2(L(g) + g - 2)}{L(g)}$$

Proof. Because H is an excluded minor for a surface of genus g, then H has genus at least g. Let Π_H be an embedding of H. Therefore, by Euler formula,

$$\begin{aligned} |V(H)| - |E(H)| + |F(H, \Pi_H)| &= 2 - g(H) \\ &\leq 2 - g \end{aligned}$$

Moreover, we have $2|E(H)| = \sum_{f \in F(H, \Pi_H)} |f| \geq \Delta_F(H, \Pi_H)|F(H, \Pi_H)|$.

Consequently,

$$\begin{aligned} |V(H)| - |E(H)| + \frac{2}{\Delta_F(H,\Pi_H)}|E(H)| &\leq 2-g \\ |V(H)| - \left(1 - \frac{2}{\Delta_F(H,\Pi_H)}\right)|E(H)| &\leq 2-g \\ \left(1 - \frac{2}{\Delta_F(H,\Pi_H)}\right)|E(H)| &\geq |V(H)| + g - 2 \end{aligned}$$

Finally, we have

$$\begin{aligned} |E(H)| &\geq \frac{1}{1 - \frac{2}{\Delta_F(H,\Pi_H)}}(|V(H)| + g - 2) \\ &\geq |V(H)| + g - 2 \\ &\geq L(g) + g - 2 \end{aligned}$$

To conclude, we indeed find that $|E(H)| \geq L(g) + g - 2$.

Now, let us prove the lower bound on the maximum degree of H. The average degree of H is

$$\frac{2|E(H)|}{|V(H)|} = \frac{2(L(g) + g - 2)}{L(g)}$$

Therefore, there is a vertex of H that has degree at least the average degree $\frac{2(L(g)+g-2)}{L(g)}$. This concludes the proof. □

A consequence of Proposition 6.4.2 is that there exists a vertex of H with degree at least $\frac{2(L(g)+g-2)}{L(g)} = \Theta(\sqrt{g})$, which is of the same asymptotic order as $|V(H)| = L(g)$.

We can also ask whether it is possible to obtain a bound on the maximum degree that is asymptotically smaller than the order of the excluded minor:

Problem 2. *Let $N(g)$ be the maximum order over all the excluded minors for surfaces of genus g. Does there exist a function f such that $\frac{f(g)}{N(g)} = o(1)$ and, for every excluded minor H for a given surface of genus g, $\Delta(H) \leq f(g)$?*

It could be the case that excluded minors of large order do not have a dominating vertex and even have maximum degree that is small compared with their order, whereas excluded minors of small order have maximum degree of the same asymptotic order as their order.

For $k \geq 0$, let H_k be the graph obtained from $k+1$ copies of K_5 by gluing them together at a common vertex (see Figure 6.7 for an illustration). It is then easy to show that H_k is an excluded minor for a surface of genus k, with a proof similar to that of Proposition 2.2.10. Moreover, H_k has a dominating vertex and therefore has maximum degree $4k$.

This gives the following lower bound on the function f defined in Problem 2:

$$f(g) \geq 4g$$

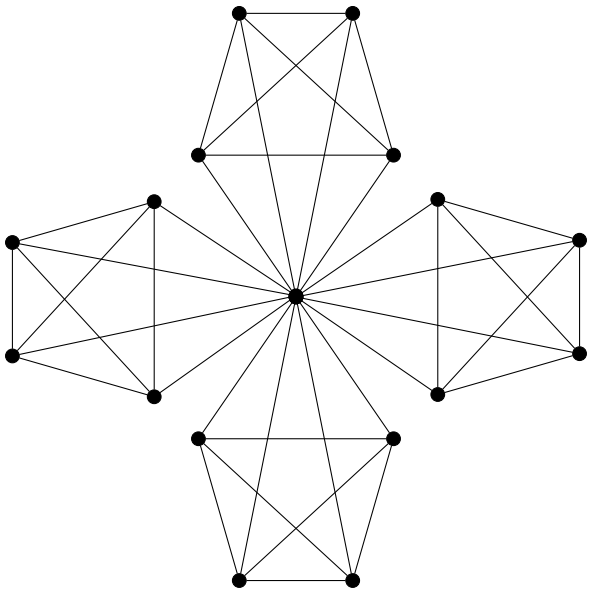


Figure 6.7: The graph H_4

Chapter 7

Bounds on the order of the excluded minors for surfaces

7.1 Balanced separator decompositions 126

7.2 Finding a subgraph G_0 of (G, Π) with Π-contractible bridges 129

7.3 Bounding the order of a 2-connected Π-contractible subgraph of (G, Π) 130

7.4 A quasi-polynomial bound for the order of the excluded minors for surfaces 132

7.5 An almost quadratic bound for the order of the excluded minors for surfaces 135

In this chapter, we finally establish bounds for the order of the excluded minors for surfaces. We prove both a quasi-polynomial and an almost quadratic bound in g for the excluded minors for surfaces.

Even though the almost quadratic bound is a strict improvement on the quasi-polynomial, we chose to present both results. Presenting both results gives a better understanding of the reasoning process that we followed to establish the current state-of-the-art result and it also allows us to compare the methods used in each proof.

The content presented here is mainly extracted from our second paper [HK26a]: in this paper, we presented a new method to split the graph G into several subgraphs of balanced order using tree decomposition, called balanced separator decomposition. It relies on the celebrated balanced separator theorem from Robertson and Seymour [RS86a]. The former method to split G, presented in [HK26b] was less elaborated and consisted in cutting equal size paths in the tree of the tree decomposition of G.

Compared to the former method from [HK26b], this method is more efficient because it produces tighter bounds for the order of G.

In Section 7.1, we first present a few results about balanced separator decompositions that directly build on the work on balanced separators from Robertson and Seymour [RS86a] and tailor this tool to our needs.

Then, in Section 7.2, we first show that there is a subgraph G_0 of G that is separated from the rest of G by a separator A of size $O(g \log^2 g)$ and so that every bridge on A in G_0 is Π-contractible. Moreover, the order of G is bounded by $|V(G_0)| \times Q(g)$ with Q some nearly linear function of g.

Next, in Section 7.3, we show that the order of a 2-connected Π-contractible subgraph G' of G is bounded by some function $P_{G'}(g)$ which depends on the bounds on the maximum degree and maximum size of a face in (G', Π), on the number of Π-well-nested cycles in (G', Π) and on the separator of G' from the rest of G.

In Section 7.4, we establish the quasi-polynomial bound on the order of G by combining Propositions 7.2.1 and 7.3.1 and conclude for all excluded minors for surfaces by using Corollary 4.2.4.

Finally, in Section 7.5, we establish the almost quadratic bound on the order of G by combining Propositions 7.2.1, 7.3.1 and 7.5.1 and conclude for all excluded minors for surfaces by using Corollary 4.2.4.

7.1 Balanced separator decompositions

In this section, we make use of the celebrated balanced separator theorem from Robertson and Seymour [RS86a] to derive a balanced separator decomposition for any graph so that the intersection of a set of the decomposition with the other sets is small (at most the product of the treewidth of the graph and a function logarithmic in the size of the decomposition).

For a tree T, we define a 1-*separation sequence* of T to be a collection $(T_1, ..., T_m)$ of subtrees of T ($m \geq 2$) such that $T_1 \cup ... \cup T_m = T$ and $|T_i \cap T_j| \leq 1$ for all $i \neq j$. In [RS86a], Robertson and Seymour proved the following theorem, known as the balanced separator theorem:

Theorem 7.1.1 (Balanced separator theorem, [RS86a, (2.6)])**.** *Let H be a graph and let $(T, (V_t)_{t \in T})$ be a tree decomposition of H. Then, there is a 1-separation (T_1, T_2) of T such that*

$$\frac{1}{3}|V(H) - V_{t_0}| \leq |\bigcup_{t \in T_1} V_t - V_{t_0}| \leq \frac{2}{3}|V(H) - V_{t_0}|$$

and

$$\frac{1}{3}|V(H) - V_{t_0}| \leq |\bigcup_{t \in T_2} V_t - V_{t_0}| \leq \frac{2}{3}|V(H) - V_{t_0}|$$

with $T_1 \cap T_2 = \{t_0\}$.

We first prove the following variant of Theorem 7.1.1:

Corollary 7.1.2 ([HK26a, Corollary 7.2])**.** *Let H be a graph and let $(T, (V_t)_{t \in T})$ be a tree decomposition of H. Suppose that $\max_{t \in T} |V_t| \leq \frac{1}{A}|V(H)|$ for some constant $A > 1$. Then, there is a 1-separation (T_1, T_2) of T such that*

$$\frac{1}{3}|V(H)| \leq |\bigcup_{t \in T_1} V_t| \leq (\frac{2}{3} + \frac{1}{3A})|V(H)|$$

and

$$\frac{1}{3}|V(H)| \leq |\bigcup_{t \in T_2} V_t| \leq (\frac{2}{3} + \frac{1}{3A})|V(H)|$$

Proof. By Theorem 7.1.1, we easily get that $|\bigcup_{t\in T_i} V_t| \geq \frac{1}{3}|V(H)|$ for $i = 1, 2$. Indeed, for $i = 1, 2$ and $T_1 \cap T_2 = \{t_0\}$, we have

$$\begin{aligned}
\frac{1}{3}|V(H) - V_{t_0}| &\leq |\bigcup_{t\in T_i} V_t - V_{t_0}| \\
\frac{1}{3}(|V(H)| - |V_{t_0}|) &\leq |\bigcup_{t\in T_i} V_t| - |V_{t_0}| \\
\frac{1}{3}|V(H)| + \frac{2}{3}|V_{t_0}| &\leq |\bigcup_{t\in T_i} V_t| \\
\frac{1}{3}|V(H)| &\leq |\bigcup_{t\in T_i} V_t|
\end{aligned}$$

Moreover, by using again Theorem 7.1.1, we get that $|\bigcup_{t\in T_i} V_t| \leq (\frac{2}{3}+\frac{1}{3A})|V(H)|$ for $i = 1, 2$. Indeed, for $i = 1, 2$ and $T_1 \cap T_2 = \{t_0\}$, we have

$$\begin{aligned}
|\bigcup_{t\in T_i} V_t - V_{t_0}| &\leq \frac{2}{3}|V(H) - V_{t_0}| \\
|\bigcup_{t\in T_i} V_t| - |V_{t_0}| &\leq \frac{2}{3}(|V(H)| - |V_{t_0}|) \\
|\bigcup_{t\in T_i} V_t| &\leq \frac{2}{3}|V(H)| + \frac{1}{3}|V_{t_0}| \\
|\bigcup_{t\in T_i} V_t| &\leq \frac{2}{3}|V(H)| + \frac{1}{3A}|V(H)| \\
|\bigcup_{t\in T_i} V_t| &\leq (\frac{2}{3} + \frac{1}{3A})|V(H)|
\end{aligned}$$

□

Then, we extend Corollary 7.1.2 to a 1-separation of T of arbitrary size:

Lemma 7.1.3 ([HK26a, Lemma 7.3])**.** *Let H be a graph and let $(T, (V_t)_{t\in T})$ be a tree decomposition of H. Suppose that $\max_{t\in T} |V_t| \leq \frac{1}{4k}|V(H)|$. For $k \in \mathbb{N}^*$, there is a 1-separation sequence $(T_1, ..., T_k)$ of T such that, for every $1 \leq i, j \leq k$,*

1. *$|\bigcup_{t\in T_i} V_t| \leq 3|\bigcup_{t\in T_j} V_t|$,*
2. *$V(T_i) \cap (\bigcup_{1\leq j\neq i\leq k} V(T_j))$ has size at most $\lfloor \log_{\frac{4}{3}} 3k \rfloor$.*

We call $(\bigcup_{t\in T_1} V_t, ..., \bigcup_{t\in T_k} V_t)$ a balanced separator decomposition of H.

