



HAL
open science

Exploration of the space of the colorings of a graph

Clément Legrand-Duchesne

► **To cite this version:**

Clément Legrand-Duchesne. Exploration of the space of the colorings of a graph. Computer Science [cs]. Université de Bordeaux, 2024. English. NNT : 2024BORD0117 . tel-04681426

HAL Id: tel-04681426

<https://theses.hal.science/tel-04681426v1>

Submitted on 29 Aug 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE

PRÉSENTÉE À

L'UNIVERSITÉ DE BORDEAUX

ÉCOLE DOCTORALE
DE MATHÉMATIQUES ET D'INFORMATIQUE

par **Clément LEGRAND-DUCHESNE**

POUR OBTENIR LE GRADE DE

DOCTEUR

SPÉCIALITÉ : INFORMATIQUE

Exploration de l'espace des colorations d'un graphe

Soutenue le 2 juillet 2024

Membres du jury :

Mme Marie ALBENQUE	Directrice de recherche CNRS, Université Paris Cité	Examinatrice
Mme Marthe BONAMY	Chargée de recherche CNRS, Université de Bordeaux	Directrice
M. Vincent DELECROIX	Chargé de recherche CNRS, Université de Bordeaux	Directeur
M. Hugo PARLIER	Professeur, University of Luxembourg	Rapporteur
M. Alexander SCOTT ..	Professeur, University of Oxford	Rapporteur
M. Stéphan THOMASSÉ	Professeur, ENS de Lyon	Examineur

Membre invité :

M. Robert CORI Professeur émérite, Université de Bordeaux Invité

Titre Exploration de l'espace des colorations d'un graphe

Résumé vulgarisé pour le grand public Cette thèse porte sur la reconfiguration combinatoire et des problèmes structurels ou géométriques. Nous associons à un espace de configuration et à une opération élémentaire un graphe de reconfiguration, dont les sommets sont les configurations et dans lequel deux configurations sont adjacentes si elles diffèrent par une opération. La reconfiguration est motivée par l'échantillonnage aléatoire et des applications en mécanique statistique. Cela soulève des questions algorithmiques, probabilistes et structurelles, notamment sur la complexité du problème d'accessibilité et la géométrie du graphe de reconfiguration.

La principale opération de reconfiguration étudiée ici est appelée changement de Kempe et opère sur les colorations propres d'un graphe. Nous étudions ensuite d'autres opérations de reconfiguration motivées par la topologie en basse dimension. Nous présentons plusieurs bornes polynomiales sur le diamètre du graphe de reconfiguration et obtenons divers résultats algorithmiques.

Mots-clés graphe, coloration, reconfiguration, mineur de graphe, surface

Laboratoire d'accueil LaBRI, 351 Cours de la Libération, 33405 Talence

Title Exploring the space of colorings of a graph

Short abstract This thesis focuses on combinatorial reconfiguration and structural problems with a geometric flavor. To a configuration space and given an elementary operation we associate their reconfiguration graph, whose vertices are the configurations and in which two configurations are adjacent if they differ by one operation. Reconfiguration is motivated by random sampling and practical applications in statistical mechanics. It raises algorithmic, probabilistic as well as structural questions, for example by looking at the geometry of the reconfiguration graph or the complexity of the reachability problem in it.

The main reconfiguration operation studied in this thesis is called a Kempe change and operates on the proper colorings of a graph. We also study other reconfiguration operations motivated by low-dimensional topology. We present several polynomial upper bounds on the diameter of the reconfiguration graph and obtain algorithmic results in various settings.

Keywords graph, coloring, reconfiguration, graph minor, surface

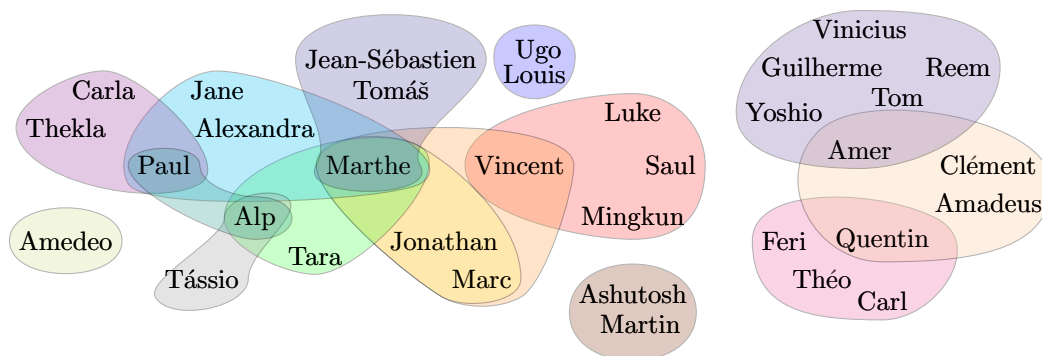
Hosting Laboratory LaBRI, 351 Cours de la Libération, 33405 Talence

Remerciements

Tout d'abord, merci à mes encadrants, Marthe et Vincent. Petit élève de master, on ne se rend pas compte de l'importance de bien choisir ses encadrants de thèse. J'ai réalisé au fur et à mesure de mon stage de M2 et de ma thèse à quel point j'ai eu beaucoup de chance dans ce choix. Au delà de l'aspect scientifique, merci pour votre accueil chaleureux en ces temps de covid. Nos innombrables discussions autour d'un goûter ou au détour d'une balade à vélo font partie de mes meilleurs souvenirs de thèse. Merci de m'avoir donné autonomie et liberté, de m'avoir encouragé à travailler à plusieurs. Enfin, merci pour le temps et les précieux conseils que vous m'avez donnés pendant ces trois années. Je vous dois énormément.

I would like to thank the members of my jury, Marie Albenque, Hugo Parlier, Alex Scott, Stéphan Thomassé and Robert Cori, for accepting to be part of it. It meant a lot to me as I look up to each of you. I hope that I will have the chance of working with you one day.

This PhD would have been very different without all the amazing people that I had the chance to work with. I learned a lot, from each and everyone of you



and it is a real pleasure to be part of such a kind and welcoming community. In particular, thank you to Ugo and Louis for their warm welcome in my favorite mountains, to Amer for showing us Beyrouth, to my former supervisor Martin and to Alp and Paul, who made my dream workshop come true (désolé, il devait pas pleuvoir comme ça, plus qu'un virage et on arrive !).

La recherche est une grande famille, je tiens à remercier la mienne : mes aînés et modèles Théo, Jonathan, Claire (ton éclat de rire manque à nos couloirs depuis ton départ) et Alex ; Hoang pour m'avoir si gentiment pris sous ton aile ; Vania et mes petits frères Paul, Timothé et Kai. Je tiens aussi à remercier tous ceux qui rendent mon quotidien au labo si vivant, à commencer par les habitants du bureau 267, Arthur, Ludovic, Paul, Zoé et bientôt Raquel; mes cobureaux Sarah et Timothée, ainsi que Théo, Pierre et Antonio. Je tiens plus généralement à

remercier les membres de l'AFoDIB et des équipes *Combinatoire et Intéractions* et *Graphes et Optimisation* qui grâce à leur implication dans l'animation de la vie de ce labo, accueillent les doctorants de manière sûre et conviviale. Merci aussi aux personnes que j'ai pu cotoyer au cours de mes enseignements, je pense en particulier à Samuel Thibault et Isabelle Dutour, qui m'ont accompagné pour mes premiers grands bains, devant une classe puis dans le monde du réseau. Merci aussi à mes élèves qui ont été adorables.

Le bon fonctionnement de ce labo tient aussi au personnel d'entretien et au personnel administratif, que je tiens à remercier. En particulier, merci à Isabelle Garcia, tu m'as toujours réservé un accueil rassurant, gentil et efficace face aux diverses complications administratives que j'ai pu rencontrer. Merci à l'ANR CODYS, ainsi qu'à l'ANR GrR et au projet CEDRE "Prix de la localité pour la reconfiguration" tous les deux portés par Nicolas Bousquet, dont les financements m'ont permis de rencontrer la communauté de la reconfiguration et certains de mes collaborateurs.

Merci à mes mentors, Laurent Feuilloley, Loïc Paulevé et Lionel Eyraud-Dubois d'avoir sù apaiser mes inquiétudes en partageant votre expérience. Merci à mon mentor officieux Nicolas Bousquet pour tes conseils et pour nos discussions de reconfiguration et de vie en recherche. Merci aussi à mon comité de suivi de thèse, Nathanaël Fijalkow et Jean-François Marckert, j'ai beaucoup apprécié nos réunions annuelles pour le regard et les conseils que vous apportiez.

Ceux qui me connaissent savent combien je peux être grincheux sans ma dose quotidienne de sport. Merci aux Aigles de Bègles de m'avoir fait découvrir ce magnifique sport qu'est l'Ultimate Frisbee. Zoé, Kévin, Nico, Léopold, Léonard, Marine, Pierre, Lucille, Seb, Mousse, Thomas et Sophie, merci de m'avoir tant appris et d'avoir égayé mes soirées et mes week-ends, sur et en dehors des terrains. Merci aussi à tous mes fidèles compagnons de grimpe: Théo, Sarah, Raquel, Nina, Gabriel, Alp, Paul, Louis, Alexandra, Géraud et Nicolas (grâce à toi ma liste de spots d'escalade à essayer concurrence dangeureusement ma liste de livres à lire), et surtout Benoît (j'attends avec impatience nos prochaines sorties falaises, en essayant de passer plus de temps sur un mur que perdus en forêt).

Enfin, je ne serai pas là sans mes proches. J'ai la chance d'avoir grandi et encore aujourd'hui d'évoluer dans un petit havre tranquille. Merci à Zoé et à toute ma famille pour la bonne humeur et la quiétude que vous diffusez dans mon quotidien. Merci d'avoir cru en moi.

Contents

Extended abstracts	1
Résumé en français	1
Abstract in English	3
Introduction	5
Introduction en français	5
Résultats et organisation de cette thèse	9
Introduction in English	14
Results and organization of this thesis	17
What is not in this thesis	21
1 Preliminaries	23
1.1 Graphs, surfaces and colorings	23
1.1.1 Graph relations	24
1.1.2 Special graphs	25
1.1.3 Colorings	25
1.1.4 Surfaces	26
1.1.5 Graph operations	27
1.1.6 Graph classes	28
1.1.7 Tree-decompositions	29
1.2 Groups and infinite graphs	29
1.2.1 Infinite graphs	30
1.2.2 Groups and Cayley graphs	31
1.3 Markov chains and random walks	34
1.3.1 Mixing time	35
1.3.2 Couplings	36
1.3.3 Bottleneck ratio and expander families	36
1.4 (Parameterized) complexity classes	37
1.4.1 Complexity of decision problems	37
1.4.2 Parameterized complexity	38
1.4.3 Counting problems	39

1.5	Knots	39
2	Recoloring with Kempe changes	41
	Introduction	41
2.1	Connectedness of $\mathcal{K}_k(G)$	46
	2.1.1 Negative results	47
	2.1.2 Positive results in sparse graphs classes	55
	2.1.3 Beyond frozen colorings: open questions	60
2.2	The diameter of $\mathcal{K}_k(G)$	71
2.3	Counting and complexity	74
2.4	Wang-Swendsen-Kotecký dynamics	75
	2.4.1 Introduction to the Potts model	75
	2.4.2 Sampling with Markov Chain Monte Carlo algorithms	77
	2.4.3 Ergodicity and mixing time of the Glauber dynamics and the WSK algorithm	78
2.5	Conclusion and perspectives	81
	2.5.1 Non-recolorability without frozen colorings	81
	2.5.2 Bounding the Kempe distance	83
3	Kempe changes in sparse graphs	85
	Introduction	85
	Recoloring version of Hadwiger’s conjecture	89
3.1	Upper bounds on $\text{diam}(\mathcal{K}_k(G))$ in sparse graph classes	90
	3.1.1 Planar graphs	90
	3.1.2 Technical improvements of Theorem 2.1.15	98
	3.1.3 Graphs of bounded degree	98
	3.1.4 Graphs of bounded maximum average degree	99
	3.1.5 Graphs of bounded treewidth	105
3.2	On a recoloring version of Hadwiger’s conjecture	109
	3.2.1 Construction	110
	3.2.2 No large simple minor	111
	3.2.3 No large double minor	111
	3.2.4 No huge triple minor	112
	Conclusion and open questions	112
4	Optimal algorithms for reconfiguration	115
	Introduction	115
	Untangling knots	115
	Playing Rubik’s Cube on a square-tiled surface	120
4.1	Parameterized complexity of untangling knots	125
	4.1.1 A reorganization result	125

4.1.2	W[P]-algorithm	126
4.1.3	Proof of Theorem 4.1.1	128
4.2	Reconfiguration of square-tiled surfaces	132
4.2.1	Background and more	132
4.2.2	Connecting to path-like configurations	142
4.2.3	Reconfiguring path-like configurations	151
4.2.4	Hyperelliptic square-tiled surfaces	159
	Conclusion and perspectives	168
	Equivalence via cylinder shears	168
	Beyond Conjecture 4.0.7	170
5	A structure theorem for locally finite quasi-transitive graphs avoiding a minor	171
	Introduction	171
	A structure theorem	173
	Hadwiger number	174
	Accessibility in graphs	174
	Accessibility in groups	175
	The Domino Problem	176
	Overview of the proof of Theorems 5.0.1 and 5.0.2	177
	Organization of this chapter	178
5.1	Tree-decompositions and tangles	178
5.1.1	Separations and canonical tree-decompositions	178
5.1.2	Tangles	185
5.1.3	An example	188
5.1.4	Tangles of order 4: orthogonality and crossing-lemma	189
5.1.5	Contracting a single crossed edge	195
5.1.6	Contracting all the crossed edges	197
5.1.7	Planarity after uncontracting crossed edges	202
5.2	The structure of quasi-transitive graphs avoiding a minor	206
5.2.1	Main results	206
5.2.2	Tools	207
5.2.3	Proof of Theorems 5.2.1 and 5.2.3	208
5.3	Applications	211
5.3.1	The Hadwiger number of quasi-transitive graphs	211
5.3.2	The Domino Problem in minor-excluded groups	211
5.4	Open problems	213
5.4.1	Related work	215

Conclusion and Perspectives	217
Recoloring with Kempe changes	217
Obstacles to Kempe recolorability	217
Recoloring at and below the degeneracy threshold	218
Reconfiguration of square-tiled surfaces	219
References	221
Publications	237

Extended abstracts

Résumé en français

Cette thèse porte sur la reconfiguration combinatoire et des problèmes structurels avec une coloration géométrique. La reconfiguration est l'étude des espaces de configuration vis-à-vis d'une opération élémentaire transformant une configuration en une autre. Les espaces de configuration sont modélisés par des graphes de reconfiguration, dont les sommets sont les configurations et dans lesquels deux configurations sont adjacentes si elles diffèrent par une opération élémentaire. La reconfiguration est motivée par des applications pratiques en mécanique statistique et soulève des questions structurelles, algorithmiques et probabilistes. Sous quelles conditions structurelles le graphe de reconfiguration est-il connexe ; quel est son diamètre ? Existe-t-il une configuration "optimale" dans toutes les composantes connexes du graphe de reconfiguration ; existe-t-il un algorithme efficace pour la trouver ? Peut-on échantillonner efficacement une configuration aléatoire en effectuant une marche aléatoire sur le graphe de reconfiguration ?

La principale opération de reconfiguration étudiée dans cette thèse est appelée changement de Kempe et opère sur les colorations propres d'un graphe. Le premier chapitre consiste en un état de l'art de la recoloration avec les changements de Kempe, avec une vue d'ensemble des principaux outils et des différentes idées de preuve. Le deuxième chapitre présente des bornes polynomiales sur le diamètre du graphe de reconfiguration de Kempe dans différentes classes de graphes creux. Nous y réfutons également une version de recoloration de la conjecture de Hadwiger : tous les graphes sans K_t -minor ont un graphe de reconfiguration de Kempe connexe lorsque le nombre de couleurs est au moins t . Dans le troisième chapitre, nous présentons d'autres opérations de reconfiguration provenant de la topologie en basse dimension, telles que les mouvements de cisaillement sur les surfaces à petits carreaux ou les mouvements de Reidemeister sur les diagrammes de nœuds. Nous obtenons là encore des bornes sur les diamètres des graphes de reconfigurations correspondants. Enfin, nous montrons comment dériver des algorithmes optimaux et des résultats de complexité paramétrée à partir de résultats de commutation sur les opérations de reconfiguration.

Dans le dernier chapitre, nous laissons la reconfiguration de côté pour présenter un résultat utilisant des notions transverses dans cette thèse. Nous prouvons un théorème de structure pour les graphes infinis ayant de nombreuses symétries et évitant un mineur, qui rappelle le théorème de structure des mineurs de graphes de Robertson et Seymour. Parmi les applications de ce résultat, nous obtenons des résultats structurels sur le nombre de Hadwiger de ces graphes et nous caractérisons la décidabilité du problème du domino sur les groupes dont les graphes de Cayley évitent un mineur.

Abstract in English

This thesis focuses on combinatorial reconfiguration and structural problems with a geometric flavor. Reconfiguration is the study of configuration spaces with respect to an elementary operation transforming a configuration into another. Configuration spaces are formalized through the reconfiguration graph, whose vertices are the configurations and in which two configurations are adjacent if they differ by one operation. Reconfiguration is motivated by practical applications in statistical mechanics and raises algorithmic, probabilistic as well as structural questions. Under which structural conditions is the reconfiguration graph connected; what is its diameter? Is there a “best” configuration in all the connected components of the reconfiguration graph; is there an efficient algorithm to find it? Can we efficiently sample a random configuration by performing a random walk on the reconfiguration graph?

The main reconfiguration operation studied in this thesis is called a Kempe change and operates on the proper colorings of a graph. The first chapter consists in a survey on recoloring with Kempe changes and illustrates the state of the art with an overview of the main tools and the different proof ideas. The second chapter presents polynomial bounds on the diameter of the Kempe reconfiguration graph in various sparse graph classes. We also disprove a recoloring version of Hadwiger’s conjecture: all K_t -minor-free graphs have a connected Kempe reconfiguration graph when the number of colors is at least t . In the third chapter, we present other reconfiguration operations arising from low-dimensional topology, such as shearing moves on squared-tiled surfaces or Reidemeister moves on knot diagrams. Once again, we prove polynomial bounds on the diameter of the reconfiguration graph in different settings. Moreover, we give general methods to obtain optimal algorithms and parameterized complexity results from some commutation properties of the elementary operations.

In the last chapter, we move away from reconfiguration to present a result that intertwines several transverse notions of this thesis. We prove a structure theorem for infinite graphs that have a lot of symmetries and avoid a minor, which is reminiscent of the Graph Minor Structure Theorem of Robertson and Seymour. As applications of this result, we obtain structural results on the Hadwiger number of these graphs and characterize the decidability domino problem on the groups with a Cayley graph avoiding a minor.

Introduction

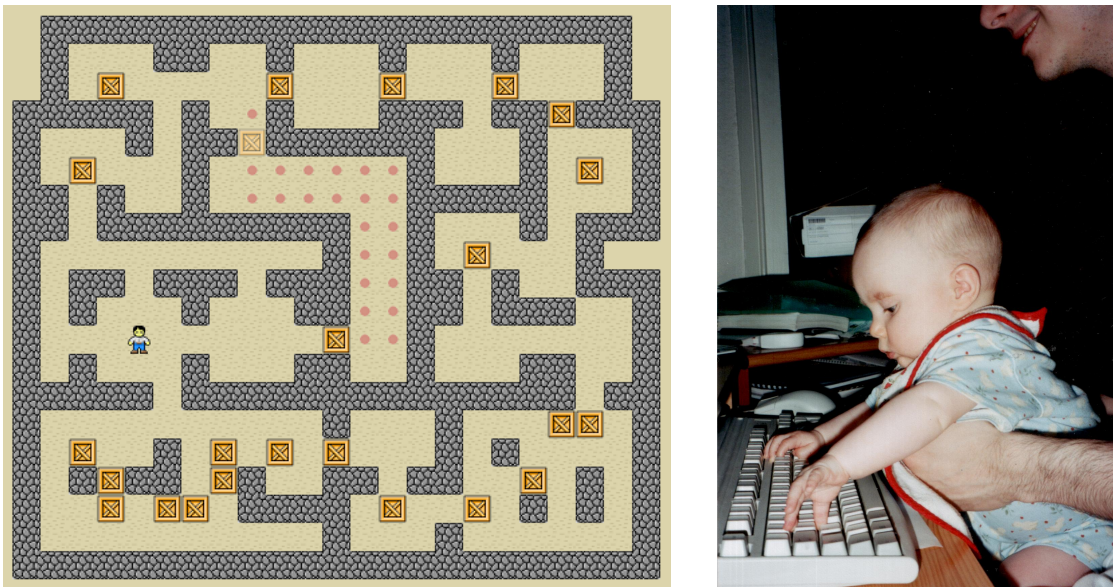
Introduction en français

Les puzzles combinatoires sont des casse-têtes dont l'état est modifié par des mathématiciens en herbe, dans le but d'atteindre une configuration spécifique. L'exemple le plus connu d'un tel puzzle est le *Rubik's Cube*, qui a connu un succès mondial dans les années quatre-vingt et reste aujourd'hui le jouet le plus vendu au monde, avec plus de 500 millions d'exemplaires vendus depuis sa création, sans compter les nombreuses variantes du cube $3 \times 3 \times 3$ classique.

Enfant, j'appréciais particulièrement un autre puzzle combinatoire du nom de *Sokoban*. Sokoban est un jeu vidéo sorti en 1982, dans lequel le joueur pousse des caisses dans un entrepôt afin de les amener sur des positions cibles (voir Figure 1). Le joueur peut bouger dans l'entrepôt et pousser les caisses, à moins qu'il n'y ait un mur ou une autre caisse derrière celle-ci. En revanche, le joueur ne peut pas se déplacer à travers une caisse ou un mur, ni tirer les caisses.

Il existe plein d'autres exemples de puzzles combinatoires. On pourrait par exemple penser aux *Tours de Hanoi* (voir Figure 2), inventées en 1883 par le mathématicien Édouard Lucas, qui consistent en des disques de diamètres variables enfilés en pyramide sur trois pics : aucun disque ne peut supporter un disque plus large. Le joueur peut alors bouger le disque au sommet d'une des trois pyramides, tout en préservant la contrainte d'empilement pyramidal. Les configurations initiale et objectif sont souvent celles composées de deux pics vides et le pic de droite ou de gauche avec les n disques. Ce puzzle est un classique des ateliers de vulgarisation scientifique ou d'informatique débranchée et est notoirement connu pour nécessiter $2^n - 1$ déplacements de disques entre ses configurations initiale et objectif.

Le charme des puzzles combinatoires tient principalement à la complexité de leur résolution et au nombre considérable de configurations (43×10^{18} configurations pour le Rubik's Cube), comparé à la simplicité apparente de ces objets. Pour les résoudre efficacement, les joueurs ont conçus des algorithmes, avec pour objectif le *God's algorithm* : un algorithme résolvant le puzzle en un minimum de mouvements. Cela est parfois impossible efficacement : Culberson a prouvé



(a) Les positions cibles sont représentées par des points rouges. Les caisses semi-transparents sont déjà sur une position cible.

(b) Depuis tout jeune, Sokoban est un de mes jeux favoris

Figure 1: Un niveau de Sokoban

en 1997 que décider si une instance de Sokoban est résoluble est un problème PSPACE-complet [Cul97]. En revanche, il y a une vaste histoire d'algorithmes triant des Rubik's Cube de taille arbitraire, et de bornes supérieures sur le nombre maximum de mouvements nécessaires pour cela. En 2014, Rokicki, Kociemba, Davidson et Dethridge [RKDD14] ont prouvé grâce à une ingénieuse exploration par ordinateur que toute configuration du Rubik's cube classique pouvait être résolue en au plus 20 rotations de faces, ce qui est optimal. Il existe plusieurs compétitions de résolution de Rubik's Cube, avec différentes épreuves, comme la résolution en le moins de mouvements possibles, ou aussi rapide que possible¹.

Dans un contexte plus sérieux, il est parfois nécessaire de modifier une configuration tout en maintenant certaines contraintes. Pour illustrer cela, nous présentons une application concrète en la conception d'ordinateurs quantiques. Pour stocker de l'information quantique, il est souvent nécessaire d'assembler des particules quantiques suivant un motif prédéfini. Supposons que nous ayons pour cela

¹Les records du monde actuels ont été réalisés en 2023 par Wong Chong Wen avec 20 mouvements en moyenne et par Yiheng Wang en 4,48 secondes en moyenne. Pour la résolution en le moins de mouvements possibles, le joueur dispose d'une heure pour écrire sa solution. Pour chacun des records, la moyenne est prise sur les trois essais médians sur cinq essais sur des configurations différents, en enlevant le meilleur et le pire résultat.



Figure 2: Une configuration intermédiaire des Tours de Hanoi

une grille de pièges optiques, chacun permettant de capturer un atome pour une courte période de temps. Les atomes sont initialement disposés au hasard sur cette grille et peuvent être relâchés par les pièges optiques ou déplacés un par un vers un piège optique vide voisin, jusqu'à obtenir la configuration désirée. À cause de leur grande instabilité, les atomes peuvent être perdus pendant l'opération et les réarranger le plus vite possible permet d'augmenter substantiellement le taux de réussite. Deux facteurs jouent sur ce temps de préparation et le taux de réussite: le temps de calcul nécessaire pour déterminer la séquence de déplacements et la distance totale parcourue par les atomes. La forte contrainte de temps fait du réarrangement efficace d'atomes un véritable défi technologique et algorithmique (voir [CESB+23, ESBD+23] pour une explication plus détaillée de ce problème et de ses motivations).

Le besoin de modifier une configuration peut être motivé par plusieurs raisons et est parfois simplement inévitable, comme nous l'avons vu précédemment. Une première motivation est d'éviter une panne pendant le changement de configuration. Cela est souvent le cas dans des contextes de réseaux de télécommunication par exemple. En second lieu, trouver une solution à partir de zéro s'avère souvent très coûteux, d'autant plus que les configurations sont souvent des solutions de problèmes combinatoires compliqués, tandis que modifier légèrement une configuration est bien plus facile. La *reconfiguration* est l'étude de ce type de problèmes,

centrés sur la structure de l'espace de configuration vis-à-vis d'une opération de reconfiguration élémentaire. Un problème de reconfiguration est la donnée d'un ensemble de configuration et d'un ensemble restreint d'opérations élémentaires transformant une configuration en une autre.

Pour modéliser les espaces de configurations ainsi que les configurations elles-mêmes, nous utiliserons souvent des *graphes*, qui sont des objets combinatoires composés de *sommets* reliés par des *arêtes* (voir le Chapitre 1 pour toutes les définitions). Par exemple, les configurations d'atomes peuvent être modélisées par une grille dont les sommets sont les pièges optiques et les arêtes joignent les sites adjacents. Les atomes sont alors représentés par des pions placés sur les sommets et autorisés à glisser le long des arêtes. Pour représenter les espaces de configurations, on définit le *graphe de reconfiguration*, dont les sommets sont les configurations et les arêtes lient les configurations qui diffèrent par une opération élémentaire. La taille de ce graphe étant le nombre de configurations, celui-ci est souvent trop large pour être analysé directement.

Les questions récurrentes en reconfiguration sont celles que nous avons illustrées jusqu'ici. Est-ce que deux configurations données A et B sont *équivalentes*, c'est-à-dire joignable par une séquence d'opérations élémentaires ? Par exemple, dans le cas du Rubik's Cube, tourner manuellement un coin du cube résulte en une configuration non-équivalente (via les rotations de faces) au cube trié. Dans le langage des graphes, cette question peut être reformulée ainsi : Y a-t-il un chemin allant de A à B dans le graphe de reconfiguration ? Il est souvent intéressant d'obtenir des algorithmes efficaces pour déterminer si cela est possible et donnant une séquence de reconfiguration le cas échéant. Plus généralement, on peut se demander si toutes les configurations sont équivalentes et quelles sont les configurations à distance maximale dans le graphe de reconfiguration, et quelle est la difficulté algorithmique de ces questions.

La reconfiguration est aussi motivée par l'échantillonnage aléatoire. L'ensemble des configurations est souvent trop large pour être énuméré en pratique, ce qui motive l'utilisation de techniques plus convoluées. L'une d'elles, appelée méthode de Monte-Carlo par chaînes de Markov, est une application directe de la reconfiguration. Ces algorithmes consistent à commencer avec une configuration arbitraire et à appliquer des opérations de reconfiguration aléatoirement. Sous certaines conditions, la distribution aléatoire des configurations obtenue après un grand nombre d'étapes est proche de la distribution que l'on essaye d'échantillonner. De plus, la vitesse de convergence de cette distribution est intimement liée à des paramètres structurels du graphe de reconfiguration. Ces méthodes sont monnaies courantes en physique statistique, où les modèles de particules ont des distributions d'états trop compliquées à analyser par d'autres moyens. Enfin, l'échantillonnage aléatoire efficace est souvent équivalent à l'approximation du nombre de configurations, ce

qui motive encore le développement de ces méthodes.

Résultats et organisation de cette thèse

Chapitre 1: Préliminaires

Nous donnons dans le Chapitre 1 toutes les définitions et propriétés élémentaires des objets que nous utiliserons dans ce manuscrit. Nous présenterons ici certaines d'entre elles, nécessaires à l'énoncé de nos résultats. Une k -coloration d'un graphe est une affectation de couleurs aux sommets du graphe, de sorte à ce que les sommets adjacents reçoivent des couleurs différentes. Ces couleurs sont généralement numérotées de 1 à k . Le nombre minimum de couleurs nécessaires pour colorier un graphe est appelé *nombre chromatique*. Un graphe est *planaire* s'il peut être dessiné dans le plan sans que ses arêtes se croisent (voir Chapitre 1 pour une définition plus précise de dessins de graphes). Les *arbres* sont les graphes *connexes* et sans *cycles*. Autrement dit, tels qu'il existe une unique manière de se rendre d'un sommet à un autre en suivant les arêtes, pour toute paire de sommets.

Étant donné un graph G fixé, un changement de Kempe est une opération sur les colorations de G , initialement introduite par Kempe pour prouver la Conjecture des Quatre Couleurs : tous les graphes planaires sont 4-coloriables. Un changement de Kempe revient à premièrement choisir un sommet u colorié a et une couleur b , avant de d'inverser récursivement les couleurs a et b , de proche en proche autour de u (voir Figure 3), de sorte à obtenir une autre coloration valide.




Figure 3: Deux 3-colorations du graphe de Petersen qui diffèrent d'un changement de Kempe. Les sommets et arêtes impactés par le changement sont en gras.

Un graphe H est un *mineur* d'un graphe G , si H peut être obtenu à partir de G en supprimant des sommets et des arêtes, et en contractant des arêtes. Dans une suite de vingt articles phares, Robertson et Seymour ont prouvé que les graphes ordonnés par la relation de mineur forment un bel ordre. Ce résultat est une des plus grosse avancée en théorie des graphes et peut être reformulé ainsi : toute famille de graphe close par mineur peut être caractérisée par un nombre

fini de mineurs interdits. Par exemple, les graphes planaires sont exactement les graphes évitant $K_{3,3}$ et K_5 comme mineurs (voir Figure 4), bien que ce résultat soit antérieur à celui de Robertson et Seymour. Une application directe de ce résultat est que le graphe de Petersen (voir Figure 3) n'est pas planaire, puisqu'il contient K_5 comme mineur.




Figure 4: Les mineurs interdits des graphes planaires : $K_{3,3}$ et K_5

Le dernier concept dont nous aurons besoin est celui de *largeur arborescente*, qui est un paramètre de graphe quantifiant la ressemblance à un arbre. Informellement, un graphe G a largeur arborescente au plus k s'il peut être couvert par un arbre, dont les sommets sont des sacs contenant au plus $k + 1$ sommets de G . On appelle *décomposition arborescente* une telle couverture (voir Chapitre 1 pour une définition plus précise).

Chapitre 2 : Recoloration avec des changements de Kempe

Une majeure partie de cette thèse est consacrée à la recoloration de graphes via les changements de Kempe. Ceux-ci sont déterminants pour obtenir des bornes sur le nombre chromatique de plusieurs classes de graphes, ce qui illustre notamment le fait que la reconfiguration peut servir à prouver l'existence de configurations solutions du problème original. Le Chapitre 2 est un état de l'art de la reconfiguration avec les changements de Kempe. Nous y présentons les avancées récentes de ce domaine, ainsi qu'une sélection de conjectures ouvertes et de directions de recherche que nous jugeons pertinentes.

Chapitre 3 : Changements de Kempe dans des graphes creux

Dans le Chapitre 3, nous présentons nos résultats sur la recoloration via les changements de Kempe dans des classes de graphes creux, c'est-à-dire contenant peu d'arêtes. Plus précisément, nous étudions le graphe de reconfiguration de Kempe dans plusieurs contextes, comme celui des graphes planaires et des graphes de largeur arborescente ou de degré borné. Nous prouvons plusieurs bornes polynomiales sur le diamètre du graphe de reconfiguration dans ces classes de graphes. Dans un second temps, nous réfutons le pendant recoloration de la conjecture de Hadwiger. La conjecture de Hadwiger est un problème ouvert central en théorie des graphes, qui est une généralisation du Théorème des Quatre Couleurs. Bien que

la variante de recoloration que nous étudions n'implique pas ni ne soit impliquée par la conjecture de Hadwiger, les deux conjectures présentent des connexions intéressantes. De plus, notre résultat soulève des questions portant sur une notion relâchant la relation de mineur, que nous appelons *quasi-mineur*.

Le contenu de ce chapitre est indépendant de celui du Chapitre 2. Nous conseillons toutefois au lecteur non-familier avec les changements de Kempe de lire le Chapitre 3 après l'introduction du Chapitre 2.

Chapitre 4 : Algorithmes optimaux en reconfiguration

Dans le Chapitre 4, nous étudions d'autres espaces de configurations et opérations, provenant de problèmes de topologie en basse dimension. Dans une première partie, nous étudions la complexité paramétrée d'un problème de théorie des nœuds. Un nœud peut être vu comme un bout de corde dont les extrémités ont été collées ensemble. L'étude des nœuds a débuté en 1860, motivée par de potentielles applications en électromagnétisme puis pour expliquer l'existence de différents éléments chimiques. Un des problèmes centraux de la théorie des nœuds est de les classer. Mis à plat, les nœuds forment des diagrammes, avec des croisements où deux brins de la corde se chevauchent. Reidemeister, Alexander et Briggs [AB26, Rei27] ont montré que deux diagrammes proviennent d'un même nœud si et seulement si ils sont équivalents via une séquence de modifications locales appelées *mouvements de Reidemeister* (voir Figure 5). Il est naturel de se demander combien de mouvements de Reidemeister sont pour cela nécessaires, notamment dans le cas restreint du nœud trivial, un simple cercle, en cherchant à atteindre le diagramme sans croisement. Ce problème est NP-complet et nous déterminons sa classe de complexité paramétrée, en utilisant une méthode s'appliquant à d'autres problèmes de reconfiguration et permettant d'obtenir des algorithmes optimaux.




Figure 5: Les trois mouvements de Reidemeister

Dans une seconde partie, nous étudions la reconfiguration des surfaces à petits carreaux, qui sont un cas particulier de surfaces de translation et plus particulièrement de quadrangulations. Les surfaces à petits carreaux sont composées de carrés collés par leurs côtés, de telle sorte à ce que les côtés horizontaux soient collés entre eux, et de même pour les côtés verticaux. Les surfaces à petits carreaux peuvent être vues comme les points entiers des espaces de modules des différentielles quadratiques d'une surface de Riemann [Zor02] et sont aussi motivées par

l'étude des billards polygonaux. Elles peuvent être classifiées en strates en fonction des angles autour des coins des carrés, aussi appelés *singularités coniques*. Les strates fixent entre autre le genre et la taille des quadrangulations. L'opération de reconfiguration que nous considérons est appelée *cisaillement* et ressemble aux rotations de faces du Rubik's Cube, bien que celui-ci ne soit pas une surface à petits carreaux. Ces cisaillements consistent à couper la surface le long d'un cercle, afin de faire glisser les deux côtés l'un contre l'autre avant de les recoller. Dans un sous-ensemble de strates, nous prouvons que les surfaces à petits carreaux sont toutes équivalentes via $O(n)$ cisaillements, où n est le nombre de carrés de la surface à petits carreaux.

Chapitre 5 : Un théorème de structure sur les graphes localement finis quasi-transitifs évitant un mineur

Dans le Chapitre 5, nous laissons de côté la reconfiguration pour présenter un résultat entremêlant largeur arborescente, mineurs et surfaces, qui sont des notions transverses dans cette thèse.

En 1961, Wang [Wan61] étudie le problème de pavage qui prendra son nom, demandant si un ensemble fini de tuiles carrées coloriées sur leurs côtés peut paver le plan, de sorte que les tuiles adjacentes aient la même couleur sur leurs côtés mitoyens. En 1966, Berger [Ber66] montre que le problème du pavage de Wang est undécidable, grâce à une réduction au problème de l'arrêt et en exhibant un ensemble de tuiles ne pouvant paver le plan que de manière apériodique. Le problème du domino est une généralisation de ce problème de pavage et de la coloration de graphe aux graphes de Cayley de groupes plus complexe que le plan $(\mathbb{Z}^2, +)$. En 2018, Ballier et Stein [BS18] conjecturent que le problème du domino est décidable uniquement dans les groupes virtuellement libres, c'est à dire ceux admettant un graphe de Cayley de largeur arborescente bornée.

Pour attaquer ce problème, une idée naturelle est de décomposer le graphe de Cayley en sous-graphes bien compris et de vérifier si le problème du domino est décidable sur chacun d'entre eux avant de recoller les morceaux ensemble. Dans le dix-septième article de leur série, Robertson et Seymour prouvent un autre résultat central de la théorie moderne des graphes en donnant un théorème de décomposition des graphes évitant un mineur : Tout graphe G évitant un mineur H fixé peut être "décomposé en torsos", qui sont des graphes presque plongeables dans des surfaces de genre borné, recollés entre eux sur de petits ensembles de sommets. Au delà de leur intérêt structurel, les décompositions arborescentes sont très utiles pour concevoir des algorithmes efficaces. Cependant, dans les graphes infinis quasi-transitifs (c'est à dire ceux dont le groupe d'automorphismes contient un nombre fini d'orbites), la décomposition arborescente de Robertson et Seymour peut contenir un nombre infini de torsos différents et ainsi perdre les symétries

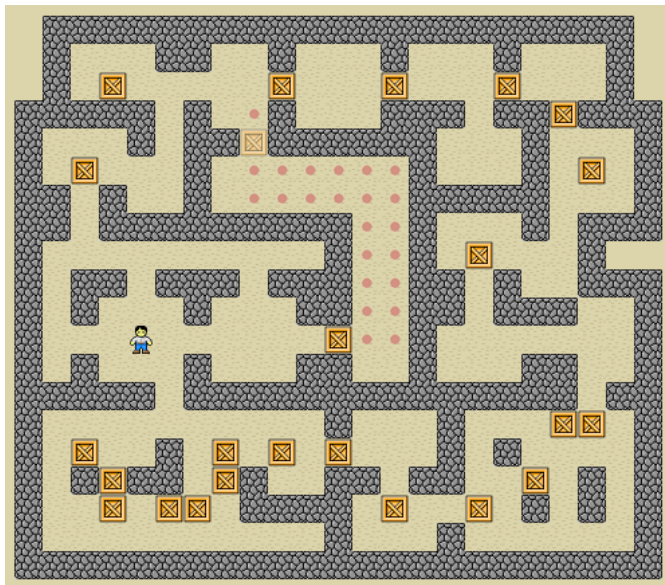
du graphe d'origine. Par conséquent, ce théorème de structure ne semble être d'aucune utilité pour prouver des résultats de décidabilité.

Nous donnons une décomposition arborescente évitant cet écueil, s'appliquant à tout graphe G localement fini quasi-transitifs qui évite un mineur. Plus précisément, celle-ci est invariante sous l'action du groupe d'automorphismes de G et ses torsos sont planaires ou finis. Parmi plusieurs applications de ce résultat, nous prouvons le cas des groupes évitant un mineur de la conjecture de Ballier et Stein.

Introduction in English

Combinatorial puzzles are recreational brain teasers in which a budding mathematician modifies the state of the puzzle to reach a specific configuration. Perhaps the most popular example of combinatorial puzzle is the *Rubik's Cube*, that received a worldwide interest in the eighties and remains the bestselling toy as of today, with over 500 millions of cubes sold, on top of the several existing variants of the original $3 \times 3 \times 3$ cube.

As a kid, I particularly enjoyed another combinatorial puzzle called *Sokoban*. Sokoban is a video game released in 1982, in which the player pushes boxes in a warehouse to bring them to their designated location (see Figure 6). The player can move in the warehouse and push the boxes, unless there is a wall or another box behind it. However, the player cannot travel through a box or a wall, nor pull the boxes.



(a) The designated locations of the boxes are represented by the red dots. The semi-transparent box is already over a designated location.



(b) Me showing an interest in Sokoban at a young age.

Figure 6: A level of Sokoban

There are many more famous examples of combinatorial puzzles. One could think for example of the *Tower of Hanoi* (see Figure 7), invented in 1883 by the mathematician Édouard Lucas, where disks of various diameter are stacked on three rods in a pyramid fashion: no disk can be above a smaller one. The player is allowed to move the top disk of any stack, while preserving the pyramidal stack

constraint. The initial and target configurations consist usually of two empty rods and either the left or the right rod with n disks stacked on it. This puzzle is a classic of science fairs and Computer science unplugged activities, and is well known for requiring $2^n - 1$ moves to go from the initial configuration to the target configuration.



Figure 7: A intermediate configuration of the Tower of Hanoi

Part of the charm of combinatorial puzzles comes from the complexity of their resolution and the tremendous number of configurations (43 quintillions different configurations in the case of the Rubik’s Cube), contrasting with the apparent simplicity of the objects. To solve them efficiently, players began designing algorithms, aiming for *God’s algorithm*: an algorithm solving the puzzle in a minimum number of moves. This is sometimes not possible efficiently: Culberson proved in 1997 that determining whether a Sokoban instance is solvable is PSPACE-complete [Cul97]. However, there is a vast history of algorithms sorting the Rubik’s Cube and its variants, and of upper bounds on the number of necessary moves to do so. In 2014, Rokicki, Kociemba, Davidson and Dethridge [RKDD14] proved with a clever computer search that any scrambled configuration of the Rubik’s Cube could be solved in at most 20 face turns, which is tight. There exist various Rubik’s Cube competitions, with trials such as solving the Rubik’s Cube in the fewest possible moves, or as fast as possible².

²The current world records were set in 2023 by Wong Chong Wen with 20 moves in average

In more serious contexts, we sometimes need to modify a configuration while maintaining some constraints. To illustrate this, we present a real-life application in the design of quantum computers. Quantum information processing often requires assembling quantum particles in a predefined pattern. Assume that you dispose of an array of optical traps, each of them able to trap an atom for a short time. The atoms are initially sprayed on this lattice in random manner and can be released by the traps or moved one by one to an adjacent empty optical trap, until reaching the target configuration. Due to their high instability, atoms may be lost in the process and rearranging the atoms as fast as possible improves significantly the chances of success. Two factors impact this preparation time and success rate: the processing time to compute the sequence of moves, and the total distance traveled by the atom in the lattice. The number of potential configurations scales exponentially with the size of the lattice. Along with the strong time constraint, this makes the design of fast algorithms to reconfigure the atoms efficiently a real challenge (see [CESB⁺23, ESD⁺23] for a more detailed explanation of this problem and its motivations).

Modifying a configuration can be motivated by several reasons and is sometimes simply unavoidable, as we have seen in the last example. The first motivation is to avoid a service failure while the configuration changes. This often occurs in networks problems for example. The second reason is that finding a solution from scratch may be very expensive, as the configurations are often solutions of hard combinatorial problems, while performing small modifications is much cheaper. *Reconfiguration* is the study of such problems, where we are interested in the structure of a configuration space, with respect to some elementary operation. A reconfiguration problem is a set of configurations, along with a small set of (elementary) operations transforming a configuration into another.

To model both the space of configurations and the configurations themselves, we will often use *graphs*, which are combinatorial objects composed of *vertices* connected by *edges* (see Chapter 1 for all definitions). For example, an atom configuration can be modeled by a grid whose vertices are the optical traps, and whose edges connect the adjacent optical traps. The atoms are then represented by a set of tokens placed on the vertices of the graph and allowed to slide one by one along the edges. To represent the space of configurations, we will often use the *reconfiguration graph*, whose vertices are all the possible configurations and in which two configurations are connected by an edge if they differ by one of the admissible elementary moves. Note that the size of this graph being the number of configurations, it is often very large and thus hard to analyze.

and Yiheng Wang in 4.48 seconds in average. For the fewest possible move competition, the contestant is given one hour to write down his solution. For both records, the average is taken from three tries out of five on different initial configurations, removing the best and the worst case.