Proof. We prove the first property by induction on k.

For $k = 1$, this is clear.

Now, suppose that the property is true for some $k \in \mathbb{N}^*$, let us prove that it is true for $k + 1$. Let H be a graph and let $(T, (V_t)_{t\in T})$ be a tree decomposition of H. By induction

hypothesis, there is a 1-separation sequence $(T_1, ..., T_k)$ of T such that, for every $1 \leq i, j \leq k$, $|\bigcup_{t\in T_i} V_t| \leq 3|\bigcup_{t\in T_j} V_t|$ and such that, for $1 \leq i \leq k$, $V(T_i) \cap (\bigcup_{1\leq j\neq i\leq k} V(T_j))$ has size at most $\lfloor \log_{\frac{4}{3}} 3k \rfloor$. Suppose that the order $T_1, ..., T_k$ is chosen so that $|\bigcup_{t\in T_i} V_t| \leq |\bigcup_{t\in T_j} V_t|$ if $i \leq j$.

Then, remark that $\max_{t\in T_k} |V_t| \leq \max_{t\in T} |V_t| \leq \frac{1}{4k}|V(H)| \leq \frac{1}{4}|\bigcup_{t\in T_k} V_t|$. Consequently, by Corollary 7.1.2, there is a 1-separation (T'_k, T''_k) of T_k such that

$$\frac{1}{3}|\bigcup_{t\in T_k} V_t| \leq |\bigcup_{t\in T'_k} V_t| \leq \frac{3}{4}|\bigcup_{t\in T_k} V_t| \quad \text{and} \quad \frac{1}{3}|\bigcup_{t\in T_k} V_t| \leq |\bigcup_{t\in T''_k} V_t| \leq \frac{3}{4}|\bigcup_{t\in T_k} V_t|$$

Let us show that the 1-separation sequence $(T_1, T_2, ..., T_{k-1}, T'_k, T''_k)$ of T has the needed property. It suffices to verify it for T_i $(1 \leq i \leq k-1)$ and T'_k without loss of generality. We indeed have that

$$|\bigcup_{t\in T_i} V_t| \leq |\bigcup_{t\in T_k} V_t| \leq 3|\bigcup_{t\in T'_k} V_t| \quad \text{and} \quad |\bigcup_{t\in T'_k} V_t| \leq |\bigcup_{t\in T_k} V_t| \leq 3|\bigcup_{t\in T_i} V_t|$$

This concludes the proof of the first property.

Now, we prove the second property. Suppose that there exists $1 \leq i \leq k$ such that

$$V(T_i) \cap (\bigcup_{1\leq j\neq i\leq k} V(T_j)) \text{ has size } n > \lfloor \log_{\frac{4}{3}} 3k \rfloor$$

Remark that by construction, T_i is obtained after a sequence of n 1-separations:

$$T, U_1, U_2, U_3, ..., U_n = T_i$$

where, for all $1 \leq i \leq n-1$, $U_{i+1} \subseteq U_i$ and $|\bigcup_{t\in U_{i+1}} V_t| \leq \frac{3}{4}|\bigcup_{t\in U_i} V_t|$ (according to Corollary 7.1.2).

Therefore, we get that

$$|\bigcup_{t\in T_i} V_t| \leq \left(\frac{3}{4}\right)^n |\bigcup_{t\in T} V_t| = \left(\frac{3}{4}\right)^n |V(H)|.$$

Moreover, by the first property, we have that

$$|\bigcup_{t\in T_i} V_t| \geq \frac{1}{3}|\bigcup_{t\in T_k} V_t| \quad \text{and} \quad k \times |\bigcup_{t\in T_k} V_t| \geq \sum_{j=1}^{k} |\bigcup_{t\in T_j} V_t| \geq |V(H)|$$

Therefore,

$$|\bigcup_{t\in T_i} V_t| \geq \frac{|V(H)|}{3k}$$

Finally, we have $(\frac{3}{4})^n \geq \frac{1}{3k}$, which is a contradiction because $n > \lfloor \log_{\frac{4}{3}} 3k \rfloor$, and therefore $(\frac{3}{4})^n < \frac{1}{3k}$. □

7.2 Finding a subgraph G_0 of (G, Π) with Π-contractible bridges

We first prove that there is a subgraph G_0 of G that is separated from the rest of G by a separator A of size $O(g \log^2 g)$ and so that every bridge on A in G_0 is Π-contractible. Moreover, the order of G is bounded by $|V(G_0)| \times Q(g)$ with Q some nearly linear function of g.

Let q and m be the values defined in Corollary 5.3.15, that is to say $q = \frac{9073}{9072}$ and $m = 2(\lfloor \log_q(3g+4) \rfloor + 2)$. Moreover, let $T(g) = 264(g+2)(m+1) - 1$ be as defined in Section 5.4.

Let $m' = \lfloor \log_{\frac{4}{3}} 3(4m^2(3g+3)+1) \rfloor$.

Proposition 7.2.1 ([HK26a, Proposition 7.6])**.** *Suppose that* $|V(G)| \geq 4(4m^2(3g+3)+1)(tw(G)+1)$*. Then, there exists a separation* (G_0, G_1) *of* G *such that* $A = V(G_0) \cap V(G_1)$ *has size at most* $(tw(G)+1) \times m'$*, every bridge on* A *in* G_0 *is* Π*-contractible and*

$$|V(G)| \leq Q(g) \times |V(G_0)|$$

with $Q(g) = 3(4m^2(3g+3)+1)$.

Proof. Let $(T, (V_t)_{t \in T})$ be a tree decomposition of G of width $tw(G)$. By Lemma 7.1.3, as $tw(G)+1 \leq \frac{1}{4(4m^2(3g+3)+1)}|V(G)|$, there is a separation $(T_1, ..., T_k)$ of T with:

- $k = 4m^2(3g+3)+1$,
- for every $1 \leq i, j \leq k$, $|\bigcup_{t \in T_i} V_t| \leq 3|\bigcup_{t \in T_j} V_t|$,
- for every $1 \leq i \leq k$, $V(T_i) \cap (\bigcup_{1 \leq j \neq i \leq k} V(T_j))$ has size at most $\lfloor \log_{\frac{4}{3}} 3k \rfloor = m'$.

For $1 \leq i \leq k$, let $G_i = \bigcup_{t \in T_i} V_t$ and let $A_i = \bigcup_{j \neq i} V(G_i) \cap V(G_j)$. Moreover, for $1 \leq i \leq k$, let G_i' be the graph with vertex set $V(G_i)$ and edge set all the edges in $E(G_i)$ with at least one endpoint in $V(G_i) - A_i$.

Let $1 \leq i \leq k$. We define $(H_i^j)_{1 \leq j \leq \ell_i}$ to be the connected components of $G_i - A_i$. Moreover, for $1 \leq j \leq \ell_i$, we define $\overline{H}_i^j$ to be the graph induced by H_i^j and the vertices in A_i that have a neighbor in H_i^j.

Let us first prove the following claim on $\overline{H}_i^j$:

Claim 7.2.1.1. *Let* $1 \leq i \leq k$ *and* $1 \leq j \leq \ell_i$*. If the graph* $\overline{H}_i^j$ *contains a* Π*-noncontractible cycle, then it contains a* Π*-noncontractible cycle with at most one vertex in* A_i*.*

Proof of the claim. Let $1 \leq i \leq k$ and $1 \leq j \leq \ell_i$. Suppose that $\overline{H}_i^j$ contains a Π-noncontractible cycle C. Suppose, moreover, that we selected C so that it has the fewest vertices in A_i.

If C contains at most one vertex in A_i, then there is nothing to prove.

Suppose therefore that C contains at least two distinct vertices u, v in A_i. Let P and P' be the two distinct subpaths of C from u to v.

As H_i^j is connected, there is a path P'' from the interior of P to the interior of P' in H_i^j. Then $P'' \cup C$ generates two cycles C' and C'' distinct from C and, by Lemma 2.1.14, at least one of them (say C') is Π-noncontractible. The cycle C' has fewer vertices in A_i than C, a contradiction. ■

Let us prove that there exists $1 \leq i \leq k$ so that, for every $1 \leq j \leq \ell_i$, $\overline{H}_i^j$ is Π-contractible. By contradiction, suppose that, for $1 \leq i \leq k$, there exists $1 \leq j_i \leq \ell_i$ so that the graph $\overline{H}_i^{j_i}$

contains a Π-noncontractible cycle. Then, by Claim 7.2.1.1, there exists a cycle C_i of $\overline{H}_i^{j_i}$ that is Π-noncontractible and contains at most one vertex in A_i. Remark that the cycles $(C_i)_{1\leq i\leq k}$ are almost disjoint.

By Corollary 5.4.6, there are at most $4m^2(3g+3)$ almost disjoint Π-noncontractible cycles in (G,Π), which leads to a contradiction. Therefore, there exists $1 \leq i_0 \leq k$ so that, for every $1 \leq j \leq \ell_{i_0}$, $\overline{H}_{i_0}^j$ is Π-contractible.

Remark that, for $1 \leq j \leq \ell_{i_0}$, the subgraph $\overline{H}_{i_0}^j$ is separated from the rest of the graph by a subset of vertices in A_{i_0} and that, by definition, A_{i_0} has size at most $(tw(G)+1) \times m'$.

Finally, we obtain

$$\begin{aligned}
|V(G)| &\leq \sum_{j=1}^{k} |\bigcup_{t\in T_j} V_t| \\
&\leq 3k|\bigcup_{t\in T_{i_0}} V_t| \\
&= 3(4m^2(3g+3)+1) \times |V(G'_{i_0})| \\
&= Q(g) \times |V(G'_{i_0})|
\end{aligned}$$

By choosing $G_0 = G'_{i_0}$ and $G_1 = G - G'_{i_0}$, we indeed get a separation (G_0, G_1) of G such that $A = A_i = V(G_0) \cap V(G_1)$ has size at most $(tw(G)+1) \times m'$, and every bridge on A in G_0 is Π-contractible with the needed relationship between $|V(G)|$ and $|V(G_0)|$. □

7.3 Bounding the order of a 2-connected Π-contractible subgraph of (G,Π)

Let G' be a 2-connected Π-contractible subgraph of (G,Π). We define the following values with respect to G':

- let $\Delta_{G'}(g)$ be a bound on the degree of G' and on the maximum size of a face in (G',Π),
- let $m_{G'}$ be a bound on the number of Π-well-nested cycles in (G',Π),
- and let $A_{G'}(g)$ be a bound on the size of the separator of G' from the rest of the graph G.

We will show that the order of G' is bounded by the function

$$P_{G'}(g) = \frac{\Delta_{G'}(g)(\Delta_{G'}(g)^{2m_{G'}}-1)}{\Delta_{G'}(g)-1} \times A_{G'}(g)$$

This result will be used to prove both the quasi-polynomial bound and the almost quadratic bound on the order of the excluded minors for surfaces. It consists of a generalization of the two versions of this result that appear in [HK26b] and [HK26a].

To prove the quasi-polynomial bound, we will use this result with $G' = D$ being a 2-connected component of a bridge on A in G_0, as defined in Proposition 7.4.1. Therefore, in that case, we have $\Delta_{G'}(g) = \Delta(g)$, $m_{G'} = m$ and $A_{G'}(g) = 6(T(g)+1) \times m'$.

Finally, the order of D is bounded by:

$$P_D(g) = \frac{\Delta(g)(\Delta(g)^{2m}-1)}{\Delta(g)-1} \times 6(T(g)+1) \times m'$$

To prove the almost quadratic bound, we will use this result with $G' = G_1$ being a 2-connected Π-contractible subgraph of G with a separator of size $A(g) = O(\log g)$, as defined in Chapter 6. Therefore, in that case, we have $\Delta_{G'}(g) = \Delta'(g)$, $m_{G'} = \tilde{m}$ and $A_{G'}(g) = A(g)$.

Finally, the order of G_1 is bounded by:

$$P_{G_1}(g) = \frac{\Delta'(g)(\Delta'(g)^{2\tilde{m}} - 1)}{\Delta'(g) - 1} \times A(g)$$

Let C be a Π-contractible cycle. We define $\mathcal{B}(C)$ to be the faces in $\mathrm{Int}(C, \Pi)$ that touch C.

Proposition 7.3.1. *We have*

$$|V(G')| \leq P_{G'}(g)$$

Proof. For a cycle C in G', let $E(C)^+$ be the set of edges in $\mathrm{int}(C, \Pi)$ incident to C. By hypothesis, every vertex of C is adjacent to at most $\Delta_{G'}(g)$ edges in G', therefore $|E(C)^+| \leq \Delta_{G'}(g) \times |V(C)|$. Let C be a cycle in G' and let $\mathcal{C}(C)$ be the cycles so that $\mathcal{B}(C)$ contains exactly the faces inside $\mathrm{Int}(C, \Pi) - \bigcup_{C' \in \mathcal{C}(C)} \mathrm{Int}(C', \Pi)$.

Claim 7.3.1.1. *Let C be a cycle in G', then*

$$\sum_{C' \in \mathcal{C}(C)} |V(C')| \leq \Delta_{G'}(g)^2 \times |V(C)|$$

Proof of the claim. Let H_C be the graph containing all the bridges inside $\mathrm{int}(C, \Pi) - \bigcup_{C' \in \mathcal{C}(C)} \mathrm{Int}(C', \Pi)$. Let $A(C)$ be the set of vertices in H_C that have attachments on some cycle in $\mathcal{C}(C)$.

Let us first prove that $A(C)$ is of order at most $\Delta_{G'}(g) \times |V(C)|$. Remark that, by construction, there are fewer edges with an endpoint in $\bigcup_{C' \in \mathcal{C}(C)} C'$ in H_C than to C in H_C. Hence, we have $|A(C)| \leq |E(C)^+| \leq \Delta_{G'}(g) \times |V(C)|$.

Let P be a path on a cycle C' of $\mathcal{C}(C)$ between two consecutive attached vertices $a, a' \in A(C')$ in H_C. Then, a Π-facial walk contains P. Therefore, by hypothesis, $|P| \leq \Delta_{G'}(g)$.

Moreover, as $|A(C)| \leq \Delta_{G'}(g) \times |V(C)|$, $\bigcup_{C' \in \mathcal{C}(C)} C'$ contains at most $\Delta_{G'}(g) \times |V(C)|$ such paths. Finally, $\sum_{C' \in \mathcal{C}(C)} |V(C')| \leq \Delta_{G'}(g)^2 \times |V(C)|$. ■

By hypothesis, the radius of G' with embedding $\Pi(G')$ is at most $m_{G'}$. Let us partition the faces of $(G', \Pi(G'))$ into their radius class with respect to the outer face boundary C_1: for $0 \leq i \leq m_{G'} - 1$, let $\mathcal{F}_i$ be the set of faces of $(G', \Pi(G'))$ of radius $i+1$. Let us define inductively a set of cycles $\mathcal{C}_i$ so that $\bigcup_{C \in \mathcal{C}_i} \mathcal{B}(C) = \mathcal{F}_i$ and $\sum_{C \in \mathcal{C}_i} |V(C)| \leq |V(C_1)| \times \Delta_{G'}(g)^{2i}$.

We define $\mathcal{C}_0 = \{C_1\}$. We have indeed that $\mathcal{B}(C_1) = \mathcal{F}_0$ and $|V(C_1)| = |V(C_1)| \times 1$. Now, suppose that for $0 \leq i \leq m_{G'} - 1$, $\mathcal{C}_i$ is defined and is so that $\bigcup_{C \in \mathcal{C}_i} \mathcal{B}(C) = \mathcal{F}_i$ and $\sum_{C \in \mathcal{C}_i} |V(C)| \leq |V(C_1)| \times \Delta_{G'}(g)^{2i}$. Let us define $\mathcal{C}_{i+1}$. For $C \in \mathcal{C}_i$, let $\mathcal{C}(C)$ be the set of cycles so that $\mathcal{B}(C)$ contains exactly the faces inside $\mathrm{Int}(C, \Pi) - \bigcup_{C' \in \mathcal{C}(C)} \mathrm{Int}(C', \Pi)$. Then, by the Claim 7.3.1.1, $\sum_{C' \in \mathcal{C}(C)} |V(C')| \leq \Delta_{G'}(g)^2 \times |V(C)|$. We define $\mathcal{C}_{i+1} = \bigcup_{C \in \mathcal{C}_i} \bigcup_{C' \in \mathcal{C}(C)} C'$.

Then, $\sum_{C \in \mathcal{C}_{i+1}} |V(C)| = \sum_{C \in \mathcal{C}_i} \sum_{C' \in \mathcal{C}(C)} |V(C')| \leq \sum_{C \in \mathcal{C}_i} \Delta_{G'}(g)^2 \times |V(C)| \leq |V(C_1)| \times \Delta_{G'}(g)^{2i+2}$. Moreover, it is easy to verify that $\bigcup_{C \in C_{i+1}} \mathcal{B}(C)$ is exactly $\mathcal{F}_{i+1}$.

In Figure 5.10, the graph represented has radius 3, and the faces of radius 1, 2, and 3 are respectively in blue, pink, and green. Moreover, we have $\mathcal{C}_0 = \{C_0\}$, $\mathcal{C}_1 = \{C_{1,1}, C_{1,2}, C_{1,3}\}$ and $\mathcal{C}_2 = \{C_{2,1}, C_{2,2}, C_{2,3}\}$.

Let $0 \leq i \leq m_{G'} - 2$. Remark that, by construction, the set $\mathcal{B}(C) \subseteq \mathcal{F}_i$ for $C \in \mathcal{C}_i$ induces a set of bridges $\mathcal{H}_C$ on C and the cycles of $\mathcal{C}_{i+1}$. By construction, each bridge in $\mathcal{H}_C$ is a tree whose leaves are on C and whose roots are on some cycles in $\mathcal{C}_{i+1}$ (two distinct leaves or roots of these trees can lie on the same vertex of C). For $i = m_{G'} - 1$ and $C \in \mathcal{C}_{m_{G'}-1}$, $\mathcal{B}(C)$ induces a set of bridges $\mathcal{H}_C$, and these bridges are trees whose leaves are on C (two distinct leaves of these trees can lie on the same vertex of C).

Claim 7.3.1.2. *For $0 \leq i \leq m_{G'} - 1$ and $C \in \mathcal{C}_i$, $V(C) \cup \bigcup_{B \in \mathcal{H}_C} V(B)$ has size at most $(\Delta_{G'}(g) + 1)|V(C)|$.*

Proof of the claim. The trees in $\mathcal{H}_C$ have in total at most $|E(C)^+|$ leaves. Moreover, by minimality of G, there is no induced path of length 2 in any tree in $\mathcal{H}_C$. Therefore, these trees have more leaves than internal vertices by Lemma 2.1.3.

Finally,

$$\sum_{B \in \mathcal{H}_C} |V(B)| \leq |E(C)^+|$$

and

$$|V(C)| + \sum_{B \in \mathcal{H}_C} |V(B)| \leq |V(C)| + |E(C)^+| \leq (\Delta_{G'}(g) + 1) \times |V(C)|$$

■

Remark that the set of vertices of C_1 on which edges of $\mathrm{Ext}(C_1, \Pi)$ have an endpoint is exactly A_1. Let P be a path on C_1 between two consecutive attached vertices $a, a' \in A_1$, then there is a Π-facial walk that contains P. Therefore, by hypothesis, $|P| \leq \Delta_{G'}(g)$. Moreover, as $|A_1| \leq A_{G'}(g)$, C_1 contains at most $A_{G'}(g)$ such paths. Finally, C_1 has size at most $A_{G'}(g) \times \Delta_{G'}(g)$.

By using Claim 7.3.1.2 and the inequality $\sum_{C \in \mathcal{C}_i} |V(C)| \leq |V(C_1)| \times \Delta_{G'}(g)^{2i}$ for $0 \leq i \leq m_{G'} - 1$, we obtain that the order of G' is bounded by

$$\begin{aligned}
\sum_{i=0}^{m_{G'}-1} \sum_{C \in \mathcal{C}_i} |V(C \cup \bigcup_{B \in \mathcal{H}_C} B)| &= \sum_{i=0}^{m_{G'}-1} \sum_{C \in \mathcal{C}_i} (\Delta_{G'}(g) + 1)|V(C)| \\
&= \sum_{i=0}^{m_{G'}-1} (\Delta_{G'}(g) + 1)\Delta_{G'}(g)^{2i}|V(C_1)| \\
&\leq (\Delta_{G'}(g) + 1)\frac{\Delta_{G'}(g)^{2m_{G'}} - 1}{\Delta_{G'}(g)^2 - 1}|V(C_1)| \\
&\leq \frac{\Delta_{G'}(g)^{2m_{G'}} - 1}{\Delta_{G'}(g) - 1}|V(C_1)| \\
&\leq \frac{\Delta_{G'}(g)(\Delta_{G'}(g)^{2m_{G'}} - 1)}{\Delta_{G'}(g) - 1} \times A_{G'}(g).
\end{aligned}$$

Finally, G' is of order at most $\frac{\Delta_{G'}(g)(\Delta_{G'}(g)^{2m_{G'}}-1)}{\Delta_{G'}(g)-1} \times A_{G'}(g) = P_{G'}(g)$. □

7.4 A quasi-polynomial bound for the order of the excluded minors for surfaces

In this section, we show a first bound on the order of the excluded minors for surfaces which is quasi-polynomial in g. The method used here to prove the quasi-polynomial bound is adapted

from [HK26a] and was not the one originally used in [HK26b]. This new method leads to the same asymptotic bound but has a less intricate proof.