The recurring questions in reconfiguration are the ones we previously illustrated. Are two given configurations A and B *equivalent*, that is can one go from A to B with a sequence of reconfiguration operations? For example in the case of the Rubik's Cube, manually rotating a corner of the cube yields a configuration that is not equivalent to the unscrambled configuration via face turns. In the world of graphs, this question can be restated as follows: is there a path from A to B in the reconfiguration graph? One is often interested in efficient algorithms to determine whether it is the case, or yielding a reconfiguration sequence as short as possible. More generally, one can wonder whether all configurations are equivalent and what are the configurations at maximal distance in the reconfiguration graph, and how hard these questions are for a computer.

Reconfiguration is also motivated by random sampling. The set of configurations is often too large and complicated to be enumerated in practice, which motivates the use of more advanced methods. One of them, called Markov Chain Monte Carlo algorithms, is a direct application of reconfiguration. These algorithms consist in starting from an arbitrary configuration, and performing random reconfiguration operations. Under certain conditions, the distribution of the configuration obtained after a large number of steps approximates the desired distribution on the configurations. Moreover, the rate of convergence of this distribution is closely related to some structural parameters of the reconfiguration graph. These methods are particularly useful in statistical physics, where some models of particles have a distribution of states difficult to analyze by other means. Finally, efficient random sampling is in many cases equivalent to approximating the number of configurations, which further motivates the development of these methods.

Results and organization of this thesis

Chapter 1: Preliminaries

We give in Chapter 1 all the definitions and basic properties of the objects we will use in this manuscript. We will present here some of those, needed to describe our results. A k -*coloring* of a graph is an assignment of colors to the vertices of the graph, such that adjacent vertices receive different colors. These colors are usually numbered from 1 to k . The least number of colors needed to color a graph is called its *chromatic number*. A graph is said to be *planar* if it can be drawn in the plane without crossing edges (see Chapter 1 for a more precise definition of graph drawing). The *trees* are the *connected* and *acyclic* graphs. In other words, the graphs such that there exists a unique way of going from one vertex to another by moving along the edges, for every pair of vertices.

Given a fixed graph G , a *Kempe change* is an operation on colorings of G

initially introduced in an attempt to prove the Four Color Theorem: all planar graphs are 4-colorable. First, one chooses a vertex u colored say a and a color b , then the Kempe change consists in changing the color of u from a to b and, in order to obtain a coloring, to recursively switch the colors a and b , around u , whenever an edge is not properly colored (see Figure 8).




Figure 8: Two 3-colorings of the Petersen graph that differ from one Kempe change. The vertices and edges impacted by the Kempe change are thickened.

A graph H is a *minor* of G , if H can be obtained from G by deleting vertices and edges, and by contracting edges. In a series of twenty seminal articles, Robertson and Seymour proved that the graph minor relation is a well-quasi-ordering. This result is one of the biggest breakthrough in structural graph theory and can be restated as follows: Every minor-closed family of graphs can be characterized by a finite number of forbidden minors. For example, although this was known before, planar graphs are exactly the graphs that forbid $K_{3,3}$ and the clique K_5 as a minor (see Figure 9). For example, the Petersen graph (see Figure 8) is not planar as it contains K_5 as a minor.




Figure 9: The forbidden minors for planar graphs: $K_{3,3}$ and K_5

The last concept we will need is that of *treewidth*, which is a graph parameter that measures how much a graph looks like a tree. Informally, a graph G has treewidth at most k if it can be covered by a tree, whose vertices are bags containing at most $k + 1$ vertices of G . Such a covering is called a *tree-decomposition* (see Chapter 1 for a more precise definition).

Chapter 2: Recoloring with Kempe changes

A significant part of this thesis deals with the reconfiguration of graph colorings with Kempe changes. Kempe changes are decisive to determine the chromatic number of several graph classes, which illustrates that reconfiguration can also be useful to prove the existence of configurations solving a problem. Chapter 2 is a survey of reconfiguration with Kempe changes. It presents recent advances on Kempe recoloring as well as a selection of open conjectures and directions of research that we find relevant.

Chapter 3: Kempe changes in sparse graphs

In Chapter 3, we present our results on reconfiguration with Kempe changes in sparse graph classes, that is graphs with few edges. More precisely, we study the Kempe reconfiguration graph in various settings, such as planar graphs, graphs of bounded degree or bounded treewidth, by proving several polynomial upper bounds on its diameter. In a second part, we disprove a recoloring version of Hadwiger's conjecture. Hadwiger's conjecture is a central open problem in graph theory and is a generalization of the Four Color Theorem. It conjectures that graphs with no K_t minor are $(t-1)$ -colorable. Although the reconfiguration variant we study neither implies nor is implied by Hadwiger's Conjecture, it presents interesting connections with it. Furthermore, our results raise questions on a substructure weakening minors, that we call *quasi-minors*.

The content of this chapter is independent from Chapter 2, but we advise the reader unfamiliar with Kempe changes to read Chapter 3 after the introduction of Chapter 2.

Chapter 4: Optimal algorithms for reconfiguration

In Chapter 4, we study other configuration spaces and operations, arising from low-dimensional topology. In a first part we study the parameterized complexity of a knot theory problem. A knot can be thought of as a piece of rope whose ends are joined together. The study of knots traces back to 1860, motivated by applications in electromagnetism, then to explain the existence of different chemical elements. One of the main problems in knot theory is to classify the knots. When laid flat, the knots form diagrams, with crossings at which a piece of the rope overpasses or underpasses another. Reidemeister, Alexander et Briggs [AB26, Rei27] showed that two diagrams correspond to a same knot if and only if they are equivalent up to a sequence of local modifications called *Reidemeister moves* (see Figure 10). When the knot is the trivial knot, a simple circle, one can wonder how many Reidemeister moves are needed to go from an arbitrary diagram to the diagram with no crossings. This problem is known to be NP-complete and we determine its

parameterized complexity, using a method producing optimal algorithms for other reconfiguration problems.




Figure 10: Reidemeister moves

In a second part, we study the reconfiguration of square-tiled surfaces, which are a special type of translation surfaces and more specifically quadrangulations. Square-tiled surfaces are composed of squares glued along their edges, such that no horizontal side is glued to a vertical side. Square-tiled surfaces can be seen as the integer points of the moduli space of quadratic differentials of a Riemann Surface [Zor02] and are also motivated by applications to polygonal billiards. They can be classified into different strata according to the total angles around the corners of the squares, also known as *conical singularities*. Among other parameters, strata fix the genus and the size of the quadrangulation. The reconfiguration operation we consider, called *shear*, is similar to the face turns on a Rubik's Cube, even though the Rubik's Cube is not a square-tiled surface. It consists in cutting the surface along a circle and sliding both sides before gluing them together again. In a subset of strata, we prove that the square tiled-surfaces are equivalent up to $O(n)$ shears, where n is the number of squares.

Chapter 5: A structure theorem for locally finite quasi-transitive graphs

In Chapter 5, we move away from reconfiguration to present a result that intertwines with treewidth, minors and surfaces, which are notions studied throughout this thesis.

In 1961, Wang [Wan61] studies the tiling problem named after him, which asks whether a given finite set of square tiles colored on their edges can tile the plane, such that adjacent tiles have the same color on their shared side. In 1966, Berger [Ber66] showed that the Wang tiling problem is undecidable, by a reduction from the halting problem and by exhibiting a set of tiles that can only tile the plane aperiodically. The domino problem is a generalization of this tiling problem and of graph colorings to Cayley graphs of groups more complicated than the plane $(\mathbb{Z}^2, +)$. In 2018, Ballier and Stein [BS18] conjectured that the Domino Problem is decidable if and only if the group is virtually-free, that is if it has a Cayley graph of bounded treewidth.

To tackle this problem, an idea is to decompose the Cayley graph in well-understood subgraphs and check if the Domino Problem is decidable in each of

them, before putting the pieces together. In the seventeenth article of their series, Robertson and Seymour proved another central result in modern graph theory by giving a structure theorem on graphs avoiding a fixed minor: Any H -minor-free graph G can be “decomposed into torsos”, that are graphs almost embeddable on surfaces of bounded genus, glued on small subsets of vertices. Beyond their structural interest, tree-decompositions are very useful to design efficient algorithms. However, in quasi-transitive graphs (infinite graphs with many symmetries (namely whose automorphism group has finitely many orbits), this tree-decomposition can contain infinitely many different torsos and hence does not reflect these symmetries. Thus this structure theorem cannot be used to prove decidability results.

For any quasi-transitive graph G avoiding a minor, we give a tree-decomposition invariant under the automorphism group of G , with planar or finite torsos. Among other applications of this result, we prove the minor-excluded case of the conjecture of Ballier and Stein.

What is not in this thesis

Polytope reconstruction

With Jane Tan, we worked on a reconstruction conjecture of Kalai [Kal94, Conjecture 8]. The term reconstruction covers all questions of the type “What partial information uniquely characterizes an object in a collection?”. Here we are interested in reconstructing the combinatorial type of simplicial d -polytopes. A natural piece of information for that is their k -skeleton, which contains the lists of all faces of dimension at most k . Dancis proved that one can recover the combinatorial type of any simplicial d -polytope from its $\lfloor d/2 \rfloor$ -skeleton [Dan84]. However, the $(\lfloor d/2 \rfloor - 1)$ -skeleton does not provide enough information. In fact the class of neighborly d -polytopes on n vertices, which we will not define here, has a unique $(\lfloor d/2 \rfloor - 1)$ -skeleton [She82] and there are $2^{\Theta(n \log n)}$ such polytopes [Pad13].

Kalai conjectured that the $(k - 1)$ -skeleton together with the space of affine k -stresses uniquely determine the face lattice of simplicial d -polytopes. The space of affine k -stresses of a polytope P is the space of affine dependencies on subsets of at most k vertices of P . For $k = 2$, these stresses can be thought of as an assignment of compensating forces applied along the edges of the 1-skeleton, that is the underlying graph of P . Novik and Zheng [NZ23] confirmed Kalai’s conjecture for $k = 2$ and for neighborly polytopes. With Jane Tan, we observed that the proof of Novik and Zheng seems to generalize to higher k , provided that a generalization of Balinski’s theorem holds.

Random embeddings of bounded degree trees with optimal spread

With Paul Bastide and Alp Müyesser, we worked on a problem connected to recurring questions in extremal combinatorics. Given a fixed graph H on n vertices, what is the infimum δ_n such that any n -vertex graph G of minimum degree δ_n contains H as a subgraph? The first example of such threshold was given in 1952 by the seminal result of Dirac [Dir52]: any n -vertex graph of minimum degree at least $n/2$ is Hamiltonian. Since then, an important branch of extremal combinatorics has been dedicated to finding such thresholds and studying the properties of the graph around this threshold. Above this minimum degree threshold, G tends to contain many different copies of H and counting them is also an extensively studied question.

This counting question is related to the following random graph question: For which *threshold* value p does the random graph $G(n, p)$ contain H as a subgraph with probability $1/2$? This third question can be considered from the perspective of random p -sparsification: removing each edge of a clique on n vertices with probability $1-p$ yields $G(n, p)$. This brings us to the last question which generalizes the previous ones: For which value of p does the p -sparsification of any n -vertex graph of minimum degree δ_n contain H with probability $1/2$?

These four questions have been studied separately for various graphs H , including Hamiltonian paths and Hamiltonian cycles, K_r -factors, perfect matchings and spanning trees of bounded degree. Park and Pham [PP24] recently proved Kahn-Kalai conjecture, that relates the threshold to the so-called *expectation threshold* p at which the expectation of the number of copies of H in $G(n, p)$ is 1. Using this result, Kelly, Müyesser and Pokrovskiy [KMP23] showed that the existence of certain probability distributions on the space of embeddings of H into any graph n -vertex graph G with minimum degree at least δ_n , called $O(1/n)$ -*spread* distributions, provides a unified approach to the three last questions (counting the embeddings, embedding in a random graph and in a random sparsification). One way to obtain these $O(1/n)$ -spread distributions is to build a random algorithm that embeds H in G progressively, such that each vertex has linearly many options among the remaining vertices at each step. Kelly, Müyesser and Pokrovskiy [KMP23] proved the existence of a $O(1/n)$ -spread distribution on embeddings of Hamiltonian cycles and Pham, Sah, Sawhney and Simkin [PSSS22] for K_r -factors, perfect matchings and spanning trees of bounded degree. With Paul Bastide and Alp Müyesser, we give another $O(1/n)$ -spread distribution on embeddings of bounded degree spanning trees. Our framework is flexible and unlike the proof of Pham, Sah, Sawhney and Simkin, it avoids the use of Szemerédi's regularity lemma. As a result, it generalizes immediately to hypertrees and digraphs, and we obtain better constants.

Chapter 1

Preliminaries

For every $n \in \mathbb{N}^*$, we let $[n] := \{1, 2, \dots, n\}$ and S_n denote the symmetric group over n elements. A graph is a structure encoding a symmetric relation within a set. More precisely, an undirected simple *graph* G consists in a set $V(G)$ of *vertices* and a set $E(G) \subseteq V(G)^2$ of edges, such that $(u, v) \in E(G)$ if and only if $(v, u) \in E(G)$. When that is the case, we say that uv is an edge of G or that u and v are *adjacent*.

We say that two edges are *incident* if they share a vertex. Several variations of the definition of graph can be considered: in a *multi-graph*, two vertices can be adjacent via several edges, in a *directed graph*, the symmetric condition $(u, v) \in E(G) \Leftrightarrow (v, u) \in E(G)$ is dropped. A *loopless* graph is a graph in which no vertex is adjacent to itself. Note that reconfiguration graphs defined are often multigraphs with loops. Unless stated otherwise, all other graphs will be loopless and simple.

The equivalence classes of the adjacency relation are called the *connected components* of a graph and a graph is *connected* if it has only one connected component. A graph is *connected* if there is a path between any two vertices. A graph is *k-connected* if it contains at least k vertices and if the removal of any $k - 1$ vertices yields a connected graph. A *k-connected-component* is a maximal k -connected subgraph and 2-connected components are referred to as *blocks*. A graph is said to be *quasi-4-connected* if it is 3-connected and for every set $S \subseteq V(G)$ of size 3 such that $G - S$ is not connected, $G - S$ has exactly two connected components and one of them consists of a single vertex. A *separator* in a graph is a subset of vertices that if removed disconnects the graph. An *s, t-separator* of G is a subset X of vertices such that s and t lie in different connected components of $G \setminus X$.

1.1 Graphs, surfaces and colorings

The notions in this section are used in all the chapters of this thesis.

Let G be an n -vertex graph. The following definitions hold even if G is infinite,

with multiple edges or loops. The *open neighborhood* of a vertex u is the set of vertices adjacent to u and is denoted $N_G(u)$. These vertices are called *neighbors* of u in G . We denote $N_G(X)$ the set of vertices that are adjacent to some vertex in $X \subseteq V(G)$. When the graph G is clear from the context we will drop the subscript and write $N(X)$ instead of $N_G(X)$. The *close neighborhood* of u is $N[u] = n(u) \cup \{u\}$. We will often refer to the open neighborhood simply as the neighborhood. The *degree* of a vertex is the size of its neighborhood. The minimum and maximum degree of a graph are denoted $\delta(G)$ and $\Delta(G)$. A graph is r -*regular* if all its vertices have degree r and 3-regular graphs are called *cubic*.

1.1.1 Graph relations

A graph *homomorphism* is a mapping $f : G \rightarrow H$ that preserves the edges, that is $\forall uv \in E(G), \phi(u)\phi(v) \in E(H)$.

We say that H is an *induced subgraph* of G if H can be obtained from G by deleting vertices (see Figure 1.1 for an example of the different relations defined in this subsection). If the set of remaining set of vertices is X , we denote $G[X]$ the corresponding graph, and call it the graph *induced by X* . If we also allow edge deletions, we say that H is a *subgraph* of G . Note that this is equivalent to saying that there is a injective homomorphism from H to G . We call H -*free* graphs the class of graphs that do not contain H as an induced subgraph.

Finally, if we allow edge contractions as well, that is identification of adjacent vertices, we say that H is a *minor* of G . Minors are especially important in the graph theory landscape, thanks to the celebrated Graph Minor Theorem of Robertson and Seymour, that states that the minor relation is a well-quasi-ordering [RS04]. In other words, there exists no infinite family of graphs incomparable for the minor relation, and equivalently, every minor-closed family is characterized by a finite set of forbidden minors.




Figure 1.1: The different graph relations

We call H -*minor-free* graphs the class of graphs that do not contain H as a minor.

A *model* of H in G is a family $(V_v)_{v \in V(H)}$ of pairwise disjoint vertex subsets of G such that each V_v induces a connected subgraph of G , and for each $uv \in E(H)$, there exists $u' \in V_u, v' \in V_v$ such that $u'v' \in E(G)$. The sets V_v are often referred to as bags, but we will call them *blobs* as the name bag interferes with the definition of tree-decompositions that we will present later. Note that H is a minor of G if and only if there is a model of H in G . When $V(H) \subseteq V(G)$, a model $(V_v)_{v \in V(H)}$ of H in G is said to be *faithful* if for each $v \in V(H), v \in V_v$. H is a *faithful minor* of G if it admits a faithful model in G .

1.1.2 Special graphs

The *independent set* I_t is a set of t vertices in which no vertices are adjacent. On the contrary, the *clique* K_t is a set of t vertices that are all pairwise adjacent. We call *triangle* the clique on three vertices. The *path* of length t is a set of $t + 1$ vertices w_0, \dots, w_t such that w_i and w_{i+1} are adjacent for all i . We denote P_t a path on t vertices (and thus $t - 1$ edges) and call *extremities of the path* the vertices w_0 and w_t .

A *cycle* is a cyclic sequence w_1, \dots, w_p such that for all i , w_i and w_{i+1} are adjacent, and w_1 is adjacent to w_p . We denote C_t a cycle on t vertices (and edges).

We say that a graph G *contains* a cycle, a clique or a path if it is a subgraph of G . For independent sets however, we require it to be an induced subgraph. For any of these graphs, we say that it is an independent set (resp. clique ...) of G if G contains it. The independence number $\alpha(G)$ and the clique number $\omega(G)$ denote the size of the largest independent set and clique of G .

There is a *path between u and v* if G contains a path, whose extremities are u and v . The *distance* between u and v is the length of the shortest path between them. The *diameter* of a graph is the maximum distance between any two vertices of a graph. We will denote it $\text{diam}(G)$.

A graph is *acyclic* if it contains no cycle. A *tree* is a connected and acyclic graph. A tree on n vertices contains exactly $n - 1$ edges. The *girth* of a graph is the length of its shortest cycle.

1.1.3 Colorings

Let $k > 0$, a *k -coloring* of a graph is a map $\alpha : V(G) \rightarrow [k]$, where $[k] = \{1, \dots, k\}$ is the set of *colors*, such that adjacent vertices have distinct colors. In particular, a k -coloring is also a $k + 1$ -coloring. When the number of colors is irrelevant or clear from the context, we will simply say coloring. The *chromatic number* of a graph G , denoted $\chi(G)$, is the minimum number of colors k such that a k -coloring of G exists. This notion can be generalized as follows. A *list assignment* is a function $L : V(G) \rightarrow \mathcal{P}(\mathbb{N})$, and an *L -coloring* of G is a coloring such that $\alpha(u) \in L(u)$

for all $u \in V(G)$. A graph is k -choosable if it admits an L -coloring, for all list assignments L such that $|L(u)| \geq k$ for all u . The *list-chromatic number* is denoted $\chi_\ell(G)$ and is the smallest k such that G is k -choosable. A k -edge-coloring is a map $\alpha : E(G) \rightarrow [k]$, such that incident edges have distinct colors. The *chromatic index* of a graph is the smallest number of colors needed to edge-color it and is denoted $\chi'(G)$. Note that homomorphisms are another generalization of graph colorings as a k -coloring is simply an homomorphism to the clique on k -vertices.

Let G be an n -vertex graph, we have $\chi(G) \leq \chi_\ell(G)$ and $\omega(G) \leq \chi(G) \leq \Delta(G) + 1$. Since all color classes are independent sets, we also have $\frac{\alpha(G)}{n} \leq \chi(G)$. As we will see in Chapter 2, Vizing proved that the chromatic index of a graph is always between its maximum degree and its maximum degree plus one [Viz64].

A graph is k -critical if it has chromatic number k and all its proper subgraphs are $(k - 1)$ -colorable.

1.1.4 Surfaces

We assume here basic knowledge of topology.

A *surface* is a manifold of dimension two, that is a topological space in which the neighborhood of all points is homeomorphic to the open disk. In more simple words, the surface looks locally flat around all points. A surface is *orientable* if no close simple curve going in one direction can be continuously deformed to itself, but going in the other direction. Most of the surfaces we will consider in this thesis are connected, closed and orientable.

The *genus* of an orientable surface is the maximum number of non-intersecting simple closed curves that can fit in the surface without disconnecting it.

A *drawing* of a graph on a surface is a map of its vertices to distinct points of the surface, together with a map of the edges to arcs on the surface such that

- for any edge $uv \in E(G)$, the endpoints of the arc corresponding to uv are the images of u and v ,
- no image of a vertex is in the interior of an arc.

If furthermore the interior of the arcs are disjoint, then we call it a *graph embedding*. If the embedding is fixed, we call the embedded graph a *combinatorial map*. Each connected component of its complement is called a *face*. Each edge of the graph is adjacent to two faces (possibly twice the same). Each edge on the combinatorial map has two *sides* an edge together one of these two adjacencies. The *degree* of a face is the number vertices (or edges) of its sides.

The *genus* of a graph is the minimum genus of a closed orientable surface it can be embedded in. The sphere has genus zero and the closed orientable surface

of genus one is called a *torus*. A *plane drawing* or *plane embedding* is a drawing or embedding on the sphere.

The *crossing number* is the lowest number of crossings in a plane drawing of G . A graph is *planar* if it has crossing number zero, that is if it can be embedded on the sphere (or the plane), such that no two edges cross. We call it a *plane graph* if this embedding is fixed.

Euler's formula then relates the number of vertices N , of edges E and of faces F with the genus, and holds even for multi-graphs:

$$N - E + F = 2 - 2g.$$

1.1.5 Graph operations

The *join* of two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ is denoted $G_1 \nabla G_2$ is the graph on the vertex set $V_1 \sqcup V_2$, such that V_i induces G_i for $i \in \{1, 2\}$ and all vertices of V_1 are adjacent to all vertices of V_2 .

There exist several graphs products: the *categorical product* \times , the *Cartesian product* \square and the *strong product* \boxtimes . Given two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, all three yield a graph defined on the vertex set $V_1 \times V_2$. Two vertices (u_1, v_1) and (u_2, v_2) in $V_1 \times V_2$ are adjacent

- in $G_1 \times G_2$ if u_1 is adjacent to u_2 and v_1 is adjacent to v_2 ,
- in $G_1 \square G_2$ if $u_1 = u_2$ and v_1 is adjacent to v_2 or if u_1 is adjacent to u_2 and $v_1 = v_2$,
- in $G_1 \boxtimes G_2$ if one of the above conditions is verified.

Note that the symbols \times , \square and \boxtimes correspond to the graph obtained by taking the product of two edges for each operation (see Figure 1.2 for an example).




Figure 1.2: The different graphs products of P_2 by P_3

Let G be a graph embedded on a surface. Its *medial graph* $\text{Med}(G)$ is the graph whose vertices are the midpoints of the edges of G and in which two vertices are

adjacent if their corresponding edges appear consecutively around one of the faces of G . Note that this construction may depend on the embedding of G .

1.1.6 Graph classes

A graph is *bipartite* if it admits a 2-coloring. Equivalently, bipartite graphs are exactly the graphs that contain no induced odd-cycle.

The *complete bipartite graph* $K_{s,t}$ is the join of two independent sets of size s and t . The renowned Wagner's theorem [Wag37] characterizes planar graphs as those that do not contain K_5 or $K_{3,3}$ as a minor. The celebrated Four Color Theorem states that all planar graphs are 4-colorable [AH89, RSST97].

The maximum average degree, denoted mad , is defined as the maximum of the average degree of its subgraphs, that is

$$\text{mad } G = \max_{H \subseteq G} \frac{|E(H)|}{|V(H)|}.$$

As a immediate consequence of Euler's formula, planar graphs have maximum average degree less than six.

A graph is *perfect* if its chromatic number equals its trivial lower bound: the clique number. A *hole* in a graph G is an induced subgraph isomorphic to a cycle on at least four vertices and an *anti-hole* is an induced subgraph whose complement is a hole in the complement of G . The celebrated Strong Perfect Graph Theorem, whose proof is due to Chudnovsky, Robertson, Seymour, and Thomas [CRST06], characterizes perfect graphs as the graphs that have no odd-holes and no odd-anti-holes.

A graph is *chordal* if it contains no induced cycle of length at least four. Chordal graphs are exactly the graphs that admit a perfect elimination ordering, that is an ordering v_1, \dots, v_n of the vertices, such that for all i , $N(v_i) \cap \{v_{i+1}, \dots, v_n\}$ induces a clique. Using this ordering, chordal graphs can be greedily colored, which proves that chordal graphs are perfect.

Cographs are the P_4 -free graphs. Equivalently, cographs can be defined inductively as follows: a single vertex is a cograph, the disjoint union of two cographs is a cograph and the join of two cographs is a cograph. Cographs are also a special case of perfect graphs.

A graph is *d -degenerate* if all its subgraphs contain a vertex of degree at most d . Alternatively, a graph is d -degenerate if and only if it admits a *degeneracy ordering*: an ordering v_1, \dots, v_n such that each v_i has at most d neighbors among $\{v_{i+1}, \dots, v_n\}$. In particular, d -degenerate graphs are $d + 1$ -colorable. Note that planar graphs are 5-degenerate (by Euler's formula), but 4-colorable, so the converse does not hold. Note also that graphs of maximum average degree at most m are $\lceil m - 1 \rceil$ -degenerate.

1.1.7 Tree-decompositions

A *tree-decomposition* of a graph G is a pair (T, \mathcal{V}) where T is a tree and $\mathcal{V} = (V_t)_{t \in V(T)}$ is a family of subsets V_t of $V(G)$ such that:

- $V(G) = \bigcup_{t \in V(T)} V_t$;
- for all nodes t, t', t'' such that t' is on the unique path of T from t to t'' , $V_t \cap V_{t''} \subseteq V_{t'}$;
- every edge $e \in E(G)$ is contained in an induced subgraph $G[V_t]$ for some $t \in V(T)$.

Note that in our definition of tree-decompositions, we allow T to have vertices of infinite degree if G is infinite. The sets V_t for every $t \in V(T)$ are called the *bags* of (T, \mathcal{V}) , and the induced subgraphs $G[V_t]$ the *parts* of (T, \mathcal{V}) . The *width* of (T, \mathcal{V}) is the supremum of $|V_t| - 1$, for $t \in V(T)$. Note that the width of a tree-decomposition can be infinite. The *treewidth* $\text{tw}(G)$ of a graph G is the minimum width of a tree-decomposition of G . The sets $V_t \cap V_{t'}$ for every $tt' \in E(T)$ are called the *adhesion sets* of (T, \mathcal{V}) and the *adhesion* of (T, \mathcal{V}) is the supremum of the sizes of its adhesion sets (possibly infinite). We also let $V_\infty(T) \subseteq V(T)$ denote the set of nodes $t \in V(T)$ such that V_t is infinite.

For connectivity reasons, we are sometimes more interested in the *torsos* of a graph (as defined below) than in its parts. Given a set X of vertices, the graph with vertex set X whose edge set consists of all the pairs uv such that $uv \in E(G)$ or there exists a connected component C of $G - X$ such that $\{u, v\} \subseteq N(C)$ is a *torso* of G and is denoted $G[[X]]$. Consider for instance a graph G with a subset of vertices $X \subset V(G)$ such that $G - X$ is connected and only two vertices of X (call them x and y) have a neighbor in $G - X$. Then the torso $G[[X]]$, is a faithful minor of G and it consists of the graph $G[X]$ with the addition of the edge xy (if it is not already present in G).

As having treewidth at most k is a minor-closed property, the Graph Minor Theorem applies and graphs of treewidth at most k are characterized by finitely many forbidden minors. For example, graphs of treewidth 1 are the forests, which are exactly the set of K_3 -minor-free graphs and graphs of treewidth at most 2 are the K_4 -minor-free graphs. On the other hand, not all minor-closed graphs classes have bounded treewidth: planar graphs have unbounded treewidth, as the (k, ℓ) -*grid*, defined as $P_k \square P_\ell$, has treewidth $\min(k, \ell)$ and is planar.

1.2 Groups and infinite graphs

The notions in this section are used in Chapter 5 only.

1.2.1 Infinite graphs

By default, the graph considered in this thesis are finite and connected and we will denote n the size of $V(G)$. In Chapter 5, we will consider *locally finite* graphs: graphs on finite or countable vertex sets, in which all the vertices have finite degree. The only graphs that will not satisfy this additional requirement will be trees of tree-decompositions.

We let K_∞ be the *countable clique*, that is the infinite complete graph with vertex set \mathbb{N} . This graph is sometimes also denoted by K_{\aleph_0} , which is less ambiguous, but we prefer to keep the notation K_∞ as used by Thomassen in [Tho92], since we reuse a number of results proved in his paper.

Rays and ends

A *ray* in a graph G is an infinite simple one-way path $P = (v_1, v_2, \dots)$. A *subray* P' of P is a ray of the form $P' = (v_i, v_{i+1}, \dots)$ for some $i \geq 1$. We say that a ray *lives* in a set $X \subseteq V(G)$ if one of its subrays is included in X . We define an equivalence relation \sim over the set of rays $\mathcal{R}(G)$ by letting $P \sim P'$ if and only if for every finite set of vertices $S \subseteq V(G)$, there is a component of $G - S$ that contains infinitely many vertices from both P and P' . When G is infinite, this is equivalent to saying that for any finite set $S \subseteq V(G)$, P and P' are living in the same component of $G - S$. The *ends* of G are the elements of $\mathcal{R}(G)/\sim$, the equivalence classes of rays under \sim . For every $X \subseteq V(G)$, we say that an end ω *lives in* X if one of its rays lives in X .

When there is a set X of vertices of G , two distinct components C_1, C_2 of $G - X$, and two distinct ends ω_1, ω_2 of G such that for each $i = 1, 2$, ω_i lives in C_i , we say that X *separates* ω_1 and ω_2 . A graph G is *vertex-accessible* if there is an integer k such that for any two distinct ends ω_1, ω_2 in G , there is a set of at most k vertices that separates ω_1 and ω_2 . The *degree* of an end ω is the supremum number $k \in \mathbb{N} \cup \{\infty\}$ of pairwise disjoint rays that belong to ω . By a result of Halin [Hal65], this supremum is a maximum, i.e., if an end ω has infinite degree then there exists an infinite countable family of pairwise disjoint rays belonging to ω . An end is *thin* if it has finite degree, and *thick* otherwise. It is an easy exercise to check that for every end ω of finite degree k and every end $\omega' \neq \omega$, there is a set of size at most k that separates ω from ω' .

For example, the infinite grid has one thick end, whereas the infinite two-way path has two thin ends (of degree 1). The interested reader is referred to Chapter 8 in [Die17] for more background and important results in infinite graph theory.

1.2.2 Groups and Cayley graphs

An *automorphism* of a graph G is a graph isomorphism from G to itself (i.e., a bijection from $V(G)$ to $V(G)$ that maps edges to edges and non-edges to non-edges). The set of automorphisms of G has a natural group structure (as a subgroup of the symmetric group over $V(G)$); the group of automorphisms of G is denoted by $\text{Aut}(G)$.

For a graph G and a group Γ , we will say that Γ *acts by automorphisms on* G (or simply that Γ *acts on* G when the context is clear) if every element of Γ induces an automorphism g of G , such that the induced application $\Gamma \rightarrow \text{Aut}(G)$ is a group morphism. We will usually use the right multiplicative notation $x \cdot g$ instead of $g(x)$ for $g \in \Gamma$, $x \in V(G)$. For every $X \subseteq V(G)$, $\Gamma' \subseteq \Gamma$ and $g \in \Gamma$, we let $X \cdot g := g(X) = \{x \cdot g : x \in X\}$ and $X \cdot \Gamma' := \bigcup_{g \in \Gamma'} X \cdot g$. We denote the set of orbits of $V(G)$ under the action of Γ by G/Γ (Γ naturally induces an equivalence relation on $V(G)$, relating elements in the same orbit of Γ). For every subset $X \subseteq V(G)$ we let $\text{Stab}_\Gamma(X) := \{g \in \Gamma : X \cdot g = X\}$ denote the *stabilizer* of X , which is always a subgroup of Γ . For each $x \in X$, we let $\Gamma_x := \text{Stab}_\Gamma(\{x\})$.

Quasi-transitive graphs

The action of a group Γ on a graph G is said to be *vertex-transitive* (or simply *transitive*) when there is only one orbit in G/Γ , i.e. when for every two vertices $u, v \in V(G)$ there exists an element $g \in \Gamma$ such that $u \cdot g = v$. The action of Γ on G is said to be *quasi-transitive* if there is only a finite number of orbits in G/Γ . We say that G is *transitive* (resp. *quasi-transitive*) if it admits a transitive (resp. quasi-transitive) group action, that is if $\text{Aut}(G)$ is transitive (resp. quasi-transitive).

It was proved, first for finitely generated groups and then in the more general graph-theoretic context, that the number of ends of a quasi-transitive graph is either 0, 1, 2 or ∞ [Fre44, Hop43, DJM93]. A graph with a single end is said to be *one-ended*.

Finitely presented groups

A *group presentation* is a pair $\langle S|R \rangle$ where S is a set of letters called *generators* and R a set of finite words over the alphabet S called *relators*. We will always assume that S is finite and closed under taking inverse, i.e. that every generator $a \in S$ comes with an associated *inverse* $a^{-1} \in S$ (which may be equal to a) such that $(a^{-1})^{-1} = a$, and we assume that $aa^{-1} \in R$. We say that $\langle S|R \rangle$ is *finite* when both S and R are finite. The group associated to $\langle S|R \rangle$ is the group $F(S)/N(R)$, where $F(S)$ is the free group over S and $N(R)$ is the normal closure of R in $F(S)$,

i.e. the set of elements of $F(S)$ of the form $(w_1 \cdot r_1 \cdot w_1^{-1}) \cdots (w_\ell \cdot r_\ell \cdot w_\ell^{-1})$ for any $\ell \in \mathbb{N}$, $r_1, \dots, r_\ell \in R$ and $w_1, \dots, w_\ell \in F(S)$. For simplicity, we will also denote this group by $\langle S|R \rangle$. Note that by definition of $\langle S|R \rangle$, since $aa^{-1} \in R$ for every element $a \in S$, the (formal) inverse a^{-1} of a in S is indeed the inverse of a in the group $\langle S|R \rangle$. A group is *finitely presented* if it is finitely generated and admits a finite presentation.

Cayley graphs

The *Cayley graph* of a finitely generated group Γ , with respect to a finite set of generators S , is the edge-labelled graph $\text{Cay}(\Gamma, S)$ whose vertex set is the set of elements of Γ and where for every two $g, h \in \Gamma$ we add an arc (g, h) labelled by $a \in S$ when $h = a \cdot g$. Note that $\text{Cay}(\Gamma, S)$ is always locally finite (as we assume S to be finite), connected and that the group Γ acts transitively by right multiplication on $\text{Cay}(\Gamma, S)$. As mentioned above we will always assume that the set S of generators is symmetric, so whenever there is an arc (u, v) labelled with $a \in S$, the graph also contained the arc (v, u) labelled a^{-1} . In this case we can consider the non-labelled version of $\text{Cay}(\Gamma, S)$ as an undirected simple graph, with a single edge uv instead of each pair of arcs (u, v) and (v, u) .

We say that a finitely generated group Γ is *planar* (resp. *minor-excluded*) if it admits a finite generating set S such that $\text{Cay}(\Gamma, S)$ is planar (resp. does not contain every finite graph as a minor). Similarly, we say that Γ is K_∞ -minor-free if it admits a finite generating set S such that $\text{Cay}(\Gamma, S)$ is K_∞ -minor-free.

Accessibility

In order to give a precise definition of accessibility in groups, we first need to define the notion of a graph of groups. The reader is referred to [Ser80, Chapter 1] for more details on graphs of groups and Bass-Serre theory.

A *graph of groups* consists of a pair (G, \mathcal{G}) such that $G = (V, E)$ is a connected graph (possibly having loops and multi-edges), and \mathcal{G} is a family of *vertex-groups* Γ_v for each $v \in V$, *edge-groups* Γ_{uv}, Γ_{vu} for each edge $uv \in E$, and of group isomorphisms $\phi_{u,v} : \Gamma_{uv} \rightarrow \Gamma_{vu}$ for each $uv \in E(G)$, such that Γ_{uv} and Γ_{vu} are respectively subgroups of Γ_u and Γ_v for each $uv \in E$.

Let T be a spanning tree of G and $\langle S_v | R_v \rangle$ be a presentation of Γ_v for each $v \in V$. The *fundamental group* $\Gamma := \pi_1(G, \mathcal{G})$ of the graph of groups (G, \mathcal{G}) is defined as the group having as generators the set:

$$S := \left(\bigsqcup_{v \in V} S_v \right) \sqcup \left(\bigsqcup_{e \in E \setminus E(T)} \{t_e\} \right)$$

and as relations:

- the relations of each R_v ;
- for every edge $uv \in E(T)$ and every $g \in \Gamma_{uv} \subseteq \Gamma_u$, the relation $\phi_{uv}(g) = g$;
- for every edge $e = uv \in E \setminus E(T)$ and every $g \in \Gamma_{uv}$, the relation $t_e \phi_{uv}(g) t_e^{-1} = g$.

It can be shown that for a given graph of groups (G, \mathcal{G}) , the definition of its fundamental group does not depend of the choice of the spanning tree T (see for example [Ser80, Section I.5.1]).

Remark 1.2.1. At the beginning of this section we defined Γ_v as the stabilizer of $\{v\}$ by the action of Γ on a graph G and the reader might be worried about a possible confusion with the notation Γ_v of vertex-groups above. On the one hand the vertex-group Γ_v has a close connection with the stabilizer of $\{v\}$ in this context, so the objects are not completely unrelated, and on the other hand vertex-groups will only be used at the very end (in Subsection 5.3.2, where stabilizers will not be used at all), so hopefully there should not be any risk of confusion.

A group Γ is said to be *accessible* if it is the fundamental group of a finite graph of groups G with finite vertex set $V(G)$ such that:

- the vertex-groups have at most one end, and
- the edge-groups are finite.

By Bass-Serre theory [Ser80, DD89], accessible groups are exactly those groups acting on trees without edge-inversion, such that the vertex-stabilizers have at most one end and the edge-stabilizers of the action are finite.

As mentioned in the introduction, Thomassen and Woess [TW93] obtained the following connection between accessibility in groups and vertex-accessibility in graphs.

Theorem 2 [TW93]

A finitely generated group is accessible if and only if it admits a Cayley graph which is vertex-accessible.

It was open for a long time whether there exist finitely generated groups which are not accessible, and Dunwoody provided a construction of such group in [Dun93]. On the other hand he also proved the following result.

Theorem 3 [Dun85]

Every finitely presented group is accessible.

In particular, planar groups form a proper subclass of accessible groups.

Theorem 4 [Dro06]

Every finitely generated planar group is finitely presented, and thus accessible.

Virtually free groups

A group is *virtually free* if it contains a finitely generated free group as a subgroup of finite index. For graphs of groups, this property can be related to the structure of the vertex groups as follows.

Theorem 5 [KPS73]

A finitely generated group is virtually free if and only if it is the fundamental group of a finite graph of groups in which all vertex-groups are finite.

1.3 Markov chains and random walks

The notions presented in this section are used in Chapter 2, and mentioned in passing in Chapters 3 and 4.

We follow here the notations and definitions of [LP17]. A finite Markov chain is a memory-less random process describing the evolution of a state in a finite set Λ , such that at each step, the next state is determined by a probability distribution $\mathbb{P}(x, \cdot)$ that depends only on the current state x . More precisely, a sequence of random variables $(X_t)_{t \in \mathbb{N}}$ is a *Markov chain with state space Λ and transition matrix (or kernel) P* if it verifies the *Markov property*: $\forall n, \forall x_0, \dots, x_n \in \Lambda$,

$$\mathbb{P}(X_n = x_n | X_{n-1} = x_{n-1}, \dots, X_0 = x_0) = \mathbb{P}(X_n = x_n | X_{n-1} = x_{n-1}) = P(x_{n-1}, x_n) \quad (1.1)$$

Its *ground graph* is the directed graph on the vertex set Λ , whose arcs (x, y) corresponds to the non-zero transitions and have weight $P(x, y)$. Conversely, given a directed graph G , possibly with multi-edges and loops, the *random walk* on G is the Markov chain with transition matrix P such that $P(x, y)$ is the proportion of arcs going out from x that lead to y .

A Markov chain on a kernel P is *irreducible* if for any two states $x, y \in \Lambda$, there is a integer t such that $P^t(x, y) > 0$, that is if the ground graph is connected. Let $\mathcal{T}(x) = \{t \geq 1 : P^t(x, x) > 0\}$, the *period* of x is the greatest common divisor of $\mathcal{T}(x)$. A chain is *aperiodic* if all its states have period 1. A distribution π on Λ is *stationary* if $\pi = \pi P$.

Theorem 6 [LP17]

Any irreducible Markov chain admits a unique stationary distribution. Moreover, the uniform distribution is stationary if P is symmetric.

We say that a Markov chain is *ergodic* if it converges towards its stationary distribution.

1.3.1 Mixing time

The total variation distance on two probability distributions μ and ν on Λ is defined as:

$$\|\mu - \nu\|_{TV} = \frac{1}{2} \sum_{x \in \Lambda} |\mu(x) - \nu(x)| = \max_{A \subset \Lambda} \mu(A) - \nu(A).$$

Let $d(t) := \max_{x \in \Lambda} \|P^t(x, \cdot) - \pi\|_{TV}$. The *mixing time* of P is

$$t_{mix} = \min\{t : d(t) \leq \frac{1}{4}\}. \quad (1.2)$$

Conceptually, the mixing time measures how fast the Markov chain equidistributes. In (1.2), there are two arbitrary choices: the choice of the constant $1/4$ and the choice of the total variation distance. We say that a Markov chain on a space of configurations is *rapidly mixing* if $t_{mix} = O(n^c)$ for some $c > 0$, where n is the size of the configurations. On the other hand, a Markov chain is *torpidly mixing* if $t_{mix} = \Omega(\exp(n^\varepsilon))$ for some $\varepsilon > 0$.

The following theorem give sufficient conditions for the ergodicity of a Markov Chain, as well as a quantitative rate of convergence.

Theorem 7 (Convergence theorem, Theorem 4.9 in [LP17])

Let P be an aperiodic and irreducible Markov chain with stationary distribution π . There exist C and $1 > \alpha > 0$ such that

$$d(t) \leq C\alpha^t$$

The proof of Theorem 7 follows from elementary linear algebra. Transforming a Markov chain with kernel P into an aperiodic one can be easily done without modifying the stationary distribution, by considering the *lazy* version of it, that has kernel $\varepsilon \text{Id} + (1 - \varepsilon)P$ for any $0 < \varepsilon < 1$. We will call *lazy random walk* on G any lazy version of the random walk on G .

A *coupling* of two probability distributions μ and ν is a couple of random variables (X, Y) defined on a single probability space, such that $\mathbb{P}(X = x) = \mu(x)$ and $\mathbb{P}(Y = y) = \nu(y)$, even though X and Y are in general far from being independent.

Lemma 8 [LP17]

Given two probability distributions μ and ν , we have

$$\|\mu - \nu\|_{TV} = \min \mathbb{P}(X \neq Y)$$

where the min is taken over all couplings (X, Y) of μ and ν .

1.3.2 Couplings

A *coupling* for P is a pair of (not necessarily independent) Markov chains $(X_t, Y_t)_{t \geq 0}$ both with kernel P . The *coalescence time* of the coupling is

$$\tau_{\text{couple}} := \min\{t : X_s = Y_s, \text{ for all } s \geq t\}.$$

A direct application of Lemma 8 gives

$$\forall t \geq 0, \quad \|P^t(x, \cdot) - P^t(y, \cdot)\|_{TV} = \min \mathbb{P}_{x,y}(X_t \neq Y_t) = \min \mathbb{P}_{x,y}(\tau_{\text{couple}} > t)$$

where the minimum is taken over all couplings $(X_t, Y_t)_{t \geq 0}$ of P and τ_{couple} denote their coalescence time.