First, we show that the order of G, which is an arbitrary 2-connected excluded minor for the surface S' of genus g, is bounded by a quasi-polynomial in g:

Proposition 7.4.1. *We have*

$$|V(G)| \leq U(g) = Q(g) \times S(g)$$

with Q defined in Section 7.2 and

$$S(g) = 10((T(g)+1) \times m' + g)((T(g)+1) \times m')P_D(g)$$

Proof. By Proposition 7.2.1, we obtain that either $|V(G)| \leq 4(4m^2(3g+3)+1) \leq U(g)$ or that there exists a subgraph G_0 of G such that G_0 is separated from the rest of the graph by a separator A of size $tw(G) \times m'$, such that every bridge on A in G_0 is Π-contractible and such that $|V(G)| \leq Q(g) \times |V(G_0)|$.

In that second case, we first bound the number of bridges and the number of 2-connected components in each bridge.

Let $(B_i)_{1 \leq i \leq \ell}$ be the bridges on A in G_0 together with their attaches in A. We bound the number of bridges ℓ:

Claim 7.4.1.1.

$$\ell \leq 2((tw(G)+1) \times m' + g)$$

Proof of the claim. Let us associate to each $1 \leq i \leq \ell$ three vertices A_i in A so that $A_i \subseteq B_i$, which is always possible because B_i has at least three attaches in A.

Let us consider the bipartite graph B with vertex set $A \cup [1, \ell]$ and so that, for $a \in A$ and $1 \leq i \leq \ell$, we put an edge between a and i if and only if $a \in A_i$.

Remark that B is a minor of $G_0 \subseteq G$, therefore B has genus at most $g(G) \leq g+2$. Let $\mathcal{F}$ be the set $F(B, \Pi(B))$ of faces of B in $\Pi(B)$. Remark that $2|E(B)| = \sum_{f \in \mathcal{F}} |f| \geq 4|\mathcal{F}|$ because B is bipartite.

By Euler's formula, we get:

$$\begin{aligned} |V(B)| - |E(B)| + |F| &\geq 2 - (g+2) \\ |V(B)| - |E(B)| + \frac{1}{2}|E(B)| &\geq -g \\ |E(B)| &\leq 2(|V(B)| + g) \\ 3\ell &\leq 2(|A| + \ell + g) \\ \ell &\leq 2((tw(G)+1) \times m' + g) \end{aligned}$$

■

Let $1 \leq i \leq \ell$, we now bound the number of 2-connected components $(D^i_j)_{1 \leq j \leq p_i}$ of B_i and, for every $1 \leq j \leq p_i$, we establish a bound on the size of the separator of D^i_j from the rest of the graph:

Claim 7.4.1.2. *$p \leq 5|A|$ and, for every $1 \leq j \leq p_i$, D^i_j has a separator of size at most $6|A|$.*

Proof of the claim. Let F_i be the block-cutvertex forest of B_i, that is to say $V(F_i)$ consists of the cutvertices of B_i together with the 2-connected components $(D_j^i)_{1\leq j\leq p_i}$ of B_i and a connected component is incident to all the cutvertices it contains.

For $1 \leq j \leq p_i$, D_j^i is separated from the rest of G by attachments in $A \cap D_j^i$ and by the cutvertices that separate it from the other subgraphs in $(D_j^i)_{1\leq j\leq p_i}$.

In particular, each subgraph D_j^i corresponding to a leaf of F_i is separated from the rest of G by attachments in $A \cap D_j^i$ and by at most one cutvertex. Because G is 2-connected, D_j^i contains at least one attachment of A that is distinct from the cutvertex.

Moreover, if there is a vertex v in F_i corresponding to a cutvertex of B_i that has degree 2, then let j and j' be its neighbors in F_i and let D_j^i and $D_{j'}$ be the corresponding 2-connected components. Then, $D_j^i \cup D_{j'}^i$ contains an attach of A distinct from the 2-separator of $D_j^i \cup D_{j'}^i$. Indeed, if that is not the case, $D_j^i \cup D_{j'}^i$ is a subgraph with at least one internal vertex (the cutvertex corresponding to v in B_i) separated from the rest of the graph by a 2-separator, which is in contradiction with Lemma 4.2.2.

For $k \geq 0$, let v_k be the number of vertices of F_i of degree k, and let v'_k and v''_k be the number of vertices of F_i of degree k which correspond respectively to cutvertices and 2-connected components of B_i. Remark that $v_k = v'_k + v''_k$. Moreover, remark that $v'_0 = v'_1 = 0$.

Because the average degree of a forest is strictly less than 2 by Lemma 2.1.2, we have

$$\begin{aligned}\sum_{k\geq 0} kv_k &< 2\sum_{k\geq 0} v_k\\ \sum_{k\geq 3}(k-2)v_k &< v_1 = v''_1\end{aligned}$$

Moreover, remark that $\sum_{k\geq 3} kv_k = \sum_{k\geq 3}(k-2)v_k < v''_1$ and $\sum_{k\geq 3} kv_k = \sum_{k\geq 3}(k-2)v_k + 2\sum_{k\geq 3} v_k < 3v''_1$.

By definition of the block-cutvertices forest F_i, we have $2v'_2 + \sum_{k\geq 3} kv'_k = v''_1 + \sum_{k\geq 2} kv''_k$. Furthermore, because of the above, $v''_0 + v''_1 \leq |A|$ and $v'_2 \leq |A|$.

Finally,

$$\begin{aligned}p_i &= v''_0 + v''_1 + \sum_{k\geq 2} v''_k\\ &\leq v''_0 + v''_1 + \sum_{k\geq 2} kv''_k\\ &= v''_0 + 2v'_2 + \sum_{k\geq 3} kv'_k\\ &\leq v''_0 + 2v'_2 + \sum_{k\geq 3} kv_k\\ &\leq v_0 + 2|A| + 3v''_1\\ &\leq 5|A|\end{aligned}$$

Let $1 \leq j \leq p_i$, D_j^i has a separator which is contained in the union of the at most $5|A|$ cutvertices of D_j^i in B_i and A, which has therefore a total size of at most $6|A|$ vertices. ■

Then, we can apply Proposition 7.3.1 to each 2-connected component D of each bridge on A in G_0 and we obtain that $|V(D)| \leq P_D(g)$.

Finally, we get

$$\begin{aligned}|V(G)| &\leq Q(g) \times |V(G_0)| \\ &\leq Q(g) \times \left(\sum_{i=1}^{\ell} |V(B_i)|\right) \\ &\leq Q(g) \times \left(\sum_{i=1}^{\ell} \sum_{j=1}^{p_i} |V(D_j^i)|\right) \\ &\leq Q(g) \times \left(\sum_{i=1}^{\ell} \sum_{j=1}^{p_i} P_D(g)\right)\end{aligned}$$

and

$$\begin{aligned}|V(G)| &\leq Q(g) \times \ell p_i P_D(g) \\ &\leq Q(g) \times 2((tw(G)+1) \times m' + g) \times 5|A| P_D(g) \\ &\leq Q(g) \times 10((T(g)+1) \times m' + g)((T(g)+1) \times m') P_D(g) \\ &\leq Q(g) \times S(g)\end{aligned}$$

Finally, in every case, $|V(G)| \leq Q(g) \times S(g) = U(g)$. □

Then, we extend this result to all the excluded minors for a given surface:

Theorem 3.1.1. *Let S be a given surface of Euler genus g. Every excluded minor for S is of order at most $U(g) = g^{O(\log^3 g)}$.*

Proof. By Proposition 7.4.1, we get that any 2-connected excluded minor for a surface S of genus g ($g \geq 0$) is of order at most $U(g)$. Moreover, as U is an increasing function with the property that $U(g_1 + g_2) \geq U(g_1) + U(g_2)$, then by Corollary 4.2.4 we get that any excluded minor for a surface S of genus g is of order bounded by the same function $U(g)$. □

7.5 An almost quadratic bound for the order of the excluded minors for surfaces

In this section, we show that the order of G is bounded by a function

$$U'(g) = ag^2 \times (\log\log g)^{b(\log\log g)^3}$$

of the genus g for some constants a and b.

Let G_0 and A be the subgraph of G and its separator in G which were defined in Proposition 7.2.1. Now, we prove a bound on the order of each bridge B on A in G_0 with respect to g and the size of its separator $A_B \subseteq A$ in G_0.

Proposition 7.5.1 ([HK26a, Proposition 7.7]). *Let (G_0, G_1) be a separation of G such that $A = V(G_0) \cap V(G_1)$ has size at least a and at most $2a-1$ with $3 \leq a \leq (tw(G)+1) \times m'$, G_0 is Π-contractible and $|V(G_0)| \geq 4((T(g)+1) \times m' + 1)(12m+8)$. Then,*

$$|V(G_0)| \leq R(g, a)$$

with

$$A(g) = 6\lfloor \log_{\frac{4}{3}} 3((T(g)+1) \times m' + 1) \rfloor \times (12m+8)$$

and

$$R(g, a) = 3 \times 2a \times \left(\frac{5}{6} A(g) P(g) + \frac{1}{6} A(g) \right)$$

Proof. By hypothesis, G_0 is contained in a disk in Π. Let C_0 be the cycle bounding G_0 in Π. Remark that $G_0 \subseteq \mathrm{Int}(C_0, \Pi)$ and let $G_0' = \mathrm{Int}(C_0, \Pi)$. Remark that, like G_0, G_0' is separated from the rest of G by A. We will show that $|V(G_0')| \leq R(g)$, therefore we will have $|V(G_0)| \leq |V(G_0')| \leq R(g)$.

According to a result of Robertson, Seymour, and Thomas [RST94, (6.2)], every planar graph with no $k \times k$ grid minor has treewidth at most $6k-5$. Remark that, by Corollary 5.3.15, G_0' does not contain $m+1$ disjoint Π-nested cycles. Therefore, G_0' cannot contain a $2(m+1) \times 2(m+1)$ grid minor, and it has treewidth at most $12m+7$ by the result stated above.

Let $(T, (V_t)_{t \in T})$ be a tree decomposition of G_0' of width $tw(G_0') \leq 12m+7$.

By Lemma 7.1.3, as

$$tw(G_0') + 1 \leq 12m+8 \leq \frac{1}{4((T(g)+1) \times m' + 1)} |V(G_0)| \leq \frac{1}{4((T(g)+1) \times m' + 1)} |V(G_0')|$$

there is a separation $(T_1, ..., T_k)$ of T with:

- $k = |A| + 1$,
- for every $1 \leq i, j \leq k$, $|\bigcup_{t \in T_i} V_t| \leq 3|\bigcup_{t \in T_j} V_t|$,
- for every $1 \leq i \leq k$, $V(T_i) \cap (\bigcup_{1 \leq j \neq i \leq k} V(T_j))$ has size at most

$$\lfloor \log_{\frac{4}{3}} 3k \rfloor = \lfloor \log_{\frac{4}{3}} 3((tw(G)+1) \times m' + 1) \rfloor$$

For $1 \leq i \leq k$, let $G_i = \bigcup_{t \in T_i} V_t$ and let $A_i = \bigcup_{j \neq i} V(G_i) \cap V(G_j)$. For $1 \leq i \leq k$, let $G_i' = G_i - A_i$. Remark that the subgraphs $(G_i')_{1 \leq i \leq k}$ of G_0' are pairwise disjoint.

Remark that, by the pigeonhole principle, there exists $1 \leq i_0 \leq k$ such that, for all $1 \leq j \leq k_i$, G_i' contains no vertex from A.

Now, we will bound the order of G_i'. Remark that G_i' is separated from the rest of the graph by the separator A_i of size at most $|A_i| \leq \lfloor \log_{\frac{4}{3}} 3((tw(G)+1) \times m' + 1) \rfloor \times (12m+8) \leq A(g)/6$.

Let $(D_j)_{1 \leq j \leq p}$ be the 2-connected components of G_i'. For $1 \leq j \leq p$, let C_j be the cycle so that $D_j = \mathrm{Int}(C_j, \Pi)$, we call it the exterior cycle of D_j.

We first bound the number of 2-connected components $(D_j)_{1 \leq j \leq p}$ and, for every $1 \leq j \leq p$, we establish a bound on the size of the separator of D_j from the rest of the graph:

Claim 7.5.1.1. *$p \leq 5|A_i|$ and, for every $1 \leq j \leq p$, D_j has a separator of size at most $6|A_i|$.*

Proof of the claim. Let F_i be the block-cutvertex forest of G'_i, that is to say $V(F_i)$ consists of the cutvertices of G'_i together with the 2-connected components $(D_j)_{1\leq j\leq p}$ of G'_i and a connected component is incident to all the cutvertices it contains.

For $1 \leq j \leq p$, D_j is separated from the rest of G by attachments in $A_i \cap D_j$ and by the cutvertices that separate it from the other subgraphs in $(D_j)_{1\leq j\leq p}$.

In particular, each subgraph D_j corresponding to a leaf of F_i is separated from the rest of G by attachments in $A_i \cap D_j$ and by at most one cutvertex. Because G is 2-connected, D_j contains at least one attachment of A_i that is distinct from the cutvertex.

Moreover, if there is a vertex v in F_i corresponding to a cutvertex of G'_i that has degree 2, then let j and j' be its neighbors in F_i and let D_j and $D_{j'}$ be the corresponding 2-connected components. Then, $D_j \cup D_{j'}$ contains an attach of A_i distinct from the 2-separator of $D_j \cup D_{j'}$. Indeed, if that is not the case, $D_j \cup D_{j'}$ is a subgraph with at least one internal vertex (the cutvertex corresponding to v in G'_i) separated from the rest of the graph by a 2-separator, which is in contradiction with Lemma 4.2.2.

For $k \geq 0$, let v_k be the number of vertices of F_i of degree k, and let v'_k and v''_k be the number of vertices of F_i of degree k which correspond respectively to cutvertices and 2-connected components of G'_i. Remark that $v_k = v'_k + v''_k$. Moreover, remark that $v'_0 = v'_1 = 0$.

Because the average degree of a forest is strictly less than 2 by Lemma 2.1.2, we have

$$\begin{aligned}\sum_{k\geq 0} kv_k &< 2\sum_{k\geq 0} v_k\\ \sum_{k\geq 3}(k-2)v_k &< v_1 = v''_1\end{aligned}$$

Moreover, remark that $\sum_{k\geq 3} kv_k = \sum_{k\geq 3}(k-2)v_k < v''_1$ and $\sum_{k\geq 3} kv_k = \sum_{k\geq 3}(k-2)v_k + 2\sum_{k\geq 3} v_k < 3v''_1$.

By definition of the block-cutvertices forest F_i, we have $2v'_2 + \sum_{k\geq 3} kv'_k = v''_1 + \sum_{k\geq 2} kv''_k$. Furthermore, because of the above, $v''_0 + v''_1 \leq |A_i|$ and $v'_2 \leq |A_i|$.

Finally,

$$\begin{aligned}p &= v''_0 + v''_1 + \sum_{k\geq 2} v''_k\\ &\leq v''_0 + v''_1 + \sum_{k\geq 2} kv''_k\\ &= v''_0 + 2v'_2 + \sum_{k\geq 3} kv'_k\\ &\leq v''_0 + 2v'_2 + \sum_{k\geq 3} kv_k\\ &\leq v_0 + 2|A_i| + 3v''_1\\ &\leq 5|A_i|\end{aligned}$$

Let $1 \leq j \leq p$, D_j has a separator which is contained in the union of the at most $5|A_i|$ cutvertices of D_j in G'_i and A_i, which has therefore a total of at most $6|A_i|$ vertices. ■

For $1 \leq j \leq p$, remark that D_j is Π-contractible, 2-connected and has a separator of size at most $6|A_i| \leq A(g)$. We can thus use Proposition 7.3.1 for $G_1 = D_j$. We get $|V(D_j)| \leq P_{G_1}(g)$. Therefore,

$$V(G'_i) = \sum_{j=1}^{p} |V(D_j)| \leq 5|A_i| \times P_{G_1}(g) \leq \frac{5}{6} A(g) P_{G_1}(g)$$

Finally,

$$\begin{aligned} |V(G'_0)| &= \sum_{j=1}^{k} |V(G_j)| \\ &\leq 3k \times |V(G_i)| \\ &= 3k \times (|V(G'_i)| + |A_i|) \\ &\leq 3 \times 2a \times \left(\frac{5}{6} A(g) P_{G_1}(g) + \frac{1}{6} A(g) \right) \leq R(g,a) \end{aligned}$$

Hence, $|V(G_0)| \leq |V(G'_0)| \leq R(g,a)$ as announced. □

We conclude by combining Propositions 7.2.1 and 7.5.1 to find a bound on the order of G. Moreover, we extend the bound on the order of G to every excluded minor for S (even those which are not 2-connected), by using Corollary 4.2.4.

Thereafter, we show that the order of G, which is an arbitrary 2-connected excluded minor for the surface S' of genus g, is bounded by an almost quadratic in g:

Proposition 7.5.2 ([HK26a, Proposition 7.8])**.** *We have*

$$|V(G)| \leq U'(g) = Q(g) \times S'(g)$$

with Q defined in Proposition 7.2.1 and

$$S'(g) = 42(m''+1)((tw(G)+1) \times m' + g) \times \left(\frac{5}{6} A(g) P(g) + \frac{1}{6} A(g) \right)$$

Proof. By Proposition 7.2.1, we obtain that either $|V(G)| \leq 4(4m^2(3g+3)+1)(tw(G)+1) \leq 4(4m^2(3g+3)+1)(T(g)+1) \leq U'(g)$ or that there exists a subgraph G_0 of G such that G_0 is separated from the rest of the graph by a separator A of size $(tw(G)+1) \times m'$, every bridge on A in G_0 is Π-contractible and such that $|V(G)| \leq Q(g) \times |V(G_0)|$.

Let $\mathcal{H}$ be the set of bridges on A in G_0 together with their attaches. We partition $\mathcal{H}$ into $m'' = \lfloor \log_2 \left(\frac{1}{3}(tw(G)+1) \times m' + \frac{1}{3} \right) \rfloor$ sets $(\mathcal{H}_i)_{0 \leq i \leq m''}$ so that the subgraphs in $\mathcal{H}_i$ for some $0 \leq i \leq m''$ have between $a_i = 3 \times 2^i$ and $2a_i - 1 = 3 \times 2^{i+1} - 1$ attaches in A. Remark that $(\mathcal{H}_i)_{0 \leq i \leq m''}$ is indeed a partition of $\mathcal{H}$: every subgraph H of $\mathcal{H}$ has at least 3 attaches because H is Π-contractible and therefore does not have a 2-separator by Lemma 4.2.2, and every subgraph H of $\mathcal{H}$ has at most $3 \times 2^{m''+1} - 1 \geq 3 \times 2^{\log_2\left(\frac{1}{3}(tw(G)+1) \times m' + \frac{1}{3}\right)} - 1 = (tw(G)+1) \times m' \geq |A|$ attaches.

Let us first determine the size ℓ_i of $\mathcal{H}_i$ for every $0 \leq i \leq m''$.

Claim 7.5.2.1. *For* $1 \leq i \leq m''$,

$$\ell_i \leq \frac{2}{a_i - 2}((tw(G)+1) \times m' + g)$$

Proof of the claim. Let us associate to each $1 \leq j \leq \ell_i$ a_i vertices A_j in A so that $A_j \subseteq H_j$, which is always possible because H_j has at least a_i attaches in A.

Let us consider the bipartite graph B with vertex set $A \cup [1, \ell_i]$ and so that, for $a \in A$ and $1 \leq j \leq \ell_i$, we put an edge between a and j if and only if $a \in A_j$.

Remark that B is a minor of $G_0 \subseteq G$, B therefore has genus at most $g(G) \leq g + 2$. Let $\mathcal{F}$ be the set $F(B, \Pi(B))$ of faces of B in $\Pi(B)$. Remark that $2|E(B)| = \sum_{f \in \mathcal{F}} |f| \geq 4|\mathcal{F}|$ because B is bipartite.

By Euler's formula, we get:

$$\begin{aligned}
|V(B)| - |E(B)| + |F| &\geq 2 - (g+2) \\
|V(B)| - |E(B)| + \frac{1}{2}|E(B)| &\geq -g \\
|E(B)| &\leq 2(|V(B)| + g) \\
a_i \ell_i &\leq 2(|A| + \ell_i + g) \\
(a_i - 2)\ell_i &\leq 2((tw(G)+1) \times m' + g) \\
\ell_i &\leq \frac{2}{a_i - 2}((tw(G)+1) \times m' + g)
\end{aligned}$$

■

For $0 \leq i \leq m''$, let $H_i \in \mathcal{H}_i$ be so that $|V(H_i)| = \max_{H \in \mathcal{H}_i} |V(H)|$.

Then, we have

$$\begin{aligned}
|V(G_0)| &\leq \sum_{i=0}^{m''} |\bigcup_{H \in \mathcal{H}_i} V(H)| \\
&\leq \sum_{i=0}^{m''} \ell_i |V(H_i)|
\end{aligned}$$

For $0 \leq i \leq m''$, we can apply Proposition 7.5.1 to the subgraph H_i and we obtain that either $|V(H_i)| \leq 4((T(g)+1) \times m' + 1)(12m + 8)$ or $|V(H_i)| \leq R(g, a_i)$. In both cases, we get that $|V(H_i)| \leq R(g, a_i)$.