Remark 1.3.9. In [LP17] this is Theorem 5.2 but they only mention the upper bound. The equality does hold.

As a consequence, for any coalescence time τ_{couple} of a coupling $(X_t, Y_t)_t$ of P we have

$$t_{\text{mix}}(P) \leq 4 \max_{x,y} \mathbb{E}_{x,y}(\tau_{\text{couple}}) \quad (1.3)$$

1.3.3 Bottleneck ratio and expander families

The rapid mixing of a Markov chain is strongly related to the *expansion* (or bottleneck ratio) of the underlying graph, defined as

$$\Phi_{\star}(G) = \min_{S: \pi(S) \leq 1/2} \frac{Q(S, \Lambda \setminus S)}{\pi(S)}$$

where $Q(A, B)$ is the probability of going from A to B , starting from the stationary distribution. If all transitions have the same weight, $\Phi_{\star}(G)$ is also known as the *Cheeger constant*, denoted $h(G)$. Intuitively, the Cheeger constant of a graph is large if and only if any splitting of the vertices in two sets induces a large number of edges between the two parts, compared to the size of the sets: there is no bottleneck. If a Markov chain has a small expansion, then it is unlikely to get in or out of the set S of states for which the expansion is attained. Formally,

Lemma 10 [LP17]

The expansion provides a lower bound on the mixing time t_{mix} :

$$t_{\text{mix}} \geq \frac{1}{4\Phi_{\star}}$$

Because of this, any family of graphs with good mixing properties must avoid bottlenecks. To that end, we call (d, α) -*expander family* any infinite family \mathcal{F} of

graphs, such that all $G \in \mathcal{F}$ are d -regular and have edge expansion $h(G) \geq \alpha$. The existence of such families is not trivial, but true for all $d \geq 3$ for some $\alpha > 0$ [Pin73]. Lazy random walks in these families have the fastest mixing time among graphs of bounded degree:

Lemma 11 (Expander mixing lemma, Proposition 13.35 in [LP17])

Let \mathcal{F} be a (d, α) -expander family. The lazy random walk on $G \in \mathcal{F}$ has mixing time $t_{mix} = O_{d,\alpha}(|G|)$.

1.4 (Parameterized) complexity classes

The notions presented in this section are used in Chapter 4 only.

1.4.1 Complexity of decision problems

We assume basic knowledge of Turing machine and refer the reader to [CLRS22, Chapter 34]. Given an alphabet Σ , the set of finite words on Σ is denoted Σ^* . A *decision problem* is a subset $P \subset \Sigma^*$. The *instances* of the problem are the words in Σ^* and the *positive instances* those in P . We will often define problem using a property characterizing the positive instances. A Turing machine *decides* a decision problem P if for any instance x , the Turing machine outputs YES if P is positive and NO otherwise. In this case, we also say that P is *decidable*.

P is the set of problems that can be decided in polynomial time by a deterministic Turing machine. For example, the set of bipartite graphs, that is 2-colorability, is in P as odd cycles can be detected in polynomial time. NP is the set of decision problems that can be decided by a non-deterministic Turing machine in time polynomial in the size of the input instance. Alternatively, NP can be defined as the problems that can be checked in polynomial time by a deterministic Turing machine, provided with a certificate. For example, k -colorability is in NP for all k , as one can check in polynomial time if a k -coloring of a graph is proper. $co-NP$ is the set of problems P whose complement $\Sigma^* \setminus P$ is in NP . Alternatively $co-NP$ is the set of problems for which there is a certificate of negative instances that can be checked in polynomial time by a deterministic Turing machine. $PSPACE$ is the set of decision problems that can be decided by a Turing machine using a polynomial memory space. Note that allowing this machine to be non-deterministic does not grant any extra power by Savitch's theorem [Sav70]. The class P is contained in both NP and $co-NP$ and both of these classes are contained in $PSPACE$. Each of these inclusions is believed to be strict, but it remains a major open problem in complexity theory.

Given two problems Π and Π' , if an algorithm for Π can be used to solve Π' , then Π is harder than Π' and we say that Π *reduces to* Π' . The exact meaning

of “used to solve” depends on the complexity class of the problems Π and Π' . In classical complexity, a (*polynomial*) *reduction* is a polynomial time algorithm that transforms an instance x of Π into an instance x' of Π' , such that x is a positive instance if and only if x' is. Given a complexity class C , we say that a problem is *C-hard* if it reduces in polynomial time to all problems of C , and that it is *C-complete* if it also belongs to C .

We will see in Chapter 2 that the reconfiguration analog of many NP-complete problems is PSPACE-complete.

1.4.2 Parameterized complexity

We refer the reader to the books of Flum and Grohe[FG06] and of Downey and Fellows [DF13]. Hard problems are hard! Though, they might become easier when restricted to a subset of instances. To that end, parameterized complexity studies the complexity with respect to the size n of the instance but also with respect to some other parameter, often denoted k . A *parameterized problem* is a subset of $\Sigma^* \times \mathbb{N}$, where Σ^* is the set of strings over a finite alphabet Σ and the input strings are of the form (x, k) . Here the integer k is called *the parameter*.

A problem is *fixed-parameter tractable (FPT)* if it can be solved by a deterministic Turing machine in time $O(f(k) \cdot n^c)$ for some $c > 0$ and a computable function f . In other words, if we fix k , then the problem can be solved in polynomial time while the degree of the polynomial does not depend on k . There is also a wider class XP of problems, that can be solved in time $O(n^{f(k)})$ by a deterministic Turing machine. The problems in XP are still polynomial time solvable for fixed k , however, at the cost that the degree of the polynomial depends on k .

Somewhere in between FPT and XP there is an interesting class W[P]. A parameterized problem belongs to the class W[P] if it can be solved by a non-deterministic Turing machine in time $h(k) \cdot n^c$ provided that this machine makes only $O(f(k) \log n)$ non-deterministic choices where f, h are computable functions and $c > 0$ is some constant. Given an algorithm for a computational problem Π , we say that this algorithm is a *W[P]-algorithm* if it is represented by a Turing machine satisfying the conditions above. The class W[P] is often considered as the parameterized analog of NP [FG06].

It is obvious that a problem in FPT belongs to W[P]—we use the same algorithm with 0 non-deterministic choices. On the other hand, given a W[P]-algorithm, it can be converted to a deterministic algorithm by trying all options for each non-deterministic choice. The running time of the new algorithm is $h(k) \cdot n^c \cdot (c')^{O(f(k) \log n)}$ for some $c' > 0$ which can be easily manipulated into a formula showing XP-membership.

Given two parameterized problems Π and Π' , we say that Π reduces to Π' via an *FPT-reduction* if there exist computable functions $f: \Sigma^* \times \mathbb{N} \rightarrow \Sigma^*$ and

$g: \Sigma^* \times \mathbb{N} \rightarrow \mathbb{N}$ such that

- $(x, k) \in \Pi$ if and only if $(f(x, k), g(x, k)) \in \Pi'$ for every $(x, k) \in \Sigma^* \times \mathbb{N}$;
- $g(x, k) \leq g'(k)$ for some computable function g' ; and
- there exist a computable function h and a fixed constant $c > 0$ such that for all input strings, $f(x, k)$ can be computed by a deterministic Turing-machine in $O(h(k)n^c)$ steps.

The classes FPT, W[P] and XP are closed under FPT-reductions. We define the notions of hardness and completeness for parameterized problems as for classical complexity but with respect to FPT-reductions.

Another way of defining FPT problems is through kernelization, a preprocessing reduction. A parameterized problem admits an FPT-reduction to a set called a *kernel* of instances of size at most $f(k)$. This kernel is said to be linear or polynomial if f is linear or polynomial. As it can easily be checked a problem is FPT if and only if it admits a kernel. In practice, this FPT-reduction is generally expressed as a set of reduction rules and this equivalent definition of FPT problems is often easier to work with.

1.4.3 Counting problems

Given a decision problem, one can wonder how many instances of size n are positive. This question is called a *counting problem*. Counting problems are not decision problems. A Turing machine *solves* a counting problem if it outputs the correct amount of positive instances of size n , for every input n . The class #P is the class of counting problems associated to an NP-problem. The reductions that preserve the complexity of a counting problem are the *parsimonious reductions* that preserve the number of positive instances of a given size. We will see in Chapter 2 that counting the number of configurations of a reconfiguration problem is a typical #P-complete problem.

1.5 Knots

The notions presented in this section are used in Chapter 4 only.

A *diagram* of a knot is a piecewise linear map $D: S^1 \rightarrow \mathbb{R}^2$ in general position; for such a map, every point in \mathbb{R}^2 has at most two preimages, and there are finitely many points in \mathbb{R}^2 with exactly two preimages (called *crossings*). Locally at a crossing two arcs cross each other transversely, and the diagram contains the information of which arc passes ‘over’ and which ‘under’. We usually depict this by interrupting the arc that passes under.

(A diagram usually arises as a composition of a (piecewise linear) knot $\kappa: S^1 \rightarrow \mathbb{R}^3$ and a generic projection $\pi: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ which also induces ‘over’ and ‘under’ information.) We usually identify a diagram D with its image in \mathbb{R}^2 together with the information about underpasses/overpasses at crossings.¹ Diagrams are always considered up-to ambient isotopy of the plane, that is up to continuous distrotion of the plane. The unique diagram without crossings is denoted U (for untangled). See Figure 1.3 for an example of a diagram.




Figure 1.3: An example of a diagram of the unknot.

Diagrams may be encoded purely combinatorially as 4-regular plane graphs (with some additional combinatorial information) and the size of this encoding is comparable to the number of crossings.

¹It would also be meaningful to consider diagrams in S^2 rather than \mathbb{R}^2 . This would make the Reidemeister moves passable through the outer face. In general, this setting would not be a big difference for our results. Working in S^2 would even make some parts of our analysis slightly easier because we would not have to distinguish the outer face as a special case.

Chapter 2

Recoloring with Kempe changes

This chapter is a survey on recoloring with Kempe changes. It contains ongoing work with Marthe Bonamy and Tomáš Kaiser; published results obtained with Quentin Deschamps, Carl Feghali, František Kardoš and Théo Pierron [8], with Marthe Bonamy and Vincent Delecroix [6], with Marthe Bonamy, Marc Heinrich and Jonathan Narboni [7]; as well as personal work in progress.

Introduction

In 1852, Guthrie observed, while coloring maps of counties of England, that it was possible to color them using no more than four colors, ensuring that no adjacent counties shared the same color [FF98]. He wondered if this phenomenon could be extended to all planar maps, giving rise to the renowned Four Color Problem: are all planar graphs 4-colorable?

In 1879, Alfred Kempe introduced an elementary modification of a graph coloring as a tool towards solving this problem [Kem79]. To change the color of a vertex u from a to b and obtain a proper coloring, the neighbors of u colored b need to be recolored. By recoloring them with a and propagating the swapping of a and b recursively, we obtain another proper coloring, with u colored b . Formally, given a k -colored graph G and two colours a and b , an $\{a, b\}$ -Kempe chain is a connected component in the subgraph induced in G by the vertices colored a or b , and Kempe change consists in swapping the two colors within a Kempe chain, thereby obtaining a new k -coloring of the graph (see Figure 2.1). We will denote $K_{u,b}(\alpha, G)$ the Kempe change on the coloring α of G that recolors u to b , or simply $K_{u,b}$ when there is no ambiguity.

Using this notion, we can easily prove that all planar graphs are 5-colorable:

Lemma 2.0.1 (Folklore)

Every planar graph is 5-colorable.




Figure 2.1: Two 3-colorings of the Petersen graph that differ by one Kempe change. The corresponding $\{\bullet, \bullet\}$ -Kempe chain is thickened.

Proof. We prove the result by induction on G . A direct consequence of Euler's formula is that planar graphs are 5-degenerate. Let G be a plane graph and v be a vertex of degree at most 5 in G .

By induction, $G - v$ is 5-colorable, let α be such a 5-coloring of $G - v$. If the neighborhood of v uses only 4 colors, then G is trivially 5-colorable. Otherwise, every color is used exactly once by α in $N(v)$. Let $u_1 \dots u_5$ be the neighbors of v in cyclic order, with $\alpha(u_i) = i$ up to color permutation. Consider the $\{1, 3\}$ -Kempe chain K containing u_1 . If $u_3 \notin K$, then performing this Kempe change results in a coloring that does not use the color 1 in $N(v)$, hence v can be colored 1. If $u_3 \in K$, then by the Jordan curve theorem, the cycle $K \cup \{v\}$ separates the sphere in two (see Figure 2.2), thus u_4 cannot belong to the $\{2, 4\}$ -Kempe chain containing u_2 , so the same reasoning applies. \square




Figure 2.2: The vertices u_1 and u_3 belong to the same $\{\bullet, \bullet\}$ -Kempe chain, represented by a dashed line, so u_2 and u_4 cannot be in the same $\{\bullet, \bullet\}$ -Kempe chain.

While Kempe's proof of the Four Color Theorem turned out to be false,

Kempe changes are a crucial tool in the computer assisted proof of Appel and Haken [AH89]. Although this proof was since shortened by Robertson, Sanders, Seymour and Thomas [RSST97], it still uses Kempe changes and computer-checks. Finding a shorter proof of the Four Color Theorem is still an active field of research.

When considering edge-colorings instead of vertex-colorings, Kempe chains are a lot more structured as they are either paths or cycles (see Figure 2.3). A trivial




Figure 2.3: A $\{\bullet, \bullet\}$ -Kempe chain in an 3-edge-coloring of the Petersen graph

lower bound on the chromatic index of a graph is the maximum degree Δ and Vizing proved in [Viz64] that all graphs are $(\Delta+1)$ -edge-colorable. Vizing used the particular structure of Kempe changes on edge-colorings to design a new recoloring operation known as Vizing fans. Kempe changes are decisive in his proof as no proof avoiding this tool is known as of today. We give a sketch of his proof to illustrate these fundamental ideas:

Theorem 2.0.2 [Viz64]

Let G be a graph of maximum degree Δ and α be a k -edge-coloring of G . There exists a $(\Delta+1)$ -edge-coloring β of G that is Kempe equivalent to α . In particular, G is $(\Delta+1)$ -edge-colorable.

Proof. We prove the result by induction on G . Let uv be an edge of G and $G' = G - uv$. Let α be a $(\Delta(G') + 1)$ -edge-coloring of G' . We say that a vertex w misses a color a if for all $x \in N(w)$, $a \neq \alpha(wx)$. Every vertex of G misses at least one color in $[\Delta(G) + 1]$.

If u and v miss a common color a , then coloring uv with a results in a $(\Delta(G)+1)$ -edge-coloring of G .

Otherwise, we build a sequence of edges (vw_0, \dots, vw_p) such that $w_0 = u$ and for each i , w_i is missing some color c_i and vw_{i+1} is colored c_i . We stop this construction either if v is missing c_p or if c_p is already used by some vw_i for $i \geq 1$. This sequence of edges (vw_0, \dots, vw_p) forms the so-called *Vizing fan*. We will represent this Vizing fan by a directed graph whose vertices are the edges vw_i , with arcs starting from vw_i and pointing to the edge incident to v that is colored

c_i for each i (see Figure 2.4). The first kind of Vizing fan, when both v and w_p are missing the same color, corresponds to a directed path (see Figure 2.4 (left)). The second kind, when w_p misses a color already used by some vw_i , corresponds to a *comet*: a directed path attached to a directed cycle (see Figure 2.4 (right)).




Figure 2.4: The two kinds of Vizing fans. The missing colors are represented by the decoration $\times \bullet$. The arcs between the edges vw_i of the Vizing fan point towards the edge colored with the missing color at w_i .

In the first case, vw_p can be recolored with c_p , then vw_{p-1} with c_{p-1} and so on, until both v and $w_0 = u$ miss a common color (see Figure 2.4 (left)).

In the second case, let vw_i be the edge of the Vizing fan that uses c_p and denote b the color missing at v . Consider the $\{b, c_p\}$ -Kempe chain containing vw_i . Because v misses b , this Kempe chain is a path P , hence it cannot contain both w_{i-1} and w_p , which are both missing c_p (see Figure 2.5). If P does not contain w_{i-1} , performing the corresponding Kempe chain results in an edge-coloring in which v and w_{i-1} miss c_p (see Figure 2.6 (left)), thus we are done by reducing to the first case. If P contains w_{i-1} , then it does not contain w_p , then performing the corresponding Kempe chain results in an edge-coloring with v and w_p missing c_p (see Figure 2.6 (right)), which reduces once again to the first case. \square

This proves that one can reach, using Kempe changes, a $(\Delta + 1)$ -edge-coloring of G , starting from any k -edge-coloring of G . Two k -colorings of a graph are *Kempe equivalent* if one can be obtained from the other through a series of Kempe changes. In 1965, Vizing conjectured the following:

Conjecture 2.0.3 [Viz65]

All edge-colorings are Kempe equivalent to an optimal edge-coloring. That is, every k -edge-coloring of any graph G is Kempe equivalent to a $\chi'(G)$ -edge-coloring of G .

In particular, this conjecture implies that all $(\chi'(G) + 1)$ -edge-colorings of G are Kempe equivalent. Recently, this conjecture was confirmed, first by Bonamy, Defrain, Klimošová, Lagoutte and Narboni on triangle-free graphs [BDK⁺23], and




Figure 2.5: The $\{\bullet, \bullet\}$ -Kempe chain P , represented by a dashed line, cannot contain both w_1 and w_4 .




Figure 2.6: Possible results of the $\{\bullet, \bullet\}$ -Kempe change P . A dashed line between w_2 and w_i indicates that w_i is contained in P , while a crossed dashed line represents the contrary. In both cases, the resulting Vizing fan is of the first kind.

then by Narboni on all graphs [Nar23]. The two proofs reduce to the case of Δ -regular graphs, before doing an induction on the degree by coloring a perfect matching identically and removing it. This is done thanks to an intricate use of Vizing's fans.

The same questions can be asked on vertex-colorings: under which conditions are all the k -colorings of a graph Kempe equivalent? Is there an optimal coloring in all the Kempe equivalence classes; is there an efficient algorithm to find it? Our main interest is the study of the reconfiguration graph $\mathcal{K}_k(G)$, whose vertices are the k -colorings of G and in which two colorings are adjacent if and only if they differ by one Kempe change. The former question can be restated as follows: in which setting is the reconfiguration graph connected, that is, any two k -colorings are Kempe equivalent (see Section 2.1)? When this is the case, can we bound the length of the shortest sequence of Kempe changes between any two colorings, i.e. the diameter of the reconfiguration graph (see Section 2.2)? As these questions can be studied independently in each of the connected components of G , we will consider only connected graphs from now on. Beyond the intrinsic interest of theoretical results, Kempe changes are motivated by practical applications in statistical physics, and closer to graph theory, Kempe equivalence can be studied with a goal of obtaining a random coloring. Random sampling is also deeply related to approximate counting the colorings of a graph (see Section 2.3 for a discussion on the complexity of these problems). Efficient random sampling can sometimes be achieved by applying random walks and rapidly mixing Markov chains. In particular, the Wang-Swendsen-Kotecký algorithm performs random Kempe changes at each step and is studied in statistical mechanics to understand the transition phase of the Potts model (see Section 2.4). Finally, we summarize in Section 2.5 the open questions and conjectures that we believe to be of special interest in the Kempe landscape.

2.1 Connectedness of $\mathcal{K}_k(G)$

The goal of this section is to present sufficient or necessary conditions under which a graph is k -recolorable, meaning that all its k -colorings are Kempe equivalent. As a warm-up, we can easily settle the case of bipartite graphs:

Lemma 2.1.1 (Folklore)

All bipartite graphs are k -recolorable, for any value of k .

Proof. Let $A \sqcup B$ be a partition of the vertices of G such that A and B are independent sets. The result follows from the fact that any coloring α of G is Kempe equivalent to a canonical coloring of G using only two colours, say a and b . For all vertices of A that are not colored a , perform the Kempe change with color

a. By a parity argument this operation does not recolor any vertex of A that was already colored a . The same procedure in B with color b then yields the desired 2-coloring. \square

Note that this recoloring sequence has length $O(|V(G)|)$, hence the Kempe reconfiguration graph of any bipartite graph G has diameter linear in the number of vertices of G .

2.1.1 Negative results

We first present the different methods to prove that some graph is not k -recolorable for some k . Two arguments are known as of today. The first and most general one consists in exhibiting a frozen coloring, while the second is a refinement of a topological argument of Fisk, introduced by Mohar and Salas.

2.1.1.1 Frozen colorings

A k -coloring is said to be *frozen* if for any pair of colors a and b , the subgraph induced by the vertices colored a or b is connected. Being frozen is invariant under Kempe changes, as swapping colors a and b does not modify the partition induced by the color classes if there is only one $\{a, b\}$ -Kempe chain in G .

Therefore an simple witness that G is not k -recolorable consists in exhibiting two k -colorings α and β of a graph G , whose color classes induce different partitions and such that α is frozen. Two colorings are *similar* if their color classes induce the same partition. For example, the 3-prism $K_2 \square K_3$ admits two non-similar frozen 3-colorings (see Figure 2.7) and hence is not 3-recolorable. This is the easiest and most common way to prove the non-recolorability of a graph.



Figure 2.7: Two non-similar frozen 3-colorings of the 3-prism $K_2 \square K_3$

A direct illustration of this is that the connectedness of $\mathcal{K}_k(G)$ is a non-monotone property (in the number of colors k):

Proposition 2.1.2 (Folklore)

Let $\ell > k > 2$, there exists a graph G such that G is k -recolorable but not ℓ -recolorable.

Proof. Let $G_{k,\ell}$ be the categorical product $K_k \times K_\ell$ (see Chapter 1): the vertices of $G_{k,\ell}$ are indexed by $[k] \times [\ell]$ and (u, v) is adjacent to (w, x) if $u \neq w$ and $v \neq x$ (see Figure 2.8).




Figure 2.8: The complement of $G_{3,5}$: the dashed edges represent the non-edges of $G_{3,5}$

The graph $G_{k,\ell}$ admits a unique k -coloring up to color permutation: $\alpha((u, v)) := u$, hence $\mathcal{K}_k(G_{k,\ell})$ is connected. On the other hand, $G_{k,\ell}$ admits a frozen ℓ -coloring that is non-similar to α , defined as $\beta((u, v)) := u$. Thus $G_{k,\ell}$ is not ℓ -recolorable. \square

The chromatic number of a graph is trivially lower bounded by its clique number ω . However, the chromatic number can be in general arbitrarily larger than the clique number: Erdős [Erd59] gave a random construction proving that there exist graphs with arbitrarily large chromatic number and girth, and thus that contain no triangles. An important topic in graph coloring consists in studying the graph classes for which there exist some relation between χ and ω . A family of graphs \mathcal{F} is χ -bounded if there exists a function f such that $\chi(G) \leq f(\omega(G))$ for all graphs $G \in \mathcal{F}$. It is tempting to ask similar questions in the context of Kempe changes. Namely,

Question 2.1.3

For which hereditary graph classes \mathcal{F} is there a χ -bounding recoloring function f , such that any graph $G \in \mathcal{F}$ is k -recolorable for $k \geq f(\chi(G))$?

One could go even further and ask which hereditary graph classes admit a function f such that any graph $G \in \mathcal{F}$ is k -recolorable for $k \geq f(\omega(G))$. The example above guarantees that both questions admit no straightforward answer. Indeed, we note that the graphs $K_k \times K_\ell$ are perfect [RP77], that is have chromatic number equal to their clique number, yet admit no χ -bounding recoloring function.

The Gyárfás-Sumner conjecture [Gyá75, Sum81] asks whether for any tree T , the graphs that do not contain T as an induced subgraph are χ -bounded. A recoloring version of this conjecture cannot hold for trees of large diameter, as $K_3 \times$

K_ℓ does not contain long induced paths. However, it is possible that forbidding other induced subgraphs yields a class admitting a χ -bounding recoloring function.

For example, cographs are the graphs that contain no induced P_4 and are known to be k -recolorable for all $k \geq \chi = \omega$ (see [BHI⁺20]).

In an other direction, graphs forbidding any fixed star S admit a χ -bounding recoloring function. Indeed, by Ramsey's theorem, any vertex of large enough degree contains either a large clique or a large independent set in its neighborhood. Hence, the graphs of bounded clique number that also forbid a fixed star have bounded degree. Since d -degenerate graphs are k -recolorable for $k \geq d+1$ [VM81], this implies that the existence of a ω -bounding recoloring function for graphs that forbid S as an induced subgraph.

Finally, if we allow ourselves to forbid induced structures other than trees, chordal graphs are a special case of perfect graphs, and are k -recolorable for all $k \geq \chi = \omega$.

Moving away from χ -bounding recoloring functions, when studying the recolorability of a graph class \mathcal{F} , a reasonable first step is to check whether or not \mathcal{F} can admit frozen colorings. An easy counting argument provides a necessary condition for a graph G to admit a frozen k -coloring:

Observation 2.1.4. If G admits a frozen k -coloring, then the number of edges in G is at least $(k-1)(n-k/2)$.

Indeed, each pair of color induces a connected subgraph, hence the number of edges with endpoints colored i and j is at least $n_i + n_j - 1$, where n_i and n_j denote the number of vertices colored i and j respectively. Therefore, the number of edges in G is at least

$$\sum_{\substack{i,j \in [k] \\ i \neq j}} (n_i + n_j - 1) = (k-1)(n-k/2).$$

We say that a graph G is k -freezable if it has a frozen k -coloring and that G is *critically k -freezable* if it is k -freezable and has exactly $(k-1)(n-k/2)$ edges.

Sometimes a graph is not k -recolorable despite having no frozen k -colorings, but the reason for its non-recolorability boils down to not being recolorable or even frozen “up to quotienting”. With Marthe Bonamy, we give an example of a construction illustrating this phenomenon:

Proposition 2.1.5

Let G_1 be a graph that is not k_1 -recolorable. For any graph G_2 , the strong product $G_1 \boxtimes G_2$ is not $(k_1 \cdot \chi(G_2))$ -recolorable (see Figure 2.10). However, $G_1 \boxtimes G_2$ is not necessarily $(k_1 \cdot \chi(G_2))$ -freezable.


More precisely, the following statement is true, as it will be clear from the proof of Proposition 2.1.5. Replacing each vertex v of G_1 by a graph G_v of fixed

chromatic number χ_2 yields a graph that is not $(k_1 \cdot \chi_2)$ -recolorable, even if the graphs $(G_v)_{v \in V(G)}$ are distinct. The intuition is that there exist some $(k_1 \cdot \chi_2)$ -colorings of the constructed graph G that can be viewed as colorings of G_1 . Indeed, given a coloring α of G , the map that associates to $v \in V(G_1)$ the set of colors used in G_v is sometimes a proper k_1 -coloring of G_1 . Moreover, we will see that this property is invariant under Kempe changes, thus the non-equivalence of colorings of G_1 directly implies the non-equivalence of some colorings of G .


To formalize this, we introduce preliminary definitions useful to compare reconfiguration graphs, before proving Proposition 2.1.5. Let G and G' be two graphs. Recall from Subsection 1.1.1 that a *homomorphism* from G to G' is a mapping of the vertices of G to the vertices of G' such that $\forall uv \in E(G), \phi(u)\phi(v) \in E(G')$. We call *projection* a surjective homomorphism in which edges of G can also be contracted: $\forall uv \in E(G), \phi(u)\phi(v) \in E(G')$ or $\phi(u) = \phi(v)$.

Given a homomorphism $\phi : G \rightarrow G'$, a k -coloring of G is ϕ -*compatible* if it induces a coloring of G' . That is, there is a k -coloring α' of G' such that the preimage by ϕ of every color class of α' is a color class of α . Note that in this case, $\phi^{-1}(v)$ is monochromatic for every $v \in V(G')$ and α' is unique up to color permutation (see Subfigure 2.9 (a)). We will say that α is ϕ -*compatible with* α' , when stressing additional information on α' .

Given an ℓ -coloring β of G' and a projection $\psi : G \rightarrow G'$, a k -coloring α of G *refines* β if every color class of α is mapped to a color class of β by ψ (see Subfigure 2.9 (b)).



(a) α is ϕ -compatible



(b) α refines β

Figure 2.9: Commutative diagrams defining the notions of compatibility and refinement. The maps whose existence is claimed are represented by dashed arrows.

Homomorphisms and projections can encode classical graph operations such as subgraphs and minors for example, but also identifications of pairs of vertices. These comparisons have been commonly used in various settings to adapt reconfiguration sequences from a graph to another graph. Throughout this survey, we will try to provide a general framework to compare the reconfiguration graphs, based on how their ground graphs and colorings relate to one another.

Proof of Proposition 2.1.5. Let $n_i = |V(G_i)|$ and label the vertices of G_i by $[n_i]$.

Let H be the strong product $G_1 \boxtimes G_2$: the vertices of H are indexed by $[n_1] \times [n_2]$ and (u, v) is adjacent to (w, x) if u is adjacent or equal to w in G_1 , or if v is adjacent or equal to x in G_2 . For convenience we will use colors in $[k_1] \times [k_2]$ to color H , instead of $[k_1 \cdot k_2]$.

Let C be a maximal subset of equivalent k_1 -colorings of G_1 . Let ϕ_i be the projection from H to G_i that maps (u, v) to u if $i = 1$ and to v if $i = 2$. Given a k_1 -coloring α of G_1 and a k_2 -coloring β of G_2 , we denote $\alpha \boxtimes \beta$ the $k_1 k_2$ -coloring of $G_1 \boxtimes G_2$ defined by $(\alpha \boxtimes \beta)(u, v) = (\alpha(u), \beta(v))$ (see Figure 2.10). The coloring $\alpha \boxtimes \beta$ refines α if β uses all the colors in $[k_2]$, however note that not all refinements of α can be decomposed as $\alpha \boxtimes \beta$ for some coloring β of G_2 .

We claim that refining a coloring $\alpha \in C$ is a property invariant under Kempe changes if $k_2 = \chi(G_2)$. Let $\alpha \in C$ and γ be a refinement of α . Let $a = (a_1, a_2)$ and $b = (b_1, b_2)$ in $[k_1] \times [k_2]$ and δ be a coloring obtained after an $\{a, b\}$ -Kempe change. Let S_a and S_b be the vertices of H colored a and b respectively by γ . Since γ refines α , the sets $\phi_1(S_a)$ and $\phi_1(S_b)$ are color classes of α , thus each copy of G_2 is colored with exactly $k_2 = \chi(G_2)$ colors. If $\phi_1(S_a) = \phi_1(S_b)$, then γ and δ differ only in one of the copies of G_2 . In this copy, both a and b still appear in δ because G_2 has chromatic number $k_2 = \chi(G_2)$, hence δ still refines α .

If $\phi_1(S_a) \neq \phi_1(S_b)$, then the graph K induced in G_1 by $\phi_1(S_a) \cup \phi_1(S_b)$ is a Kempe chain of α in G_1 , because γ refines α . For every edge $(u, v) \in E(K)$, all vertices (u, w) colored a are adjacent to all the vertices (v, x) colored b , hence the vertices colored a or b in $\phi_1^{-1}(K)$ are exactly $S_a \cup S_b$. As a result, δ refines the coloring of α' obtained from α after a Kempe change on K and hence still in C .

Given a k_1 -coloring α of G_1 and a k_2 -coloring β of G_2 , the coloring $\alpha \boxtimes \beta$ refines α . The coloring $\alpha' \boxtimes \beta$ cannot be a refinement of α for any α' that is not similar to α . Hence, if $\alpha' \notin C$, the coloring $\alpha' \boxtimes \beta$ is not equivalent to any $\alpha \boxtimes \beta$ with $\alpha \in C$. This proves that $G_1 \boxtimes G_2$ is not $(k_1 \cdot \chi(G_2))$ -recolorable.

Let us assume now that G_1 is critically k_1 -freezable with $k_1 \geq 2n_1 - 1$, and that G_2 is a clique. The number of edges in $G_1 \boxtimes G_2$ is $n_1|E(G_2)| + n_2|E(G_1)| + 2|E(G_1)||E(G_2)|$. By assumption, $|E(G_1)| = (k_1 - 1)(n_1 - k_1/2)$ and $|E(G_2)| = n_2(n_2 - 1)/2$. We have

$$\begin{aligned} |E(G_1 \boxtimes G_2)| &= \frac{1}{2} [n_1 n_2 (n_2 - 1) + n_2 (k_1 - 1) (2n_1 - k_1) \\ &\quad + n_2 (n_2 - 1) (k_1 - 1) (2n_1 - k_1)] \\ &= \frac{1}{2} [-n_1 n_2 - k_1 (k_1 - 1) n_2^2 + (2k_1 - 1) n_1 n_2^2] \end{aligned}$$

By Observation 2.1.4, note that $G_1 \boxtimes G_2$ needs to contain at least

$$\frac{1}{2} [k_1 n_2 - 2n_1 n_2 - k_1^2 n_2^2 + 2k_1 n_1 n_2^2]$$

edges to admit a frozen $k_1 n_2$ -coloring. The difference between the second term and the first is

$$\frac{1}{2} [k_1 n_2 - n_1 n_2 - k_1 n_2^2 + n_1 n_2^2] = \frac{1}{2} (n_1 - k_1)(n_2^2 - n_2) > 0,$$

hence $G_1 \boxtimes G_2$ does not admit a frozen $k_1 n_2$ -coloring. \square




Figure 2.10: The strong product of the 3-prism and a triangle: $(K_3 \square K_2) \boxtimes K_3$

2.1.1.2 Invariant from algebraic topology

We now move away from frozen colorings to present the second known obstruction to recolorability. Fisk initiated the study of Kempe equivalence using tools from algebraic topology by proving the following result:

Theorem 2.1.6 [Fis77a]

All 3-colorable triangulations of the sphere are 4-recolorable.

To prove Theorem 2.1.6, Fisk shows that every 4-coloring of a 3-colorable triangulation T is equivalent to the 3-coloring of T , unique up to color permutation. Given a 4-coloring α , every edge uv is part of two triangles uvw and uvx and is either *singular* if $\alpha(x) = \alpha(w)$ or *non-singular* if $\alpha(x) \neq \alpha(w)$. It is easy to see that

a coloring of T uses only three colors if and only if all the edges of T are singular. The approach of Fisk consists in greedily removing the non-singular edges to reach the 3-coloring of T . Using elementary tools from algebraic topology, Fisk observes that for any pair of colors a and b , all vertices have even degree in the subgraph formed by the non-singular edges whose endpoints are colored a and b . Therefore, in a coloring using exactly four colours, one can find a cycle C of non-singular edges colored with two colors a and b . Removing this cycle disconnects the graph in two components A and B . As it can be easily checked, swapping the remaining colors in A makes all the edges of C singular and does not create any other non-singular edge.

The main notion used by Fisk is that of the *degree* of a 4-coloring. The degree is defined in algebraic topology for any continuous map between topological manifolds of the same dimension, but we give here a purely combinatorial definition for simplicial maps. A 4-coloring α of a triangulation T of a surface can be viewed as a simplicial map from T to the surface $\partial\Delta^3$ of the tetrahedron. The triangles of T are mapped to triangles of $\partial\Delta^3$, with their orientation either preserved or reversed.

The degree of α is defined as the number of triangles of T that map to a fixed triangle t of $\partial\Delta^3$ with their orientation preserved minus those that have their orientation reversed. This quantity is constant over all triangles t of $\partial\Delta^3$, thus the degree is well-defined. Tutte proved a relation between the degree of a 4-coloring of a triangulation, and the degrees of the vertices of any color class:

Theorem 2.1.7 [Tut69]

Let T be a triangulation of a closed orientable surface. For any 4-coloring α and any color $a \in [4]$,

$$\deg(\alpha) \equiv \sum_{\substack{u \in V(T) \\ \alpha(u)=a}} \deg(u) \pmod{2}$$

As a result, Mohar and Salas observed that the parity of the degree of a 4-coloring is invariant under Kempe changes. Using this and a result of Fisk [Fis73], they proved the following:

Theorem 2.1.8 [MS09]

The degree of a 4-coloring of any 3-colorable triangulation of a closed surface is invariant modulo 12 under Kempe changes. In particular, the triangular lattice $T(3L, 3M)$ (see Subfigure 2.11 (a)) with $M \geq L \geq 3$ is not 4-recolorable.

The triangular grid $T(3L, 3M)$ is the toroidal graph built from the square grid $C_{3L} \square C_{3M}$ to which parallel diagonals are inserted in each of the 4-faces (see Subfigure 2.11 (a)).

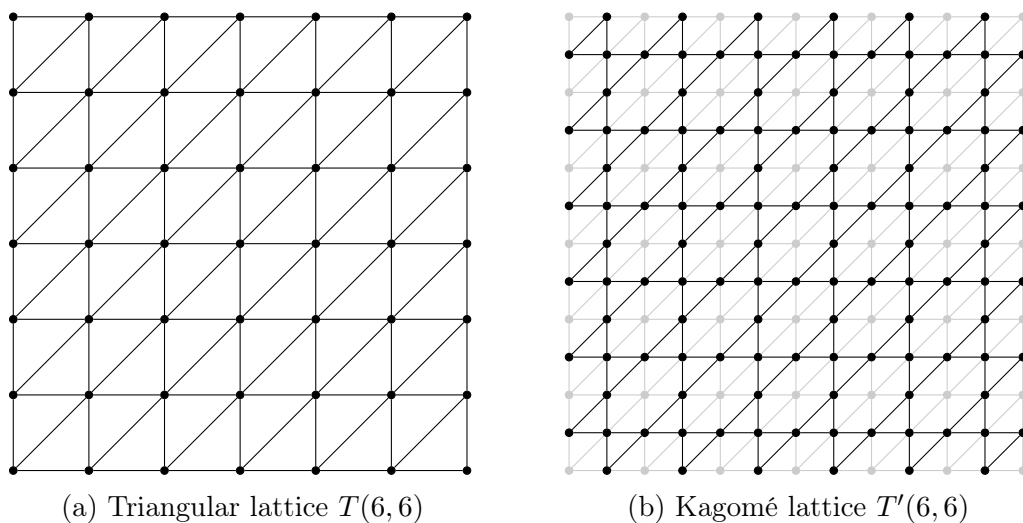


Figure 2.11: The Kagomé lattice $T'(3L, 3M)$ is defined as the medial graph of the triangular lattice $T(3L, 3M)$, but can also be viewed as a subgraph of $T(6L, 6M)$. In both subfigures, the upper and lower rank are identified, as well as leftmost and rightmost, to form a toroidal graph.

In a follow-up paper, Mohar and Salas [MS10] extended this result to the 3-colorings of the Kagomé lattice. The Kagomé lattice $T'(3L, 3M)$ (see Subfigure 2.11 (b)) is defined as the medial graph of the triangular grid $T(3M, 3L)$.

Mohar and Salas proved that the 3-edge-colorings of a triangulation T of a surface are in one-to-one correspondence with the 3-colorings of its medial graph and this bijection preserves the Kempe equivalence. The Kagomé lattice $T'(3M, 3L) = \text{Med}(T(3M, 3L))$ is also an induced subgraph of $T(6M, 6L)$, obtained by removing an independent set (see Subfigure 2.11 (b)). As a result, the 3-colorings of $T'(L, M)$ can be mapped one to one to a subset of 4-colorings of $T(6M, 6L)$, in such a way that $\mathcal{K}_3(T'(3M, 3L))$ is a subgraph of $\mathcal{K}_4(T(6M, 6L))$. Using these two correspondences and pointing out that Theorem 2.1.8 holds even for this subset of 4-colorings of $T(6M, 6L)$, Mohar and Salas were able to prove the following result:

Theorem 2.1.9 [MS10]

The Kagomé lattice $T'(3L, 3M)$ for $M \geq L \geq 3$ is not 3-recolorable.

The triangular lattice $T(3L, 3M)$ for $M \geq L \geq 3$ is not 3-edge-recolorable.

This correspondence between edge-coloring of a graph and vertex-coloring of its medial graph holds for all triangulations of a surface, but only for 3-colorings. We will see on other examples that comparing reconfiguration graphs with relations such as (induced) subgraphs or minors can sometimes ease significantly the proofs

of (non)-recolorability. The tools developed by Fisk, Mohar and Salas do not extend to a higher number of colors [MS09, Section 6] and it remains open whether there exist general invariants other than frozen colorings.

2.1.2 Positive results in sparse graphs classes

Independently of the early work of Fisk [Fis73, Fis77a], Meyniel studied the recolorability of planar graphs with a more combinatorial approach. In 1975, he proved the following:

Theorem 2.1.10 [Mey75]

All planar graphs are 5-recolorable. This is tight as there exist planar graphs that admit non-similar frozen 4-colorings (see Figure 2.12).



Figure 2.12: Two non-similar frozen 4-coloring of a planar graph

A closer look at the proof of Theorem 2.1.10 shows that it gives a reconfiguration sequence of length $O(5^n)$ between the 5-colorings of any n -vertex planar graph.

By Wagner’s celebrated characterization, planar graphs are exactly the graphs that forbid K_5 and $K_{3,3}$ as a minor. Meyniel proved with Las Vergnas that Theorem 2.1.10 is true even without the $K_{3,3}$ -minor-free assumption:

Theorem 2.1.11 [VM81]

All K_5 -minor-free graphs are 5-recolorable.

In the same paper, they draw interesting connections between recoloring with Kempe changes and Hadwiger’s conjecture, one of the most famous open problems in graph theory. Hadwiger conjectured in 1943 that for all t , every K_t -minor free graph is $(t - 1)$ -colorable. This conjecture is a generalization of the Four Color Theorem and was proven for t at most 6. We refer the curious reader to a survey of Seymour [Sey16] on Hadwiger’s conjecture. Las Vergnas and Meyniel conjectured the following, which generalizes Theorem 2.1.11 and can be seen as a reconfiguration counterpoint to Hadwiger’s conjecture, though it neither implies it nor is implied by it.

Conjecture 2.1.12 (Conjecture A in [VM81])

For every t , every K_t -minor-free graph is t -recolorable.

With Marthe Bonamy, Marc Heinrich and Jonathan Narboni [2], we disprove this conjecture by proving that for all ε , for large t , there exists a K_t -minor-free graph that admits a frozen $(\frac{3}{2} + \varepsilon)t$ -coloring and hence is not $(\frac{3}{2} + \varepsilon)t$ -recolorable. This raises the natural question:

Conjecture 2.1.13 [2]

There exists a constant c such that any K_t -minor-free graph is ct -recolorable.

Our construction shows that if such a constant exists then $c > 3/2$, but it does not seem to give any insights on how to answer Conjecture 2.1.13. However, we prove that no K_t -minor-free graph is k -freezable for any $k \geq 2t$, which motivates and supports this conjecture. This conjecture is reminiscent of the so-called “linear Hadwiger’s conjecture”: is there a constant c such that all K_t -minor-free graphs are ct -colorable? In the 1980s [Kos82, Kos84] and [Tho84] proved independently that K_t -minor-free graphs have degeneracy $O(t\sqrt{\log t})$, with current best hidden factor in [KP20]. As of today, the best bound for the linear Hadwiger’s conjecture is $O(t \log(\log(t)))$ [DP21], below the degeneracy threshold.

For Conjecture 2.1.13, the best known bound is $O(t\sqrt{\log t})$, due to the degeneracy argument combined with the following results of Las Vergnas and Meyniel:

Proposition 2.1.14 [VM81]

Let G be a graph and u be a vertex of degree less than k . If $G - u$ is k -recolorable, then G is k -recolorable.

We will give a proof of Proposition 2.1.14 as it is elementary. The following theorem is an immediate consequence of it.

Theorem 2.1.15 [VM81]

All d -degenerate graphs are k -recolorable, for all $k > d$.

Proof of Proposition 2.1.14. Let α and β be two k -colorings of G . Denote $G' := G - u$. We claim that we can apply a series of Kempe changes on α in G so as to obtain a coloring γ that agrees with β on G' . Assuming this claim, if $\gamma(u) \neq \beta(u)$, then neither $\gamma(u)$ nor $\beta(u)$ appear in $\gamma(N(u))$ since $\gamma|_{G'} = \beta|_{G'}$. As a result, performing the trivial Kempe change $K_{u,\beta(u)}(\gamma, G)$ yields β .

It remains only to prove the above claim. By hypothesis, there exists a sequence S of Kempe changes in G' leading from $\alpha|_{G'}$ to $\beta|_{G'}$. We will extend each Kempe change of S to one or two consecutive valid Kempe changes in G that result in a coloring that agrees on G' .

Let γ' and δ' be two colorings of G' such that γ' and δ' differ by one $\{a, b\}$ -Kempe chain K in G' . Let γ be an extension of γ' to G . Let L be the $\{a, b\}$ -Kempe

chain of γ that contains K . If $L \cap G' = K$, then performing the Kempe change L on γ results in a coloring δ whose restriction to G' is δ' as expected.