We obtain

$$\begin{aligned}
|V(G_0)| &\leq \sum_{i=0}^{m''} \ell_i |V(H_i)| \\
&\leq \sum_{i=0}^{m''} \frac{2}{a_i - 2}((tw(G)+1) \times m' + g) \times 3(2a_i + 1) \times \left(\frac{5}{6}A(g)P(g) + \frac{1}{6}A(g)\right) \\
&\leq \sum_{i=0}^{m''} 42((tw(G)+1) \times m' + g) \times \left(\frac{5}{6}A(g)P(g) + \frac{1}{6}A(g)\right) \\
&\leq 42(m''+1)((tw(G)+1) \times m' + g) \times \left(\frac{5}{6}A(g)P(g) + \frac{1}{6}A(g)\right) \\
&= S'(g)
\end{aligned}$$

Finally, $|V(G)| \leq Q(g) \times S'(g) = U'(g)$. □

Then, we extend this result to all the excluded minors for a given surface:

Theorem 3.2.1 ([HK26a, Theorem 1.3])**.** *Let S be a given surface of Euler genus g. Every excluded minor for S is of order at most $U'(g) = g^{2+o(1)}$.*

Proof. By Proposition 7.5.2, we get that any 2-connected excluded minor for a surface S of genus g ($g \geq 0$) is of order at most $U'(g)$. Moreover, as U' is an increasing function with the property that $U'(g_1 + g_2) \geq U'(g_1) + U'(g_2)$, then by Corollary 4.2.4 we get that any excluded minor for a surface S of genus g is of order bounded by the same function $U'(g)$. □

Chapter 8

Conclusion and open problems

To conclude this manuscript, we summarize our main contributions and discuss several open questions.

The main result of this manuscript is an improved bound U' on the order of the excluded minors for surfaces, which is almost quadratic in the genus:

Theorem 3.2.1. *Let S be a given surface of Euler genus g. Every excluded minor for S is of order at most $U'(g) = g^{2+o(1)}$.*

Moreover, as shown in Proposition 2.2.10, the graph consisting of the union of $g+1$ disjoint copies of $K_{3,3}$ is an excluded minor for every surface of genus g. Consequently, the upper bound on the order of the excluded minors for a surface of genus is at least linear in g:

$$U'(g) \geq 6(g+1)$$

The central open problem in this area of research is therefore to determine the correct asymptotic behavior of U'. While this thesis establishes the upper bound $U'(g) = g^{2+o(1)}$, the best known lower bound remains linear in g.

This naturally leads to the following questions:

Problem 3. *Is there an almost linear upper bound on the order of G?*

Problem 4. *Is the optimal bound almost quadratic?*

During private conversations, both Kawarabayashi and Mohar conjectured that the correct asymptotic behavior should be (almost) linear in g.

A related open problem is to construct a family of excluded minors for surfaces providing a superlinear lower bound on the maximum order of excluded minors as a function of the genus.

Problem 5. *For $g \geq 0$, construct an excluded minor for a surface of genus g of order $f(g)$, so that*

$$\frac{f(g)}{g} \xrightarrow[g \to +\infty]{} +\infty$$

Another promising direction for future research is to investigate excluded minors specifically for nonorientable surfaces. Several pieces of evidence suggest that excluded minors behave more tractably in the nonorientable setting than in the orientable one.

This phenomenon already appears in the historical development of the subject (see Subsection 2.2.2). The proof that nonorientable surfaces admit only finitely many excluded minors [AH89] is comparatively short and and substantially simpler than the general proof for all surfaces.

A similar contrast can be observed in the explicit determination of excluded minors. The complete list of excluded minors for the projective plane has been known since 1981 [GHW79, Arc80, Arc81], whereas no exhaustive list is currently known for any other surface. In particular, obtaining the complete list for the torus still appears to be far out of reach, and consequently also for higher-genus surfaces obtained from the torus by adding handles or crosscaps. This suggests that orientable surfaces, already beginning with the torus, present substantially greater difficulties in the study of excluded minors.

Our own experience during the proofs developed in this thesis points in the same direction. In several case analyses, the orientable configurations consistently proved more difficult to handle than the nonorientable ones.

Taken together, these observations suggest that restricting attention to nonorientable surfaces may lead to further progress. We therefore ask the following question:

Problem 6. *Is there an almost linear upper bound on the order of the excluded minors for a nonorientable surface?*

This question is closely related to the following conjecture from [MT01]:

Conjecture 3 ([MT01, Conjecture 6.6.2])**.** *Let H be a minimal excluded topological minor for the nonorientable surface $\mathbb{N}_k$. Then H can be written as the union of $k+1$ subgraphs each of which is a subdivision of K_5 or $K_{3,3}$.*

Finally, another promising research direction is the computational determination of the exact list of excluded minors for the torus and the Klein bottle.

With the steady increase in computational power over recent years, a systematic computational search for excluded minors of small-genus surfaces is becoming increasingly realistic.

The use of computational methods to investigate excluded minors is in fact not new, as discussed in Subsection 2.2.2. Over several decades, researchers have already attempted to address this problem by implementing algorithmic approaches on computers.

In particular, under the initiative of Myrvold and her collaborators, a substantial database of excluded minors for the torus has been constructed using computational methods, comprising more than 17,500 graphs [Neu93, NM97, Cha02, Woo06, MW18].

It is natural to expect that the techniques and results developed in this manuscript could be used to refine such computational approaches. In particular, they may help reduce the search space and thereby enable more efficient algorithms, potentially leading to a more complete understanding of excluded minors for the torus and related surfaces.

Appendix A

Lower bound for the excluded minors for a surface

A tight lower bound for the excluded minors for a surface can be quickly obtained by using the known results on the genus of cliques. The lower bound $L(g)$ for the excluded minors for a surface is therefore

$$L(g) = \left\lfloor \frac{7 + \sqrt{1 + 24g}}{2} \right\rfloor + 1 = \Theta(\sqrt{g})$$

as proved in Theorem 3.1.4.

A.1 Proof of the lower bound

The genus of every complete graph is precisely known, using the following formula:

Proposition A.1.1 ([MT01, Theorems 4.4.5 and 4.4.6])**.** *For $k \geq 3$,*

$$g(K_k) = \left\lceil \frac{(k-3)(k-4)}{6} \right\rceil$$

Using this result, it is also possible to find the smallest complete graph that cannot be embedded in a surface of genus g:

Lemma A.1.2 ([HK26b, Lemma A.2])**.** *Let g be an integer and let $L(g) = \min\{k \geq 3, g(K_k) > g\}$. Then,*

$$L(g) = \left\lfloor \frac{7 + \sqrt{1 + 24g}}{2} \right\rfloor + 1$$

Proof. The result is obtained by solving the following quadratic equation in k:

$$g(K_k) = \left\lceil \frac{(k-3)(k-4)}{6} \right\rceil > g$$

□

Finally, we prove the lower bound on the order of an excluded minor for a surface of genus g.

Theorem 3.1.4 ([HK26b, Theorem 1.7]). *Let S be a surface of genus g. Let $L(g)$ be the function in Lemma A.1.2. There exists an excluded minor for S of order*

$$|V(G)| = L(g)$$

and no excluded minor for S is of order less than $L(g)$.

Proof. By definition of $L(g)$, the clique $K_{L(g)}$ cannot be embedded in S. Therefore, $K_{L(g)}$ contains an excluded minor G for S as a minor with $|V(G)| \leq L(g)$.

Moreover, by definition of $L(g)$, every clique K_ℓ with $\ell < L(g)$ can be embedded in S. And, as every graph of order $\ell < L(g)$ is a minor of K_ℓ, no excluded minor for S is of order less than $L(g)$.

Finally, we proved that there exists an excluded minor G for S with $|V(G)| = L(g)$ and no excluded minor for S is of order less than $L(g)$. □

A.2 Heawood's bound

Heawood [Hea90] proved in 1890 the following analogue of the Four Color Theorem for general surfaces:

Theorem A.2.1 ([Hea90],[MT01, Theorem 8.3.1]). *Let S be a surface of genus $g > 0$ and let G be a graph embedded in S. Then*

$$\chi(G) \leq H(g)$$

with $H(g) = \left\lfloor \frac{7+\sqrt{1+24g}}{2} \right\rfloor$.

Heawood wrongfully believed that $H(g)$ was best possible for every g and that he had shown equality.

In 1960s, several groups of researchers participated to the effort to find if Heawood's bound was actually tight. Ringel and Youngs [RY68] finally showed in 1968 that this upper bound is the best possible for every surface except the Klein bottle. Moreover, remark that the chromatic number of every surface except the Klein bottle is witnessed by a complete graph $K_{H(g)}$ with g the genus of the surface.

We can remark that the lower bound we find on the order of the excluded minor for a surface of genus g is closely related to Heawood's bound. Indeed, we have $L(g) = H(g) + 1$.

Bibliography

[AGK08] Isolde Adler, Martin Grohe, and Stephan Kreutzer. Computing excluded minors. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 641–650, 2008. URL: https://dl.acm.org/doi/10.5555/1347082.1347153.

[AH77] Kenneth Appel and Wolfgang Haken. Every planar map is four colorable. Part I: Discharging. *Illinois Journal of Mathematics*, 21(3):429–490, September 1977. doi:10.1215/ijm/1256049011.

[AH89] Dan Archdeacon and Phil Huneke. A Kuratowski theorem for nonorientable surfaces. *Journal of Combinatorial Theory, Series B*, 46(2):173–231, 1989. doi:10.1016/0095-8956(89)90043-9.

[APC90] Stefan Arnborg, Andrzej Proskurowski, and Derek Corneil. Forbidden minors characterization of partial 3-trees. *Discrete Math.*, 80(1):1–19, February 1990. doi:10.1016/0012-365X(90)90292-P.

[Arc80] Dan Archdeacon. *A Kuratowski theorem for the projective plane*. PhD thesis, The Ohio State University, 1980. URL: http://rave.ohiolink.edu/etdc/view?acc_num=osu1487093435463451.

[Arc81] Dan Archdeacon. A Kuratowski theorem for the projective plane. *Journal of Graph Theory*, 5(3):243–246, 1981. doi:10.1002/jgt.3190050305.

[Arc92] Dan Archdeacon. Densely embedded graphs. *Journal of Combinatorial Theory, Series B*, 54(1):13–36, January 1992. doi:10.1016/0095-8956(92)90064-5.

[BM04] John Boyer and Wendy Myrvold. On the cutting edge: Simplified O(n) planarity by edge addition. *Journal of Graph Algorithms and Applications*, 8(3):241–273, 2004. doi:10.7155/jgaa.00091.

[BM07] Adrian Bondy and U. S. R. Murty. *Graph Theory*. Springer London, December 2007. URL: https://link.springer.com/book/9781846289699.

[BW86] Rainer Bodendiek and Klaus Wagner. A characterization of the minimal basis of the torus. *Combinatorica*, 6(3):245–260, 1986. URL: https://doi.org/10.1007/BF02579385.

[CC16] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM*, 63(5), 2016. URL: https://doi.org/10.1145/2820609.

[Cha02] John Chambers. Hunting for torus obstructions. Master's thesis, University of Victoria, 2002. URL: https://dspace.library.uvic.ca/items/760d538c-023d-45ff-8d85-57fabd1cd858.

[Chl02] Janka Chlebíková. The structure of obstructions to treewidth and pathwidth. *Discrete Appl. Math.*, 120(1â 3):61â 71, August 2002. doi:10.1016/S0166-218X(01)00281-5.

[CT21] Julia Chuzhoy and Zihan Tan. Towards tight(er) bounds for the excluded grid theorem. *Journal of Combinatorial Theory, Series B*, 146:219–265, 2021. doi:10.1016/j.jctb.2020.09.010.

[Dec78] Richard Decker. *The genus of certain graphs.* PhD thesis, The Ohio State University, 1978. URL: http://rave.ohiolink.edu/etdc/view?acc_num=osu1487077662989342.

[DH72] Richard Duke and Gary Haggard. The genus of subgraphs of K_8. *Israel Journal of Mathematics*, 11:452–455, 1972. doi:10.1007/BF02761469.

[Die24] Reinhard Diestel. *Graph Theory: 6th Edition.* Springer, July 2024. URL: https://link.springer.com/book/10.1007/978-3-662-70107-2.

[FL89] Michael Fellows and Michael Langston. An analogue of the Myhill-Nerode theorem and its use in computing finite-basis characterizations. *Annual Symposium on Foundations of Computer Science (Proceedings)*, pages 520–525, 1989. doi:10.1109/SFCS.1989.63528.

[GH78] Henry Glover and John Huneke. The set of irreducible graphs for the projective plane is finite. *Discrete Mathematics*, 22(3):243–256, 1978. doi:10.1016/0012-365X(78)90057-2.

[GHW79] Henry Glover, John Huneke, and Chin San Wang. 103 graphs that are irreducible for the projective plane. *Journal of Combinatorial Theory, Series B*, 27(3):332–370, 1979. doi:10.1016/0095-8956(79)90022-4.

[GKR13] Martin Grohe, Ken-ichi Kawarabayashi, and Bruce Reed. A simple algorithm for the graph minor decomposition – logic meets structural graph theory –. In *Proceedings of the 2013 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 414–431. Society for Industrial and Applied Mathematics, 2013. doi:10.1137/1.9781611973105.30.

[GMC09] Andrei Gagarin, Wendy Myrvold, and John Chambers. The obstructions for toroidal graphs with no $K_{3,3}$'s. *Discrete Mathematics*, 309(11):3625–3631, June 2009. doi:10.1016/j.disc.2007.12.075.

[GSW25] Maximilian Gorsky, Michał Seweryn, and Sebastian Wiederrecht. Polynomial bounds for the graph minor structure theorem. *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, 2025. doi:10.1109/FOCS63196.2025.00104.

[Hea90] Percy Heawood. Map-colour theorem. *Quarterly Journal of Mathematics*, 24:332–339, 1890.

[HK26a] Sarah Houdaigoui and Ken-ichi Kawarabayashi. An almost quadratic bound for the minimal excluded minors for a surface, September 2026. doi:10.48550/arXiv.2604.02796.

[HK26b] Sarah Houdaigoui and Ken-ichi Kawarabayashi. A quasi-polynomial bound for the minimal excluded minors for a surface. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 4602–4627. Society for Industrial and Applied Mathematics, January 2026. doi:10.1137/1.9781611978971.167.

[Hla97] Amy Hlavacek. *9-vertex irreducible graphs on the torus*. PhD thesis, The Ohio State University, 1997. URL: http://rave.ohiolink.edu/etdc/view?acc_num=osu1487944660932981.

[HT74] John Hopcroft and Robert Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, 1974. doi:10.1145/321850.321852.

[Hud04] Rick Hudson. Planar graphs : A historical perspective. Master's thesis, University of Louisville, 2004. doi:10.18297/etd/647.

[IKM+26] Yuta Inoue, Ken-ichi Kawarabayashi, Atsuyuki Miyashita, Bojan Mohar, Carsten Thomassen, and Mikkel Thorup. The Four Color Theorem with Linearly Many Reducible Configurations and Near-Linear Time Coloring, May 2026. arXiv:2603.24880, doi:10.48550/arXiv.2603.24880.

[Juv95] Martin Juvan. *Algorithms and obstructions for embedding graphs in the torus*. PhD thesis, University of Ljubljana, 1995.

[Kin92] Nancy Kinnersley. The vertex separation number of a graph equals its path-width. *Information Processing Letters*, 42(6):345–350, 1992. doi:10.1016/0020-0190(92)90234-M.

[KKR12] Ken-ichi Kawarabayashi, Yusuke Kobayashi, and Bruce Reed. The disjoint paths problem in quadratic time. *Journal of Combinatorial Theory, Series B*, 102(2):424–435, 2012. doi:10.1016/j.jctb.2011.07.004.

[KMR08] Ken-ichi Kawarabayashi, Bojan Mohar, and Bruce Reed. A simpler linear time algorithm for embedding graphs into an arbitrary surface and the genus of graphs of bounded tree-width. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 771–780, 2008. doi:10.1109/FOCS.2008.53.

[KPS24] Tuukka Korhonen, Michał Pilipczuk, and Giannos Stamoulis. Minor containment and disjoint paths in almost-linear time. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 53–61, 2024. doi:10.1109/FOCS61266.2024.00014.

[KR21] Ken-ichi Kawarabayashi and Benjamin Rossman. A polynomial excluded-minor approximation of treedepth. *Journal of the European Mathematical Society*, 24, 2021. doi:10.4171/JEMS/1133.

[KS19] Ken-ichi Kawarabayashi and Anastasios Sidiropoulos. Polylogarithmic approximation for Euler genus on bounded degree graphs. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 164–175, New York, NY, USA, 2019. Association for Computing Machinery. URL: https://doi.org/10.1145/3313276.3316409.

[KTW20] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. Quickly excluding a non-planar graph. *to appear in Memoirs of the American Mathematical Society. Soc.*, 2020. doi:10.48550/arXiv.2010.12397.

[Kur30] Kazimierz Kuratowski. Sur le problème des courbes gauches en topologie. *Fundamenta Mathematicae*, 15(1):271–283, 1930.

[Lag98] Jens Lagergren. Upper Bounds on the Size of Obstructions and Intertwines. *Journal of Combinatorial Theory, Series B*, 73(1):7–40, May 1998. doi:10.1006/jctb.1997.1788.

[Moh95] Bojan Mohar. Uniqueness and minimality of large face-width embeddings of graphs. *Combinatorica*, 15(4):541–556, 1995. doi:10.1007/BF01192526.

[Moh01] Bojan Mohar. *Graph minors and graphs on surfaces*, pages 145–164. London Mathematical Society Lecture Note Series. Cambridge University Press, 2001. doi:10.1017/CBO9780511721328.008.

[MŠ14] Bojan Mohar and Petr Škoda. Obstructions of Connectivity Two for Embedding Graphs into the Torus. *Canadian Journal of Mathematics*, 66(6):1327–1357, December 2014. doi:10.4153/CJM-2014-025-x.

[MŠ24a] Bojan Mohar and Petr Škoda. Excluded minors for the Klein bottle I. Low connectivity case. *Journal of Combinatorial Theory, Series B*, 164:299–320, January 2024. doi:10.1016/j.jctb.2023.10.002.

[MŠ24b] Bojan Mohar and Petr Škoda. Excluded minors for the Klein bottle II. Cascades. *Journal of Combinatorial Theory, Series B*, 166:80–108, May 2024. doi:10.1016/j.jctb.2023.12.006.

[MT01] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Baltimore, MD: Johns Hopkins University Press, 2001. doi:10.56021/9780801866890.

[MW18] Wendy Myrvold and Jennifer Woodcock. A Large Set of Torus Obstructions and How They Were Discovered. *The Electronic Journal of Combinatorics*, pages 1–16, January 2018. doi:10.37236/3797.

[Nd12] Jaroslav Nešetřil and Patrice Ossona de Mendez. Bounded Height Trees and Tree-Depth. In Jaroslav Nešetřil and Patrice Ossona de Mendez, editors, *Sparsity: Graphs, Structures, and Algorithms*, pages 115–144. Springer, Berlin, Heidelberg, 2012. doi:10.1007/978-3-642-27875-4_6.

[Neu93] Eugene Neufeld. *Practical toroidality testing*. PhD thesis, University of Victoria, 1993. URL: https://dspace.library.uvic.ca/server/api/core/bitstreams/e5edb8d9-ba85-44ba-9c60-7fed8759f29d/content.

[NM97] Eugene Neufeld and Wendy Myrvold. Practical toroidality testing. In *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '97, pages 574–580, USA, January 1997. Society for Industrial and Applied Mathematics. URL: https://dl.acm.org/doi/10.5555/314161.314392.

[Ram97] Siddharthan Ramachandramurthi. The structure and number of obstructions to treewidth. *SIAM Journal on Discrete Mathematics*, 10(1):146–157, 1997. doi:10.1137/S0895480195280010.

[Ram25] Clément Rambaud. Excluding a rectangular grid, January 2025. doi:10.48550/arXiv.2501.11617.

[RS86a] Neil Robertson and Paul Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, September 1986. doi:10.1016/0196-6774(86)90023-4.

[RS86b] Neil Robertson and Paul Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.

[RS90] Neil Robertson and Paul Seymour. Graph minors. VIII. A Kuratowski theorem for general surfaces. *Journal of Combinatorial Theory, Series B*, 48(2):255–288, 1990. doi:10.1016/0095-8956(90)90121-F.

[RS95] Neil Robertson and Paul Seymour. Graph minors. XIII. The Disjoint Paths Problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.

[RS04] Neil Robertson and Paul Seymour. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004. doi:10.1016/j.jctb.2004.08.001.

[RS26a] Neil Robertson and Paul Seymour. Excluding disjoint kuratowski graphs. *J. Comb. Theory Ser. B*, 178(C):294â 322, 2026. doi:10.1016/j.jctb.2026.01.002.

[RS26b] Neil Robertson and Paul Seymour. Excluding sums of kuratowski graphs. *J. Comb. Theory Ser. B*, 178(C):323–353, 2026. doi:10.1016/j.jctb.2026.01.001.

[RSST96] Neil Robertson, Daniel 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, pages 571–575, New York, NY, USA, July 1996. Association for Computing Machinery. doi:10.1145/237814.238005.

[RST94] Neil Robertson, Paul Seymour, and Robin Thomas. Quickly Excluding a Planar Graph. *Journal of Combinatorial Theory, Series B*, 62(2):323–348, November 1994. doi:10.1006/jctb.1994.1073.

[RY68] Gerhard Ringel and J. W. T. Youngs. Solution of the Heawood map-coloring problem. *Proceedings of the National Academy of Sciences of the United States of America*, 60(2):438–445, June 1968. doi:10.1073/pnas.60.2.438.

[Sey93] Paul Seymour. A bound on the excluded minors for a surface. unpublished, 1993. URL: https://web.math.princeton.edu/~pds/papers/surfacebound/bound.pdf.

[SST23] Ignasi Sau, Giannos Stamoulis, and Dimitrios Thilikos. k-apices of minor-closed graph classes. i. bounding the obstructions. *J. Comb. Theory Ser. B*, 161(C):180â 227, July 2023. doi:10.1016/j.jctb.2023.02.012.