If $L \cap G' \neq K$, the issue is that adding u to G' results in merging two connected bichromatic components of G' . This can only be the case if u and at least two of its neighbors are colored a or b . Since u has degree at most d and $k > d$, there exists a color c that is not used in $N[u]$. After the trivial Kempe change that recolors u to c , the subgraph K becomes a valid Kempe chain of G . This concludes the proof of the claim, hence of Proposition 2.1.14. \square

We note that the sequence of Kempe changes arising from this proof has length exponential in the number of vertices. Giving a recoloring sequence of polynomial length for Theorem 2.1.15 is an open problem that we will discuss in Section 2.2 and Chapter 3. Theorem 2.1.15 is a central result in the recoloring landscape, as it can be used for many sparse graph classes. However, better bounds can be obtained in several sparse graph classes using their specific structure, as we have seen with planar graphs that are 5-degenerate but still 5-recolorable by Theorem 2.1.10.

As some planar graphs admit non-similar frozen 4-colorings, this is optimal, but more can be said on recolorability of planar graphs (see Figure 2.12). First, notice that taking p copies of the graph drawn on Figure 2.12 and gluing them by performing 3-sums on triangles, creates a planar triangulation with 2^p different frozen-colorings, thereby proving that the number of Kempe equivalence classes can be exponential in the number of vertices. A similar construction was described by Mohar in [Moh85, Moh07] starting with a different graph.

Notice that the planar triangulation drawn on Figure 2.12 is not 3-colorable as it has vertices of odd degree. Using the topological arguments mentioned in the previous section, Fisk proved that this condition is necessary:

Theorem 2.1.16 [Fis77a]

Let T be a 3-colorable triangulation of the sphere, the torus or the projective plane. All the 4-colorings of T of degree divisible by 12 are Kempe equivalent. In particular, all 3-colorable triangulated planar graphs are 4-recolorable.

This result was extended by Mohar to near-triangulations of the plane and then to 3-colorable planar graphs in general:

Theorem 2.1.17 [Moh07]

All 3-colorable planar graphs are 4-recolorable.

Combined with Lemma 2.1.1 and Theorem 2.1.10, Theorem 2.1.17 refines considerably our understanding of the recolorability of planar graphs:

Corollary 2.1.18 [Moh07]

All planar graphs are k -recolorable, for all $k > \chi(G)$.

Aiming for a full characterization of k -recolorable planar graphs, Mohar conjectured that “almost 3-colorable” planar graphs should not admit frozen 4-colorings and hence should be 4-recolorable. This was confirmed by Feghali:

Theorem 2.1.19 [Feg23]

All 4-critical planar graphs are 4-recolorable.

Moving away from planar graphs, another potential extension of Theorem 2.1.15 are Δ -colorings of graphs of maximum degree Δ . All connected graphs of maximum degree Δ are Δ -degenerate and hence $(\Delta + 1)$ -colorable and $(\Delta + 1)$ -recolorable. In fact, Brooks proved most of these graphs are even Δ -colorable. Namely, all graphs but the cliques and odd-cycles are Δ -colorable [Bro87]. Today, many different proofs of Brooks’ theorem are known, some of them using Kempe changes [CR15]. This raises the following question, first asked by Mohar in [Moh07]: which graphs are Δ -recolorable? Note that Proposition 2.1.14 provides a positive answer to this question for all connected graphs that have maximum degree Δ but are not Δ -regular, as such graphs are $(\Delta - 1)$ -degenerate. Hence only the regular case is of interest. This question was first answered in the restricted case of cubic graphs by Feghali, Johnson and Paulusma.

Theorem 2.1.20 [FJP17]

All cubic graphs but the 3-prism (see Figure 2.7) are 3-recolorable.

Bonamy, Bousquet, Feghali and Johnson then gave an answer for k -regular graphs, thereby settling the conjecture in full generality.

Theorem 2.1.21 [BBFJ19]

All k -regular graphs are k -recolorable for $k \geq 4$.

Corollary 2.1.22 [BBFJ19]

Let $k \geq \Delta \geq 3$. Any graph G of maximum degree Δ is k -recolorable, unless $k = 3$ and $G = K_2 \square K_3$.

A key step in the proof of Theorem 2.1.21 consists in observing that a Δ -regular graph G that is not a clique contains a vertex u adjacent to a pair of vertices v and w , called eligible pair, that are identically colored. Merging v with w decreases the degree of u by one and results in a $(\Delta - 1)$ -degenerate graph G' that is Δ -recolorable. This proves that all the colorings of G that are identical on v and w are Kempe equivalent.

Observe that this argument applies to more complicated reduction operations than the identification of one pair of vertices, we generalize this idea to pairs of graphs related by a homomorphism, using the notion of compatible colorings:

Lemma 2.1.23

Let G and G' be two graphs, with a homomorphism $\phi : G \rightarrow G'$. Let α be a ϕ -compatible coloring of G . Let α' be the k -coloring of G' induced by α . The Kempe equivalence class of α contains all colorings ϕ -compatible with a coloring β' Kempe equivalent to α' in G' .

$$\begin{array}{ccc} \alpha & \longleftrightarrow & \beta \\ \downarrow \phi & & \downarrow \phi \\ \alpha' & \longleftrightarrow & \beta' \end{array}$$

Proof. Let α' and β' be two colorings of G' that differ by one $\{a, b\}$ -Kempe change K . The preimage of K by ϕ is a collection of disjoint $\{a, b\}$ -Kempe chains. Performing a Kempe change in each of them results in a coloring β ϕ -compatible with β' . The results follows by an immediate induction on the length of the reconfiguration sequence between α' and β' . \square

In particular, given a coloring α of G , Lemma 2.1.23 applies when ϕ is a combination of identifications of vertices colored identically in α and addition of edges between vertices colored differently.

In [BBFJ19], Bonamy, Bousquet, Feghali and Johnson cleverly use this identification argument twice: given two colorings α and β of G , the trick is to identify for each of them an eligible pair, respectively (v, w) and (x, y) , such that the pairs are sufficiently far away from one another. The above argument shows that the colorings that color v and w (resp. x and y) alike are all Kempe equivalent. Since the pairs are far away, there exists a coloring γ that colors v and w alike, as well as x and y , which concludes the proof.

Once again, this argument can be generalized to more complicated reduction operations and, in principle, the two operations do not need to be the same. Although one might prefer in practice to use this idea directly in a proof rather than apply a technical lemma, we give the following generalization that adapts the ideas of Lemma 2.1.23:

Lemma 2.1.24

Let G , H_1 and H_2 be three graphs such that G contains two disjoint copies of H_1 and H_2 . Let $\phi_i : H_i \rightarrow H'_i$ be two homomorphisms. Denote $\psi_i : G \rightarrow G_i$ the homomorphism obtained by applying ϕ_i on H_i and the identity on the other vertices and edges.

If there is a coloring ψ_1 and ψ_2 -compatible, and both G_1 and G_2 are k -recolorable, then the colorings of G that are ψ_1 -compatible or ψ_2 -compatible are all Kempe equivalent.

Proof. This is a direct application of Lemma 2.1.23 on G_1 and G_2 . \square

2.1.3 Beyond frozen colorings: open questions

The topological techniques used by Mohar and Salas to prove Theorems 2.1.8 and 2.1.9 are specific to 4-colorings and cannot be used for a higher number of colors. Therefore, we wonder if for large k , frozen colorings are the only obstacle to k -recolorability. As it is, Proposition 2.1.5 provides a negative answer to this question, hence we ask this more general question:

Question 2.1.25

Is there another obstacle to k -recolorability, beside frozen colorings and the topological argument of Fisk, Mohar and Salas?

In the remainder of this section, we present two graph classes that cannot admit frozen k -colorings, hoping that the study of their recolorability can either support this question or lead to the development of other techniques to prove that some graph is not k -recolorable.

2.1.3.1 Grötzsch recoloring

In 1959, Grötzsch proved that triangle-free planar graphs are 3-colorable [Grö59]. In 2003, Thomassen gave a shorter proof by showing that a minimum counterexample cannot contain any four-cycle and by showing that all planar graphs of girth at least 5 are 3-list-colorable [Tho03]. This approach was then further simplified by Asghar in [Asg12].

Dvořák, Král and Thomas generalized this result by proving the existence of an absolute constant d such that any planar graph whose triangles are at distance at least d from one another is 3-colorable [DKT21]. Among the other generalizations of Grötzsch theorem, we mention the following generalization of Aksionov, that is very useful in inductions:

Theorem 2.1.26 [Aks74]

Let G be a triangle-free planar graph and C be a cycle of G . If $|C| \leq 5$, any 3-coloring of C can be extended to a 3-coloring of G .

When $6 \leq |C| \leq 8$, Dvořák and Lidický gave a full characterization of the cases in which a 3-coloring of C can be extended to G [DL15].

We will here focus on the 3-recolorability of triangle-free planar graphs.

Proposition 2.1.27

No triangle-free planar graph admits a frozen 3-coloring.

Proof. By Observation 2.1.4, an n -vertex graph G that admits a frozen 3-coloring contains at least $2n - 3$ edges. By Euler's formula, the number of edges in an

n -vertex triangle-free planar graph is at most $2n - 4$:

$$\begin{aligned} |V(G)| - |E(G)| + |F(G)| &= 2 \\ 4n - 2|E(G)| + \sum_{f \in F(G)} (4 - \deg(f)) &= 8 \\ 2n - 4 &\geq |E(G)| \end{aligned}$$

because $\deg(f) \geq 4$ for all $f \in F(G)$. □

This motivates the following conjecture, independently proposed by Salas and Sokal:

Conjecture 2.1.28 (Bonamy, Legrand-Duchesne and [SS22])

All triangle-free planar graphs are 3-recolorable.

There are two natural directions to weaken Conjecture 2.1.28: either increase the number of colors or use the sparsity of planar graphs. Note that by Theorem 2.1.17, all triangle-free planar graphs are 4-recolorable, hence only the case $k = 3$ is open. On the other hand, the average degree of a planar graph, hence its degeneracy, is related to its girth via Euler's formula. A planar graph of girth at least six is 2-degenerate and thus is 3-recolorable by Theorem 2.1.15. So the first interesting case is that of planar graphs of girth 5, therefore, we conjecture the following.

Conjecture 2.1.29

All planar graphs of girth 5 are 3-recolorable.

Following an approach reminiscent of Thomassen's proof of Grötzsch's theorem, we will show that Conjecture 2.1.28 reduces to Conjecture 2.1.29:

Proposition 2.1.30

Conjecture 2.1.28 and Conjecture 2.1.29 are equivalent.

Clearly Conjecture 2.1.28 implies Conjecture 2.1.29.

In what follows, we describe general methods to prove statements of the form "In the hereditary class of graphs \mathcal{F} , all graphs are k -recolorable" and use them to prove the other direction of Proposition 2.1.30. A common approach consists in a proof by minimal counterexample. Consider S_1, \dots, S_p some well-chosen configurations such that any non-empty graph in \mathcal{F} contains at least one of them. Let G be a minimal graph of \mathcal{F} that is not k -recolorable and let i be such that G contains S_i . Using S_i , we build from G a smaller graph $G' \in \mathcal{F}$ such that the reconfiguration sequences of G' can be extended to G . By minimality of G , the graph G' is k -recolorable, so G is also k -recolorable, a contradiction.

Note that this method can also be viewed as an involved induction. When working on planar graphs, a double counting argument, called the *discharging method*, combined with Euler's formula is often particularly handy to find the configurations S_i . For example, the proofs of the Four Color Theorem are proofs by minimal counterexample that rely heavily on discharging to produce a list of unavoidable configurations (633 of them for the proof of Robertson, Sanders, Seymour and Thomas and 1,834 of them for the original proof of Appel and Haken).

To shorten the size of the reconfiguration sequences and simplify the proofs, a trick consists in considering *proper partitions* instead of proper colorings: a proper k -partition of G is an unlabelled partition of $V(G)$ into k stable sets.

Alternatively, proper k -partitions can be defined as the quotient of proper k -colorings by the similarity relation: $\alpha \sim \beta$ if α and β are similar. We will denote $\tilde{\alpha}$ the proper partition corresponding the equivalence class under \sim of a coloring α . We will then say that $\tilde{\alpha}$ and $\tilde{\beta}$ differ by a *Kempe swap*.

Quotienting $\mathcal{K}_k(G)$ by \sim results in a graph whose vertices are the proper k -partitions of G and in which two proper partitions $\tilde{\alpha}$ and $\tilde{\beta}$ are adjacent if there exists a coloring γ that differs from α by one Kempe change and such that $\gamma \sim \beta$. We will denote this quotient graph $\tilde{\mathcal{K}}_k(G)$.

Lemma 2.1.31

Let α and β be two colorings of an n -vertex graph G such that $\tilde{\alpha}$ and $\tilde{\beta}$ are at distance d in $\tilde{\mathcal{K}}_k(G)$. Then α and β are at distance at most $d + n$ in $\mathcal{K}_k(G)$. As a result, the number of connected components in $\mathcal{K}_k(G)$ and $\tilde{\mathcal{K}}_k(G)$ are equal and for any connected component C of $\mathcal{K}_k(G)$,

$$\text{diam}(C/\sim) \leq \text{diam}(C) \leq \text{diam}(C/\sim) + n$$

Proof. A straightforward induction on d shows that for any two colorings α and β such that $\tilde{\alpha}$ and $\tilde{\beta}$ are at distance d in $\tilde{\mathcal{K}}_k(G)$, α is at distance d from a coloring γ with $\gamma \sim \beta$. Starting from γ , recoloring a color class A that is not colored identically in β can be done in at most $|A|$ Kempe changes. This results in a coloring similar to β , that agrees with β on A . By induction, we can continue with the other color classes, performing in total at most n Kempe changes. \square

Given a graph G and a subgraph S , we will call k -prints of S in G the set of k -partitions of S that are restrictions of k -partitions of G . We will call k -print graph of S the graph $\tilde{\mathcal{K}}_k(S, G)$ whose vertices are the k -prints of S in G and in which two k -prints are adjacent if they differ by one Kempe swap.

Consider three connected graphs H_1 , H_2 and H along with two surjective homomorphisms $\phi_1 : H_1 \rightarrow H$ and $\phi_2 : H_2 \rightarrow H$. We say that H_1 and H_2 are *simultaneously walkable* if for all (u_1, u_2) and $(v_1, v_2) \in V(H_1) \times V(H_2)$ with $\phi_1(u_1) = \phi_2(u_2)$

and $\phi_1(v_1) = \phi_2(v_2)$, there exists a walk in $H_1 \times H_2$ between (u_1, u_2) and (v_1, v_2) such that at each step (w_1, w_2) , we have $\phi_1(w_1) = \phi_2(w_2)$.

A separator S is said to be k -nice if it satisfies the following conditions:

- (S1) The set of prints of G , G_1 and G_2 are equal,
- (S2) The minor H_i obtained from $\tilde{\mathcal{K}}_k(G_i)$ by contracting adjacent partitions with identical k -prints is connected, for all $i \in \{1, 2\}$,
- (S3) H_1 and H_2 are simultaneously walkable with respect to the k -print graph of S in G .

Proposition 2.1.32

Let G be a graph and S be a separator of G . Let G_1 and G_2 be two subgraphs of G such that $G = G_1 \cup G_2$ and $G_1 \cap G_2 = S$. If S is a nice k -separator, then $\mathcal{K}_k(G)$ is connected.

Proof. Let S be a nice k -separator in G . Let H_1 and H_2 be the minors of $\mathcal{K}_k(G_1)$ and $\mathcal{K}_k(G_2)$ respectively obtained by contracting adjacent colorings with identical prints.

Let α and β be two k -colorings of G . Since S is nice, there exists a sequence $\tilde{\gamma}_0, \dots, \tilde{\gamma}_\ell$ of proper k -partitions of G_1 , and a sequence $\tilde{\delta}_0, \dots, \tilde{\delta}_\ell$ of proper k -partitions of G_2 , such that $\tilde{\gamma}_0 = \tilde{\alpha}|_{G_1}$, $\tilde{\gamma}_\ell = \tilde{\beta}|_{G_1}$, $\tilde{\delta}_0 = \tilde{\alpha}|_{G_2}$, $\tilde{\delta}_\ell = \tilde{\beta}|_{G_2}$ and for all i , the prints induced by $\tilde{\gamma}_i$ and $\tilde{\delta}_i$ on S are equal. Furthermore, for all i , the partitions $\tilde{\gamma}_i$ and $\tilde{\gamma}_{i+1}$ (respectively $\tilde{\delta}_i$ and $\tilde{\delta}_{i+1}$) are equivalent up to a sequence of Kempe swaps that do not modify the print on S . So α and β are Kempe equivalent by Lemma 2.1.31. □

Let G be a planar graph of girth 5 and F be a 5-face of G . By Theorem 2.1.26, there are five 3-prints of F in G , each one of them is characterized by which of the five vertices of F is alone in its color class. Hence, the 3-print graph of F in G is a C_5 (see Figure 2.13). As a direct illustration of Proposition 2.1.32, we prove the following:

Observation 2.1.33. Conjecture 2.1.29 would imply the following property:

Consider H the minor obtained from $\tilde{\mathcal{K}}_3(G)$ by contracting adjacent partitions with identical 3-prints of F . For all $u \in H$, the vertex u has two neighbors of different prints. In particular, for any two adjacent vertices x and y of F , for any 3-coloring α that colors x uniquely in F , there exists a recoloring sequence that leads to a 3-coloring β that colors y uniquely in F , without recoloring F intermediately.

Proof of Observation 2.1.33. Assume that Conjecture 2.1.29 holds and let G be a planar graph of girth five. Let u and v and w be three vertices of some 5-face




Figure 2.13: The 3-print graph of F in G

F of G with u adjacent to v and w . Assume towards contradiction that there is a 3-coloring α that colors u uniquely in F , such that any recoloring sequence changing the partition induced in F recolors w uniquely in F .

Let x and y be the remaining vertices of F and consider the graph G' composed of two copies G_1 and G_2 of G glued along F in opposite directions: both copies of u are identified, while each copy of v is identified with the opposite copy of w and likewise for x and y . The graph G' is planar as well and we will argue that G' is not 3-recolorable. We label the vertices of F in G' according to their labels in G_1 . Let α' be the coloring of G' defined as $\alpha'_{|G_1} = \alpha_{|G_1}$ and $\alpha'_{|G_2 \setminus F} = \sigma \circ \alpha_{|G_2 \setminus F}$ where σ is the transposition switching $\alpha(v)$ with $\alpha(w)$ (see Figure 2.14).




Figure 2.14: The graph G'

By assumption, starting from α' , any sequence of Kempe changes that modifies the 3-print of F , must first color w uniquely in F because of G_1 , but should also first color v uniquely in F because of G_2 . Thus, α' is only Kempe equivalent to colorings that have the same 3-print of F . However, by Theorem 2.1.26, there exist colorings of G' with different 3-prints of F , which leads to a contradiction. \square

In what follows, we will prove that Conjecture 2.1.29 implies Conjecture 2.1.28. To that end, we first show that a minimal counterexample to Conjecture 2.1.28 is 3-connected (see Lemma 2.1.35) and contains no separating clique (see Lemma 2.1.34) and no separating C_4 (see Lemma 2.1.36). Finally, we prove that assuming Conjecture 2.1.29, a minimal counterexample to Conjecture 2.1.28 cannot contain a 4-face, and thus Conjecture 2.1.29 implies Conjecture 2.1.28.

Lemma 2.1.34

Let $G = G_1 \cup G_2$ with $G_1 \cap G_2 = S$, where S is a clique. For any k , if both G_1 and G_2 are k -recolorable, then S is a k -nice separator and hence G is k -recolorable.

In particular, a minimal counterexample to Conjecture 2.1.28 contains no separating clique.

Proof. Since a clique S has only one proper partition, it has a unique print and (S1) is verified. G_1 and G_2 are k -recolorable, hence H_1 and H_2 are connected, and in fact they are both reduced to a single vertex, which proves (S2). Finally, H_1 and H_2 are simultaneously walkable with respect to the k -print graph of S because all three graphs are reduced to a single vertex, which proves (S3). \square

Lemma 2.1.35

Let $G = G_1 \cup G_2$ with $G_1 \cap G_2 = S$, where $|S| \leq 2$. For any k , if both G_1 and G_2 are k -recolorable and any k -coloring of G_1 (respectively G_2) extends to a coloring of G , then S is a k -nice separator and hence G is k -recolorable.

In particular, a minimal counterexample to Conjecture 2.1.28 is 3-connected.

Proof. The case where S induces an edge is implied by Lemma 2.1.34. Assume now that S is composed of two non-adjacent vertices u and v . The assumption that any k -coloring of G_1 (respectively G_2) extends to a coloring of G ensures that the number of k -prints of S are the same in G , G_1 and G_2 . If this number is one, then the proof is identical to the proof of Lemma 2.1.34.

Otherwise, G , G_1 and G_2 have the same two k -prints of S , so (S1) is verified. Hence, H_1 and H_2 are bipartite. Since G_1 and G_2 are k -recolorable, H_1 and H_2 are connected, which proves (S2).

Let $(u_1, u_2), (v_1, v_2) \in V(H_1) \times V(H_2)$ with $\phi_1(u_1) = \phi_2(u_2)$ and $\phi_1(v_1) = \phi_2(v_2)$. Since H_1 and H_2 are connected, there exist two walks (x_1, \dots, x_p) in H_1 and (y_1, \dots, y_q) in H_2 such that $x_1 = u_1$, $x_p = v_1$, $y_1 = u_2$ and $y_2 = v_2$. If $p = q$ then for all $i \in [p]$, $\phi_1(x_i) = \phi_2(y_i)$ because H_1 and H_2 are bipartite and

$\phi_1(u_1) = \phi_2(u_2)$. If $p \neq q$, say $p < q$, alternating between x_p and x_{p-1} for $q - p$ steps at the end of (x_1, \dots, x_p) reduces to the previous case. As a result, H_1 and H_2 are simultaneously walkable with respect to H , which proves (S3).

To show that this lemma applies to a minimal counterexample to Conjecture 2.1.28, we only need to prove that any 3-coloring of G_i extends to G , when G is a triangle-free planar graph and u and v are not adjacent. This can be easily seen by adding a vertex adjacent to both u and v , as this creates a four-cycle that can be precolored arbitrarily by Theorem 2.1.26. \square

Lemma 2.1.36

Let $G = G_1 \cup G_2$ with $G_1 \cap G_2 = S$, where S is a C_4 . If both G_1 and G_2 are 3-recolorable, and any 3-coloring of G_1 (respectively G_2) extends to a coloring of G , then S is a 3-nice separator and hence G is 3-recolorable.

In particular, a minimal counterexample to Conjecture 2.1.28 cannot contain a separating C_4 .

Proof. Denote a, b, c and d the vertices of C . There are three possible 3-partitions of a four-cycle. Two of them have two opposite vertices in distinct classes, we will refer to them as *one-sided*. In the last one, each pair of opposite vertices is in one class, we will call it *contractible* (see Figure 2.15). By assumption, G_1, G_2 and G




Figure 2.15: The three 3-prints of a C_4 in G

have the same 3-print by S . If they contain at most two distinct prints, then the proof is identical to that of Lemma 2.1.34 or Lemma 2.1.35, hence we will assume that G_1, G_2 and G_3 all contain the three prints drawn on Figure 2.15. Note that this is the case when G_1 and G_2 are triangle-free planar graphs, by Theorem 2.1.26.

Since G_1 and G_2 are 3-recolorable, H_1 and H_2 are connected. No Kempe change can transform a one-sided print into the other, thus the 3-print graph of S is isomorphic to P_3 , with the contractible print being the vertex of degree two.

H_1 is connected, thus there exist $w_1 \in V(H_1)$ contractible that is adjacent to one-sided vertices corresponding to different 3-prints of S in G_1 , and likewise in H_2 . We will call *pivots* the vertices of H_1 or H_2 with this property.

We will build a simultaneous walk starting from any $(u_1, u_2) \in V(H_1) \times V(H_2)$ with $\phi_1(u_1) = \phi_2(u_2)$ that finishes at a pair of contractible pivots. After that, it will suffice to exhibit simultaneous walks between pairs of pivots.

Let $(u_1, u_2) \in V(H_1) \times V(H_2)$ with $\phi_1(u_1) = \phi_2(u_2)$. Since H_1 and H_2 are connected, there exist two walks (x_1, \dots, x_p) in H_1 and (y_1, \dots, y_q) in H_2 such that $x_1 = u_1$, $y_1 = u_2$, and with x_p and y_q the two only pivots of their respective sequence. Assume without loss of generality that x_1 and y_1 are one-sided. Since H is bipartite, for all i , the partitions x_{2i} and y_{2i} are contractible while x_{2i+1} and y_{2i+1} are one-sided. As x_p and y_q are the only pivots in the two sequences, all one-sided vertices in the sequences have the same print. Hence, up to alternating between y_1 and y_2 for $p - q$ steps at the beginning if $p > q$, there exists a simultaneous walk between (u_1, u_2) and (x_p, y_q) .

Let $(u_1, u_2), (v_1, v_2) \in V(H_1) \times V(H_2)$ be two pairs of pivots, with $\phi_1(u_1) = \phi_2(u_2)$ and $\phi_1(v_1) = \phi_2(v_2)$. As H_1 and H_2 are connected, there exist two walks (x_1, \dots, x_p) in H_1 and (y_1, \dots, y_q) in H_2 between u_1 and v_1 and between u_2 and v_2 respectively. As $u_2 = y_1$ is a pivot, there exist y_0 adjacent to y_1 that is one-sided with a print different from y_2 . Likewise, there exist x_{p+1} one-sided, adjacent to x_p with a print different from x_{p-1} .

We build the simultaneous walk $((x'_1, y'_1), \dots, (x'_{p+q}, y'_{p+q}))$ such that for $i \in [p]$,

$$x'_i = x_i \quad \text{and} \quad y'_i = \begin{cases} y_0 & \text{if } x_i \text{ is one-sided, with the same print as } y_0 \\ y_1 & \text{if } x_i \text{ is contractible} \\ y_2 & \text{if } x_i \text{ is one-sided, with the same print as } y_2 \end{cases}$$

and for all $i \in [q]$,

$$x'_{p+i} = \begin{cases} x_{p-1} & \text{if } y_i \text{ is one-sided, with the same print as } x_{p-1} \\ x_p & \text{if } y_i \text{ is contractible} \\ x_{p+1} & \text{if } y_i \text{ is one-sided, with the same print as } x_{p+1} \end{cases} \quad \text{and} \quad y'_{p+i} = y_i$$

By Theorem 2.1.26, any coloring of S extends to G , thus this lemma applies to a minimal counterexample to Conjecture 2.1.28. \square

To conclude the proof of Proposition 2.1.30, it only remains to prove the following lemma:

Lemma 2.1.37

Assuming that Conjecture 2.1.29 holds, a minimal counterexample to Conjecture 2.1.28 cannot contain a four-face.

Proof. Let G be a minimal counterexample to Conjecture 2.1.28. By Lemmas 2.1.34 and 2.1.35, G is 3-connected and has no separating C_4 . Let C be a four-face in G and denote a, b, c and d the vertices of C in cyclic order. By Theorem 2.1.26, there are exactly three different 3-prints of C in G (see Figure 2.15, taking $G_2 = G$ and $G_1 = \emptyset$). Since G is triangle-free and planar, it is impossible that $\text{dist}(a, c) \leq 3$ and $\text{dist}(b, d) \leq 3$ in the graph $G' = G \setminus E(C)$. Suppose without loss of generality that $\text{dist}_{G'}(a, c) > 3$, then identifying a and c results in a triangle-free planar graph G^* . By assumption, G^* is 3-recolorable, hence all 3-colorings of G with a and c colored identically are Kempe equivalent and it suffices to show that any 3-partition $\tilde{\alpha}$ is equivalent to a 3-partition with a and c in the same class. If we also have $\text{dist}_{G'}(b, d) > 3$, then the same reasoning holds and since G admits a partition of contractible print, all the 3-partitions of G are equivalent.

We assume now that $\text{dist}_{G'}(b, d) \leq 3$. As $\text{dist}_{G'}(b, d) = 2$ would imply that G admits a separating C_4 and hence contradict Lemma 2.1.36, we can even assume that $\text{dist}_{G'}(b, d) = 3$. Thus, G can be decomposed as $G_1 \cup G_2$ with $G_1 \cap G_2$ isomorphic to C_5 . Among all such four-faces, there is one for which G_1 is also of girth 5. Without loss of generality, assume that C is such a face. Denote $D = G_1 \cap G_2$ and e and f the vertices on the path of length 3 going from d to b in D . Assume with loss of generality that c and f are colored identically (see Figure 2.16). By Observation 2.1.33, there exists a sequence \mathcal{S} of Kempe changes




Figure 2.16: Final configuration to reduce: $G = G_1 \cup G_2$ with G_1 of girth 5

in G_1 that does not affect the coloring of D after which some Kempe change can be performed which results in f being colored uniquely in D . Let β be the coloring obtained after \mathcal{S} . In β , the vertex c does not belong to the $\{\beta(f), \beta(e)\}$ -Kempe chain containing f . So recoloring c with $\beta(e)$ does not change the color of a and results in a coloring with a and c in the same class, as desired. \square

2.1.3.2 Reed's Conjecture

The chromatic number of any graph G is at least its clique number $\omega(G)$ and a straightforward greedy algorithm can color G using at most $\Delta(G) + 1$ colors. A natural question is which of these bounds lies closer to the chromatic number. Reed conjectured that the chromatic number is upper bounded by the average of those two bounds, up to rounding:

Conjecture 2.1.38 [Ree98]

For any graph G ,

$$\chi(G) \leq \left\lceil \frac{\omega(G) + \Delta(G) + 1}{2} \right\rceil.$$

In particular, for any $\varepsilon \leq 1/2$ and any graph G ,

$$\chi(G) \leq \lceil (1 - \varepsilon)(\Delta(G) + 1) + \varepsilon\omega(G) \rceil.$$

In the same paper, Reed proved using the probabilistic method that this bound is tight if true. He also proved that there exists $\varepsilon > 0$ such that for all G , $\chi(G) \leq \lceil (1 - \varepsilon)(\Delta(G) + 1) + \varepsilon\omega(G) \rceil$. More recently, King and Reed [KR16] gave a significantly shorter proof of this result and proved that for large enough Δ , Reed's conjecture holds for $\varepsilon \leq \frac{1}{320e^6}$. In [BPP22], Bonamy, Perrett and Postle proved the same statement, for $\varepsilon \leq 1/26$. Delcourt and Postle [DP17] then proved a similar statement for the list chromatic number, improving on the way the ratio for the chromatic number: for all G with $\Delta(G)$ sufficiently large, $\chi(G) \leq \chi_\ell(G) \leq \lceil \frac{12}{13}(\Delta(G) + 1) + \frac{1}{13}\omega(G) \rceil$. This conjecture generated a lot of interest over the past years, the best bound known today is due to Hurley, de Joannis de Verclos and Kang:

Theorem 2.1.39 [HdK21]

For all graph G with sufficiently large maximum degree,

$$\chi(G) \leq \lceil 0.881(\Delta(G) + 1) + 0.119\omega(G) \rceil.$$

We consider here a recoloring version of this conjecture:

Question 2.1.40

What is the largest $\varepsilon \leq 1/2$ such that all graphs G are k -recolorable for all $k \geq \lceil (1 - \varepsilon)(\Delta(G) + 1) + \varepsilon\omega(G) \rceil$?

Note that ε is at most $1/3$. Indeed, we built with Marthe Bonamy, Marc Heinrich and Jonathan Narboni [7] a random graph with degrees concentrating around $\Delta = 3n/4$ and expected clique number $O(\log(n))$, that admits an $n/2$ -frozen coloring and is not $n/2$ -recolorable with high probability. This construction is detailed in Chapter 3.

In [Wei16], Weil proves that Reed's conjecture holds for odd-hole free graphs. With Marthe Bonamy and Tomáš Kaiser, we give an alternative proof of this result using Kempe changes:

Theorem 2.1.41 (Bonamy, Kaiser and LeGrand-Duchesne)

Let G be an odd-hole-free graph. Any k -coloring of G is Kempe equivalent to a $\left\lceil \frac{\Delta(G)+1+\omega(G)}{2} \right\rceil$ -coloring of G . In particular, Reed's conjecture holds for odd-hole-free graphs.

Proof. Let α be a k -coloring of G . Denote $q = \left\lceil \frac{\Delta(G)+\omega(G)+1}{2} \right\rceil$. If α is not a q -coloring, let u be a vertex with maximal color c . Let A be the set of vertices in $N(u)$ such that there exists at least one other vertex colored identically in $N(u)$ and B be the remaining vertices of $N(u)$. If B is not a clique, then let $v, w \in B$ such that $vw \notin E(G)$. The Kempe chain $K = K_{v,\alpha(w)}(G, \alpha)$ cannot contain w , otherwise the shortest path between v and w in K would form an odd-hole with u (see Figure 2.17). As a result, we can perform the corresponding Kempe change and add v and w to A . We can now assume that B is a clique. The number of colors used in the neighborhood of G is at most $|A|/2 + |B| \leq (|N(u)| - |B|)/2 + |B| \leq \frac{\Delta(G)+\omega(G)}{2}$. As a result there exists at least one color c of $[q]$ that is not used in $N(u)$ and u can be recolored to c with a trivial Kempe change. Notice that we never performed any Kempe change involving the color c until the recoloring of u , so the number of vertices colored with c decreases by exactly one. By repeating this process on all vertices colored c , one reaches a coloring with one fewer color and the result is obtained by induction. \square




Figure 2.17: Greedy recoloring of the neighborhood. The $\{\bullet, \bullet\}$ -Kempe chain containing v and represented by a dashed line cannot contain w .

This in particular proves that odd-hole-free graphs have no frozen k -colorings for $k > \left\lceil \frac{\Delta(G)+1+\omega(G)}{2} \right\rceil$. This motivates the following conjecture:

Conjecture 2.1.42 (Bonamy, Kaiser and Legrand-Duchesne)

All odd-hole free graphs of maximum degree Δ and clique number ω are k -recolorable for $k \geq \lceil \frac{\Delta+1+\omega}{2} \rceil$.

By Theorem 2.1.41, note that it suffices to prove that odd-hole-free graphs are $\lceil \frac{\Delta(G)+1+\omega(G)}{2} \rceil$ -recolorable to prove Conjecture 2.1.42 for higher k . We settle the special case of perfect graphs and prove a weakening of Conjecture 2.1.42 for general odd-hole-free graphs:

Theorem 2.1.43 (Bonamy, Kaiser and Legrand-Duchesne)

Let G be an odd-hole-free graph. G is k -recolorable for $k > \lceil \frac{\Delta(G)+1+\chi(G)}{2} \rceil$. In particular, any perfect graph G is k -recolorable for all $k > \lceil \frac{\Delta(G)+1+\omega(G)}{2} \rceil$, and any odd-hole-free graph G is k -recolorable for $k > \lceil \frac{3(\Delta(G)+1)+\omega(G)}{4} \rceil$.

Proof. We prove the result by induction on $\chi(G)$. Let α and β be two k -colorings of G with $k \geq q + 1 = \lceil \frac{\Delta(G)+\chi(G)+3}{2} \rceil$. Let α' and β' be the q -colorings obtained from α and β respectively by Theorem 2.1.41. Let γ be a $\chi(G)$ -coloring of G . Let A be the set of vertices of one of the color classes of γ . Recolor all the vertices of A in α' and in β' with color $q + 1$. $G[V \setminus A]$ is odd-hole-free and has chromatic number $\chi(G) - 1$, so we can apply induction on it with the induced colorings. \square

Note that the sequence of Kempe changes constructed in Theorem 2.1.41 has length $O(kn)$ which proves that $\mathcal{K}_k(G)$ has diameter $O(k^2n)$ for odd-hole-free graphs with $k > \lceil \frac{\Delta(G)+1+\chi(G)}{2} \rceil$.

2.2 The diameter of $\mathcal{K}_k(G)$

The recoloring sequences arising from most of the proofs of the results presented in Subsection 2.1.2 have length exponential in the number of vertices. In particular, the sequences constructed in Theorem 2.1.15 and Theorem 2.1.10 have length $O(2^n)$ and $O(5^n)$ and several result use Theorem 2.1.15 as a subroutine. In the context of single vertex-recoloring, all the k -colorings of a d -degenerate graph are equivalent for $k \geq d + 2$. A famous conjecture of Cereceda asks whether they are equivalent up to a quadratic number of recolorings. Cereceda's conjecture remains open, even for $d = 2$, but Bousquet and Heinrich exhibited in a breakthrough paper [BH22] a recoloring sequence of length $O(n^{d+1})$. Unfortunately, this proof does not seem to translate to Kempe changes with one fewer color, namely when $k = d + 1$. Bonamy, Bousquet, Feghali and Johnson conjectured the following:

Conjecture 2.2.1 [BBFJ19]

Let G be an n -vertex d -degenerate graph. Any two k -colorings of G are Kempe equivalent up to $O(n^2)$ Kempe changes for $k \geq d + 1$.

Proving a polynomial upper bound in Conjecture 2.2.1 would already be a significant result. With Marthe Bonamy and Vincent Delecroix, we proved in [6] three polynomial bounds on the diameter of the reconfiguration graph in closely related settings. We only mention the following strengthening of Corollary 2.1.22 and refer the reader to Chapter 3 for more details on this.

Theorem 2.2.2 [6]

Let G be an n -vertex graph of maximum degree Δ . For $k \geq \Delta$, any two k -colorings of G are equivalent up to at most $O_\Delta(n^2)$ Kempe changes, unless $k = 3$ and G is the 3-prism.

In another attempt to prove Conjecture 2.2.1, we studied with Quentin Dechamps, Carl Feghali, František Kardoš and Théo Pierron the k -recolorability of planar graphs for $k = 5$ and 6, as they are known to be 5-degenerate and even 5-recolorable by Theorem 2.1.10. We improved on Theorem 2.1.10 by proving a polynomial upper bound on the diameter of $\mathcal{K}_5(G)$ for planar graphs. Although this is not mentioned in [Moh07], we also observed that the proof of Theorem 2.1.17 gives a reconfiguration sequence of length $O(n^3)$ between any two 4-colorings of a 3-colorable planar graph on n vertices. As a result, we strengthen Corollary 2.1.18 as follows:

Theorem 2.2.3 [8]

Let G be an n -vertex planar graph, then $\mathcal{K}_k(G)$ has diameter polynomial in the number of vertices of G , for all $k > \chi(G)$.

Table 2.1 sums up the know results on the diameter of the Kempe reconfiguration graph.

While several upper bounds on the diameter of $\mathcal{K}_k(G)$ are known, proving lower bounds is significantly harder and no superlinear lower bound on the Kempe distance is known as of today.

Question 2.2.4

Is there an infinite family $(G_i, k_i, \alpha_i, \beta_i)$ such that the distance between α_i and β_i in $\mathcal{K}_{k_i}(G)$ is superlinear in $|G_i|$?

In the context of single vertex recoloring, Bonsma and Cereceda [BC09] showed that the problem of deciding whether two colorings are equivalent up to a sequence of single vertex recolorings is PSPACE-complete for $k \geq 4$. This implies, assuming that PSPACE \neq NP, that there exist pairs of k -colorings of at superpolynomial

Graph class	Number of colors	Diameter of $\mathcal{K}_k(G)$	Reference
Bipartite	$k \geq \chi$	$O(n)$	
d -degenerate	$k \geq d + 2$	$O(n^{d+1})$	[BH22]
Planar	$k \geq \chi + 1$	$O(n^{195})$ for $\chi = 4$ $O(n^3)$ for $\chi = 3$	[8] [Moh07]
Cographs	$k \geq \chi$	$O(n \log n)$	[BHI+20]
Chordal graphs	$k \geq \chi$	n	[BHI+20]
Bounded degree	$k \geq \Delta$	$O_\Delta(n^2)$ unless $G = K_2 \square K_3$	[6]
Bounded mad	$k \geq \text{mad} + \varepsilon$	$O(n^{c_{k,\varepsilon}})$	[6]
Bounded treewidth	$k \geq \text{tw}$	$O(\text{tw} \cdot n^2)$	[6]
Odd-hole-free	$k > \lfloor \frac{\Delta(G)+1+\chi(G)}{2} \rfloor$	$O_k(n)$	Subsection 2.1.3

Table 2.1: Known results on the diameter of the Kempe reconfiguration graph

single vertex recoloring distance. In the same paper, they give an explicit construction of such graphs and colorings, thereby avoiding the assumption $\text{PSPACE} \neq \text{NP}$ (see Section 1.4 for the definition of these complexity classes): for every $k \geq 4$, there is a family (G_i, α_i, β_i) such that the minimum number of single vertex recolorings between α_i and β_i is $\Omega(2^{\sqrt{n_i}})$ where $n_i = |V(G_i)|$.

In the context of Kempe changes, Bonamy, Heinrich, Ito, Kobayashi, Mizuta, Mühlenthaler, Suzuki and Wasa proved in [BHI+20] that deciding whether two colorings are Kempe equivalent is also PSPACE -complete, as soon as $k \geq 3$. This also proves, assuming that $\text{PSPACE} \neq \text{NP}$, that there exist pairs of 3-colorings of some planar graphs at superpolynomial Kempe distance (see Section 2.3 for a discussion on the matter). However, this does not give an explicit construction of colorings at even superlinear Kempe distance and the construction of [BC09] for single vertex recoloring does not seem to adapt directly to Kempe changes. A better understanding of the Kempe distance and an explicit construction could be useful in path coupling methods as we will see in Section 2.4, which further motivates Question 2.2.5.

As planar graphs are not necessarily 3-recolorable, it is possible that when the number of colors is high enough for $\mathcal{K}_k(G)$ to be connected, its diameter becomes polynomial, which also raises the following question:

Question 2.2.5

Is there an infinite family (G_i, k_i) such that for all i , $\mathcal{K}_{k_i}(G)$ is connected but has a superpolynomial diameter in $|G_i|$?

2.3 Counting and complexity

Determining if a graph is k -colorable, for $k \geq 3$, is a notoriously NP-complete problem [Kar72]. Several (decision) reconfiguration problems can be considered:

Problem k -KEMPE EQUIVALENT

INPUT A graph G and two k -coloring α and β of G .

QUESTION Are α and β Kempe equivalent?

Problem k -KEMPE DISTANCE

INPUT A graph G , two k -coloring α and β of G and an integer d .

QUESTION Are α and β at distance at most d in $\mathcal{K}_k(G)$?

Problem k -KEMPE CONNECTIVITY

INPUT A graph G .

QUESTION Is G k -recolorable?

Problem k -KEMPE DIAMETER

INPUT A graph G and an integer d .

QUESTION Does $\mathcal{K}_k(G)$ have diameter at most d ?

The single-vertex recoloring version of k -KEMPE EQUIVALENT and k -KEMPE CONNECTIVITY have been extensively studied: k -SINGLE VERTEX RECOLORING EQUIVALENT is polynomially solvable for $k \leq 3$ and PSPACE-complete for $k \geq 4$, while k -SINGLE VERTEX RECOLORING CONNECTIVITY is polynomially solvable for $k = 2$ and co-NP-complete for $k = 3$. See the 2013 survey [vdH13] of van den Heuvel and the references therein for a comprehensive overview.

For Kempe changes, Bonamy, Heinrich, Ito, Kobayashi, Mizuta, Mühlenthaler, Suzuki and Wasa proved that k -KEMPE EQUIVALENT is PSPACE-complete for any $k \geq 3$, and remains PSPACE-complete even for $k = 3$ and restricted to planar graphs of maximum degree 6 [BHI⁺20]. This implies that k -KEMPE DISTANCE is also PSPACE-hard. However, while k -KEMPE EQUIVALENT and k -KEMPE CONNECTIVITY become polynomially solvable on k -recolorable graphs classes, it is possible that k -KEMPE EQUIVALENT and k -KEMPE DIAMETER remain hard even then.

As mentioned in [BHI⁺20], the PSPACE-hardness of k -KEMPE DISTANCE implies that there exist colorings at super polynomial Kempe-distance, under the assumption $\text{PSPACE} \neq \text{P}$.

Bubley, Dyer, Greenhill and Jerrum proved that the exact counting of the number of k -colorings of a graph is a #P-hard problem [BDGJ99]. Jerrum, Valiant and Vazirani showed that approximately counting the number of k -colorings of a graph up to an ε -factor is in difficulty equivalent to the approximate random uniform generation of a k -coloring [JVV86]. More precisely, a *randomized approximation*

scheme for k -coloring is an algorithm that given a graph G and $\varepsilon > 0$ outputs a number X such that

$$\mathbb{P}[(1 - \varepsilon)|\Omega_k(G)| \leq X \leq (1 + \varepsilon)|\Omega_k(G)|] \geq \frac{3}{4}$$

where the constant $\frac{3}{4}$ is arbitrary and can be replaced by any constant strictly between $1/2$ and 1 . Such an algorithm is said to be *fully polynomial* (FPRAS) if it runs in time polynomial in $|V(G)|$ and ε^{-1} . An *almost uniform generator* for k -colorings is an algorithm M that given a tolerance $\varepsilon > 0$ and a graph G , outputs a random k -coloring of G such that there exists $c > 0$ for which,

$$(1 - \varepsilon)c \leq \mathbb{P}[M \text{ outputs } \alpha] \leq (1 + \varepsilon)c$$

for all coloring α of G . An almost uniform generator is fully polynomial if it runs in time polynomial in $|V(G)|$ and $\log(\varepsilon^{-1})$.

The existence of a fully polynomial randomized approximation scheme for k -colorings on a class of graphs \mathcal{C} is equivalent to the existence of a fully-polynomial almost uniform generator for k -colorings of graphs in \mathcal{C} , see [JVV86] for a detailed study of the relations between (approximate) counting and generation problems. This motivates the study of approximate uniform generation (see Section 2.4).

Another facet of this counting problem is the extremal problem consisting in finding the n -vertex graphs that maximize or minimize the number of k -colorings for a given k . This question is of small interest without any further restriction, as an independent set of size n has k^n colorings. However, when restricted to r -regular graphs, the answer is far from obvious and presents interesting connections with the Potts model (see Section 2.4). Disjoint copies of the bipartite graph $K_{d,d}$ were conjectured to maximize the number of k -colorings, which was confirmed by Sah, Sawhney, Stoner and Zhao [SSSZ20] after numerous progressive improvements on this question.

2.4 Wang-Swendsen-Kotecký dynamics

We now outline the connections between the so-called *Potts model* and the geometry of the Kempe reconfiguration graph. We will first briefly present the Potts Model and how random walks on the Kempe reconfiguration graphs can be used to sample colorings.

2.4.1 Introduction to the Potts model

The Potts model is a model of statistical mechanics that represents the interactions between the spins of atoms in a crystal structure. The atoms are represented by

the vertices of a graph, while the interactions between them are modeled by edges. These “crystal” graphs are often highly regular (such as grids or infinite d -valent trees for example). Each atom can take a range of states $1, \dots, q$ called *spins* that graph theorists call colors. When $q = 2$, we call it the *Ising model*. Depending on the matter, the nature of the interactions varies and the edges tend to take identical spins (ferromagnetic case) or different spins (antiferromagnetic case). The Potts model gives a probability distribution on the (non necessarily proper) colorings of G .

Given a (non necessarily proper) coloring $\sigma \in [q]^{V(G)}$, the interaction energy of an edge uv is given by a coefficient J_{ij} , where $\sigma(u) = i$ and $\sigma(v) = j$. For the Ising model, the coefficient $J_{i,j}$ usually equals one if $i = j = 1$ and zero otherwise. Here we will consider the simpler case where the interaction energy at each edge can only take values: J if the endpoints of e are colored identically and 0 otherwise, where J is a parameter called *coupling constant*¹. If $J > 0$ the model is in the *ferromagnetic* regime that favors monochromatic edges and large monochrome clusters, while $J < 0$ corresponds to the *antiferromagnetic* regime that favors bichromatic edges and proper colorings. Let M_σ be the number of monochromatic edges in σ . The energy of a configuration σ is given by the Hamiltonian $H_J(\sigma) = -JM_\sigma$.

The *partition function* is then defined as

$$Z_G^{\text{Potts}}(q, J\beta) = \sum_{\sigma: V \rightarrow [q]} e^{-\beta H_J(\sigma)}$$

where β is a parameter called the *inverse temperature* and is equal to $1/kT$ where k is Boltzmann’s constant and T the temperature of the system. The probability for the system to be in the configuration σ is given by the so-called *Gibbs measure*

$$\pi(\sigma) = \frac{e^{-\beta H_J(\sigma)}}{Z_G^{\text{Potts}}(q, J\beta)}.$$

The Gibbs measure may drastically change at thresholds called the *phase transitions*. Understanding the Gibbs measure and the phase transitions, that is what a random configuration looks like and the different regimes, is one of the main objective of statistical physicists and probabilists.

Surprisingly, the partition function of the Potts model is a polynomial in q :

Theorem 2.4.1 Fortuin–Kasteleyn representation of the Potts Model

For every $q \geq 1$,

$$Z_G^{\text{Potts}}(q, J\beta) = Z_G(q, v) = \sum_{A \subseteq E(G)} q^{k(A)} v^{|A|}$$

where $k(A)$ is the number of connected component in $(V(G), A)$ and $v = e^{J\beta} - 1$.

¹Despite its misleading name, J has nothing to do with the couplings from probability theory.

Note that $\sum_{A \subseteq E(G)} q^{k(A)} v^{|A|}$ is an equivalent definition of the famous Tutte polynomial, after a change of variable.

In the following, we will focus on the antiferromagnetic case, taking $J < 0$. Note that the proper configurations have probability $\frac{1}{Z_G^{\text{Potts}}(q, J\beta)}$ while the non proper ones have probability $\frac{e^{J\beta M_\sigma}}{Z_G^{\text{Potts}}(q, J\beta)}$. So at zero temperature, that is $\beta \rightarrow \infty$, the constraints become *hard*: adjacent vertices cannot have the same color, whereas they are *soft* when $\beta < \infty$ (monochromatic edges have a finite cost). In this regime called the *ground state* of the Potts model, only the proper colorings have non-zero probability. Therefore, the number of proper q -colorings of G is given by $\lim_{\beta \rightarrow \infty} Z_G^{\text{Potts}}(q, J\beta) = P_G(q)$ which is the *chromatic polynomial*.

On the other hand, when $\beta \rightarrow 0$, all configurations have the same weight and $Z_G^{\text{Potts}}(q, 0) = q^n$.

2.4.2 Sampling with Markov Chain Monte Carlo algorithms

Markov Chain Monte Carlo algorithms are a common tool used to sample from complex probability distributions such as the Gibbs measure (see Section 1.3 for all definitions). They consist in performing a Markovian random walk in the configuration space, starting from an arbitrary configuration. If the corresponding Markov chain is irreducible and aperiodic then it converges towards its stationary distribution (see Theorem 7 and more generally Section 1.3 for all prerequisites on Markov chains). By design, this stationary distribution will be the Gibbs measure in the examples we will consider. Note that if the reconfiguration operations can be reversed, then the reconfiguration graph undirected, and the Markov Chain is symmetric. By Theorem 6, this implies that its stationary distribution is the uniform one. Note also that the mixing time of a Markovian random walk is strongly connected to the expansion of its ground graph: if the ground graph has a Cheeger constant exponentially small in n^ε for some $\varepsilon > 0$, where n is the size of the ground graph, then the Markov chain is torpidly mixing [LV05], while expander families have the fastest mixing time among graphs of bounded degree [LP17]. Thus, a structural interpretation of how fast some Markov Chain Monte Carlo algorithm mixes can be found in the expansion of the corresponding reconfiguration graph. As explained in Section 2.3, on top of the motivations from statistical physics, designing efficient approximate uniform generators for q -colorings provides approximate bounds on the number of q -colorings of a graph.

To sample from the antiferromagnetic Potts model, two Markov chains have received a lot of attention: the Glauber dynamics and the Wang-Swendsen-Kotecký (WSK) algorithm. At each iteration of the Glauber dynamics, a vertex v is chosen at random and the configuration σ is updated by sampling among the configurations whose restriction to $G \setminus v$ agrees with σ .

An iteration of the original WSK algorithm consists in choosing uniformly at random two colors $a, b \in [q]$ and fixing the colors in $G \setminus G_{a,b}$, where $G_{a,b}$ is the subgraph induced by colors a and b . The WSK algorithm then samples a new 2-coloring on $G_{a,b}$. A variant of this chain, first considered by Wolff in [Wol89], consists in choosing a vertex v and a color a at random, and with probability $1/2$ swap the colors a and $\sigma(v)$ in the connected component colored a and $\sigma(v)$ containing v . This alternative definition is equivalent in terms of ergodicity and mixing times up to a constant factor. We refer to it as the WSK algorithm as it is the most common version of the WSK algorithm in the graph theory literature.

At zero temperature, only proper colorings are considered, thus the Glauber dynamics performs random single vertex recolorings and the WSK algorithm performs random Kempe changes. In both cases, an iteration consist in picking uniformly at random a vertex v and a color c . Then the Glauber dynamics recolors v with c with probability $1/2$, if $c \notin \sigma(N(v))$; while the WSK dynamics performs the Kempe change $K_{v,\sigma(c)}$ with probability $1/2$. Both dynamics are symmetric, thus when the corresponding reconfiguration graphs is connected they converge towards the uniform distribution.

2.4.3 Ergodicity and mixing time of the Glauber dynamics and the WSK algorithm

The ergodicity of the Glauber dynamics and the WSK dynamics are immediate at positive temperature, however at zero temperature, their ergodicity is equivalent to the connectedness of respectively the single vertex recoloring reconfiguration graph and the Kempe reconfiguration graph. From now on, we will restrict ourselves to the study at zero temperature. Denote Λ_q be the set of all the proper q -colorings of a graph G .

The WSK algorithm introduces non-local updates and generalizes the Glauber dynamics. Hence, the ergodicity or the rapid mixing of the Glauber dynamics implies that of the WSK algorithm [DS93]. However, the converse is false: a simple obstacle to the (good) mixing of the Glauber dynamics is the existence of *locally frozen* colorings: colorings in which all vertices see all the colors in their neighborhood. The non-locality of the WSK algorithm allows more flexibility by avoiding the pitfall of locally frozen colorings. For example, on a star with n vertices and $n^{1-\varepsilon}$ colors, most of the colorings are locally frozen which makes recolorings of the central vertex highly unlikely. As a result, the Glauber dynamics is torpidly mixing, while the WSK algorithm mixes rapidly [LV05]. However, this non-locality also results in a much harder analysis of both the ergodicity and the mixing time.

We know that the Glauber dynamics is ergodic on any graph G provided that

the number k of colors is at least $\Delta(G) + 2$. Salas and Sokal [SS97] and Jerrum [Jer95], independently proved that the Glauber dynamics is rapidly mixing when the number of colors is larger than twice the maximum degree Δ of the graph. Jerrum's proof became a textbook example of a coupling argument and extends to the case $k = 2\Delta$, while Salas and Sokal's proof uses Dobrushin uniqueness theorem and extends to the Potts model at non-zero temperature. This raises the following question:

Question 2.4.2

What is the smallest c such that the WSK algorithm mixes rapidly on any graph of maximum degree Δ , when the number of colors is at least $c\Delta$?

Vigoda broke the barrier of 2Δ by proving that the Glauber dynamics mixes in $O(n^2 \log n)$ when $q \geq \frac{11\Delta}{6}$ [Vig00]. In fact, Vigoda proves that the *flip Dynamics* mixes rapidly in this setting by using a path coupling argument and derives from it the rapid mixing of Glauber dynamics. The updates of the flip dynamics are bounded size Kempe changes, placing it in between the Glauber dynamics and the WSK algorithm. Recently, Delcourt, Perarnau and Postle [DPP18] improved this result by showing that these dynamics mix rapidly for $k \geq (\frac{11}{6} - \eta)\Delta$ for some small constant $\eta > 0$, even in the more general setting of list-coloring. As of today, this bound has not been improved and it remains open whether the Glauber dynamics or the WSK dynamics mix rapidly for $k \geq \Delta + 2$.

Question 2.4.2 can also be studied on restricted graphs classes. In this direction, Łuczak and Vigoda [LV05] proved that the number of colors at which the WSK algorithm mixes rapidly depends on Δ even in strongly structured graph classes such as bipartite graphs and planar graphs, despite their low chromatic number and the polynomial diameter of their reconfiguration graph with $O(1)$ colors. Their argument mimics the obstruction of locally frozen colorings for single vertex coloring: let G_n be the a 5-cycle in which two non adjacent vertices are blown-up into equal size independent set A and B (see Figure 2.18). They prove that the WSK algorithm is torpidly mixing on G_n with $n^{1-\varepsilon}$ colors for any fixed $\varepsilon > 0$. We give the proof of Łuczak and Vigoda [LV05] of this result, as it is a short and classical method to prove that a Markov Chain Monte Carlo algorithm mixes torpidly.

Proof [LV05]. Denote w the vertex of maximal degree and u and v the two other vertices of the original 5-cycle. Let S be all the k -colorings of G_n that color u and w identically. By symmetry, $|S| \leq \frac{\binom{\Delta_k}{2}}{2}$. To exit S , one must perform a Kempe change on v or w that does not affect the vertices in A , hence A must use at most $k - 2$ colors among the $k - 1$ available. Hence the Cheeger constant of the




Figure 2.18: Planar graph on which the WSK mixes torpidly for $k \leq n^{1-\varepsilon}$

reconfiguration graph is at most

$$\frac{\partial S}{S} \leq \frac{(k-1)(k-2)^n}{(k-1)^n} \leq \exp\left(-\frac{n-1}{k-1}\right)$$

For $k \leq n^{1-\varepsilon}$ for some fixed ε , this gives $h(\mathcal{K}_k(G)) \leq \exp(-n^\varepsilon)$ and hence $t_{\text{mix}} = \Omega(\exp(n^\varepsilon))$. \square

Using a similar proof, Łuczak and Vigoda [LV05] prove the torpid mixing of a family of bipartite graphs for $k = \Omega(\Delta/\log(\Delta))$. These results highlight the large gap in which the reconfiguration graph can have polynomial diameter but bad expansion.

2.4.3.1 Upper bounding the mixing time using couplings

Using Eq. (1.3), a natural method to bound the mixing time of a Markov chain is to design efficient couplings. That is, a coupling (X_1, Y_1) between each pair of states, along with a metric $\rho : \Lambda_k \times \Lambda_k \rightarrow \mathbb{R}_+$, such that ρ is in average contracted by the coupling:

$$\forall x, y \in \Lambda_k, \quad \mathbb{E}_{x,y}(\rho(X_1, Y_1)) \leq e^{-\varepsilon} \rho(x, y) \quad (2.1)$$

By iterating Eq. (2.1), we obtain after t steps $\mathbb{E}_{x,y}(\rho(X_t, Y_t)) \leq e^{-\varepsilon t} \rho(x, y)$, which gives an immediate upper bound on the coalescence time and hence on the mixing time:

$$t_{\text{mix}} \leq \frac{1}{\varepsilon} \log(\text{diam}_\rho(\Lambda_k))$$

where $\text{diam}_\rho(\Lambda_k)$ is defined as the maximum over all $x, y \in \Lambda_k$ of $\rho(x, y)$.

In practice, designing couplings on all pairs of states is no easy task: it is often unclear how to modify two arbitrary colorings to make them closer with respect to ρ . When ρ is the shortest path distance on the reconfiguration graph, then the *path coupling* technique eases the design of couplings, by extending a coupling defined only on pairs of adjacent states to a coupling between all pairs of states. This technique originally developed by Bubley and Dyer [BD97], detailed in

Section 14.2 of [LP17] or in the comprehensive survey of Frieze and Vigoda [FV07], is at the core of Jerrum’s coupling [Jer95] and Vigoda’s breakthrough [Vig00].

Another regularly used trick is to use the path coupling technique with the Hamming distance $\rho(\alpha, \beta) = \sum_{v \in V(G)} \mathbb{1}_{\alpha(v) \neq \beta(v)}$. As the path coupling method requires the distance to be a shortest path metric, this approach requires to work on $[k]^V$ instead of Λ_k by also considering non-proper colorings. Since the Glauber dynamics does not create non-proper colorings from proper ones, proving rapid mixing of this Markov chain on the enlarged space results in rapid mixing of the standard Glauber dynamics on proper colorings. There are several advantages to the use of the Hamming distance: the Hamming distance between any two colorings is at most n , it is usually easier to analyze than the recoloring distance and it is close to the single vertex recoloring distance for nearly identical colorings, as the Hamming distance increases by at most one when recoloring a single vertex, making Eq. (2.1) easy to check on pairs of colorings that differ at a single vertex. However, the Hamming distance is of little usefulness to analyze the mixing of the WSK algorithm, as it can increase tremendously when performing one Kempe change. We hope that a better understanding of the Kempe distance transfers to the analysis of path coupling for the WSK algorithm.

2.5 Conclusion and perspectives

2.5.1 Non-recolorability without frozen colorings

As of today, only two arguments are known to prove that a graph admits several Kempe equivalence classes. As the topological argument of Fisk, Mohar and Salas is specific to 4-colorings of graphs embedded on surfaces, we conjecture that frozen colorings are the only general obstacle to recolorability:

Question 2.1.25

Is there another obstacle to k -recolorability, beside frozen colorings and the topological argument of Fisk, Mohar and Salas?

To support this conjecture or to develop new Kempe invariants, we advertise the study of the recolorability of three natural graphs classes that admit no frozen colorings: colorings of K_t -minor-free graphs, 3-colorings of triangle-free planar graphs and k -colorings of odd-hole-free graphs for $k > \left\lceil \frac{\Delta(G)+1+\chi(G)}{2} \right\rceil$.

2.5.1.1 Recoloring version of the Linear Hadwiger’s conjecture

K_t -minor-free graphs can admit frozen $(2/3 - \varepsilon)t$ -colorings, hence are not t -recolorable (see [2]), but have no frozen k -colorings for $k \geq 2t$, hence we conjecture the following:

Conjecture 2.1.13 [2]

There exists a constant c such that any K_t -minor-free graph is ct -recolorable.

Progress on this recoloring version of the Linear Hadwiger's conjecture would require to go below the degeneracy threshold of $O(t \log t)$, which is the best bound for Conjecture 2.1.13 as of today.

2.5.1.2 Recoloring version of Grötzsch theorem

Grötzsch's theorem states that triangle-free planar graphs are 3-colorable. A counting argument shows that triangle-free planar graphs cannot have frozen 3-colorings, so we conjecture the following:

Conjecture 2.1.28 (Bonamy, Legrand-Duchesne and [SS22])

All triangle-free planar graphs are 3-recolorable.

We prove that this conjecture is equivalent to the following:

Conjecture 2.1.29

All planar graphs of girth 5 are 3-recolorable.

Note that Conjecture 2.1.29 implies that in a triangle-free planar graph G , any Kempe recoloring sequence of a 5-face F can be extended to recoloring sequence of G (see Observation 2.1.33): for any coloring α of G and any coloring β' of F that differs from $\alpha|_F$ by one Kempe change, there is a recoloring sequence that leads to a 3-coloring β of G with $\beta|_G = \beta'$, such that F is not affected by the intermediate changes. This property appears to be a key ingredient in showing that a minimal counterexample to Conjecture 2.1.29 contains no separating 5-cycle and would constitute a promising first step towards proving Conjectures 2.1.28 and 2.1.29.

2.5.1.3 Recoloring version of Reed's conjecture

In an independent attempt to progress on Question 2.1.25, we study with Marthe Bonamy and Tomáš Kaiser a recoloring version of Reed's coloring conjecture on χ , ω and Δ :

Question 2.1.40

What is the largest $\varepsilon \leq 1/2$ such that all graphs G are k -recolorable for all $k \geq [(1 - \varepsilon)(\Delta(G) + 1) + \varepsilon\omega(G)]$?

We note that for general graphs, a construction in [7] shows that ε is at most $1/3$. We show that odd-hole-free graphs satisfy Reed's coloring conjecture and that they admit no frozen k -colorings for $k > \left\lceil \frac{\Delta(G)+1+\omega(G)}{2} \right\rceil$, which motivates the study of Question 2.1.40 in the restricted class of odd-hole free graphs.

We prove that for odd-hole-free graphs, ε is at least $1/4$ and that all the Kempe equivalence classes contain a $\lceil \frac{\Delta(G)+1+\omega(G)}{2} \rceil$ -coloring. This motivates the following conjecture, that we prove in the special case of perfect graphs (using one extra color):

Conjecture 2.1.42 (Bonamy, Kaiser and Legrand-Duchesne)

All odd-hole free graphs of maximum degree Δ and clique number ω are k -recolorable for $k \geq \lceil \frac{\Delta+1+\omega}{2} \rceil$.

2.5.2 Bounding the Kempe distance

The second direction of work we promote is the study of the structure of the Kempe reconfiguration graph, more precisely its diameter and its expansion. Regarding upper bounds, we recall the conjecture of Bonamy, Bousquet, Feghali and Johnson on the polynomial Kempe recoloring of degenerate graphs:

Conjecture 2.2.1 [BBFJ19]

Let G be an n -vertex d -degenerate graph. Any two k -colorings of G are Kempe equivalent up to $O(n^2)$ Kempe changes for $k \geq d + 1$.

This conjecture, analog to the central conjecture of Cereceda's in the context of single vertex recoloring, remains widely open as we do not even know a polynomial bound in this setting. Although polynomial bounds can be obtained using one extra color or in various related graphs classes, these methods do not seem to apply here.

The Kempe distance is very poorly understood. In particular, we do not know of any non-trivial lower bounds:

Question 2.2.4

Is there an infinite family $(G_i, k_i, \alpha_i, \beta_i)$ such that the distance between α_i and β_i in $\mathcal{K}_{k_i}(G)$ is superlinear in $|G_i|$?

Although there are 3-colorings at superpolynomial Kempe distance in planar graphs under the complexity assumption that PSPACE \neq NP, no explicit construction is known for Question 2.2.4 despite its apparent simplicity. Designing superlinear lower bounds on the Kempe distance could translate into improvements in sampling methods via the Wang-Swendsen-Kotecký algorithm and the use of path coupling methods.

Regardless of the answer to Question 2.2.4, it is possible that when the number of colors becomes sufficiently high for the graph to be k -recolorable, all k -colorings become at polynomial distance in the Kempe reconfiguration graph. This raises the following extremal question:

Question 2.2.5

Is there an infinite family (G_i, k_i) such that for all i , $\mathcal{K}_{k_i}(G)$ is connected but has a superpolynomial diameter in $|G_i|$?

A reasonable first step towards answering Questions 2.2.4 and 2.2.5 could be to study the geometry of the Kempe reconfiguration graphs around its connectivity threshold: For what values of p and k_p does the random Erdős-Rényi graph $G(n, p)$ become k -recolorable with high probability for all $k \geq k_p$? For slightly smaller values of k , does the Kempe reconfiguration graph of $G(n, p)$ contain a giant component? If so, do the smaller components correspond to frozen colorings? In graph classes that we know to be k -recolorable but not $(k - 1)$ -recolorable, do we observe similar phenomena when the number of colors is only $k - 1$? Answers to these questions could also shed light on Question 2.1.25.

Chapter 3

Kempe changes in sparse graphs

This chapter focuses on Kempe recoloring in sparse graph classes. It contains published results obtained with Quentin Deschamps, Carl Feghali, František Kardoš and Théo Pierron [8], with Marthe Bonamy and Vincent Delecroix [6] and with Marthe Bonamy, Marc Heinrich and Jonathan Narboni [7].

Introduction

Alongside the early results of Fisk on the 4-recolorability of graphs embedded on surfaces, the groundwork of Las Vergnas and Meyniel on the recolorability of degenerate graphs and planar graphs initiated the study of Kempe equivalence. We recall these results:

Theorem 2.1.15 [VM81]

All d -degenerate graphs are k -recolorable, for all $k > d$.

Applied to planar graphs, Theorem 2.1.15 proves that planar graphs are 6-recolorable. In [Mey75], Meyniel proved that this is true even for $k = 5$:

Theorem 2.1.10 [Mey75]

All planar graphs are 5-recolorable. This is tight as there exist planar graphs that admit non-similar frozen 4-colorings (see Figure 2.12).

By Wagner's theorem, planar graphs are exactly the graphs that do not contain $K_{3,3}$ or K_5 as a minor. Las Vergnas and Meyniel strengthened Theorem 2.1.10 by removing the assumption of being $K_{3,3}$ -minor-free:

Theorem 2.1.11 [VM81]

All K_5 -minor-free graphs are 5-recolorable.

These results raise two natural questions: what is the diameter of the reconfiguration graph in these settings and can Theorem 2.1.11 be generalized to the k -recolorability of K_k -minor-free graphs for higher k ?

The original proof of Theorem 2.1.15 in [VM81] implies that $\mathcal{K}_k(G)$ has diameter $O(2^{|V(G)|})$. Similarly, the proofs of Theorems 2.1.10 and 2.1.11 in [Mey75, VM81] imply that $\mathcal{K}_5(G)$ has diameter $O(5^{|V(G)|})$. Bonamy, Bousquet, Feghali and Johnson conjectured that the former can be significantly improved.

Conjecture 2.2.1 [BBFJ19]

Let G be an n -vertex d -degenerate graph. Any two k -colorings of G are Kempe equivalent up to $O(n^2)$ Kempe changes for $k \geq d + 1$.

For $k > d + 1$ in Conjecture 2.2.1, a result of Bousquet and Heinrich [BH22] regarding the reconfiguration graph $\mathcal{C}_k(G)$ for the single vertex recoloring operation on colorings of graphs with bounded degeneracy implies the following:

Theorem 3.0.1

For all d -degenerate n -vertex graph G and all $k > d + 1$, the single vertex recoloring reconfiguration graph $\mathcal{C}_k(G)$ has diameter $O(n^{d+1})$, thus so does $\mathcal{K}_k(G)$.

The case $k = d + 1$ in Conjecture 2.2.1 seems much more challenging. The object of Section 3.1 is to break the $k > d + 1$ (in fact, $k = d + 1$) barrier for various subclasses of degenerate graphs.

With Quentin Deschamps, Carl Feghali, František Kardoš and Théo Pierron we addressed the case of planar graphs (see Subsection 3.1.1).

Theorem 2.2.3 [8]

Let G be an n -vertex planar graph, then $\mathcal{K}_k(G)$ has diameter polynomial in the number of vertices of G , for all $k > \chi(G)$.

With Marthe Bonamy and Vincent Delecroix, we addressed the case of graphs of bounded treewidth, graphs with bounded maximum average degree and $(\Delta - 1)$ -degenerate graphs, where Δ denotes the maximum degree.

A common feature of planar graphs and graphs of bounded maximum average degree is that they admit a degeneracy ordering by “layers of linear size”. In each setting, Euler’s formula or the maximum average degree bound, combined with a counting argument, shows that one can find an independent set S of linear size containing only vertices of low degree. The idea is then to recolor all the vertices in S , using similar techniques as the ones used by Las Vergnas and Meyniel in [VM81] and apply induction. This step is repeated a logarithmic number of times, because the independent set S has linear size, hence, even if each step doubles the length of the reconfiguration sequence, the final sequence has polynomial size. This argument cannot be used in degenerate graphs, as they can have a unique degeneracy ordering, and no such set of linear size.

In the case of bounded degree graphs and graphs of bounded treewidth, the polynomial bounds we obtain rely on different techniques. For graphs of maximum degree Δ , it suffices to obtain a polynomial bound in Theorem 2.1.15 under an additional assumption on the degree of the vertices to immediately obtain a polynomial bound of $O(\Delta^3 n^2)$ in Corollary 2.1.22, but a more careful analysis of the proof of Corollary 2.1.22 can bring it down to $O(\Delta n^2)$. For graphs of bounded treewidth, we rely on the tree structure and reduce the problem to the recolorability of chordal graphs, which are polynomially recolorable [BHI+20].

Kempe changes in planar graphs

The following theorem is due to Mohar:

Theorem 2.1.17 [Moh07]

All 3-colorable planar graphs are 4-recolorable.

By making a few simple observations, we show that the proof of Theorem 2.1.17 recolors each vertex at most $O(n^2)$ times. For planar graphs of chromatic number four, we explain how to reduce to the recoloring of a planar graph of size $(1 - \varepsilon)n$, by using the recolorability of 3-colorable planar graphs as a black box along with a sparsity argument. By applying this reduction step a logarithmic number of times, we are able to recolor any planar graph using a polynomial number of Kempe changes:

Theorem 2.2.3 [8]

Let G be an n -vertex planar graph, then $\mathcal{K}_k(G)$ has diameter polynomial in the number of vertices of G , for all $k > \chi(G)$.

Kempe equivalence of Δ -colorings

Any graph G can be greedily colored with $(\Delta + 1)$ colors, where Δ is the maximum degree of G . Brooks' theorem states that if G is not a clique or an odd-cycle, then Δ colors suffice. Many different proofs of this theorem exist, see [CR15] for a collection of proofs of Brooks' theorem using various techniques — for that matter, some of them using Kempe changes.

Mohar [Moh07] conjectured that all the k -colorings of a graph are Kempe equivalent for $k \geq \Delta$. Note that the result of Las Vergnas and Meyniel [VM81] settles the case of non-regular graphs. Feghali, Johnson and Paulusma [FJP17] proved that the conjecture holds for all cubic graphs but the 3-prism (see Fig. 2.7). Indeed, the 3-prism $K_2 \square K_3$ admits non-similar frozen 3-colorings and hence is not 3-recolorable.

Bonamy, Bousquet, Feghali and Johnson [BBFJ19] later showed that the conjecture also holds for Δ -regular graphs with $\Delta \geq 4$. Both papers heavily rely

on the lemma of Las Vergnas and Meyniel [VM81] and provide sequences that possibly have exponential length.

In Subsection 3.1.3, we give a polynomial upper bound on the diameter of the reconfiguration graph in this setting.

Theorem 2.2.2 [6]

Let G be an n -vertex graph of maximum degree Δ . For $k \geq \Delta$, any two k -colorings of G are equivalent up to at most $O_\Delta(n^2)$ Kempe changes, unless $k = 3$ and G is the 3-prism.

The main idea of the proof of Theorem 2.2.2 is to improve the result of Las Vergnas and Meyniel [VM81]. We show that there exists a sequence of Kempe changes of length $O(dn^2)$ between any two k -colorings of a d -degenerate graph when $k \geq d + 1$ under the additional assumption that all the vertices but one have degree at most $d + 1$ (see Subsection 3.1.2).

Kempe equivalence in graphs of bounded maximum average degree

For all $d \geq 1$, if a graph has mad less than d , then it is $(d - 1)$ -degenerate: all its subgraphs have average degree less than d , so admit a vertex of degree at most $d - 1$. In Subsection 3.1.4, we show that the mad of a graph, while related to its degeneracy, is easier to work with in this setting.

Theorem 3.0.2 [6]

Let k be a positive integer and $\varepsilon > 0$. There exists a constant $c_\varepsilon > 0$ such that for any n -vertex graph G with $\text{mad}(G) \leq k - \varepsilon$, any two k -colorings of G are Kempe equivalent up to $O(n^{c_\varepsilon})$ Kempe changes.

We prove the above theorem by giving an upper bound on the number of Kempe changes needed to reconfigure list-colorings (see Subsection 3.1.2) and adapting ideas developed by Bousquet and Perarnau in [BP16] in the setting of single vertex recoloring (see Subsection 3.1.4).

Kempe equivalence in graphs of bounded treewidth

Another way to strengthen the degeneracy assumption involves the treewidth of a graph. A graph of treewidth k is k -degenerate, while there are 2-degenerate graphs with arbitrarily large treewidth. Bonamy and Bousquet [BB18] confirmed Cereceda’s conjecture for graphs of treewidth tw , by proving that all their $(\text{tw} + 2)$ -colorings are equivalent up to $O(n^2)$ vertex recolorings. In Subsection 3.1.5, we extend this result to non-trivial Kempe changes with one fewer color.

Theorem 3.0.3 [6]

Let tw be a positive integer and G be a n -vertex graph of treewidth tw . For $k \geq tw + 1$, any two k -colorings of G are equivalent up to $O(tw n^2)$ Kempe changes.

Additionally, the proof of Theorem 3.0.3 is constructive and yields an algorithm to compute such a sequence in time $f(k) \cdot n^c$ for some fixed $c > 0$. Given a witness that the graph has treewidth at most k , the complexity drops to $k \cdot n^c$.

Recoloring version of Hadwiger's conjecture

We now discuss the second question of interest of this chapter: can Theorem 2.1.11 be generalized to higher k ? In [VM81], Las Vergnas and Meyniel conjectured the following, which can be seen as a reconfiguration counterpoint to Hadwiger's conjecture, though it neither implies it nor is implied by it.

Conjecture 2.1.12 (Conjecture A in [VM81])

For every t , every K_t -minor-free graph is t -recolorable.

Note that their proof of Theorem 2.1.11 settles the case $k = 5$ and that the conjecture is also true for smaller values of k : K_3 -minor-free graphs are the forest which are k -recolorable for $k \geq 2$ and K_4 -minor-free graphs are exactly the graphs of treewidth at most 2 and thus are k -recolorable for $k \geq 3$. Las Vergnas and Meyniel also proposed a related conjecture that is weaker assuming Hadwiger's conjecture holds.

Conjecture 3.0.4 (Conjecture A' in [VM81])

For every t and every graph with no K_t -minor, every equivalence class of t -colorings contains some $(t - 1)$ -coloring.

With Marthe Bonamy, Marc Heinrich and Jonathan Narboni, we present in Section 3.2 a random construction that disproves both Conjectures 3.0.4 and 2.1.12, as follows.

Theorem 3.0.5 [7]

For every $\varepsilon > 0$ and for any large enough t , there is a graph with no K_t -minor, whose $(\frac{3}{2} - \varepsilon)t$ -colorings are not all Kempe equivalent.

In fact, we prove that for every $\varepsilon > 0$ and for any large enough t , there is a graph that does not admit a K_t -minor but admits a frozen $(\frac{3}{2} - \varepsilon)t$ -coloring. To obtain Theorem 3.0.5, we then argue that the graph admits a coloring with a different color partition.

The notion of frozen k -coloring is related to that of *quasi- K_p -minor*, introduced in [VM81]. Recall that a graph G admits a K_p -minor if it admits p non-empty,

pairwise disjoint and connected blobs $B_1, \dots, B_p \subset V(G)$ such that for any $i \neq j$, there is an edge between some vertex in B_i and some vertex in B_j . For the notion of quasi- K_p -minor, we drop the restriction that each B_i should induce a connected subgraph of G , and replace it with the condition that for any $i \neq j$, the set $B_i \cup B_j$ induces a connected subgraph of G . If the graph G admits a frozen p -coloring, then it trivially admits a quasi- K_p -minor¹, while the converse may not be true.

If all p -colorings of a graph form a single equivalence class, then either there is no frozen p -coloring or there is a unique p -coloring of the graph up to color permutation. The latter situation in a graph with no K_p -minor would disprove Hadwiger's conjecture, so Las Vergnas and Meyniel conjectured that there is no frozen p -coloring in that case. Namely, they conjectured the following.

Conjecture 3.0.6 Conjecture C in [VM81]

For any t , any graph that admits a quasi- K_t -minor admits a K_t -minor.

Conjecture 3.0.6 was proved for increasing values of t , and is now known to hold for $t \leq 10$ [Jør94, Son05, Kri21]. As discussed above, we strongly disprove Conjecture 3.0.6 for large t . It is unclear how large t needs to be for a counterexample.

3.1 Upper bounds on $\text{diam}(\mathcal{K}_k(G))$ in sparse graph classes

3.1.1 Planar graphs

The results in this subsection are joint work with Quentin Deschamps, Carl Feghali, František Kardoš and Théo Pierron [8]. Theorem 2.2.3 is our main result and strengthens Theorem 2.1.10 and Corollary 2.1.18.

The proof of Theorem 2.2.3 is based on a proof method introduced in [Feg19] and some ideas from [EF20], but of course has many of its own features. We prove the specific case of $k = 5$ but the proof can be adapted for a larger number of colors. We first analyze a proof from [Moh07] to get an essential estimate (Proposition 3.1.1) that we then use to prove Theorem 2.2.3.

3.1.1.1 The case of 3-colorable planar graphs

We recall the following theorem of Mohar:

Theorem 2.1.17 [Moh07]

All 3-colorable planar graphs are 4-recolorable.

¹One blob for each color class.

By making a few simple observations, we show that the proof of Theorem 2.1.17 in fact establishes the following stronger fact that is crucial to Theorem 2.2.3.

Proposition 3.1.1

Let G be a 3-colorable planar graph on n vertices. Then for every 4-colorings α and β of G , there exists a sequence of Kempe changes from α to β that changes the color of each vertex at most $O(n^2)$ times.

In particular, note that Proposition 3.1.1 implies that $\mathcal{K}_4(G)$ has diameter $O(n^3)$ when G is a 3-colorable planar graph on n vertices. The rest of this section is devoted to the proof of Proposition 3.1.1. We basically follow the same steps as the proof of Theorem 2.1.17, except that we add some complexity estimates to Theorem 2.1.6 of Fisk [Fis77a] that handles the case of 3-colorable planar triangulations, and a reduction to this case as done by Mohar [Moh07].

We recall the following theorem of Fisk discussed in Chapter 2:

Theorem 2.1.6 [Fis77a]

All 3-colorable triangulations of the sphere are 4-recolorable.

We show that the proof of Theorem 2.1.6, word for word, gives us the following estimate.

Lemma 3.1.2 [Fis77a]

Let G be a 3-colorable triangulation of the plane on n vertices. Then for every 4-colorings α and β of G , there exists a sequence of Kempe changes from α to β that changes the color of each vertex at most $O(n^2)$ times.

Let α be a 4-coloring of a triangulation G of the plane, and let $e = uv$ be an edge of G . We denote by $\alpha(e) = \{\alpha(u), \alpha(v)\}$ the color of e under α . If uvw and uvx are the two triangles containing an edge uv , we recall that uv is *singular* under α if $\alpha(w) = \alpha(x)$. Note that, for a triangulation G of the plane, if G has a 3-coloring, then this coloring is unique up to permutations of colors, and all edges are singular.

As explained in Subsubsection 2.1.1.2, the proof of Theorem 2.1.6 and by extension of Lemma 3.1.2 relies on the following key structural lemma extracted from [Fis77a].

Lemma 3.1.3

Let G be a 3-colorable triangulation of the plane, and let α be a 4-coloring of G . Then every monochromatic set of non-singular edges of G contains a cycle that bounds some region of the plane.

Using this lemma, we can conclude the proof of Lemma 3.1.2.

Proof of Lemma 3.1.2. We shall show, by exhibiting at most $n \cdot |E(G)|$ Kempe changes, that any 4-coloring of G is Kempe equivalent to the (unique) 3-coloring of G , which will prove the lemma. To do so, given a 4-coloring α of G , we show how to obtain, via at most n Kempe changes, a new coloring β that is Kempe equivalent to α and with fewer non-singular edges. As no edge in the 3-coloring of G is non-singular, by iterating this argument at most $|E(G)|$ times the result will follow.

Let e be a non-singular edge of G . By Lemma 3.1.3, there is a cycle in G whose edges have the same color as e and bounding some region D of the plane. By interchanging the two colors in $\{1, 2, 3, 4\} \setminus \alpha(e)$ in the interior of D , we obtain a new coloring β with fewer non-singular edges than α (singular edges in the interior of D stay singular, while the edges of the cycle on the boundary of D change from non-singular to singular). \square

We now explain how to reduce to the triangulation case in Proposition 3.1.1. We first restate Proposition 4.3 from Mohar [Moh07] except that we add some observations about the number of vertices of the resulting triangulation and the number of Kempe changes involved – these directly follow from Mohar’s proof.

Proposition 3.1.4 [Moh07]

Let G be a planar graph with a facial cycle C and two 4-colorings α, β . Then there exists a graph H formed from G by adding a near-triangulation of size $O(|C|)$ inside C and two 4-colorings α', β' of H such that $\alpha'|_{V(G)}$ and $\beta'|_{V(G)}$ are obtained from α, β using $O(1)$ Kempe changes. Moreover, if the restriction of α to C is a 3-coloring, then α' is a 3-coloring of H that agrees with α on $V(G)$.

We may now prove Proposition 3.1.1.

Proof of Proposition 3.1.1. The proof follows the same steps as [Moh07, Theorem 4.4]. We apply Proposition 3.1.4 to each face of G (instead [Moh07, Proposition 4.3]). We thus made $O(n)$ Kempe changes, and the resulting triangulation T has $O(n)$ vertices. We then apply Lemma 3.1.2 (instead [Moh07, Theorem 4.1]) to obtain a sequence of Kempe changes between the two colorings of T that changes the color of each vertex at most $O(n^2)$ times. \square

3.1.1.2 Strategy to 5-recolor planar graphs

We now prove Theorem 2.2.3. Thanks to the celebrated Four Color Theorem, it suffices to prove the following result.

Theorem 3.1.5

Let G be a plane graph with n vertices. For every 5-coloring α of G and every 4-coloring β of G , there is a sequence of Kempe changes from α to β where each vertex is recolored polynomially many times.

In the remainder of this section, we prove Theorem 3.1.5. Let us briefly sketch the details of the approach. The proof proceeds by induction on the number of vertices. Our aim is to describe a sequence of Kempe changes from α to β such that each vertex is recolored at most $f(n)$ times, where f will satisfy a recurrence relation given at the end of the section. To establish this, we roughly adopt the following strategy:

1. We find a ‘large’ independent set I that is monochromatic in both α and β and that contains vertices of degree at most 6 in G (that I is ‘large’ will enable us to show that f is a polynomial function).
2. We introduce an operation at a vertex that we call *collapsing*, which when applied to each vertex of I gives a new graph H where the degree of each vertex of I is at most 4 in H and such that $F = H - I$ is planar. We use these to show that any sequence of Kempe changes in F extends to a sequence of Kempe changes in H and, in turn, in G .
3. We apply induction to find a sequence of Kempe changes in F from any 5-coloring of F to a 4-coloring of F avoiding the color $\alpha(I)$. Applying Step 2, this sequence extends to a sequence in G ending at a 5-coloring, where color 5 may appear only on I .
4. By definition, $I \subset B$ for some color class B of β ; so we can recolor each vertex of B to color 5. Finally, noting that $G - B$ is a 3-colorable planar graph, we then apply Proposition 3.1.1 to recolor the remaining vertices in $G - B$ to their color in β .

In the rest of this section, we give the details and conclude with a small analysis of the maximum number of times a vertex changes its color.

Step 1: Constructing I .

We prove that the required independent set I exists.

Lemma 3.1.6

There exists an independent set I of G such that:

- all the vertices of I have degree at most 6,
- I is contained in a color class of α and of β ,
- $|I| \geq \frac{n}{140}$.

Proof. Let S be the set of vertices of degree at most 6 in G . Then $|S| > n/7$ since otherwise

$$\sum_{v \in V(G)} d(v) \geq \sum_{v \in V(G) - S} d(v) \geq 7(n - \frac{n}{7}) = 6n,$$

which contradicts Euler's formula.

For $i \in \{1, \dots, 5\}$ and $j \in \{1, \dots, 4\}$ define the set

$$S_{i,j} = S \cap \alpha^{-1}(i) \cap \gamma^{-1}(j).$$

Note that each $S_{i,j}$ satisfies all the criteria from the lemma, except maybe the last. However, by the pigeonhole principle, there exists i and j such that $S_{i,j}$ contains at least $|S|/(5 \times 4) \geq n/140$ vertices, which concludes the proof. \square

From now on, we fix a set I satisfying the hypotheses of Lemma 3.1.6.

Step 2: Constructing H and extending recoloring sequences.

In order to construct H , we want to identify vertices in $N(I)$ that are colored alike so that vertices of I end up with degree 4 and the resulting graph with I excluded is planar. We show that we can modify the coloring α so that such identification becomes possible.

Let P be a plane graph. For a 5-coloring φ of P and a vertex v of K with $d(v) = 6$, we say that v is φ -good if, in φ , the vertex v has three neighbors colored alike or two pairs (a, b) and (c, d) of neighbors colored alike that are *non-overlapping*, i.e. such that $P - v + ab + cd$ is planar (see Figure 3.1). A sequence of Kempe changes is said to *avoid* color a if no vertex involved in some Kempe change in the sequence changes its color to a .




Figure 3.1: Examples of non-overlapping neighborhoods: removing the central vertex and merging the vertices colored identically results in a planar graph.

Lemma 3.1.7

Let P be a plane graph, φ a 5-coloring of K , and $v \in V(P)$ such that $d(v) = 6$. There exists a sequence of at most three Kempe changes avoiding $\varphi(v)$ that transforms φ into a 5-coloring β of P such that v is β -good.