[ST96] Paul Seymour and Robin Thomas. Uniqueness of highly representative surface embeddings. *Journal of Graph Theory*, 23(4):337–349, 1996. doi:10.1002/(SICI)1097-0118(199612)23:4<337::AID-JGT2>3.0.CO;2-S.

[Tho90] Carsten Thomassen. Embeddings of graphs with no short noncontractible cycles. *Journal of Combinatorial Theory, Series B*, 48(2):155–177, April 1990. doi:10.1016/0095-8956(90)90115-G.

[Tho97] Carsten Thomassen. A simpler proof of the excluded minor theorem for higher surfaces. *Journal of Combinatorial Theory, Series B*, 70(2):306–311, 1997. doi:10.1006/jctb.1997.1761.

[TUK94] Atsushi Takahashi, Shuichi Ueno, and Yoji Kajitani. Minimal acyclic forbidden minors for the family of graphs with bounded path-width. *Discrete Mathematics*, 127(1):293–304, 1994. doi:10.1016/0012-365X(94)90092-2.

[Wag37] Klaus Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114(1):570–590, December 1937. doi:10.1007/BF01594196.

[Whi92] Hassler Whitney. 2-Isomorphic Graphs. In James Eells and Domingo Toledo, editors, *Hassler Whitney Collected Papers*, pages 125–134. Birkhäuser, Boston, MA, 1992. doi:10.1007/978-1-4612-2972-8_8.

[Woo06] Jennifer Woodcock. A faster algorithm for torus embedding. Master's thesis, University of Victoria, 2006. URL: https://dspace.library.uvic.ca/items/25c5ba80-a9c0-4111-9d93-f6586dfdb47e/full.

Publication List

[HK26a] Sarah Houdaigoui and Ken-ichi Kawarabayashi. An almost quadratic bound for the minimal excluded minors for a surface, September 2026. doi:10.48550/arXiv.2604.02796.

[HK26b] Sarah Houdaigoui and Ken-ichi Kawarabayashi. A quasi-polynomial bound for the minimal excluded minors for a surface. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 4602–4627. Society for Industrial and Applied Mathematics, January 2026. doi:10.1137/1.9781611978971.167.

Index

arch, 104, 105
  empty above, 105
  empty below, 105
  half-arches, 110
  intersection point, 105
  maximal arches, 105
  noncontractible, 105
  size, 105

bipartite graph, 28
block, 30
  2-connected block, 30
  block-cutvertex graph, 30
boundary, 39
boundary of C, 75
boundary of D, 86
bridge, 30
  attach, 30

chromatic number, 30
closed under taking, 39
colorable, 30
  k-colorable, 30
coloring, 30
  k-coloring, 30
  color, 30
  color classes, 30
  proper coloring, 30
complete graph, 28
connected
  k-connected, 27
connected components, 28
crosscap, 32
  add a crosscap, 32
curve, 31
  simple closed curve, 31
cutvertex, 27
  block-cutvertex graph, 30
cycle, 27
  almost disjoint cycles, 57
  boundary, 39
  bounding an empty disk, 39
  closest, 75, 86
  contractible, 34
  cutting along a cycle, 37, 38
  cycles having the same relative orientation, 60
  cycles in this order, 75
  cycles on a spanning tree, 59
  homotopic, 38
  left graph, 34
  nested, 35
  one-sided, 33
  right graph, 34
  separating, 34
  signature, 33
  two-sided, 33
  well homotopic, 85
  well nested, 73
cylinder
  boundaries, 39
  bounding an empty cylinder, 39

degeneracy, 31
disk, 31
  bounds a disk, 31

edge, 26
  contraction, 27
  endpoint, 26
  incident, 26
edge-sharing components, 76
embedding, 33
  dual embedding, 121

equivalent, 33
face, 33
face traversal procedure, 33
facewidth, 95
facial walk, 33
genus, 34
induced embedding, 33
local change, 33
nonorientable, 33, 34
orientable, 33, 34
planar embedding, 35
signature, 33
Euler formula, 34

face, 33
arched face, 104
arches, 104, 105
face of an induced embedding, 102
facewidth, 96
facial walk, 33
maximum face degree, 33
outer face, 35
radius, 96
size, 33
face travesal procedure, 33
facewidth, 95
maximum facewidth, 96
fan, 102
arch path, 103
arched face, 104
arches, 104, 105
column, 103
fan with an arch, 103
horizontal path, 102, 105
intersection point of the arch path, 103
size, 103
vertical paths, 102, 105
flipping, 36
forest, 28

genus, 33
Euler genus, 33
nonorientable genus, 33
orientable genus, 33
graph, 26
average degree, 26
bipartite graph, 28
block-cutvertex graph, 30
bridge, 30
chromatic number, 30
circuit, 27
closed walk, 27
complete graph, 28
connected, 27
connected components, 28
cycle, 27
disconnected, 28
edge, 26
edge set, 26
embeddable, 34
embedding, 33
graph parameter, 40
induced subgraph, 26
maximum degree, 26
minor, 27
order, 26
path, 27
planar graph, 35
radius, 96
subdivision, 27
subgraph, 26
topological minor, 27
vertex, 26
vertex set, 26
walk, 27
graph parameter, 40
minor-monotone, 40
grid, 28
hexagonal grid, 93

handle, 32
add a handle, 32
homotopic, 38
(f,f')-well-homotopic, 86
(f,free)-well-homotopic, 85
(v,f)-well-homotopic, 86
(v,v')-well-homotopic, 85
(v,free)-well-homotopic, 85
(free,free)-well-homotopic, 85
freely well homotopic, 85
well homotopic, 85
well homotopic pinched on a face, 85
well homotopic pinched on a piece, 86
well homotopic pinched on a vertex, 85
well homotopic pinched on a vertex and a face, 86

well homotopic pinched on two faces, 86
well homotopic pinched on two pieces, 86
well homotopic pinched on two vertices, 85

independent set, 30
isolated paths, 69
almost disjoint, 70
disjoint, 70
joint, 70

loop, 26

minor, 27
excluded minor, 40
excluded topological minor, 40
minor ordering, 41
proper minor, 27
topological minor, 27
multigraph, 26
dual multigraph, 121
edge set, 26
geometric dual multigraph, 121
loop, 26
multiple edge, 26
parallel edge, 26
vertex set, 26
multiple edge, 26

nested, 35
(f,f')-well-nested, 73
(f,free)-well-nested, 73
(v,f)-well-nested, 73
(v,v')-well-nested, 73
(v,free)-well-nested, 73
(free,free)-well-nested, 73
freely well nested, 73
well nested, 73, 75
well nested pinched on a face, 73
well nested pinched on a piece, 73
well nested pinched on a vertex, 73
well nested pinched on a vertex and a face, 73
well nested pinched on two faces, 73
well nested pinched on two piece, 73
well nested pinched on two vertices, 73
nonorientable, 32

parallel edge, 26
path, 27
interior, 27
isolated paths, 69
path decomposition, 29
width, 29
pathwidth, 29
piece, 52
arch, 106
arched faces, 105
column, 106
degree, 69
horizontal path, 105
vertical paths, 105

radius, 96
representativity, *see* facewidth

separation, 28
k-separation, 28
separator, 27
k-separator, 27
cutvertex, 27
square
bad with respect to e and Π_e, 76, 88
contractible square, 75
full, 75, 87
good with respect to e and Π_e, 76, 88
homotopic square, 87
pinched on a piece, 75, 87
pinched on two pieces, 75, 87
subdivision, 27
subgraph, 26
independent subgraphs, 107
induced subgraph, 26
surface, 31
double torus, 32
genus, 33
Klein bottle, 32
orientable, 32
projective plane, 32
torus, 32

torus, 32
double torus, 32
tree, 28
1-separation sequence, 126
diameter, 28

internal vertex, 28
leaf, 28
root, 28
rooted, 28
spanning tree, 28
tree decomposition, 29
minimal, 29
width, 29
treewidth, 29

vertex, 26
adjacent, 26
degree, 26
dominating vertex, 122
incident, 26
neighbor, 26

Symbol index

$\cap$, 27
$\cup$, 27
$\leq_m$, 41
$|\cdot|$, 26
$\subseteq$, 26

$A_{G'}(g)$, 130
$\mathcal{A}(H)$, 105
$\mathcal{A}_{max}(H)$, 105

$\mathcal{B}(C)$, 75, 131
$\mathcal{B}(D)$, 86
$\mathcal{B}_{\mathrm{Con}}(C)$, 75
$\mathcal{B}_{\mathrm{Con}}(D)$, 86
$\mathcal{B}_{\mathrm{N}}(C)$, 75
$\mathcal{B}_{\mathrm{N}}(D)$, 86

$\overline{C}$, 37
C_A, 103, 105
C_H, 104
$\chi(\Pi)$, 34
$\chi(G)$, 30

$\Delta(G)$, 26
$d(p)$, 69
$d(v)$, 26
$d_G(v)$, 26
$d_{(G,\Pi)}(p)$, 69
$\Delta_{G'}(g)$, 130
$\Delta_F(G,\Pi)$, 33

$E(G)$, 26
$\mathrm{Ext}(C)$, 35
$\mathrm{Ext}(C,\Pi)$, 35, 36
$\mathrm{ext}(C)$, 35
$\mathrm{ext}(C,\Pi)$, 34
$\mathrm{ext}(C,\Pi)$, 36

$\mathcal{F}(C,\Pi)$, 73
$\mathcal{F}(G,\Pi)$, 73
$\mathcal{F}(G_1,\Pi)$, 102
$\mathcal{F}(H)$, 103
$F(G,\Pi)$, 33
$\mathcal{F}_{arch}(H)$, 104
$fw(G,\Pi)$, 95

G^*, 121
$G_l(C,\Pi)$, 34
$G_r(C,\Pi)$, 34
$\overline{G}$, 37
$G-X$, 26
$G-e$, 26
$G-v$, 26
G/X, 27
G/e, 27
$g(G)$, 34
$g(\Pi)$, 34
G_e, 61

$\mathcal{I}(C)$, 75
$\mathcal{I}(D)$, 86
$\mathcal{I}_{\mathrm{N}}(C)$, 75
$\mathcal{I}_{\mathrm{N}}(D)$, 86
$\mathrm{Int}(C\cup C',\Pi)$, 38
$\mathrm{Int}(C)$, 35
$\mathrm{Int}(C,\Pi)$, 35, 36
$\mathrm{int}(C)$, 35
$\mathrm{int}(C,\Pi)$, 34
$\mathrm{int}(C,\Pi)$, 36
$\mathrm{int}(P)$, 27

J_k, 93

K_n, 28

λ, 33
$\lambda(C)$, 33

$m_{G'}$, 130
max-$fw(G, \Pi)$, 97

$\mathbb{N}_c$, 32

$pw(G)$, 29
Π^*, 121
    λ^*, 121
    π^*, 121
Π, 33
    λ, 33
    π, 33
$\Pi(H)$, 33
$\overline{\Pi}$, 37, 38
Π_e, 61

$rad(f)$, 96

$\mathbb{S}_h$, 32

$tw(G)$, 29

$V(G)$, 26