3. Kempe changes in sparse graphs

Proof. We can assume that v is not φ -good. Then the neighbors of v can be colored in four possible ways (up to permutation of colors). We present a visual proof in Figure 3.2, where the six neighbors of v are represented by circles from left to right in the cyclic ordering around v , a bold circle represents an attempt to perform a Kempe change, a curved edge between two vertices u and w represents a Kempe chain containing both u and w , and dashed arrows represent the actual Kempe changes. \square

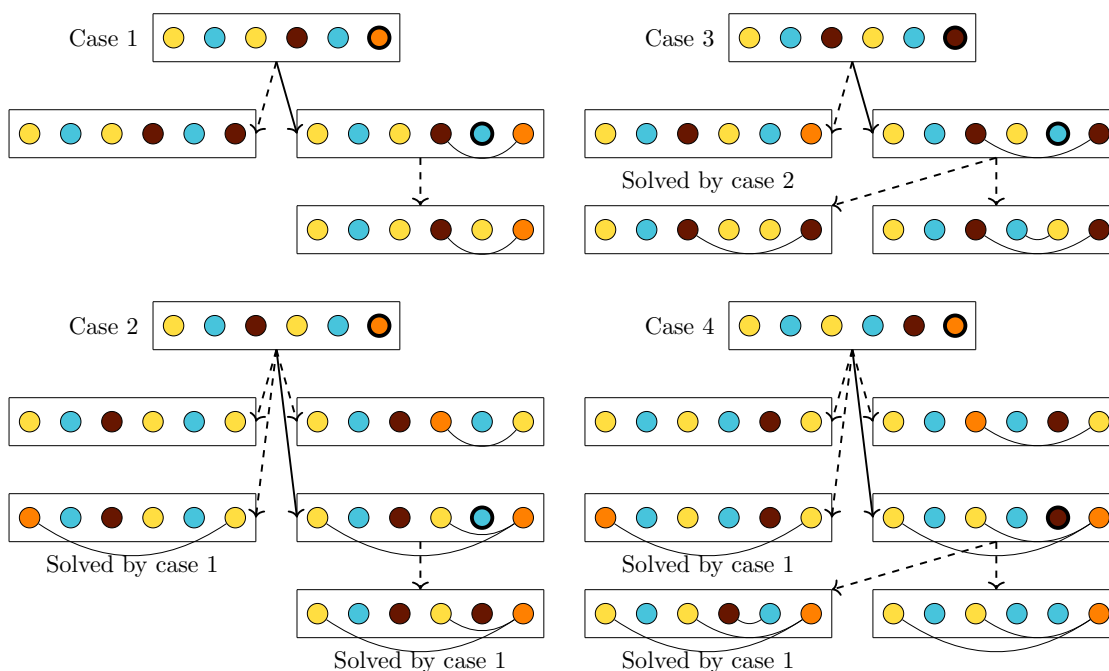


Figure 3.2: The proof of Lemma 3.1.7

We may now successively process each vertex of I : for each $v \in I$, we apply Lemma 3.1.7 to make v φ -good, then we identify vertices in $N(v)$ so that v has degree at most 4 in the resulting graph P' , and moreover, $P' - v$ is planar. Note that since we avoid the color of v in a Kempe change by Lemma 3.1.7, the vertices of I are never recolored at any stage of the construction. We now formalize this intuition.

Let P be a plane graph. Let φ be a 5-coloring of P . For a vertex v of P of degree at most 6 such that v is φ -good if $d(v) = 6$, we say that (P', φ') is the result of *collapsing* (P, v, φ) if

- in the case $1 \leq d(v) \leq 4$, $P' = P$ and $\varphi' = \varphi$;

- in the case $d(v) = 5$, P' is the graph obtained from P by identifying two neighbors u and w of v with $\varphi(u) = \varphi(w)$ into a new vertex z and φ' is the coloring obtained from φ by setting $\varphi'(z) = \varphi(u) = \varphi(w)$;
- in the case $d(v) = 6$, P' is the graph obtained from P by either
 - identifying three neighbors u, w, z of v with $\varphi(u) = \varphi(w) = \varphi(z)$ into a new vertex x . In this case, φ' is obtained from φ by setting $\varphi'(x) = \varphi(u)$, or
 - for two non-overlapping pairs (u, w) and (x, z) of neighbors of v with $\varphi(u) = \varphi(w)$ and $\varphi(x) = \varphi(z)$, identifying u, w into a new vertex s and x, z into a new vertex t , and defining φ' by setting $\varphi'(s) = \varphi(u)$ and $\varphi'(t) = \varphi(x)$.

Since v is φ -good, it should be immediate that P' and φ' are well-defined, and that $P' - v$ is planar. We now focus on showing that we can extend a given recoloring sequence of the collapsed graph to a recoloring sequence of the original graph.

Lemma 3.1.8

Let P be a plane graph with a 5-coloring φ . Let v be a vertex of P of degree at most 6 such that v is φ -good if $d(v) = 6$. Let (P', φ') be the result of collapsing (P, v, φ) . Then every sequence of Kempe changes in $P' - v$ starting from $\varphi'|_{P' - v}$ extends to a sequence of Kempe changes in P starting from φ . Moreover, each vertex of $P - v$ changes its color as many times as in $P' - v$, and v changes its color at most once every time one of its neighbors in P' changes its color.

Proof. Each time a neighbor w of v is recolored in $P' - v$, we may use the same Kempe change in P' unless it involves the color of v and there is another neighbor u of v of the same color as w . In this case, since at most 3 colors appear on $N_{P'}(v)$ (as v has degree 4 in P') we can precede this Kempe change by first recoloring v to a color not appearing in its neighborhood. This shows that any sequence of Kempe changes in $P' - v$ extends to a sequence in P' . To extend the sequence to P , observe that we can simulate in P a Kempe change in P' at a vertex w by performing a Kempe change at each vertex that was identified to form w . Clearly, each vertex of $P - v$ changes its color as many times as in $P' - v$, and v changes its color at most once every time one of its neighbors in P' changes its color, which concludes. \square

We now successively apply Lemma 3.1.8 to each vertex of I in G , so that every vertex of I becomes φ -good.

Lemma 3.1.9

Let P be a plane graph with a 5-coloring φ . If I is an φ -monochromatic independent

set of vertices of degree at most 6 in P , then there is a 5-coloring ψ of P for which I is ψ -good and a sequence of Kempe changes from φ to ψ that changes the color of each vertex at most $3|I|$ times.

Proof. We shall prove by induction on $|V(P)|$ the stronger claim that there is such a sequence that avoids the color of I .

By Lemma 3.1.7, we can assume, up to at most three Kempe changes, that I contains a vertex v of degree at most 5 or a vertex of degree 6 that is φ -good. So we can let (P', φ') be the result of collapsing (P, v, φ) . Since $P'' = P' - v$ is planar, we can apply our induction hypothesis to P'' (with $I \setminus \{v\}$ instead of I and $\varphi'_{|P''}$ instead of φ) to find a sequence of Kempe changes in P'' that avoids the color of v and that transforms $\varphi'_{|P''}$ to some 5-coloring φ'' of P'' so that the following holds:

- $I \setminus \{v\}$ is φ'' -good, and
- each vertex changes its color at most $3(|I| - 1)$ times.

By Lemma 3.1.8, this sequence extends to a sequence in G . Moreover, v does not change its color (as the sequences avoids $\varphi(v)$), and every other vertex changes its color at most $3(|I| - 1) + 3 = 3|I|$ times. This completes the proof. \square

Step 3: Induction.

By Lemma 3.1.9, we can assume that I is α -good. Let $n = |V(G)|$ and $f(n)$ be the maximum number of times a vertex in G is involved in a Kempe change. Write $I = \{v_1, \dots, v_m\}$, set $G_1 = G$, $\psi_1 = \alpha$ and, for $i = 2, \dots, m + 1$, let $H_{i-1} = G_i - \{v_1, \dots, v_{i-1}\}$ where (G_i, ψ_i) is the result of collapsing $(G_{i-1}, v_{i-1}, \psi_{i-1})$. Then the final graph H_m is planar. Hence, by the induction hypothesis combined with Lemma 3.1.6 (with H_m in place of G), there is a sequence of Kempe changes from $(\psi_{m+1})_{|H_m}$ to some 4-coloring γ' of H_m on colors $\{1, \dots, 5\} \setminus \alpha(I)$ where each vertex changes its color at most $f(n - |I|)$ times. By successively applying Lemma 3.1.8 to H_m etc. up until H_1 this same sequence extends to a sequence in G from α to some 5-coloring ψ of G where $\psi_{|G-I}$ uses only colors $\{1, \dots, 4\}$. Now, recalling Step 4 verbatim, by definition, $I \subset B$ for some color class B of β ; so we can recolor each vertex of B to color 5. Finally, noting that $G - B$ is a 3-colorable planar graph, we then apply Proposition 3.1.1 to recolor the remaining vertices in $G - B$ to their color in β .

Complexity analysis.

By Lemma 3.1.9, the color of each vertex is changed at most $3|I|$ times in order to reach a 5-coloring ψ in which I is ψ -good. During the induction step, by Lemma 3.1.8, the color of each vertex in I is changed at most $4f(n - |I|)$ times,

while the color of the other vertices changes at most $f(n - |I|)$ times. By Proposition 3.1.1, the final step requires at most $O(n^2)$ changes of color per vertex. We deduce that $f(n)$ satisfies the recurrence

$$f(n) \leq 3|I| + 4f(n - |I|) + O(n^2) \leq 4f\left(n - \frac{n}{140}\right) + O(n^2).$$

The master theorem then yields that each vertex changes its color $O(n^{\log_{139} 140(4)}) = O(n^{194})$ times, hence the sequence has length at most $O(n^{195})$, which concludes the proof of Theorem 3.1.5.

3.1.2 Technical improvements of Theorem 2.1.15

The remaining of this section is joint work with Marthe Bonamy and Vincent Delecroix. In this subsection, we strengthen Theorem 2.1.15 under two different additional assumptions:

Proposition 3.1.10 [6]

Let G be a graph and let $v_1 \prec \dots \prec v_n$ be an ordering of $V(G)$. If the ordering yields a $(d - 1)$ -degeneracy sequence and $\deg(v) \leq d$ for every vertex v but v_n , then any two k -colorings of G are Kempe equivalent up to $O(dn^2)$ Kempe changes, for $k \geq d$. More precisely, each vertex is affected by $O(dn)$ Kempe changes.

Proposition 3.1.11 [6]

Let G be a graph and let $v_1 \prec \dots \prec v_n$ be an ordering of $V(G)$. Given a list assignment L of G , if $|L(v)| \geq \deg(v) + 1$ for all vertices v but v_n , then any two L -colorings of G are equivalent up to $O(n)$ Kempe changes.

The proofs of Propositions 3.1.10 and 3.1.11 are very similar. Given two colorings α and β , the main idea is to recolor progressively the vertices starting from v_n to v_1 , such that at the i -th step, the colorings agree on $\{v_{n-i}, \dots, v_n\}$ and the color of v_{n-i} will never be changed in further steps. This recoloring step is done recursively using at most $O(n)$ Kempe changes in Proposition 3.1.10 and in $O(1)$ Kempe changes in Proposition 3.1.11. We refer the reader to Section 2 of [6] for the proof of these results.

We will use Proposition 3.1.10 to prove Theorem 2.2.2 in Subsection 3.1.3, and Proposition 3.1.11 in Subsection 3.1.4 to prove Theorem 3.0.2.

3.1.3 Graphs of bounded degree

Let G be a graph of maximum degree Δ other than the 3-prism and let $k \geq \Delta$. Feghali, Johnson, Paulusma [FJP17], Bonamy, Bousquet, Feghali and Johnson [BBFJ19] proved that $\mathcal{K}_k(G)$ is connected. If $k > \Delta$ or if G is not regular,

Proposition 3.1.10 states that $\mathcal{K}_k(G)$ has diameter $O(\Delta n^2)$. Assume that $k = \Delta$ and that G is regular. Let u be a vertex of G . Given any Δ -coloring of G , there are at least two neighbors of u that are colored alike. Denote G_{v+w} the graph where two non-adjacent neighbors v and w of u are identified and $v + w$ the resulting vertex.




Figure 3.3: The graph G_{v+w} (right) obtained from the 4-regular graph G (left) is 3-degenerate with all its vertices but $v + w$ of degree at most 4

The Δ -colorings of G_{v+w} are in one-to-one correspondence with the Δ -colorings of G that color v and w alike. Hence, we have

$$V(\mathcal{K}_\Delta(G)) = \bigcup_{\substack{v,w \in N(u) \\ vw \notin E(G)}} V(\mathcal{K}_\Delta(G_{v+w}))$$

Performing a Kempe change in G_{v+w} corresponds to performing one or two Kempe changes in G , to maintain identical colors on v and w . As a result,

$$\text{diam}(\mathcal{K}_\Delta(G)) \leq 2 \sum_{\substack{v,w \in N(u) \\ vw \notin E(G)}} \text{diam}(\mathcal{K}_\Delta(G_{v+w})). \quad (3.1)$$

Note that G_{v+w} is $(\Delta - 1)$ -degenerate and all its vertices but $v + w$ are of degree at most Δ (see Figure 3.3). By Proposition 3.1.10, $\mathcal{K}_\Delta(G_{v+w})$ has diameter at most $O(\Delta n^2)$. Together with Eq. (3.1), this proves that $\text{diam}(\mathcal{K}_\Delta(G)) = O(\Delta^3 n^2)$. As a result, we obtain Theorem 2.2.2 with $O_\Delta(n^2) = O(\Delta^3 n^2)$.

One can improve this bound to $O(\Delta n^2)$ by replacing the use of Theorem 2.1.15 with Proposition 3.1.10 in the proof of Bonamy, Bousquet, Feghali and Johnson [BBFJ19] and counting the number of Kempe changes performed. This is tedious however and we will not develop it here.

3.1.4 Graphs of bounded maximum average degree

A key ingredient to prove Theorem 3.0.2 is the notion of t -layering of bounded degree, introduced by Bousquet and Perarnau in [BP16]. A t -layering of a graph

G is an ordered partition $V = V_1 \sqcup \dots \sqcup V_t$ of its vertices into t subsets. We call *layers* the atoms V_i of the partition. Given a t -layering, we denote $G_i = G[\bigcup_{j \geq i} V_j]$ for $1 \leq i \leq t$ and we define the *level* $\ell(v)$ of any vertex v as the index i of the subset V_i it belongs to. We say that a t -layering has *degree* k if v has degree at most k in $G_{\ell(v)}$ for all $v \in V$. A t -layering of degree k can be thought of as a k -degeneracy ordering in which vertices are removed layer by layer instead of one by one.

The proof of Theorem 3.0.2 can be decomposed in three main steps. First, proving that if G has mad less than $k - \varepsilon$, then it admits a t -layering of degree $k - 1$ with t being logarithmic in n . This is achieved by Proposition 3.1.12, proved by Bousquet and Perarnau in [BP16]. We do not include the proof and until the end of this subsection, we consider a t -layering of G verifying those properties.

Proposition 3.1.12 [BP16]

For every $k \geq 1$ and every $\varepsilon > 0$, there exists a constant $C = C(k, \varepsilon) > 0$ such that every graph G on n vertices that satisfies $\text{mad}(G) \leq k - \varepsilon$ admits a $(C \log_k n)$ -layering of degree $k - 1$.

The second step consists in proving that if α_i and β_i are two colorings of G_i that differ by only one Kempe change, then any extension of α_i to G differs by a polynomial number of Kempe changes from an extension of β_i to G , see Proposition 3.1.13.

Proposition 3.1.13

Let α be a k -coloring of G and K a Kempe chain in G_i with respect to $\alpha|_{G_i}$ that contains only vertices in V_i . Let γ be the coloring of G_i obtained from $\alpha|_{G_i}$ by performing the Kempe change on K in G_i .

There exists a k -coloring β of G within $n^2 \cdot (2k)^i$ Kempe changes of α such that $\beta|_{G_i} = \gamma$.

Finally, by Proposition 3.1.11, the layers $G[V_i]$ can be recolored one by one with $O(|V_i|)$ Kempe changes for each i , starting from $G[V_t]$ down to $G[V_1]$.

3.1.4.1 Technical tools

Before we start proving Proposition 3.1.13, we state a few definitions that will be used in the rest of this subsection. Consider a graph G and a t -layering of degree $k - 1$ of it. Consider an arbitrary total order \prec on the vertices that is consistent with the layering, that is:

$$\forall i < j, \forall (u, v) \in V_i \times V_j, u \prec v.$$

Note that every vertex has at most $k - 1$ greater neighbors. We say a sequence S_1 of vertices is *lexicographically smaller* than another sequence S_2 if

- S_1 is non-empty and S_2 is empty, or
- the first vertex of S_1 is smaller than the first vertex of S_2 , or
- $S_1 = x \oplus S'_1$ and $S_2 = x \oplus S'_2$ for some vertex x , and S'_1 is lexicographically smaller than S'_2 , where \oplus denotes the sequence concatenation.

We then denote $S_1 \prec_{lex} S_2$. Note that in particular, the empty sequence is the biggest element for this order. A sequence of vertices $S = (v_1, \dots, v_p)$ is said to be *level-decreasing* if the sequence of levels $(\ell(v_1), \dots, \ell(v_p))$ is decreasing.

Finally, we adapt the ideas of Subsection 3.1.2 to the context of t -layerings. Given a coloring α , a vertex v and a color c , we say that a vertex u is *problematic* for the pair (v, c) if there exists a non-trivial level-decreasing path of vertices in $K_{v,c}(\alpha, G)$ going from v to u , and such that u has at least two neighbors in $K_{v,c}(\alpha, G_{\ell(u)})$. We give sufficient conditions to extend a Kempe change from G_i to G :

Lemma 3.1.14

Let α be a coloring and K a Kempe chain of G_i for some i . If for all $v \in K$, v has no problematic vertices x for the other color of K with $\ell(x) < i$, then the Kempe chain K' of G containing K verifies $G_i \cap K' = K$.

Proof. We have $K \subseteq G_i \cap K'$. Assume towards contradiction that there exist $w \in (G_i \cap K') \setminus K$. Let $P = (w, u_1, u_2, \dots, u_m)$ be a shortest path in K' from w to some vertex u_m of K . Note the u_{m-1} has level less than i , because u_{m-1} belongs to $K' \setminus K$ and K is a Kempe chain of G_i . Let j be the smallest index such that (u_j, \dots, u_m) is a decreasing path. We have $j \neq m$ because u_{m-1} has level less than i . Therefore, u_j is problematic for (u_m, c) , where c is the color of K unused by u_m , which is a contradiction. \square

3.1.4.2 Removing problematic vertices for a single vertex

We will now design Algorithm 1 whose goal is to remove all problematic vertices for a pair (v, c) , without modifying the colors of the vertices of level at least $\ell(v)$.

Algorithm 1 works recursively on inputs of the form (α, S, c) where α is a k -coloring, S a level-decreasing sequence of vertices and c a color. Even though we will always call Algorithm 1 on sequences of only one vertex, it is useful to keep track of the recursive problematic vertices in the sequence S to bound the total number of Kempe changes performed. Lemma 3.1.15 states the correctness and the complexity of Algorithm 1, when called with a sequence of only one vertex.

Lemma 3.1.15

Let α be a k -coloring of G and v be a vertex of G . For any color c , Algorithm 1 on

Algorithm 1:

Input : A k -coloring α of G , a level-decreasing sequence of vertices S , a color c .

Output: A k -coloring β of G which agrees with α on $V(G_{\ell(v)})$ where v is the last vertex of S . Moreover, the pair (v, c) admits no problematic vertices with respect to β .

Let v be the last vertex of S ;
 Let $\beta = \alpha$;
while the pair (v, c) admits problematic vertices with respect to the coloring β **do**
 1 | Let u be the largest problematic vertex for (v, c) with respect to β ;
 | Let c_u be a color in $[k] \setminus \beta(N_{G_{\ell(u)}}[u])$;
 | Let β be the result of Algorithm 1 on input $(\beta, S \oplus u, c_u)$;
 2 | Perform $K_{u, c_u}(\beta, G)$ in β ;
end
 Return β ;

input $(\alpha, (v), c)$ yields a k -coloring β of G within $n(2(k-1))^{\ell(v)}$ Kempe changes of α , such that α agrees with β on $G_{\ell(v)}$, and the pair (v, c) admits no problematic vertices with respect to β .

To prove Lemma 3.1.15, we start by proving the correctness of Algorithm 1.

Proof of correctness of Algorithm 1. We prove correctness of the algorithm by induction on $\ell(v)$. If $\ell(v) = 1$, then $V(G_{\ell(v)}) = V$, so (v, c) does not admit any problematic vertex. For $\ell(v) > 1$, at each iteration of the while loop, line 1 is possible since u has degree at most $k-1$ in $G_{\ell(u)}$ and has at least two neighbors in $G_{\ell(u)}$ that are colored identically. By the induction hypothesis, the Kempe change in line 2 does not modify the color of the vertices in $G_{\ell(u)} \setminus \{u\}$. Thus, $\ell(v)$ decreases at each iteration of the loop and at most n calls are generated by the current call. \square

Now that we have established Algorithm 1 to be correct, let us bound the number of Kempe changes performed by it. This will yield Lemma 3.1.15 and is the focus of the remaining of this subsection.

Given a call C of Algorithm 1, we will denote S_C the sequence provided in input, and u_C its last vertex.

Observation 3.1.16. If Algorithm 1 is called on (α, S, c) where S is a level-decreasing sequence, and makes some recursive call C , then the sequence S_C is longer than S and is of the form $S \oplus u_C$. Moreover, S_C is also a level-decreasing sequence by construction.

Claim 3.1.17 analog to [BP16]

If a call D is initiated after a call C , then $S_D \prec_{lex} S_C$.

Proof. If D is called by C , then by Observation 3.1.16, S_D is of the form $S_C \oplus u_C$ thus $S_D \prec_{lex} S_C$. By applying this argument inductively, we also have $S_D \prec_{lex} S_C$ if the call D is generated by C .

Now, assume that D is not generated by C . Denote I the initial call of Algorithm 1 and recall that the recursive calls generated by I have a natural tree structure, rooted at I . There exists a unique sequence $\mathcal{S}_C = C_1, C_2, \dots, C_{t_1}$ such that C_i calls C_{i+1} for each $i < t_1$, with $C_1 = I$ and $C_{t_1} = C$; and a unique sequence $\mathcal{S}_D = D_1, D_2, \dots, D_{t_2}$ such that D_i calls D_{i+1} for each $i < t_2$, with $D_1 = I$ and $D_{t_2} = D$. Denote B the last common ancestor of C and D . Since neither C nor D is generated by the other, B is distinct from C and D and thus is not the last element of the sequences \mathcal{S}_C and \mathcal{S}_D . Let B_C be the call after B in \mathcal{S}_C and B_D be the call after B in \mathcal{S}_D . We have:

- B_C and B_D are both called by B and B_C is initiated before B_D ,
- $B_C = C$ or B_C generates C ,
- $B_D = D$ or B_D generates D .

It follows from Observation 3.1.16 that $S_D \prec_{lex} S_{B_D}$, so we just need to show that $S_{B_D} \prec_{lex} S_C$. By Observation 3.1.16 applied inductively, S_C can be written as $S_{B_C} \oplus T_C$ with $S_{B_C} = S_B \oplus u_{B_C}$. Furthermore, S_{B_D} can be written as $S_B \oplus u_{B_D}$. Since both B_D and B_C are called by B and B_C is called before B_D , we have $u_{B_D} \prec u_{B_C}$, so $S_{B_D} = S_B \oplus u_{B_D} \prec_{lex} S_B \oplus u_{B_C} \oplus T_C = S_C$. \square

Claim 3.1.18 [BP16]

Given that the t -layering $V = V_1 \sqcup \dots \sqcup V_t$ has degree at most $k - 1$, the number of level-decreasing paths between two vertices v and w in different levels is at most $(k - 1)^{i-1}$ where $i = |\ell(v) - \ell(w)|$.

We can now finish the proof of Lemma 3.1.15 by proving that the number of Kempe changes performed by Algorithm 1 on input $(\alpha, (v), c)$ is bounded by $n(2(k - 1))^{\ell(v)}$.

Proof of the complexity of Algorithm 1. Observe that the sequences of vertices considered in the recursive calls are subsequences of level-decreasing paths. Each level-decreasing path between v and w has at most $2^{\ell(v) - \ell(w) - 1}$ subsequences that contain v and w . As a result, each vertex w with level $\ell(w)$ less than $\ell(v)$ is the origin of at most $2^{\ell(v) - \ell(w) - 1} (k - 1)^{\ell(v) - \ell(w) - 1}$ Kempe changes. \square

3.1.4.3 Emulate a Kempe change in G_i

In this subsection, we design Algorithm 2, whose goal is to emulate a Kempe change in some G_i by a sequence of Kempe changes in G .

Algorithm 2:

Input : A k -coloring α of G , a level i , a Kempe chain K of G_i recoloring only vertices in V_i .

Output: A k -coloring β of G , whose restriction to G_i is the coloring resulting from the Kempe change K in $\alpha|_{G_i}$.

Let c_1, c_2 be the colors involved in K ;

Let $\beta = \alpha$;

while there exists a vertex $v \in K$ such that (v, c_1) or (v, c_2) admits problematic vertices in β **do**

Among all problematic vertices for some (v, c_1) or (v, c_2) with $v \in K$, let u be the largest one;

Let $v \in K$ and $j \in \{1, 2\}$ be such that u is problematic for (v, c_j) in β ;

Update β with the result of Algorithm 1 on input $(\beta, (v), c_j)$.

end

Let K' be the Kempe chain with respect to G and β containing K ;

Perform the Kempe change K' on β ;

Return β ;

This algorithm will be used to prove Proposition 3.1.13 that we restate:

Proposition 3.1.13

Let α be a k -coloring of G and K a Kempe chain in G_i with respect to $\alpha|_{G_i}$ that contains only vertices in V_i . Let γ be the coloring of G_i obtained from $\alpha|_{G_i}$ by performing the Kempe change on K in G_i .

There exists a k -coloring β of G within $n^2 \cdot (2k)^i$ Kempe changes of α such that $\beta|_{G_i} = \gamma$.

Proof of Proposition 3.1.13. The coloring β is obtained by applying Algorithm 2. The correctness of Algorithm 2 follows from the correctness of Algorithm 1, which holds by Lemma 3.1.15. After each iteration of the loop the largest possible vertex that is problematic for a pair of vertex and a color of K decreases. Therefore, the loop is executed at most n times and at the end, none of the vertices of the Kempe chain K of G_i admit problematic vertices. By Lemma 3.1.14, the Kempe chain K' of G containing K verifies $K' \cap G_i = K$. Thus the coloring β returned by the algorithm is such that $\beta|_{G_i}$ is the coloring of G_i resulting from the Kempe change K applied on $\alpha|_{G_i}$.

Combined with the cost analysis of Algorithm 1 proved in Subsubsection 3.1.4.2, this proves that Algorithm 2 performs at most $n^2(2(k-1))^i + 1$ Kempe changes and is correct. \square

3.1.4.4 Combining the arguments

We are now ready to prove Theorem 3.0.2.

of Theorem 3.0.2. Let $V_1 \sqcup \dots \sqcup V_t$ be a t -layering of G of degree $(k-1)$, with $t = C \log_k n$ (see Proposition 3.1.12).

We claim that G can be recolored layer by layer, starting from $G[V_t]$, with a polynomial number of Kempe changes. We prove this by a downward induction on the level of the layer. Let $1 \leq i \leq t$, let α and β be two k -colorings of G and assume that α and β agree on all vertices of level more than i . For $v \in V_i$, let $L(v) = [k] \setminus \alpha(N(v) \cap G_{i+1})$. We have for all $v \in V_i$, $|L(v)| \geq \deg_{G[V_i]}(v) + 1$, so by applying Proposition 3.1.11, there exists a sequence S of Kempe changes in $G[V_i]$ of size $O(|V_i|)$ leading from $\alpha|_{V_i}$ to $\beta|_{V_i}$. Moreover, each of the Kempe change in S is a proper Kempe change in G_i . By Proposition 3.1.13, each of these Kempe changes can be performed in G , after $O(n^{c_\varepsilon})$ Kempe changes affecting only the vertices of level less than i , for some $c_\varepsilon > 0$ that depends on ε . \square

3.1.5 Graphs of bounded treewidth

This subsection consists in the proof of Theorem 3.0.3. In [BHI⁺20], Bonamy, Heinrich, Ito, Kobayashi, Mizuta, Mühlenthaler, Suzuki and Wasa present a linear time algorithm that recolors chordal graphs:

Proposition 3.1.19 [BHI⁺20]

Given an n -vertex chordal graph G and an integer k , any two k -colorings of G are equivalent up to at most n Kempe changes.

Let G be an n -vertex graph of treewidth tw . Let H be a chordal graph such that G is a (not necessarily induced) subgraph of H , with $\omega(H) = \chi(H) \leq \text{tw} + 1$ and $V(H) = V(G)$. Computing H is equivalent to computing a so-called tree decomposition of G , which can be done in time $f(\text{tw}) \cdot n$ [Bod96].

Since G and H are defined on the same vertex set, there may be confusion when discussing neighborhoods and other notions. When useful, we write G or H in index to specify. There is an ordering $v_1 \prec \dots \prec v_n$ of the vertices of G such that for all $v \in V(G)$, $N_H^+[v]$ induces a clique in H . The ordering can be computed from H in $O(n)$ time using Lex-BFS [RTL76].

The core of the proof of Theorem 3.0.3 lies in the following proposition:

Proposition 3.1.20

Given any k -coloring α of G with $k \geq \text{tw} + 1$, there exists a k -coloring α' of H and hence of G too, that is equivalent to α up to $O(\text{tw} \cdot n^2)$ Kempe changes in G . Algorithm 3 computes α' and a sequence of Kempe changes leading to it.

Proof of Theorem 3.0.3 assuming Proposition 3.1.20. Let α and β be two k -colorings of G . By Proposition 3.1.20, there exists a k -coloring α' (resp. β') that is equivalent to α (resp. β) up to $O(\text{tw} \cdot n^2)$ Kempe changes. Additionally, both α' and β' yield k -colorings of H . Since H is chordal, by Proposition 3.1.19, there exists a sequence of at most n Kempe changes in H from α' to β' . Each of these Kempe changes in H can be simulated by at most n Kempe changes in G , which results in a sequence of length $O(\text{tw} \cdot n^2)$ between α and β . \square

To prove Proposition 3.1.20 and obtain a k -coloring of H , we gradually “add” to G the edges in $E(H) \setminus E(G)$. To add an edge, we first reach a k -coloring where the extremities have distinct colors. Then, we pretend the edge exists, by propagating any Kempe change involving one extremity to the other extremity immediately after. We formalize this process through Algorithm 3. Let $v_1w_1 \prec \dots \prec v_qw_q$ be the edges in $E(H) \setminus E(G)$ in the lexicographic order, where $v_i \prec w_i$ for every i .

We will prove the following three lemmas. Note that Proposition 3.1.20 follows from Lemmas 3.1.21 and 3.1.22, while Lemma 3.1.23 simply guarantees that the proof of Theorem 3.0.3 is indeed constructive.

Lemma 3.1.21

Algorithm 3 outputs a k -coloring α' of G that is Kempe equivalent to α and such that $\alpha'(u) \neq \alpha'(v)$ for all $uv \in E(H)$.

Lemma 3.1.22

Algorithm 3 performs $O(\text{tw} \cdot n^2)$ Kempe changes in G to obtain α' from α .

Lemma 3.1.23

Algorithm 3 runs in $O(\text{tw}^2 \cdot n^3)$ time.

In Algorithm 3, the variable G' keeps track of how close we are to a k -coloring of H . Before the computations start, $G' = G$. When the algorithm terminates, $G' = H$. At every step, G is a subgraph of G' . To refer to G' or α' at some step of the algorithm, we may say the *current* graph or *current* coloring. The Kempe changes that we discuss are performed in G' . Consequently, the corresponding set of vertices might be disconnected in G , and every Kempe change in G' may correspond to between 1 and n Kempe changes in G .

We first show that Algorithm 3 effectively “adds” every edge in H .

Algorithm 3: Reaching a k -coloring of the underlying chordal graph

Input : G a graph of treewidth tw , a k -coloring α of G with $k \geq \text{tw} + 1$
and $v_1w_1 \prec \dots \prec v_qw_q$ such that $G + \{v_1w_1, \dots, v_qw_q\}$ is a
 $\text{tw} + 1$ -colorable chordal graph H

Output: A k -coloring α' of H , that is Kempe equivalent to α

Let $G' = G$;
Let $\alpha' = \alpha$;

1 **for** j from 1 to q **do**

2 **if** $\alpha'(v_j) = \alpha'(w_j)$ **then**

Let $c \in [k] \setminus \alpha'(N_H^+[v_j])$;
// Possible because $\alpha'(v_j) = \alpha'(w_j)$ and $|N_H^+[v_j]| \leq \text{tw} + 1 \leq k$.
Let $U = N_{G'}^-(v_j) \cap N_{G'}^-(w_j) = \{u_1 \prec \dots \prec u_p\}$;

3 **for** $i = p$ down to 1 **do**

4 **if** $\alpha'(u_i) = c$ **then**

Let $c_i \in [k] \setminus \alpha'(N_{G'}^+[u_i])$;
// Possible because $\alpha'(v_j) = \alpha'(w_j)$ and $|N_{G'}^+[u_i]| \leq \text{tw} + 1 \leq k$.
Perform $K_{u_i, c_i}(\alpha', G')$;

5 **end**

6 **end**

// Now $c \notin \alpha'(\{x \in N_{G'}(v_j) \mid x \succeq u_i\})$ (see Claim 3.1.24)

7 **end**

// Now $c \notin \alpha'(N_{G'}(v_j))$ (see Claim 3.1.24)

8 Perform $K_{v_j, c}(\alpha', G')$;

9 **end**

Let $G' = G' \cup \{v_jw_j\}$;

10 **end**

Proof of Lemma 3.1.21. By construction, at every step α' is Kempe equivalent to α . We prove the following loop invariant: at every step, α' is a k -coloring of G' . Since $G' = H$ at the end of the algorithm, proving the loop invariant will yield the desired conclusion.

The invariant holds at the beginning of the algorithm, when $G' = G$.

Assume that at the beginning of the j -th iteration of the loop 1, α' is a proper coloring of G' . All the Kempe changes performed inside the loop are performed in G' , so we only need to prove that at the end of the iteration, we have $\alpha'(v_j) \neq \alpha'(w_j)$.

This is a direct consequence of the following claim:

Claim 3.1.24

The two following statements hold:

- At the end of the if condition of line 3 we have $c \notin \alpha'(\{x \in N_{G'}(v_j) \mid x \succeq u_i\})$,

- At the end of the for loop of line 2, we have $c \notin \alpha'(N_{G'}(v_j))$.

Proof of Claim 3.1.24. The latter item is a direct consequence of the former, so we focus on arguing that after the step i of the inner loop 2, we have $c \notin \alpha'(\{x \in N_{G'}(v_j) \mid x \succeq u_i\})$. The key observation is that at the i -th step of the inner loop 2, the Kempe changes performed at line 4 does not involve vertices greater than u_i . We prove by induction the stronger statement: the Kempe chain T involved in the Kempe change $K_{u_i, c_i}(\alpha', G')$ is a tree rooted at u_i in which all the nodes are smaller than their parent. We prove this by induction on the distance of the nodes to u_i .

First note that $c_i \notin \alpha'(N_{G'}^+(u_i))$ so all the neighbors of u_i in T are smaller than u_i . Let x be a vertex of T , other than u_i . If x has a neighbor $y \in T$ such that $x \prec y \preceq u_i$, then all neighbors w of x different from y must be smaller than x . Otherwise, $w, y \in N^+(x)$ and since H is chordal, this implies that $wy \in E(H)$. Since $y \prec u_i \prec v_j$, we have $wy \in E(G')$. But a Kempe chain is bichromatic and cannot contain the triangle xwy , raising a contradiction.

As a result, the Kempe change of line 4 does not recolor any vertex larger than u_i , and the first item is true. As explained before, the second item follows directly from the first one. \square

Therefore at the beginning of line 5, $\alpha'(v_j) = \alpha'(w_j)$ and $c \notin \alpha(N(v_j))$. At the end of line 5, v_j and w_j are colored differently, the loop invariant holds. \square

We now show that Algorithm 3 “adds” all the edges of H using $O(\text{tw} \cdot n^2)$ Kempe changes.

Proof of Lemma 3.1.22. We now prove that the number of Kempe changes in G performed by the algorithm is $O(\text{tw} \cdot n^2)$.

We first prove that for each vertex x , there exists at most one step j of the loop 1 for which $v_j = x$ and we enter the conditional statement $\alpha'(v_j) = \alpha'(w_j)$. Indeed, the first time we enter the conditional statement, the vertex x is recolored with a color c not in $N_H^+(x)$ at line 5. Note that all the edges xy with $y \succ x$ are consecutive in the ordering of $E(H) \setminus E(G)$. Therefore, once the vertex x is recolored, all the remaining edges xy are handled without Kempe change, as the conditional statement is not satisfied. This implies directly that line 5 is executed at most n times.

Now, we bound the number of times x plays the role of u_i in the Kempe change at line 4. For each step j of loop 1 for which it happens, we have $v_j, w_j \in N_H^+(x)$. Since $|N_H^+(x)| \leq \text{tw}$ and each v_j is involved at most once by the above argument, we obtain that x plays this role at most tw times.

Consequently the overall number of Kempe changes performed in G' by the algorithm is $O(\text{tw} \cdot n)$. Performing a Kempe change in G' is equivalent to performing a Kempe change in all the connected component of G of the Kempe chain of

G' . Therefore, the number of Kempe changes performed in G by the algorithm is $O(\text{tw} \cdot n^2)$. \square

Finally, we argue that Algorithm 3 runs in time $O(\text{tw}^2 \cdot n^3)$.

Proof of Lemma 3.1.23. In total, the loop 2 is executed at most once for every pair of vertices in $N_H^+(u)$ for each $u \in V(G)$, that is $O(\text{tw}^2 n)$ times. However, we also take into account the number of Kempe changes that need to be performed. By Lemma 3.1.22, only $O(\text{tw} \cdot n^2)$ Kempe changes are performed in G . As a result, the total complexity of the algorithm is $O(\text{tw}^2 \cdot n + \text{tw}^2 \cdot n^3)$, that is $O(\text{tw}^2 \cdot n^3)$, since computing a Kempe chain $K_{u,c}(\alpha', G')$ can be done in a naive way in time $O(|V(G')| + |E(G')|) = O(\text{tw} \cdot n)$ by doing a depth-first-search from u . \square

3.2 On a recoloring version of Hadwiger's conjecture

We recall the three conjectures of Las Vergnas and Meyniel:

Conjecture 2.1.12 (Conjecture A in [VM81])

For every t , every K_t -minor-free graph is t -recolorable.

Conjecture 3.0.4 (Conjecture A' in [VM81])

For every t and every graph with no K_t -minor, every equivalence class of t -colorings contains some $(t - 1)$ -coloring.

Conjecture 3.0.6 Conjecture C in [VM81]

For any t , any graph that admits a quasi- K_t -minor admits a K_t -minor.

Here, we disprove both Conjectures 3.0.4, 3.0.6 and 2.1.12, as follows.

Theorem 3.2.1

For every $\varepsilon > 0$ and for any large enough t , there is a graph that admits a frozen t -coloring (and hence a quasi- K_t -minor) but does not admit a $K_{(\frac{2}{3}+\varepsilon)t}$ -minor. This graph admits other non-similar colorings and thus is not t -recolorable.

Following the submission and advertisement of this result, we learned that the exact same construction had been independently discovered in [Ree93] and independently in [BKT10]. While [Ree93] addresses directly the conjectures of Las Vergnas and Meyniel, this paper received very little attention and there exists no electronic version of it. On the other hand, the authors of [Ree93] and the rest of the community were seemingly unaware of the implications on quasi-minors and Kempe equivalence that we discuss here.

3.2.1 Construction

Let $n \in \mathbb{N}$ and let $\varepsilon > 0$. We build a random graph G_n on vertex set $\{a_1, b_1, a_2, b_2, \dots, a_n, b_n\}$ as follows: for every $i \neq j$ independently, we select one pair uniformly at random among $\{(a_i, a_j), (a_i, b_j), (b_i, a_j), (b_i, b_j)\}$ and add the three other pairs as edges to the graph G_n .

Note that the sets $\{a_i, b_i\}_{1 \leq i \leq n}$ form a quasi- K_n -minor, as for every $i \neq j$, the set $\{a_i, b_i, a_j, b_j\}$ induces a path on four vertices in G_n , and hence is connected.

Our goal is to argue that if n is sufficiently large then, with high probability, the graph G_n does not admit any $K_{(\frac{2}{3}+\varepsilon)n}$ -minors. This will yield Theorem 3.2.1. To additionally obtain Theorem 3.0.5, we need to argue that, with high probability, G_n admits an n -coloring with a different color partition than the natural one, where the color classes are of the form $\{a_i, b_i\}$. Informally, we can observe that each of $\{a_1, \dots, a_n\}$ and $\{b_1, \dots, b_n\}$ induces a graph behaving like a graph in $\mathcal{G}_{n, \frac{3}{4}}$ (i.e. each edge exists with probability $\frac{3}{4}$) though the two processes are not independent. This argument indicates that $\chi(G_n) = O(\frac{n}{\log n})$, but we prefer a simpler, more pedestrian approach.

Assume that for some i, j, k, ℓ , none of the edges $a_i b_j, a_j b_k, a_k b_\ell$ and $a_\ell b_i$ exist. Then the graph G_n admits an n -coloring α where $\alpha(a_p) = \alpha(b_p) = p$ for every $p \notin \{i, j, k, \ell\}$ and $\alpha(a_i) = \alpha(b_j) = i$, $\alpha(a_j) = \alpha(b_k) = j$, $\alpha(a_k) = \alpha(b_\ell) = k$ and $\alpha(a_\ell) = \alpha(b_i) = \ell$ (see Figure 3.4). Since every quadruple (i, j, k, ℓ) has a positive and constant probability of satisfying this property, G_n contains such a quadruple with very high probability when n is large.




Figure 3.4: A different n -coloring given an appropriate quadruple. Dashed edges represent the absence of an edge.

We are now ready to prove that the probability that G_n admits a $K_{(\frac{2}{3}+\varepsilon)n}$ -minor tends to 0 as n grows to infinity. We consider three types of K_p -minors in

G , depending on the size of the blobs involved. If every blob is of size 1, we say that it is a *simple* K_p -minor – in fact, it is a subgraph. If every blob is of size 2, we say it is a *double* K_p -minor. If every blob is of size at least 3, we say it is a *triple* K_p -minor. We prove three claims, as follows.

Claim 3.2.2

For any $\delta > 0$, $\mathbb{P}(G_n \text{ contains a simple } K_{\delta n}\text{-minor}) \rightarrow 0$ as $n \rightarrow \infty$.

Claim 3.2.3

For any $\delta > 0$, $\mathbb{P}(G_n \text{ contains a double } K_{\delta n}\text{-minor}) \rightarrow 0$ as $n \rightarrow \infty$.

Claim 3.2.4

G_n does not contain a triple $K_{\frac{2}{3}n+1}$ -minor.

Claims 3.2.2 to 3.2.4 are proved in Subsections 3.2.2 to 3.2.4, respectively. If a graph admits a K_p -minor, then in particular it admits a simple K_a -minor, a double K_b -minor and a triple K_c -minor such that $a + b + c \geq p$. Setting δ appropriately and combining Claims 3.2.2 to 3.2.4, we derive the desired conclusion.

3.2.2 No large simple minor

Proof of Claim 3.2.2. Let S be a subset of k vertices of G_n . The probability that S induces a clique in G_n is at most $\left(\frac{3}{4}\right)^{\binom{k}{2}}$. Indeed, if $\{a_i, b_i\} \subseteq S$ for some i , then the probability is 0. Otherwise, $|S \cap \{a_i, b_i\}| \leq 1$ for every i , so we have $G[S] \in \mathcal{G}_{k, \frac{3}{4}}$, i.e. edges exist independently with probability $\frac{3}{4}$. Therefore, the probability that S induces a clique is $\left(\frac{3}{4}\right)^{\binom{k}{2}}$.

By union-bound, the probability that some subset on k vertices induces a clique is at most $\binom{2n}{k} \cdot \left(\frac{3}{4}\right)^{\binom{k}{2}}$. For any $\delta > 0$, we note that $\binom{2n}{\delta n} \leq 2^{2n}$. Therefore, the probability that G_n contains a simple $K_{\delta n}$ -minor is at most $2^{2n} \cdot \left(\frac{3}{4}\right)^{\binom{\delta n}{2}}$, which tends to 0 as n grows to infinity. \square

3.2.3 No large double minor

Proof of Claim 3.2.3. Let S' be a subset of k pairwise disjoint pairs of vertices in G_n such that for every i , at most one of $\{a_i, b_i\}$ is involved in S' .

We consider the probability that G_n/S' induces a clique, where G_n/S' is defined as the graph obtained from G_n by considering only vertices involved in some pair of S' and identifying the vertices in each pair.

We consider two distinct pairs $(x_1, y_1), (x_2, y_2)$ of S' . Without loss of generality, $\{x_1, x_2, y_1, y_2\} = \{a_i, a_j, a_k, a_\ell\}$ for some i, j, k, ℓ . The probability that there is an edge between $\{x_1, y_1\}$ and $\{x_2, y_2\}$ is $1 - \left(\frac{1}{4}\right)^4$. In other words,

$\mathbb{P}(E((x_1, y_1), (x_2, y_2)) = \emptyset) = (\frac{1}{4})^4$ and since at most one of $\{a_i, b_i\}$ is involved in S' for all i , all such events are mutually independent. Therefore, the probability that S' yields a quasi- $K_{|S'|}$ -minor is $(1 - (\frac{1}{4})^4)^{\binom{|S'|}{2}}$.

For any $\delta' > 0$, the number of candidates for S' is at most $\binom{2n}{2\delta'n}$ (the number of choices for a ground set of $2\delta'n$ vertices) times $(2\delta'n)!$ (a rough upper bound on the number of ways to pair them). Note that $\binom{2n}{2\delta'n} \cdot (2\delta'n)! \leq (2n)^{2\delta'n}$. We derive that the probability that there is a set S' of size $\delta'n$ such that $G_n/S' = K_{|S'|}$ is at most $(2n)^{2\delta'n} \cdot (1 - (\frac{1}{4})^4)^{\binom{\delta'n}{2}}$, which tends to 0 as n grows large.

Consider a double K_k -minor S of G_n . Note that no pair in S is equal to $\{a_i, b_i\}$ (for any i), as every blob induces a connected subgraph in G_n . We build greedily a maximal subset $S' \subseteq S$ such that S' involves at most one vertex out of every set of type $\{a_i, b_i\}$. Note that $|S'| \geq \frac{|S|}{3}$. Therefore, by taking $\delta' = \frac{\delta}{3}$ in the above analysis, we obtain that the probability that there is a set S of δn pairs that induces a quasi- $K_{|S|}$ -minor tends to 0 as n grows large. \square

3.2.4 No huge triple minor

Proof of Claim 3.2.4. The graph G_n has $2n$ vertices, and a triple K_k -minor involves at least $3k$ vertices. It follows that if G_n contains a triple K_k -minor then $k \leq \frac{2n}{3}$. \lrcorner

Conclusion and open questions

We leave four open questions. The first one, already discussed in Chapter 2, is the Kempe version of the Cereceda's conjecture:

Conjecture 2.2.1 [BBFJ19]

Let G be an n -vertex d -degenerate graph. Any two k -colorings of G are Kempe equivalent up to $O(n^2)$ Kempe changes for $k \geq d + 1$.

The three others are motivated by our result on quasi-minors and investigate the links between minors and quasi-minors. More precisely, the first two ask whether the conjectures of Las Vergnas and Meyniel are true up to a linear factor. Theorem 3.2.1 proves that there is no equivalence between clique minors and clique quasi-minors, but we can wonder if a weaker version of Conjecture 3.0.6 holds:

Question 3.2.5

What is the infimum c such that for any large enough t , there is a graph that admits a quasi- K_t -minor but no K_{ct} -minor?

Trivially, every graph that admits a quasi- K_{2t} -minor admits a K_t -minor, thus $\frac{2}{3} \geq c \geq \frac{1}{2}$. It is unclear what should be the value of c . For graphs with a quasi- K_t -minor such that all blobs have size at most three, $c = 2/3$, but the proof does not generalize easily to more complicated quasi- K_t -minor.

Our proofs shows that not all K_t -minor-free graphs are t -recolorable, hence the recoloring version of Hadwiger's conjecture is false. However, it does not give any insight on the following relaxed question, that recalls the Linear Hadwiger conjecture:

Conjecture 2.1.13 [2]

There exists a constant c such that any K_t -minor-free graph is ct -recolorable.

Our proof shows that $c \geq \frac{3}{2}$ if it exists. In the 1980s, [Kos82, Kos84] and [Tho84] proved independently that a graph with no K_t -minor has degeneracy $O(t\sqrt{\log t})$, with current best hidden factor in [KP20]. Since all the k -colorings of d -degenerate graphs are equivalent for $k > d$ [VM81], this proves that K_t -minor-free graphs are $O(t\sqrt{\log t})$ -(re)colorable. This is the best bound known for Conjecture 2.1.13, unlike the Linear Hadwiger conjecture for which the best bound is as of today $O(t \log(\log(t)))$, below the degeneracy threshold.

Finally, note that if Hadwiger's conjecture is true, then any n -vertex graph without stable sets of size 3 has chromatic number at most $n/2$ and thus admits a clique minor of size $n/2$. This second statement is known as the Seagull conjecture:

Conjecture 3.2.6 (Seagull conjecture)

Any n -vertex graph without stable sets of size 3 admits a clique minor of size $\lfloor n/2 \rfloor$.

Hadwiger's conjecture implies the Seagull conjecture, but surprisingly, the converse holds when restricting to graphs without stable sets of size 3 [PST03]. A motivation for the Seagull conjecture is that the known cases of Hadwiger's conjecture are K_t minors-free graphs for $t \leq 6$, which are sparse and close to being planar, while the Seagull conjecture attacks the problem on very dense graphs (see [CS12] for a detailed discussion on the connections between Hadwiger's and the Seagull conjecture).

Note that in an n -vertex graph G without stable sets of size 3, any P_3 dominates G . Hence packing greedily as many P_3 as possible leaves two cliques and G trivially admits a clique minor of size $\lfloor n/3 \rfloor$. Proving a quasi version of this conjecture would be a interesting first step towards the Seagull conjecture:

Conjecture 3.2.7 (Bonamy, Quasi-Seagull Conjecture)

Any n -vertex graph without stable sets of size 3 admits a clique quasi-minor of size $\lfloor n/2 \rfloor$.

Chapter 4

Optimal algorithms for reconfiguration

This chapter presents algorithms and structural results on other reconfiguration operations. It contains ongoing work with Vincent Delecroix and published results obtained with Guilherme C. M. Gomes, Reem Mahmoud, Amer E. Mouawad, Yoshio Okamoto, Vinicius F. dos Santos and Tom C. van der Zanden [3] and with Ashutosh Rai and Martin Tancer [9].

Introduction

In this Chapter, we study two reconfiguration operations arising from problems in low-dimensional topology: reconfiguration of knot diagrams with Reidemeister moves (see Section 4.1) and reconfiguration of square tiled surfaces with cylinder shears (see Section 4.2). The content of Section 4.1 comes from [9] but has been significantly shortened and reworked to highlight the key ideas.

Untangling knots

A classical and extensively studied question in algorithmic knot theory is to determine whether a given diagram of a knot is actually a diagram of the unknot. This question is known as the *Unknot Recognition Problem*. The first algorithm for this problem was given by Haken [Hak61]. Currently, it is known that the Unknot Recognition Problem belongs to $\text{NP} \cap \text{co-NP}$ but no polynomial time algorithm is known. See [HLP99] for the NP-membership and [Lac21a] for co-NP-membership (co-NP-membership modulo Generalized Riemann Hypothesis was previously established in [Kup14]). In addition, a quasi-polynomial time algorithm for Unknot Recognition has been recently announced by Lackenby [Lac21b].

One possible path for attacking the Unknot Recognition Problem is via Reidemeister moves (see Figure 4.1): if D is a diagram of the unknot, then D can be untangled to a diagram U with no crossing by a finite number of Reidemeister moves. In addition, Lackenby [Lac15] provided a polynomial bound (in the number of crossings of D) on the required number of Reidemeister moves. This is an alternative way to show that the Unknot Recognition Problem belongs to NP, because it is sufficient to guess the required Reidemeister moves for unknotting.

However, if we slightly change our viewpoint, de Mesmay, Rieck, Sedgwick, and Tancer [dRST21] showed that it is NP-hard to decide whether a given diagram can be untangled using a given number of Reidemeister moves. (An analogous result for links has been shown to be NP-hard slightly earlier by Koenig and Tsvietkova [KT21].) For more background on unknotting and unlinking problems, we also refer to Lackenby's survey [Lac17].

The main aim of Section 4.1 is to extend the line of research started in [dRST21] by determining the parameterized complexity of untangling knots via Reidemeister moves. On the one hand, it is easy to see that if we consider parameterization by the number of Reidemeister moves, then the problem is in FPT (class of *fixed parameter tractable* problems). This happens because of a somewhat trivial reason: if a diagram D can be untangled with at most k moves, then D contains at most $2k$ crossings, thus we can assume that the size of D is bounded by k and the problem admits a linear kernel.

On the other hand, we also consider parameterization with an arguably much more natural parameter called the defect (used in [dRST21]). This parameterization is also relevant from the point of view of *above guarantee parameterization* introduced by Mahajan and Raman [MR99]. Here we show that the problem is W[P]-complete with respect to this parameter.

Reidemeister moves.

Let D be a diagram of a knot (see Section 1.5 for a precise definition). *Reidemeister moves* are the local modifications of the diagram depicted in Figure 4.1. We distinguish Reidemeister moves of types I, II, and III as depicted in the figure. In addition, for types I and II, we distinguish whether the moves remove crossings (types I⁻ and II⁻) or whether they introduce new crossings (types I⁺ and II⁺). We denote $D(m)$ and $D(\mathcal{S})$ the diagrams obtained from D after one Reidemeister move m or a sequence \mathcal{S} of such moves.

A diagram D is a *diagram of the unknot* if it can be transformed into the untangled diagram U by a finite sequence of Reidemeister moves. (This is well known to be equivalent to stating that the lift of the diagram to \mathbb{R}^3 , keeping the underpasses/overpasses, is ambient isotopic to the unknot, that is, a standardly embedded S^1 in \mathbb{R}^3 [AB26, Rei27].) For example, the diagram in Figure 1.3 is a




Figure 4.1: Reidemeister moves

diagram of the unknot.

Now we discuss different parameterizations of untangling (un)knots via Reidemeister moves in more detail.

Parameterization via number of Reidemeister moves.

First, we consider the following parameterized problem.

Problem UNKNOTTING VIA NUMBER OF MOVES

INPUT	<i>A diagram D of a knot.</i>
PARAMETER	k .
QUESTION	<i>Can D be untangled to the unknot using at most k Reidemeister moves?</i>

Observation 4.0.1. UNKNOTTING VIA NUMBER OF MOVES belongs to FPT.

Proof. Each Reidemeister move removes at most 2 crossings. Thus, if D contains at least $2k + 1$ crossings, then we immediately answer NO. This can be determined in time $O(|D|)$, where $|D|$ stands for the size of the encoding of D . In the terminology of parameterized complexity, this means that we get a (linear) kernel, which immediately implies that the problem belongs to FPT. \square

Parameterization via defect.

As we discussed earlier, parameterization in the number of Reidemeister moves has the obvious disadvantage that once we fix k , the problem becomes trivial for arbitrarily large inputs (they are obviously a NO instance). We also see that if we have a diagram D with n crossings and want to minimize the number of Reidemeister moves to untangle D , presumably the most efficient way is to remove two crossings in each step, thus requiring at least $n/2$ steps. This motivates the following definition of the notion of defect.

Given a diagram D , by an *untangling* of D we mean a sequence $\mathcal{D} = (D_0, \dots, D_\ell)$ such that $D = D_0$; D_{i+1} is obtained from D_i by a Reidemeister move; and $D_\ell =$

U is the diagram with no crossings. Then we define the *defect* of an untangling \mathcal{D} as above as

$$\text{def}(\mathcal{D}) := 2\ell - n$$

where n is the number of crossings in D . Note that ℓ is just the number of Reidemeister moves in the untangling. It is easy to see that $\text{def}(\mathcal{D}) \geq 0$ and $\text{def}(\mathcal{D}) = 0$ if and only if all moves in the untangling are II^- moves. More precisely, given a Reidemeister move m , let us define the *weight* of this move $w(m)$ via the following table:

Type of the move	II^-	I^-	III	I^+	II^+
$w(m)$	0	1	2	3	4

The following lemma shows that $\text{def}(\mathcal{D})$ in some sense measures the number of “extra” moves in the untangling beyond the trivial bound.

Lemma 4.0.2

Let \mathcal{D} be an untangling of a diagram D . Then $\text{def}(\mathcal{D})$ is equal to the sum of the weights of the Reidemeister moves in \mathcal{D} .

The proof of Lemma 4.0.2 relies on a simple discharging argument and allows us to extend the definition of the defect to non-untangling sequences, as the sum of the weights of their moves.

As we have seen above, asking the question whether a diagram can be untangled with defect at most k is the same as asking if it can be untangled in $k/2$ moves above the trivial, but tight lower bound of $n/2$. In addition, it is possible to get diagrams with arbitrarily large number of crossings but with defect bounded by a constant (even for defect 0 this is possible). Finally, the defect also plays a key role in the reduction in [dRST21] which suggests that the hardness of the untangling really depends on the defect. For these reasons, we find the defect to be a more natural parameter than the number of Reidemeister moves. Therefore, we consider the following problem.

Problem UNKNOTTING VIA DEFECT

INPUT *A diagram D of a knot.*
PARAMETER *k .*
QUESTION *Can D be untangled with defect at most k ?*

With Ashutosh Rai and Martin Tancer, we prove the following result.

Theorem 4.0.3 [9]

The problem UNKNOTTING VIA DEFECT is $W[P]$ -complete.

The proof of Theorem 4.0.3 consists of two main steps: W[P]-membership and W[P]-hardness. Both of them are non-trivial. The W[P]-membership proof relies on a *reorganization* argument, and will be presented in Section 4.1, while the W[P]-hardness reduction will not be discussed here, but can be found in [9].

Reorganization arguments

We refer to results of the following form as reorganization arguments: There is a subset of X of all the reconfiguration operations, such that for any reconfiguration sequence \mathcal{S} leading from a configuration A to a configuration B , there exists a reconfiguration sequence between A and B that starts with a move in X and has the same weight as \mathcal{S} . This weight can be the length of the sequence, or more complicated parameters, such as the defect.

Reorganization results directly give the correctness of a greedy algorithm that computes a reconfiguration sequence of optimal weight, which is the core of the W[P]-membership proof, taking X to be a large subset of the II^- moves. We will give a precise statement of this reorganization argument on Reidemeister moves and its application in Section 4.1.

Finally, let us mention another situation in which we used reorganization arguments to obtain optimal algorithms. With Guilherme C. M. Gomes, Reem Mahmoud, Amer E. Mouawad, Yoshio Okamoto, Vinicius F. dos Santos and Tom C. van der Zanden, we looked at the reconfiguration of minimum s, t -separators in a graph [3]. The configurations are the sets of vertices of minimum size that are s, t -separators. Like most problems where configurations are subsets of vertices of a graph, these subsets are represented by tokens placed on the corresponding vertices, and two reconfiguration operations can be considered. The first one, called *token sliding* consists in moving a single token along an edge, and the second one, called *token jumping*, consists in moving a single token to an arbitrary location in the graph. In the specific case of minimum separator reconfiguration, by Menger's theorem there are k disjoint paths going from s to t , where k is the size of a minimum s, t -separator, and each token can only move along the path it lies on. A reorganization argument allows us to prove that if there is a sequence of moves going from A to B (those can be restricted to token slides or not), then there is one of equal length that starts with a token going towards its final destination on its path. This has several consequences. The first one is that we obtain a polynomial time algorithm to decide if two s, t -separators are equivalent up to a sequence of token slides or up to a sequence of token jumps. In the case of token slides, this algorithm also gives the length of a shortest reconfiguration sequence, and hence decides in polynomial time whether two s, t -separators are separated by at most ℓ token slides. For token jumping however, we show that this problem is NP-complete, but use our reorganization result to prove that it is FPT parameterized

by k the size of a minimum separator. Finally, we prove that the problem admits a polynomial kernel parameterized by $k + \ell$ where ℓ is the length of the reconfiguration sequence, but admits no polynomial kernel parameterized by k only, unless $\text{NP} \subseteq \text{coNP/poly}$.

Playing Rubik's Cube on a square-tiled surface

The *Quadratis puzzles*, designed by Gutiérrez, Parlier and Turner [GPT22], form a family of combinatorial puzzles similar to the Rubik's Cube in which the player moves colored square tiles on a square-tiled surface by sliding cylinders (see Figure 4.2). The aim of this puzzle is twofold: designing a sliding puzzle that unlike the Rubik's Cube, has a reconfiguration graph that can be visualized; the second goal is educational, as this puzzle is presented in the Luxembourg Science Centre and has been the object of activities in science fairs. The game is now available in a free app [GPT].




Figure 4.2: Two configurations of a Quadratis puzzle. The labels indicate the glued edges and are not displayed in the app.

At the Dagstuhl Seminar ‘Computation and Reconfiguration in Low-Dimensional Topological Spaces’ held in February 2022 [BLd⁺22], Hugo Parlier presented the Quadratis puzzles and Saul Schleimer suggested a variant of this idea, since named the *shearing block game*. In Section 4.2, we study the connectedness and the diameter of the reconfiguration graph of this variant. Before we introduce it, we overview the usual reconfiguration operations on combinatorial maps and shift progressively the focus towards the square-tiled surfaces.

Diagonal flips in polygons and embedded graphs

Consider a convex n -gon and its triangulations. The *flip* operation consists in flipping an edge: remove an edge of the triangulation and replace it with the other

diagonal in the created quadrilateral. In that context, the reconfiguration graph \mathcal{T}_n is usually called a *flip graph*. Sleator, Tarjan and Thurston [STT88] proved that the diameter of the flip graph \mathcal{T}_n is equal to $2n - 6$ for $n \geq 11$. This is one of the very few examples where the value of the diameter of the reconfiguration graph is exactly known. Molloy, Reed and Steige [MRS97] proved that the associated Markov chain, called the *flip dynamics*, mixes rapidly. Since then, the mixing time of the flip dynamics has been improved several times, although the exact rate is not known and has been conjectured by Aldous [Ald94] to be $O(n^{3/2})$, up to logarithmic factors. The current best upper bound of $O(n^3 \log^3(n))$ is due to Eppstein and Frishbreg [EF22].

Let us now move from polygons to graphs embedded in surfaces (see Subsection 1.1.4 for all definitions). Let e be an edge adjacent to two distinct faces of degree d_1 and d_2 in a combinatorial map. A *flip* of e consists in first removing e from the graph (see Figure 4.3). This merges the two adjacent faces into one of degree $d_1 + d_2 - 2$. Then, we reintroduce an edge e' inside the newly created face. Note that among the flips associated to e there are some that preserve the faces' degrees.




Figure 4.3: Two combinatorial maps that differ by one flip. The flipped edge is dashed. This flip preserves the degree of the faces.

Let $\mathcal{T}_{g,n}$ and $\mathcal{Q}_{g,n}$ be respectively the set of isomorphism classes of triangulations and of quadrangulations of genus g with n vertices. One can turn $\mathcal{T}_{g,n}$ and $\mathcal{Q}_{g,n}$ into reconfiguration graphs by putting an edge between m and m' if m and m' are related by a flip. We denote $\mathcal{T}_{g,n}^{lab}$ and $\mathcal{Q}_{g,n}^{lab}$ the triangulations and quadrangulations with vertices labeled $[n]$. Up to automorphism, a triangulation in $\mathcal{T}_{g,n}$ gives rise to $n!$ labeled triangulations in $\mathcal{T}_{g,n}^{lab}$. Finally, we denote $\mathcal{T}_{g,n}^{root}$ and $\mathcal{Q}_{g,n}^{root}$ the unlabeled triangulations and quadrangulations of genus g with n vertices and one rooted edge, that is distinguished and oriented.

Theorem 4.0.4 Disarlo-Parlier [DP19]

Both $\mathcal{T}_{g,n}$ and $\mathcal{T}_{g,n}^{lab}$ are connected. The diameter of $\mathcal{T}_{g,n}$ and $\mathcal{T}_{g,n}^{lab}$ are respectively $\Theta(g \log(g + 1) + n)$ and $\Theta(g \log(g + 1) + n \log(n))$.

The lower bound on the diameter follows from the techniques in [STT92].

Theorem 4.0.5 [Bud17, CS20]

The flip dynamics on $\mathcal{T}_{g,n}^{root}$ has mixing time $\Omega(n^{5/4})$. The flip dynamics on $\mathcal{Q}_{g,n}^{root}$ has mixing time $\Omega(n^{5/4})$ and $O(n^{13/2})$.

Note that our definitions of $\mathcal{T}_{g,n}^{root}$ and $\mathcal{Q}_{g,n}^{root}$ differ from [CS20, Bud17], in which n denotes the number of faces. However, by Euler’s formula, the number of faces is proportional to the number of vertices, up to an additive term which is linear in the genus. Hence, their results extend to our definitions of $\mathcal{T}_{g,n}^{root}$ and $\mathcal{Q}_{g,n}^{root}$.

Cylinder shears in square-tiled surfaces

A square-tiled surface is a quadrangulation of a connected surface with a 2-coloring of its edges such that all faces have an alternating boundary. One obtains a square-tiled surface geometrically by taking copies of a unit square and gluing horizontal edges to horizontal edges and vertical edges to vertical edges. We refer the reader to Subsubsection 4.2.1.1 for more detailed definitions. Square-tiled surfaces were introduced in the context of Abelian and quadratic differentials on Riemann surfaces by Zorich [Zor02]. In this whole chapter, we will actually consider a slight generalization of square-tiled surfaces where we allow *folded edges* which turns out to be crucial in our connectedness proof (see Subsubsection 4.2.1.1 for a precise definition).

A single diagonal flip operation does not preserve a square-tiled surface in general. Indeed, diagonal flips change the parity of the degree of the vertices incident to the flipped edge (see Figure 4.4), but the 2-coloring of the edges of the square-tiled surfaces imposes the vertices to have even degrees.




Figure 4.4: Diagonal flips do not preserve square-tiled surfaces.

A *horizontal cylinder* in a square-tiled surface is a cycle of squares adjacent along their vertical edges. A *horizontal cylinder shear* consists in doing all flips on the vertical edges in the cylinder (see Figure 4.5). Note that these flips commute. One defines similarly *vertical cylinders* and their shear. Cylinder shears hence preserve the set of square-tiled surfaces and the number of squares. We will see in Subsubsection 4.2.1.2, that cylinder shears present unexpected similarities with Kempe changes.





Figure 4.5: A horizontal cylinder shear.

Beyond the number of squares, it turns out that the degree of vertices is also invariant under cylinder shears. The *profile* of a square-tiled surface is a pair (μ, k) made of an integer partition $\mu = [1^{\mu_1} 2^{\mu_2} \dots]$ where μ_i is the number of vertices of degree $2\mu_i$ and k is the number of folded edges. Euler's formula relates the profile (μ, k) to the genus g of the underlying surface as follows,

$$\sum_{i \geq 1} \mu_i(i - 2) = 4g - 4 + k. \quad (4.1)$$

We let $ST(\mu, k)$ denote the set of square-tiled surfaces with given profile (μ, k) .

It turns out that $ST(\mu, k)$ might still be non-connected under cylinder shears. We introduce a first invariant of the connected components. Let τ be a square-tiled surface. We say that τ is *Abelian* if one can build this surface in such way that the top sides are only glued to the bottom sides, and the right sides to the left sides. We say that τ is *quadratic* otherwise. Let $ST_{Ab}(\mu, k)$ and $ST_{quad}(\mu, k)$ be respectively the set of Abelian and quadratic square-tiled surfaces in $ST(\mu, k)$. In Figure 4.6 we show an Abelian and a quadratic square-tiled surface in the stratum $ST([4^2], 0)$. Note also that the Quadratis puzzle on Figure 4.2 is an Abelian square-tiled surface in the stratum $ST_{Ab}([2^{21}, 6^1], 0)$. It is easy to see that the profile (μ, k) of an Abelian square-tiled surface is such that $k = 0$ and the entries in μ are even, that is $\mu_{2i+1} = 0$ for all $i \geq 0$.



(a) An Abelian square-tiled surface in $ST_{Ab}(4^2)$ (stratum $\mathcal{H}(1^2)$).

(b) A quadratic square-tiled surface in $ST_{quad}(4^2)$ (stratum $\mathcal{Q}(2^2)$).

Figure 4.6: Two square-tiled surfaces in $ST(4^2)$ (genus $g = 2$) that are not connected under cylinder shears.

The names Abelian and quadratic are borrowed from the theory of Riemann surfaces and their Abelian and quadratic differentials. It turns out that our main

conjecture for the connectedness under cylinder shears (Conjecture 4.0.7 below) is intimately linked to them. We introduce the necessary background now and refer the reader to Subsection 4.2.1 for more details. To each square-tiled surface, one can associate a Riemann surface endowed with a quadratic differential. It turns out that $ST_{Ab}(\mu)$ (respectively $ST_{quad}(\mu)$) corresponds to the quadratic differentials that are the squares of an Abelian differential (resp. not the square of an Abelian differential). The moduli space of the Abelian differentials and the quadratic differentials that are not the square of an Abelian one are stratified according to the degree of their zeros. The stratum of the Abelian differentials associated to the square-tiled surfaces in $ST_{Ab}(\mu)$ is denoted $\mathcal{H}(0^{\mu_2}1^{\mu_4} \dots)$ while the one associated to $ST_{quad}(\mu)$ is denoted $\mathcal{Q}(-1^{\mu_1+k}0^{\mu_2}1^{\mu_3} \dots)$. The strata are not necessarily connected but their connected components have been classified by M. Kontsevich and A. Zorich [KZ03] and by E. Lanneau [Lan08]. As a consequence of this classification we have that:

Theorem 4.0.6 [KZ03, Lan08]

1. Each Abelian stratum $\mathcal{H}(\kappa)$ has at most three connected components.
2. Each quadratic stratum $\mathcal{Q}(\kappa)$ has at most two connected components.
3. In genus 0, that is $\sum i\kappa_i = -4$, $\mathcal{Q}(\kappa)$ is non-empty and connected.

In this chapter we will focus on the genus 0 case and some specific connected components of $\mathcal{H}(\nu)$ in higher genera that are called *hyperelliptic components* (see Subsection 4.2.1). For now, let us just mention that for $\mu = [2^{\mu_2}, 4g - 2]$ or $\mu = [2^{\mu_2}, (2g)^2]$ where $g \geq 2$ and $\mu_2 \geq 0$, there is a subset $ST_{Ab}^{hyp}(\mu)$ of $ST_{Ab}(\mu)$, closed under cylinder shears, which correspond to the hyperelliptic connected component of $\mathcal{H}(0^{\mu_2}, 2g - 2)$ or $\mathcal{H}(0^{\mu_2}, (g - 1)^2)$ respectively.

With Vincent Delecroix, we conjecture that the connected components of the moduli space partition the square-tiled surfaces into equivalence classes for the shearing operation:

Conjecture 4.0.7 (Delecroix and Legrand-Duchesne, generalizing Conjecture 4 in [BLd+22])

Let μ be an integer partition of n and k a non-negative integer such that

$$\sum_{i \geq 1} \mu_i(i - 2) = 4g - 4 + k.$$

Let S and S' be two square-tiled surfaces in $ST(\mu, k)$. Then S and S' are equivalent via cylinder shears if and only if they belong to the same connected component of the moduli space of quadratic differentials.

According to Theorem 4.0.6, each $ST(\mu, k)$ would be made of at most 5 connected components under cylinder shears. Conjecture 4.0.7 is a direct generalization of Conjecture 4 in [BLd+22], which was supported by numerical evidence, to square-tiled surfaces with folded edges.

It is relatively straightforward to prove that being in the same connected component of the moduli space is a necessary condition:

Proposition 4.0.8

Let S be a square-tiled surface and S' obtained from S by a cylinder shear. Then the quadratic differentials associated to S and S' belong to the same connected component of the strata of the moduli space of quadratic differentials.

The underlying reason of Proposition 4.0.8 is that a cylinder shear can be realized on the moduli space as a continuous motion. Our main contribution is the two following theorems that provide a partial answer to Conjecture 4.0.7.

Theorem 4.0.9

Let $g \geq 1$ and $\mu = [2^{\mu_2}, 4g - 2]$ or $\mu = [2^{\mu_2}, (2g)^2]$. Then any two square-tiled surfaces in $ST_{Ab}^{hyp}(\mu)$ are connected by a sequence of at most $O(g)$ powers of cylinder shears.

Theorem 4.0.10

Let μ be an integer partition and k a positive integer such that $\sum \mu_i(i - 2) = k - 4$. Then $ST(\mu, k)$ is non-empty. Moreover, if $\mu_1 \leq 1$ and $(\mu, k) \neq ([1, 2^{\mu_2}, 3], 4)$, then $ST(\mu, k)$ has diameter $\Theta(k)$ with respect to powers of cylinder shears.

The proofs of these theorems are the object of Section 4.2.

4.1 Parameterized complexity of untying knots

4.1.1 A reorganization result

To obtain an efficient parameterized algorithm for UNKNOTTING VIA DEFECT, one needs to deal with (most of) the II^- moves in polynomial time, as they do not contribute to the defect, by Lemma 4.0.2. Therefore, the question we ask ourselves is: which II^- moves can be performed greedily, without increasing the defect of the diagram? In a preliminary attempt to proving Theorem 4.0.3, we showed that in a diagram of defect zero, no feasible II^- move results in a diagram of higher defect, thus diagrams of defect zero can be untangled greedily, by performing any feasible II^- at each step. However, in diagrams of higher defect, the answer is not that simple. For example, there exist diagrams in which no II^- and no I^- are feasible and in which the untying sequence of minimum defect starts by

increasing the number of crossings, by performing a II^+ move [KL12]. In the resulting diagram, the II^- move that removes the newly introduced crossings is feasible but is a step back and results in the original diagram, hence increases the defect and should not be performed. Thus, this proves that crossings that are affected by moves with non-zero weight in an untying sequence of minimal defect should not always be removed greedily with a II^- move. In fact, we prove that this restriction is essentially the only situation in which a II^- move should not be performed greedily.

Let D be a diagram of the unknot and \mathcal{S} be a feasible sequence of Reidemeister moves. We define the set S of *free crossings* with respect to \mathcal{S} as follows:

- Any crossing that is affected by a Reidemeister move of type III or I^- of \mathcal{S} is not free.
- Any crossing incident to an edge affected by a I^+ or a II^+ move of \mathcal{S} is not free.
- Among the set X of remaining crossings, any two crossings $x, y \in X$ that are removed together by a II^- move of \mathcal{S} are free.

In particular, if x is a crossing that is removed by a II^- move alongside with a crossing y affected by a move of non-zero weight in \mathcal{S} , then x is not free.

The core of our $W[P]$ -algorithm relies on the following reorganization result:

Theorem 4.1.1 (Theorem 9 in [9])

Let D be a diagram of a knot. Let $\mathcal{S} = (m_1, \dots, m_\ell)$ be a feasible sequence of Reidemeister moves for D . Assume that there are two free crossings x and z that are removed by \mathcal{S} and such that $\tilde{m} = \{x, z\}_{\text{II}^-}$ is feasible in D . Then there is a feasible sequence of defect k that starts with \tilde{m} and results in the same diagram as \mathcal{S} .

We will first explain how to obtain a $W[P]$ -algorithm using Theorem 4.1.1 before sketching its proof in Subsection 4.1.3.

4.1.2 $W[P]$ -algorithm

As we have seen, the freedom condition is crucial in Theorem 4.1.1. We spell out a naive (and wrong) algorithm anyway, that does not take the freedom condition into account, so that we can easily upgrade it afterward.

NAIVEGREEDY(D, k):

1. If $k < 0$, then output NO. If $D = U$ is a diagram without crossings and $k \geq 0$, then output YES. In all other cases continue to the next step.

2. If there is a feasible Reidemeister II^- move \tilde{m} , run $\text{NAIVEGREEDY}(D(\tilde{m}), k)$ otherwise continue to the next step.
3. (Non-deterministic step.) If there is no feasible Reidemeister II^- move, enumerate all possible Reidemeister moves m_1, \dots, m_t in D up to isotopy. Make a ‘guess’ which m_i is the first move to perform and run $\text{NAIVEGREEDY}(D(m_i), k - w(m_i))$.

This algorithm terminates as $k + |D|$ decreases at each recursive call of $\text{NAIVEGREEDY}(D, k)$. One can show that NAIVEGREEDY has $W[P]$ -running time, the key step being that the non-deterministic step is performed at most $k + 1$ times. However, not all feasible Reidemeister II^- moves can be performed in the second step without increasing the defect of the shortest untangling sequence, only those operating on free crossings.

The issue of Theorem 4.1.1 as stated is that the freedom of the crossings depends on a reconfiguration sequence, that is yet to be found by the algorithm. To circumvent this issue, we prove in [9] a different version of Theorem 4.1.1, that requires a more technical proof but can be directly applied to obtain an algorithm. In [9], the set S of non-free crossings, called *special*, is fixed in advance. Then S and \tilde{m} have to be feasible with respect to D but also to S : we require S to be the free crossings of S and \tilde{m} to operate on two free crossings. In [9], we call *greedy* these II^- moves that involve only free crossings. Finally, we prove that if a diagram admits an untangling sequence of defect k , then it contains at most $3k$ non-free crossings. Hence the predefined set S has size at most $3k$, which allows us to define the following $W[P]$ -algorithm and prove the $W[P]$ -membership part of Theorem 4.0.3:

$\text{SPECIALGREEDY}(D, k)$:

0. (Non-deterministic step.) Guess a set S of at most $3k$ crossings in D . Then run $\text{SPECIALGREEDY}(D, S, k)$.

$\text{SPECIALGREEDY}(D, S, k)$:

1. If $k < 0$, then output NO. If $D = U$ is a diagram without crossings and $k \geq 0$, then output YES. In all other cases continue to the next step.
2. If there is a feasible greedy Reidemeister II^- move m with respect to S , run $\text{SPECIALGREEDY}(D(m), S(m), k)$ otherwise continue to the next step.
3. (Non-deterministic step.) If there is no feasible greedy Reidemeister II^- move with respect to S , enumerate all possible special Reidemeister moves m_1, \dots, m_t in D with respect to S up to isotopy. If there is no such move, that is, if $t = 0$, then output NO. Otherwise, make a guess which m_i is performed first and run $\text{SPECIALGREEDY}(D(m_i), S(m_i), k - w(m_i))$.

4.1.3 Proof of Theorem 4.1.1

The goal of this section is to sketch the proof of Theorem 4.1.1. We start by two auxiliary results on rearranging consecutive moves.




Figure 4.7: Two feasible moves $m = \{x\}_{\text{I}^-}$ and $m' = \{y, z\}_{\text{II}^-}$ in a diagram D with disjoint private balls. After applying the two moves in any order, we get the same diagram.

To each feasible Reidemeister move m in a diagram D , we can assign its *private* 2-ball $B(m) \subseteq \mathbb{R}^2$ so that any modifications of D via m are performed only inside $B(m)$. Now let us assume that we have two feasible Reidemeister moves m and m' in a diagram D such that we can choose $B(m)$ and $B(m')$ so that they are disjoint. Considering the moves purely topologically, we can perform the two moves in any order yielding the same diagram (see Figure 4.7).

Applying this reasoning inductively results in the following lemmas, that give sufficient conditions to bring forward or postpone a II^- move on free crossings:

Lemma 4.1.2 (Corollary 11 in [9])

Let D be a diagram, $\mathcal{S} = (m_1, \dots, m_\ell)$ be a feasible sequence of Reidemeister moves of defect k . Assume that m_ℓ is a II^- move removing free crossings. Then there is a feasible sequence of Reidemeister moves $\widehat{\mathcal{S}} = (m_\ell, \widehat{m}_1, \dots, \widehat{m}_{\ell-1})$ such that

(i) $D(m_1, \dots, m_\ell) = D(m_\ell, \widehat{m}_1, \dots, \widehat{m}_{\ell-1})$,

(ii) $w(m_i) = w(\widehat{m}_i)$ for $i \in [\ell - 1]$,

(iii) the crossings affected by m_i and \widehat{m}_i are the same.

In particular $\widehat{\mathcal{S}}$ has defect k , produces the same diagram as \mathcal{S} , starts with m_ℓ and has the same set of free crossings as \mathcal{S} .

Lemma 4.1.3 (Corollary 13 in [9])

Let D be a diagram, $\ell \geq 2$ and $\mathcal{S} = (m_1, \dots, m_\ell)$ be a feasible sequence of moves for D . Assume that $m_1 = \{x_1, y_1\}_{\text{II}^-}$ and $m_\ell = \{x_\ell, y_\ell\}_{\text{II}^-}$ are II^- moves where x_1, y_1, x_ℓ, y_ℓ are free. Finally, assume also that $\{x_1, x_\ell\}_{\text{II}^-}$ is a feasible Reidemeister move for D . Then, there is a feasible sequence $\widehat{\mathcal{S}} = (\widehat{m}_2, \widehat{m}_3, \dots, \widehat{m}_{\ell-1}, m_1, m_\ell)$ of moves for D such that

$$(i) \ D(m_1, \dots, m_{\ell-1}, m_\ell) = D(\widehat{m}_2, \widehat{m}_3, \dots, \widehat{m}_{\ell-1}, m_1, m_\ell),$$

$$(ii) \ w(m_i) = w(\widehat{m}_i) \text{ for } i \in \{2, \dots, \ell - 1\},$$

(iii) the crossings affected by m_i and \widehat{m}_i are the same for $i \in \{2, \dots, \ell - 1\}$.

In particular, $\widehat{\mathcal{S}}$ has defect k , produces the same diagram as \mathcal{S} , ends with (m_1, m_ℓ) and has the same set of free crossings as \mathcal{S} .

Schematically, Lemmas 4.1.2 and 4.1.3 can be represented by the commutative diagrams drawn on Figure 4.8 and Figure 4.9 respectively.

$$\begin{array}{ccc} D & \xrightarrow{m_1, \dots, m_{\ell-1}} & D(m_1, \dots, m_{\ell-1}) \\ \downarrow m_\ell & & \downarrow m_\ell \\ D(\widehat{m}_\ell) & \xrightarrow{\widehat{m}_1, \dots, \widehat{m}_{\ell-1}} & D(m_1, \dots, m_\ell) \end{array}$$

Figure 4.8: Commutative diagram illustrating Lemma 4.1.2

To obtain an algorithm, one needs to work with a combinatorial description of the moves. This is much more cumbersome: to swap the order of m and m' , there are cases where one needs to change the combinatorial description of m for it to be feasible in $D(m')$ and yield the same diagram. This is the last difference with the statements of [9]: we consider here the moves purely topologically, whereas in [9], we give a combinatorial description of the move, which is necessary to encode them in an algorithm but presents the drawback of making the reconfiguration statements and proofs much more tedious. For this reason, we do not present the proofs of Lemmas 4.1.2 and 4.1.3 here.

Lemmas 4.1.2 and 4.1.3 change the order in which the moves are performed in a reconfiguration sequence, but not the set of crossings on which each move is




Figure 4.9: Commutative diagram illustrating Lemma 4.1.3

performed. To prove Theorem 4.1.1, we prove a final auxiliary lemma that allows us to rearrange consecutive Π^- moves on free crossings. The commutative diagram is drawn on Figure 4.10.




Figure 4.10: Commutative diagram illustrating Lemma 4.1.4


Lemma 4.1.4 (Lemma 14 in [9])

Let D be a diagram of a knot and let $m = \{x, y\}_{\Pi^-}$, $\tilde{m} = \{x, z\}_{\Pi^-}$ be two feasible Reidemeister moves in D where $y \neq z$. Assume also that $m' = \{w, z\}_{\Pi^-}$ is feasible in $D(m)$ with x, y, z, w mutually distinct. Then

- (i) $\tilde{m}' = \{w, y\}_{\Pi^-}$ is feasible in $D(\tilde{m})$; and
- (ii) $D(m, m') = D(\tilde{m}, \tilde{m}')$.

Proof. There are combinatorially only four options (up to a mirror image) for how x, y, z , and w can be arranged in the diagram, each of them yielding the required conclusion; see Figure 4.11. We give in [9] a direct argument as an alternative proof. □

We now have all the tools to prove Theorem 4.1.1, that we recall here for the reader's convenience.


 Figure 4.11: Four options how may x, y, z and w be arranged.

Theorem 4.1.1 (Theorem 9 in [9])

Let D be a diagram of a knot. Let $\mathcal{S} = (m_1, \dots, m_\ell)$ be a feasible sequence of Reidemeister moves for D . Assume that there are two free crossings x and z that are removed by \mathcal{S} and such that $\tilde{m} = \{x, z\}_{\text{II}^-}$ is feasible in D . Then there is a feasible sequence of defect k that starts with \tilde{m} and results in the same diagram as \mathcal{S} .

Sketch of proof of Theorem 4.1.1. There are two possible cases: either \tilde{m} appears in \mathcal{S} , or \tilde{m} removes two free crossings x and y , that are removed by different II^- moves, say $\{w, x\}_{\text{II}^-}$ and $\{y, z\}_{\text{II}^-}$.

First, let us assume that $\tilde{m} = m_j$ for some $j \in [\ell]$. If $j = 1$, there is nothing to prove, thus we may assume $j \geq 2$. Then, by using Lemma 4.1.2 on the sequence (m_1, \dots, m_j) , we get a sequence of moves $(m_j, \hat{m}_1, \dots, \hat{m}_{j-1})$ feasible for D , resulting in the same diagram. Hence, the sequence $(m_j, \hat{m}_1, \dots, \hat{m}_{j-1}, m_{j+1}, \dots, m_\ell)$ is also feasible in D and $D(\mathcal{S}) = D(m_j, \hat{m}_1, \dots, \hat{m}_{j-1}, m_{j+1}, \dots, m_\ell)$. By item (ii) of Lemma 4.1.2 and Lemma 4.0.2, the defect of this sequence equals k , and that is what we need.

Thus it remains to consider the case where $\tilde{m} \neq m_j$ for every $j \in [\ell]$. Because x and z are free with respect to \mathcal{S} , they are removed by some II^- moves $m_i = \{x, y\}_{\text{II}^-}$ and $m_j = \{w, z\}_{\text{II}^-}$. Since we assumed that x and z are free, we get that y and w are free as well. Without loss of generality, let us assume $i < j$.

Let $D' := D(m_1, \dots, m_{i-1})$. By Lemma 4.1.3, applied to D' and the sequence (m_i, \dots, m_j) feasible for D' , we get another sequence $(\hat{m}_{i+1}, \dots, \hat{m}_{j-1}, m_i)$ feasible for D' . By item (i) of Lemma 4.1.3, we get that $(\hat{m}_{i+1}, \dots, \hat{m}_{j-1}, m_i, m_j, \dots, m_\ell)$ is also feasible for D' .

Next, let $D'' := D'(\hat{m}_{i+1}, \dots, \hat{m}_{j-1})$. Because x and z are free, the edges between x and z are not used to perform I^+ or II^+ moves and the moves $m_1, \dots, m_{i-1}, \hat{m}_{i+1}, \dots, \hat{m}_{j-1}$ are performed outside of the private ball $B(\tilde{m})$. Hence \tilde{m} is fea-

sible in D'' . By Lemma 4.1.4 used in D'' , we get that $\tilde{m}' := \{w, y\}_{\Pi^-}$ is feasible in $D''(\tilde{m})$ and $D''(m_i, m_j) = D''(\tilde{m}, \tilde{m}')$. Altogether, by expanding D'' and D' , we get that

$$\tilde{\mathcal{S}} = (m_1, \dots, m_{i-1}, \hat{m}_{i+1}, \dots, \hat{m}_{j-1}, \tilde{m}, \tilde{m}', m_{j+1}, \dots, m_\ell) \quad (4.2)$$

is a feasible sequence of Reidemeister moves for D . We also get that the defect of $\tilde{\mathcal{S}}$ is equal to k by Lemma 4.0.2 and item (ii) of Lemma 4.1.3, when we used it. Note that all the moves m_i , m_j , \tilde{m} and \tilde{m}' are Π^- moves, thus they do not contribute to the weight.

The sequence $\tilde{\mathcal{S}}$ is not the desired sequence yet, because it does not start with \tilde{m} . However, it contains \tilde{m} , thus we can further modify this sequence to the desired sequence starting with \tilde{m} as in the first paragraph of this proof. \square

4.2 Reconfiguration of square-tiled surfaces

The goal of this section is to prove Theorems 4.0.9 and 4.0.10. We start by giving all definitions related to square-tiled surfaces and shears in Subsection 4.2.1. This subsection also contains background on Abelian and quadratic differentials and the definition of the weighted stable graph of a square-tiled surface, that encodes the structure of its cylinders. In Subsection 4.2.2, we focus on the square-tiled surfaces of genus 0 and show how they can be connected to some *path-like* square-tiled surface, which contains only one horizontal cylinder. Then in Subsection 4.2.3, we show that all these path-like configurations are related by cylinder shears, which concludes the proof of Theorem 4.0.10. Finally, we address in Subsection 4.2.4 the case of Abelian hyperelliptic square-tiled surfaces by reducing to square-tiled surfaces of genus zero, thereby proving Theorem 4.0.9.

4.2.1 Background and more


4.2.1.1 Square-tiled surfaces and their dual tricolored cubic graphs

A square-tiled surface is a quadrangulation of a connected surface whose edges are assigned a 2-coloring corresponding to the horizontal and vertical directions.

Let us introduce a natural encoding of a square-tiled surface that we will use as our main combinatorial definition. We fix an arbitrary labelling of the bottom-left and top-right corners of each square by the elements of $[n]$ where n is twice the number of squares. We define τ_R , τ_G , τ_B three involutions without fixed points on $[n]$, such that τ_R , τ_G and τ_B record respectively the adjacencies of the labels across the horizontal edges, inside a square and across the vertical edges. By drawing diagonals in each square we turn the square-tiled surface τ into a triangulation. In

this triangulation, the involution τ_G encodes adjacencies along the diagonal. We always use the following color convention for the edges (see Figure 4.12):

- horizontal edges (associated to τ_R) are red,
- diagonal edges (associated to τ_G) are green,
- vertical edges (associated to τ_B) are blue.



(a) A labelling of the square-tiled surface from Subfigure 4.6 (a).

(b) A labelling of the square-tiled surface from Subfigure 4.6 (b).

Figure 4.12: Labellings of the two square-tiled surfaces from Figure 4.6.

Conversely, starting from a triple of involutions $\tau = (\tau_R, \tau_G, \tau_B)$ such that they act transitively on $[n]$, one can build a square-tiled surface. If the action is not transitive, then the resulting surface is not connected.

We now introduce folded edges. Instead of using bicolored squares as building blocks we use tricolored triangles (that should be thought as “half” of a square cut along its diagonal). We pick n copies of them that we label from 1 to n . Then we consider three involutions (possibly with fixed points) τ_R , τ_G and τ_B that determine the gluings of the edges with colors respectively R , G and B . If j is a fixed point of τ_i then the edge colored i on the j -th triangle is glued to itself by a 180-degree rotation (see Figure 4.13). From now on, we will use the term square-tiled surface to refer to a triple of involutions (possibly with fixed points) that act transitively on $[n]$ and we will denote $S(\tau)$ the tricolored triangulation we have just constructed (see Subfigure 4.13 (a)). We call n the *number of triangles* in $S(\tau)$. A *red half-edge* (respectively *green half-edge* and *blue half-edge*) is a fixed point of τ_R (resp. τ_G and τ_B).


We denote $S^*(\tau)$ the *tricolored cubic graph* dual to the triangulation $S(\tau)$ (see Subfigure 4.13 (b)). The graph $S^*(\tau)$ has n vertices and is embedded in the square-tiled surface. We call *i-faces* of $S^*(\tau)$ or *3i-gons* the faces of $S^*(\tau)$ of degree $3i$. For example, there are two 2-faces or hexagons in Subfigure 4.13 (b). Note that

the i -faces are dual to the vertices of degree $3i$ in $S(\tau)$. The *profile* of τ is the pair (μ, k) , where μ is the integer partition counting the i -faces for any i , while k is the total number of half-edges. We use the standard notation $\mu = [1^{\mu_1}, 2^{\mu_2}, \dots]$ to denote 1 repeated μ_1 times, 2 repeated μ_2 times, etc. We denote $\text{ST}(\mu, k)$ the set of isomorphism classes of square-tiled surfaces with profile (μ, k) . Let us note that Euler's characteristic allows to compute the genus g of a square-tiled surface from its profile (μ, k) by

$$4g - 4 = \sum_i \mu_i \cdot (i - 2) - k. \tag{4.3}$$

Let us also note that the number of triangles n satisfies

$$n = \sum_{i \geq 1} i\mu_i. \tag{4.4}$$



(a) The triangulation $S(\tau)$ encoded by τ (b) Its dual tricolored cubic graph $S^*(\tau)$

Figure 4.13: A square-tiled surfaces of genus 0 with profile $([2, 3], 5)$ and its dual graph.

Given a triangulation $S(\tau)$ with n triangles and $\sigma \in S_n$ we denote $\tau^\sigma = (\sigma\tau_R\sigma^{-1}, \sigma\tau_G\sigma^{-1}, \sigma\tau_B\sigma^{-1})$. The triangulated square-tiled surface $S(\tau^\sigma)$ corresponds to the same underlying triangulation as $S(\tau)$ but where the labels of the triangles have been permuted by σ . We say that two square-tiled surfaces τ and τ' are *isomorphic* if they have the same number of triangles n and there exists a permutation $\sigma \in S_n$ such that $\tau' = \tau^\sigma$, in other words, $S(\tau)$ is isomorphic to $S(\tau')$.

A square-tiled surface τ is *Abelian* if there exists a non-trivial partition $A_+ \sqcup A_- = [n]$ such that $\tau_i(A_+) = A_-$ for each $i \in \{R, G, B\}$. Equivalently, it is Abelian if it has no half-edge and its dual graph $S^*(\tau)$ is bipartite. A square-tiled surface which is not Abelian is called *quadratic*. We denote $\text{ST}_{Ab}(\mu, k)$ and $\text{ST}_{quad}(\mu, k)$ respectively the Abelian and quadratic square-tiled surfaces of profile (μ, k) . By definition $\text{ST}(\mu, k)$ is the disjoint union of $\text{ST}_{Ab}(\mu, k)$ and $\text{ST}_{quad}(\mu, k)$.

We finish this introduction to square tiled-surfaces by giving some parity restrictions on the profile.

Lemma 4.2.1

Let τ be a square-tiled surface on $[n]$ with profile (μ, k) . Then we have

$$3n - k \equiv 2 \sum_{i \geq 1} \mu_{2i} \pmod{4}. \tag{4.5}$$

Proof. The faces of $S^*(\tau)$ are in bijection with the cycles of the product $\tau_R \tau_G \tau_B$ whose cycle type is μ . We claim that (4.5) follows from the signature morphism $\epsilon : (S_n, \cdot) \rightarrow (\mathbb{Z}/2\mathbb{Z}, +)$ which gives $\epsilon(\tau_R) + \epsilon(\tau_G) + \epsilon(\tau_B) = \epsilon(\mu)$. Indeed $\epsilon(\tau_i) = \frac{n-k_i}{2} \pmod{2}$ where k_i is the number of fixed points of τ_i , because τ_i is an involution. Hence $\epsilon(\tau_R) + \epsilon(\tau_G) + \epsilon(\tau_B) = \frac{3n-k}{2} \pmod{2}$. And we have $\epsilon(\mu) = \sum_{i \geq 1} \mu_{2i} \pmod{2}$ which concludes the proof of the lemma. \square

4.2.1.2 Cylinder shear

We now introduce formally the cylinder shears.

Let $\tau = (\tau_R, \tau_G, \tau_B)$ be a square-tiled surface. Let us consider a pair of colors $\{i, j\} \subset \{R, G, B\}$. A $\{i, j\}$ -component is a connected component of the subgraph of $S^*(\tau)$ in which we keep only edges of colors i and j . More combinatorially, these correspond to the orbit of the subgroup $\langle \tau_i, \tau_j \rangle$. Because the graph induced on $\{i, j\}$ has degree 2 and possibly half-edges, each $\{i, j\}$ -component is either a path or a cycle and we call them respectively $\{i, j\}$ -path or $\{i, j\}$ -cycle.

Let $c \subset [n]$ be a $\{R, G\}$ -component. Let $c_R := \tau_R|_c$ and $c_G := \tau_G|_c$ the restrictions of τ_R and τ_G to c (by definition τ_R and τ_G preserve c). The (R, G) -shear along c is the square-tiled surface $S_{c,R,G}(\tau) := (c_G c_R \tau_R, c_R c_G \tau_G, c_G \tau_B c_G)$. Note here that $c_R \tau_R$ and $c_G \tau_G$ are respectively the permutations τ_R and τ_G induced on the complement of the component c . In particular they commute with c_R and c_G because they have disjoint supports.

More graphically, the (R, G) -shear along c affects $S^*(\tau)$ by switching the colors R and G in c and sliding all adjacent edges colored B along the edge colored G (see Figure 4.14). Note that this comes down to performing a Kempe change on the 3-edge-coloring of $S^*(\tau)$, before sliding the incident edges to preserve the cyclic ordering of the colors around the faces.

The identity induces a natural bijection between the vertices of $S^*(\tau)$ and $S^*(S_{c,R,G}(\tau))$. Moreover, there is a natural bijection between the edges of $S^*(\tau)$ and $S^*(\tau') = S^*(S_{c,R,G}(\tau))$ that preserves the colors of all edges but those in the component c , which are exchanged. To ease the notations in its definition, we will use a slight abuse of the notation: $(ij) \in \sigma$ denotes that (ij) appears in the cycle decomposition of the involution σ and $(i) \in \sigma$ denotes that i is a fixed point of σ . For all i, j ,

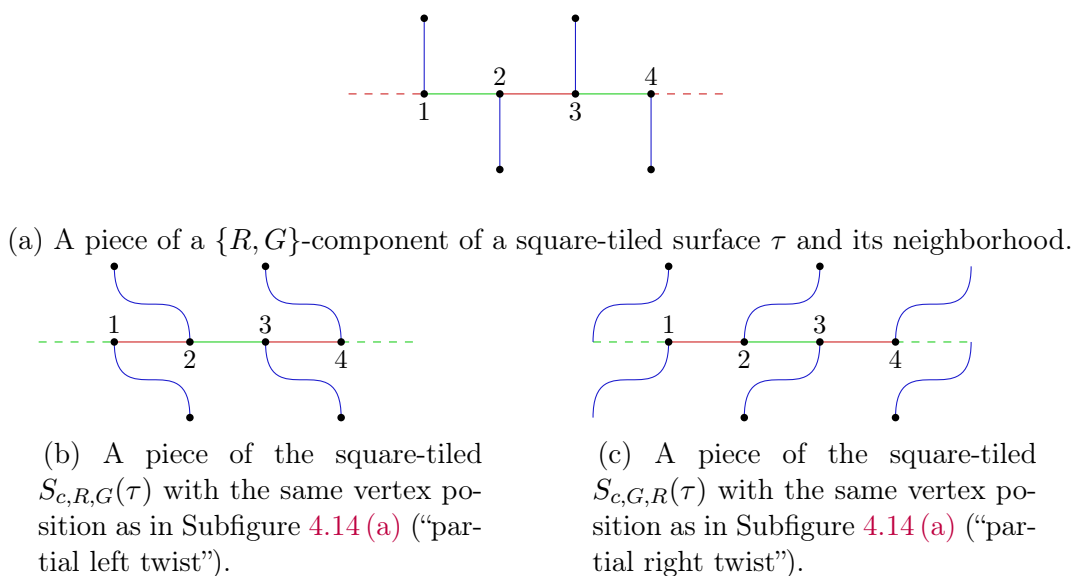


Figure 4.14: A piece of a $\{R, G\}$ -component and its neighborhood before and after a (R, G) -shear.




Figure 4.15: Two alternative labellings for left cylinder shear of a $\{R, G\}$ -cycle.

- the edge $(ij) \in \tau_R$ is mapped to $(ij) \in \tau'_R$ if $i, j \notin c$, or to $(ij) \in \tau'_G$ otherwise. Likewise, the half-edge $(i) \in \tau_R$ is mapped to $(i) \in \tau'_R$ if $i \notin c$ or to $(i) \in \tau'_G$ otherwise.
- the edge $(ij) \in \tau_G$ is mapped to $(ij) \in \tau'_G$ if $i, j \notin c$, or to $(ij) \in \tau'_R$ otherwise. Likewise, the half-edge $(i) \in \tau_G$ is mapped to $(i) \in \tau'_G$ if $i \notin c$ or to $(i) \in \tau'_R$ otherwise.
- the edge $(ij) \in \tau_B$ is mapped to $c_G(ij)c_G \in \tau'_B$, and the half-edge $(i) \in \tau_B$ is mapped to $c_G(i) \in \tau'_B$.

We define similarly (i, j) -shears for any pair of colors in $\{R, G, B\}$. Note that $S_{c,R,G}$ and $S_{c,G,R}$ are different operations (they differ on the image of τ_B) and we get 6 possible shears. We call *horizontal cylinders* the $\{G, B\}$ -components and *horizontal shears* the (B, G) and (G, B) -shears. Similarly, we call *vertical cylinders* the

$\{R, G\}$ -components and *vertical shears* the (R, G) and (G, R) -shears. Note that one could similarly define *diagonal shears* as the (R, B) and (B, R) -shears, however we will avoid them for the following reasons. The geometric interpretation of a diagonal cylinder on the quadrangulation corresponding to a square-tiled surface is unclear and there is no apparent reason to favor one diagonal over the other. Moreover, these shears can be realized by a sequence of horizontal and vertical shears whose length is linear in the size of the diagonal cylinder.

We now prove that shears preserve the stratum and stabilize the Abelian components.

Lemma 4.2.2

Let τ be a square-tiled surface, i, j two colors in $\{R, G, B\}$ and c a $\{i, j\}$ -component. The square-tiled surface $S_{c,i,j}(\tau)$ obtained after a cylinder shear in τ has the same profile as τ . Moreover if τ is Abelian then so is $S_{c,i,j}(\tau)$.

Proof. Recall that the profile of τ is the pair (μ, k) made of the integer partition μ associated to the conjugacy class of the product $\tau_R\tau_G\tau_B$ and the integer k which is the sum of the number of fixed points in τ_R, τ_G and τ_B .

Let $\tau' = (\tau'_R, \tau'_G, \tau'_B) = S_{c,R,G}(\tau)$ and (μ', k') its profile. By the commutations $[c_R, \tau_R] = [c_G, c_R\tau_R] = [c_G, \tau_G] = [c_R, c_G\tau_G] = 1$ we obtain

$$\tau'_R\tau'_G\tau'_B = (c_Gc_R\tau_R) \cdot (c_Rc_G\tau_G) \cdot (c_R\tau_Bc_R) = (c_Gc_R\tau_R) \cdot (c_G\tau_G) \cdot (\tau_Bc_R) = c_R\tau_R\tau_G\tau_Bc_R.$$

The latter is indeed conjugate to $\tau_R\tau_G\tau_B$ which shows that $\mu' = \mu$.

We also have $k' = k$ since the number of fixed points (k'_R, k'_G, k'_B) in $(\tau'_R, \tau'_G, \tau'_B)$ are respectively equal to

- the sum of the number of fixed points in $c_R\tau_R$ and c_G ,
- the sum of the number of fixed points in $c_G\tau_G$ and c_R ,
- the number of fixed points in τ_B .

Finally, if τ admits an invariant partition $A_+ \sqcup A_-$ then the same partition is invariant by τ' . □

Next, we immediately see from the definition that $S_{c,R,G}$ and $S_{c,G,R}$ are inverse of each other.

Lemma 4.2.3

Let τ be a tricolored cubic graph, i, j two colors in $\{R, G, B\}$ and c a $\{i, j\}$ -component. Then c is also a $\{i, j\}$ -component in $S_{c,i,j}(\tau)$ and $S_{c,j,i}(S_{c,i,j}(\tau)) = \tau$.

We extend the definition of the shear $S_{c,i,j}$ when c is a union of $\{i, j\}$ -components. In particular when $c = [n]$ we note $S_{i,j}$. The formula is simply

$$S_{G,R}(\tau) = (\tau_G, \tau_R, \tau_R\tau_B\tau_R)$$

Lemma 4.2.4

For any square-tiled surface $\tau = (\tau_R, \tau_G, \tau_B)$ we have

$$S_{G,R} \circ S_{B,G}(\tau) = (\tau_B, \tau_R, \tau_G)^{\tau_B}.$$

Proof. Indeed, we have

$$\begin{aligned} S_{G,R}(S_{B,G}(\tau)) &= S_{G,R}(\tau_B\tau_R\tau_B, \tau_G, \tau_B) \\ &= (\tau_B\tau_G\tau_B, \tau_B\tau_R\tau_B, \tau_B) \\ &= (\tau_G, \tau_R, \tau_B)^{\tau_B}. \end{aligned}$$

□

4.2.1.3 Abelian and quadratic differentials

We now introduce Abelian and quadratic differentials on Riemann surfaces and their associated moduli spaces. We explain how Abelian and quadratic square-tiled surfaces are particular cases of differentials and how a cylinder shear can be realized as a continuous path in the moduli space. The main consequence of this fact is the proof of Proposition 4.0.8. For a more detailed introduction to Abelian and quadratic differentials, we refer the reader to the survey [Zor06] and the book [AM24].

One often uses the term *translation surface* to refer to a compact Riemann surface endowed with a non-zero Abelian differential (or equivalently holomorphic one-form). The reason is that the most elementary way to define such object is by considering a finite collection of Euclidean polygons whose edges are identified by translations (see [Zor06, Section 1.2] and [AM24, Chapter 2]). The Abelian square-tiled surfaces we consider in this chapter are indeed particular cases of translation surfaces. More precisely, let $\tau = (\tau_R, \tau_G, \tau_B)$ be an Abelian translation surface on $[n]$ where n is twice the number of quadrilaterals. Let $A^+ \sqcup A^-$ be a τ -invariant bipartition of $[n]$. We consider $n/2$ copies of unit squares with bottom left labeled with elements of A^+ and from now on, we identify the squares with their labels in A^+ . Recall that $[n]$ corresponds to the triangles and that the “other half” of the bottom left triangle labeled $i \in A^+$ is labeled $\tau_G(i) \in A^-$ so that a square is made of two triangles. Now we glue the right side of the square labeled $i \in A^+$ to the left side of $\tau_R(\tau_G(i))$ and its top side to the bottom side of $\tau_B(\tau_G(i))$. Note that there is one ambiguity in the construction which is the choice of the bipartition. If

the partitions A^+ and A^- are swapped, we obtain a (possibly different) translation surface built from squares that are all rotated by 180 degrees. In other words, we associate to an Abelian square-tiled surface τ a translation surface up to rotation by 180 degrees that we denote by $\pm M$.

Two translation surfaces M and M' are *isomorphic* if one can obtain one from the other by a sequence of cutting and gluing operations on the polygons (see [AM24, Section 2.5.4]).

Lemma 4.2.5

Let τ and τ' be two Abelian square-tiled surfaces on $[n]$. Then τ and τ' are isomorphic (for the definition from Subsubsection 4.2.1.1) if and only if the associated translation surfaces $\pm M$ and $\pm M'$ are isomorphic.

Proof. Let us first remark that the number of triangles n of τ and τ' is twice the number of squares, which is the areas of M and M' .

It is clear that if τ and τ' are isomorphic then $\pm M$ and $\pm M'$ are isomorphic.

To prove the converse, one needs to show that there is “no choice” on where to put the squares in M . When τ has a profile with an entry different from 2 (that is $\mu_{2i} > 0$ for some $i > 1$) then the corresponding point on M must be a corner of a square. If the profile is $[2^{\mu_2}]$ then the associated translation surface is a torus and the position of the squares do not matter (since it is a group). \square

To a given translation surface, we associate the vector of positive integers $\kappa_1, \dots, \kappa_s$ such that the angle at the vertices in the surface are $(1 + \kappa_1)2\pi, \dots, (1 + \kappa_s)2\pi$. These numbers κ_i are also the order of vanishing of the Abelian differential and satisfy $\kappa_1 + \dots + \kappa_s = 2g - 2$ where g is the genus of the underlying surface. Given an Abelian square-tiled surface with profile $\mu = [2^{\mu_2}, 4^{\mu_4}, \dots]$, the vector associated to its translation surface is $\kappa = [1^{\mu_4}, 2^{\mu_6}, \dots]$.

The set of isomorphism classes of the translation surfaces with fixed vector κ form an algebraic variety, in particular a topological space which is denoted $\mathcal{H}(\kappa)$ (see [Zor06, Section 3.1] and [AM24, Chapter 3]). It follows that $\mathcal{H}(\kappa)$ has finitely many connected components, which are also path-connected. These connected components have been classified in [KZ03]. The topology of $\mathcal{H}(\kappa)$ is easy to describe: two surfaces are nearby if they are obtained by the same gluing patterns of nearby collections of polygons.

Lemma 4.2.6

Let M in $\mathcal{H}(\kappa)$ be the translation surface associated to an Abelian square-tiled surface τ . Let M' be the translation surface associated to the square-tiled surface $\tau' = S_{c,i,j}(\tau)$ obtained after a cylinder shear, where c is some $\{i, j\}$ -component of τ . Then M' belongs to $\mathcal{H}(\kappa)$ and there is a continuous path from M to M' in $\mathcal{H}(\kappa)$.

Proof. By Lemma 4.2.2, the profile of τ' is the same as the one of τ and τ' is Abelian. Hence the translation surface M' belongs to the same stratum $\mathcal{H}(\kappa)$ as M . Though, we will also obtain this fact from the explicit construction of a continuous path from M to M' .

Now, our goal is to build a continuous path of translation surfaces $\{M_t\}_{t \in [0,1]}$ inside $\mathcal{H}(\kappa)$ such that $M_0 = M$ and $M_1 = M'$. The $\{i, j\}$ -component c in τ corresponds to a horizontal cylinder in the translation surface M . More precisely, the closure of the subset of squares corresponding to this $\{i, j\}$ -component form a cylinder of height one. We define the surface M_t to be the translation surface built from an identical number of quadrilaterals where the quadrilaterals outside of c remains squares but the one in c are the “slanted” quadrilaterals with vertices $[(0,0), (1,0), (1+t,1), (t,0)]$. It is easy to notice that M_1 and M' are indeed isomorphic. \square

We now turn to quadratic differentials on Riemann surfaces which are sometimes referred as *half-translation surfaces*. One can build a half-translation surface by taking finitely many Euclidean polygons and gluing their edges using translations and 180-degree rotations (see [Zor06, Section 8.1] and [AM24, Section 2.7]). Similarly to the Abelian case, to a half-translation surface we associate a vector $\kappa = (\kappa_1, \dots, \kappa_s)$ of integers in $\{-1, 1, 2, 3, \dots\}$ such that the angle at the vertices are $(2 + \kappa_1)\pi, \dots, (2 + \kappa_s)\pi$. Euler’s formula in the quadratic case reads $\kappa_1 + \dots + \kappa_s = 4g - 4$. Similarly to the Abelian case, there exist strata $\mathcal{Q}(\kappa)$ that consist of quadratic strata that are not squares of Abelian differentials.

Given a quadratic square-tiled surface in $\text{ST}_{quad}(\mu, k)$ on $[n]$, one obtains a half-translation surface built from n right-angle triangles (which should be thought of as “half squares”). The vector associated to this half-translation surface is $\kappa = [(-1)^{k+\mu_1}, 1^{\mu_3}, 2^{\mu_4}, \dots]$. Note that both half-edges (corresponding to k) and vertices of degree 1 (corresponding to μ_1) gives rise to singularities of angle π .

Similarly to Lemma 4.2.6, we have the following.

Lemma 4.2.7

Let M in $\mathcal{Q}(\kappa)$ be the half-translation surface associated to a quadratic square-tiled surface τ . Let M' be the half-translation surface associated to the square-tiled surface $\tau' = S_{c,i,j}(\tau)$ obtained after a cylinder shear, where c is some $\{i, j\}$ -component of τ . Then M' belongs to $\mathcal{Q}(\kappa)$ and there is a continuous path from M to M' in $\mathcal{Q}(\kappa)$.

Proof. By Lemma 4.2.2, τ' is quadratic and its profile is the same as the one of τ . Hence, the half-translation surface M' belongs to the same stratum $\mathcal{Q}(\kappa)$ as M .

It is straightforward to extend the proof of Lemma 4.2.6 to the quadratic case when c is a cycle. In the case where c is a path, the union of triangles corresponding

to c in M is still a horizontal cylinder, though of height $1/2$. The same shearing construction works. \square

We can now prove Proposition 4.0.8, which states that the connected components induced by the cylinder shears refine the connected components of the strata of the moduli space of quadratic differentials.

Proof of Proposition 4.0.8. By Lemma 4.2.6 and Lemma 4.2.7 cylinder shears can be realized as a continuous path in respectively $\mathcal{H}(\kappa)$ and $\mathcal{Q}(\kappa)$. In particular, if τ belongs to a given connected component of $\mathcal{H}(\kappa)$ or $\mathcal{Q}(\kappa)$ its image under a cylinder shear belongs to the same connected component. \square

4.2.1.4 Cylinder decomposition and weighted stable graphs

We introduce now the weighted stable graph of a square-tiled surface that encodes the geometry of its $\{B, G\}$ -components. It follows closely the notions in [DGZZ20, Section 4.7] and [DGZZ21, Section 2.2]. The weighted stable graph will be the main player in Subsection 4.2.2.

Let τ be a square-tiled surface. Recall from Subsubsection 4.2.1.3 that when viewed as the gluing of right-angled triangles, $S(\tau)$ carries the geometry of an Abelian or quadratic differential. In such a surface, it makes sense to consider horizontal segments: these are continuous paths that are horizontal in every triangle. The *critical graph* of $S(\tau)$ is the graph embedded in the surface $S(\tau)$ whose vertices are the singularities of $S(\tau)$ (i.e. the union of vertices and mid-points of half-edges) and its edges are the union of all horizontal segments ending at these singularities. A connected component of the critical graph either consists entirely of red edges of $S(\tau)$ or is a saddle connection between two singularities of angle π . Such saddle connections are in bijection with the $\{B, G\}$ -paths in $S^*(\tau)$.

The *weighted stable graph* associated to τ is the multigraph $\Gamma(\tau)$ together with vertex decorations $(\mu^{(v)}, k^{(v)})$ for each vertex $v \in V(\Gamma(\tau))$ and edge decorations (w_e, h_e) for each edge $e \in E(\Gamma(\tau))$ built as follows.

- $V(\tau)$ is the set of connected components of the critical graph of $S(\tau)$; the vertex decoration $(\mu^{(v)}, k^{(v)})$ associated to a component records the singularity pattern of that component ignoring regular vertices (corresponding to μ_2),
- $E(\tau)$ is the set of cylinders, i.e. the connected components of $S(\tau)$ minus the critical graph, the two ends of an edge are the two critical graphs the boundary of the cylinder is glued to; the edge decoration $(w_e, h_e) \in \mathbb{Z}_{>0} \times \frac{1}{2}\mathbb{Z}_{>0}$ records the width and height of the cylinder. Note that this allows loops.




Figure 4.16: A square-tiled surface $S(\tau)$ and its weighted stable graphs $\Gamma(\tau)$.

By construction

- the union of the vertex decorations $(\mu^{(v)}, k^{(v)})$ is the *reduced profile* of τ , that is the profile μ of τ where the singularities corresponding to μ_2 are omitted.
- the number of triangles $n = \sum_{i \geq 0} \mu_i$ of $S(\tau)$ satisfies $\sum_{e \in E(\Gamma)} w_e \cdot h_e = \frac{n}{2}$.

Let us terminate this subsection with three remarks. First, we note that one can read the genus of $S(\tau)$ from $\Gamma(\tau)$. Namely, to each vertex $v \in V(\Gamma)$ one can associate the genus of the corresponding singular layer $g^{(v)}$ which is obtained from the vertex decoration and the degree of v in $\Gamma(\tau)$ as

$$4g^{(v)} - 4 = \sum_i \mu_i^{(v)} \cdot (i - 2) - k^{(v)} - 2 \deg(v).$$


Then the genus of the surface $S(\tau)$ is $\sum_{v \in \Gamma(\tau)} g^{(v)} + |E(\Gamma(\tau))| - |V(\Gamma(\tau))| + 1$. In particular, for a spherical square-tiled surfaces, $\Gamma(\tau)$ is always a tree.

Secondly, if $S(\tau)$ and $S(\tau')$ are two square-tiled surfaces that differ by some horizontal shears then their weighted stable graphs $\Gamma(\tau)$ and $\Gamma(\tau')$ are isomorphic.

Finally, if $S(\tau)$ is a square-tiled surface and c_1 and c_2 are two $\{B, G\}$ -components that correspond to the same edge in $\Gamma(\tau)$ then the associated cylinder shears produce isomorphic square-tiled surfaces (see Figure 4.17).

4.2.2 Connecting to path-like configurations

In this subsection, we focus on the square-tiled surfaces of genus 0, also called *spherical square-tiled surfaces*. These square-tiled surfaces are exactly those whose profile is *spherical*, i.e. satisfies $\sum_i (i - 2)\mu_i = k - 4$. We will assume that it is the case in this whole subsection.



(a) Two $\{B, G\}$ -cycles of length 6 (separated by three hexagons) whose shearing produces isomorphic square-tiled surfaces. (b) A $\{B, G\}$ -path of length 3 and a $\{B, G\}$ -cycle of length 6 whose shearing produces isomorphic square-tiled surfaces.

Figure 4.17: Equivalence of cylinder shears.

4.2.2.1 Path-like configurations

A square-tiled surface is called a *path-like configuration* if it has a single horizontal cylinder (that is, a $\{G, B\}$ -component) which furthermore is a path. The path-like square-tiled surfaces correspond to some specific Jenkins-Strebel differentials appearing in [Zor08] and to the one-cylinder square-tiled surfaces studied in [DGZ⁺20].

Note that if (μ, k) is the profile of a path-like configuration then necessarily $k \geq 2$ because of the two ends of the path. One important result of this subsection is that this condition is the only restriction on the profile for the existence of path-like square-tiled surfaces.

Theorem 4.2.8

Let (μ, k) be a spherical profile, with $k \geq 2$. Then there exists a path-like configuration with profile (μ, k) .

In order to prove Theorem 4.2.8 we construct a surjective map between path-like configurations and certain decorated plane trees. This map will be used in Subsection 4.2.3 to reconfigure path-like square-tiled surfaces.

A *decorated plane tree* is a pair (T, L) where T is a plane tree and L is a subset of the leaves of T . We say that (T, L) has *profile* (μ, ℓ) if L has size ℓ and the

vertices of T not in L have degrees given by the partition μ . The *perimeter* of a decorated plane tree is its number of edges adjacent to a leaf of L plus twice the number of other edges.

Let τ be a path-like square-tiled surface. We explain how to associate to τ its decorated tree. The vertex set of T contains the vertices of $S(\tau)$, that is the conical singularities of τ , together with the set L containing one vertex for each half-edge of $S(\tau)$. To obtain T , it suffices to keep only the red edges of $S(\tau)$, and to replace the red half-edges by an edge connecting its endpoint to a leaf of L . Note that (T, L) can also be defined by taking the dual of the red edges of $S^*(\tau)$. The main result of this subsection is the following relationship between path-like configurations and decorated plane trees.

Theorem 4.2.9

Let (μ, k) be a spherical profile with $k \geq 2$. Then the map which takes a path-like configuration with profile (μ, k) to its decorated tree is a surjection onto decorated trees with profile $(\mu, k - 2)$. Furthermore, for each decorated tree (T, L) with profile $(\mu, k - 2)$ the corresponding preimage of path-like configurations is made of a single orbit under the horizontal shears.

Proof. Let τ be a path-like configuration with profile (μ, k) . We first prove that its image is indeed a decorated tree. The only thing to prove is that T is indeed a tree. As τ is spherical and $S(\tau)$ has only one $\{B, G\}$ -component that is a path, hence not separating, the complement of T in the sphere is indeed a disk.

To obtain the surjectivity, we now explain how to build the tricolored cubic graph dual to a path-like configuration that maps to a given pair (T, L) with profile $(\mu, k - 2)$. We first consider a dual of the pair (T, L) on a disk as in Figure 4.18. Then in another disk, we put the $\{G, B\}$ -path on n vertices and put red arrows going to the boundary (see Figure 4.19), where n is the perimeter of (T, L) . Finally we make a sphere by gluing together the two disks.

We now prove that all the preimages of a decorated tree form a unique orbit under horizontal shears. Let $S^*(\tau)$ be a tricolored cubic graph dual to a given a preimage τ of (T, L) . We will mark (T, L) to track the position of the $\{G, B\}$ -path of $S^*(\tau)$ compared to the red edges. We associate a *marking site* to each endpoint of the red edges of $S^*(\tau)$ (each half-edge receives only one marking site instead of two). In (T, L) , this corresponds to placing a marking site on each edge of T adjacent to a leaf of L and one on each side of the other edges. We now need to distinguish two cases depending on the parity of the perimeter of (T, L) , that is the number of marking sites in (T, L) . Note that (T, L) has an even perimeter if and only if the half-edges at the endpoint of the $\{G, B\}$ -path of $S^*(\tau)$ (and all other preimages of (T, L)) have identical colors.

Assume that they have distinct colors, this corresponds to the case where (T, L) has odd perimeter. Mark the marking site adjacent to the unique green half-edge

of the $\{G, B\}$ -path of $S^*(\tau)$. The marking sites of (T, L) appear in a cyclic ordering on the boundary of the complement of T . Notice that performing two (G, B) -shear along the unique $\{G, B\}$ -path of $S^*(\tau)$ moves the mark to the next marking site in the anti-clockwise order. This proves that the preimages of (T, L) form a single orbit under the horizontal shears and that the number of such preimages is the perimeter of (T, L) .

Assume now that the perimeter is even, this corresponds to the case where the half-edges at the endpoints of the $\{G, B\}$ -path of $S^*(\tau)$ have identical colors. Up to performing one (G, B) -shear (which has no effect on (T, L)), we can assume that both are green. We place two marks on (T, L) , one next to each half-edge of the $\{G, B\}$ -path of $S^*(\tau)$. These marks are *antipodal*: the number of marking sites between them on each side of the boundary of the complement of T is half the perimeter of (T, L) . Like before, performing two (G, B) -shears along the unique $\{G, B\}$ -path of $S^*(\tau)$ moves both marks to the next marking site in the anti-clockwise order. Since all the markings corresponding to a preimage with two green half-edges are antipodal, this proves that these the preimages are in a single orbit under the horizontal shears. Moreover, the number of such preimages is the half perimeter of (T, L) . The same is true for preimages with two blue half-edges, and one goes from one case to the other with each (G, B) -shear. \square




Figure 4.18: A decorated plane tree (T, L) with profile $([4, 3^2, 1^2], 4)$ and its dual as a generalized arch configuration. We drawing a decorated tree, we will represent the leaves of L by a star, and the other leaves by triangular nodes.

Theorem 4.2.8 directly follows from Theorem 4.2.9 and the following lemma.




Figure 4.19: Matching the arch configuration with a $\{G, B\}$ -path.

Lemma 4.2.10

Let (μ, k) be a spherical profile with $k \geq 2$. Then there exists a decorated plane tree with profile $(\mu, k - 2)$.

Proof. It is a standard result that a tree on n vertices with a given degree sequence d_1, d_2, \dots, d_n exists if and only if $d_1 + d_2 + \dots + d_n = 2(n - 1)$. Now for a tree with profile $(\mu, k - 2)$ the number of vertices is $n = \sum \mu_i + k - 2$ and the degree sequence is $d = (1^{\mu_1+k-2}, 2^{\mu_2}, \dots)$. The planarity condition can be rewritten as

$$\sum_{i=1}^n d_i = k - 2 + \sum i\mu_i = 2 \sum \mu_i + 2k - 6 = 2(n - 1).$$

Thus there exists a tree with degree sequence d . In order to build a decorated plane tree, one just needs to pick any planar embedding and choose a subset of $k - 2$ leaves. \square

Note that the proof of Theorem 4.2.9 also provides a way of counting the number of path-like configurations in a stratum. For each decorated tree (T, L) in the stratum, the number of preimages, up to automorphism, is equal to the perimeter of (T, L) divided by its number of automorphisms.

Finally, the following proposition is an immediate consequence of the existence of path-like configurations.

Proposition 4.2.11

Let (μ, k) be a spherical profile, with $k \geq 2$. Then there are square-tiled surface with profile (μ, k) that are separated by $\Omega(k)$ cylinder shears.

Proof. A shear can only change by one the number of half-edges colored red in $S^*(\tau)$. By Theorem 4.2.8, there is a path-like square tiled-surface τ of profile (μ, k) and by symmetry, there is also a square tiled-surface τ' of profile (μ, k) that has a single vertical cylinder which furthermore is a path. Since τ and τ' have respectively k and at most two red half-edges, this concludes the proof. \square

4.2.2.2 Connecting to a path-like graph

We now state and prove a general result about reduction to path-like configurations of spherical square-tiled surface (see Lemma 4.2.12 below). Let us emphasize that it requires an assumption on the profile that is not satisfied in general by spherical square-tiled surfaces, namely $\mu_1 \leq 1$. It is a partial step towards Theorem 4.0.10.

Lemma 4.2.12

Let τ be a spherical square-tiled surface with profile (μ, k) , such that $\mu_1 \leq 1$. Then τ can be connected to a path-like square-tiled surface τ' with a sequence of $O(k)$ vertical cylinder shears and powers of horizontal cylinder shears.

The main technique to transform a spherical square-tiled surface into a path-like configuration is the fusion that we introduce now. Let τ be a square-tiled surface. A *fusion path* in τ is a vertical (or $\{R, G\}$) component of $S^*(\tau)$ that is a path and such that

- its intersection with each horizontal (or $\{B, G\}$) cycle is either empty or a single green edge,
- its intersection with each horizontal (or $\{B, G\}$) path is either empty or a single green half-edge.

Note that by definition a fusion path intersects at most two horizontal paths.

Lemma 4.2.13

Let τ be a square-tiled surface with n triangles and c be a fusion path in τ . Let U be the subset of $[n]$ that consists of elements which belong to vertical cylinders of τ which intersect c . Then in the square-tiled surfaces $S_{c,R,G}(\tau)$ and $S_{c,G,R}(\tau)$ the vertices of U form a single $\{B, G\}$ -component.

Let us note that Lemma 4.2.13 does not require planarity.

Proof. Let τ be a square-tiled surface and c a fusion path. We prove the case $S_{c,R,G}(\tau)$, the case of $S_{c,G,R}(\tau)$ being similar.

By relabelling τ , we assume that the vertices of c are labeled $1, 2, \dots, m$ in such way that $\tau_R(i)$ and $\tau_G(i)$ belong to $\{i-1, i, i+1\}$. In particular, the two half-edges of at the endpoints of c are adjacent to 1 and m .

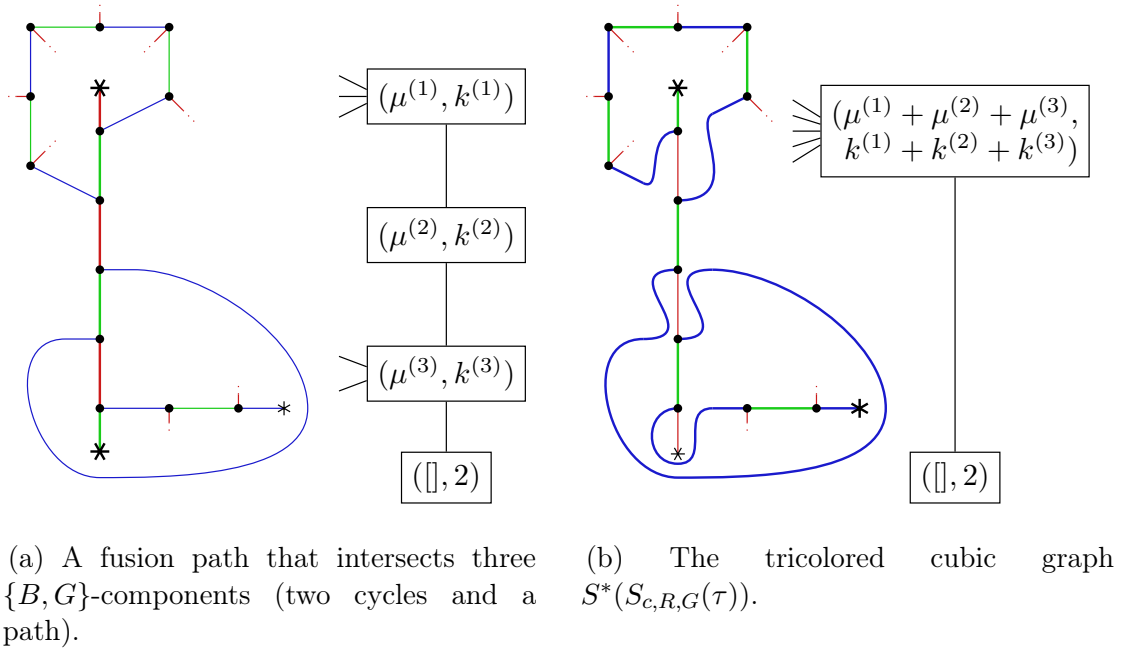


Figure 4.20: A vertical shear on a fusion path results in merging the horizontal cylinders it crosses. To the right of each square-tiled surface τ and $S_{c,R,G}(\tau)$ we draw the corresponding part of the stable graphs $\Gamma(\tau)$ and $\Gamma(S_{c,R,G}(\tau))$.

Each horizontal (or $\{B, G\}$) cycle that intersects c does it along an edge with vertices i and $\tau_G(i) = i + 1$. We let H denote the ordered list of vertices on the cycle starting at $\tau_G(i)$ and ending at i . We let H' denote the same list where the first and last elements $\tau_G(i)$ and i have been exchanged.

Let s be the number of $\{B, G\}$ -cycles intersecting the fusion path c (we have $m - 2 \leq 2s \leq m$). We label these cycles from 1 to s according to the ordering of the vertices of c induced by our choice of vertex labelling of c and denote H_1, H_2, \dots, H_s their associated ordered lists.

If the vertex 1 in c is adjacent to a green half-edge, then we denote P_1 the list of vertices of the $\{B, G\}$ -path containing 1 so that it ends with 1. If not, we let P_1 be the empty list. If the vertex m in c is adjacent to a green half-edge, then we denote P_m the list of vertices of the $\{B, G\}$ path containing m so that it starts with m . If not, we let P_m be the empty list.

The set U is the disjoint union of the elements in $P_1, P_m, H_1, \dots, H_m$. After a (R, G) -shear in c , the $\{B, G\}$ -path becomes the concatenation (in that order) of $P_1 \cdot H'_1 \cdot H'_2 \cdots H'_s \cdot P_m$. \square

Let us recall that in Subsubsection 4.2.1.4, we introduced the weighted stable graph $\Gamma(\tau)$ associated to a square-tiled surface τ . Let us note that τ is path-like

if and only if $\Gamma(\tau)$ is a graph made of two vertices, one of them having decoration $(\square, 2)$ and a unique edge between these two vertices that carries the decoration $(2n, \frac{1}{2})$.

We now reformulate Lemma 4.2.13 in terms of the weighted stable graph $\Gamma(\tau)$.

Lemma 4.2.14

Let τ be a spherical square-tiled surface and let c be a fusion path in τ . In $S^*(\tau)$, each green edge of c belongs to a unique $\{B, G\}$ -cycle and each green half-edge of c to a unique $\{B, G\}$ -path. These $\{B, G\}$ -components form a path c' in $\Gamma(\tau)$ that starts and ends at vertices with decorations $(\mu^{(start)}, k^{(start)})$ and $(\mu^{(end)}, k^{(end)})$ such that $k^{(start)} > 0$ and $k^{(end)} > 0$.

Let $\tau' = S_{c,R,G}(\tau)$. Then $\Gamma(\tau')$ is obtained from $\Gamma(\tau)$ by replacing all vertices in c' with two vertices u and v with decorations $(\mu^{(u)}, k^{(u)}) = (\square, 2)$ and $(\mu^{(v)}, k^{(v)}) = \sum_{i \in c'} (\mu^i, k^i) - (\square, 2)$ connected by a single edge and all neighbors of vertices in c in $\Gamma(\tau)$ get plugged to v .

Proof. The trace of c in $\Gamma(\tau)$ is a sequence of adjacent edges that do not repeat. Because $\Gamma(\tau)$ is a tree, this sequence of edges is an induced path. By definition of a fusion path, each endpoint of c is either a red half-edge or a green half-edge. In both cases, the associated vertex in $\Gamma(\tau)$ has a signature with a positive number k of half-edges.

The operation of fusion on $\Gamma(\tau)$ can be directly read from the description in Lemma 4.2.13. □

Finally, we state a lemma that allows us to find fusion paths in spherical square-tiled surfaces. It is a bit more general than what will be used to prove Lemma 4.2.12.

Lemma 4.2.15

Let τ be a spherical square-tiled surface made of n triangles. Let $\Gamma(\tau)$ be the associated weighted stable tree. We assume that for each leaf $v \in V(\Gamma)$ the associated partition $(\mu^{(v)}, k^{(v)})$ satisfies $k^{(v)} \geq 1$. Let $i \in [n]$ such that it is either adjacent to a red half-edge in τ (i.e. $\tau_R(i) = i$) or such that i and its red neighbor $\tau_R(i)$ belong to distinct horizontal (or $\{B, G\}$) components. Then, there exists a square-tiled surface τ' obtained from τ by doing horizontal cylinder shears such that the vertical component in τ' containing the image of this red edge or half-edge is a fusion path. Furthermore, τ' can be chosen so that this fusion path ends at a leaf of the weighted stable graph $\Gamma(\tau) = \Gamma(\tau')$.

Proof. Let us first consider the case of i being adjacent to a red half-edge, that is $i = \tau_R(i)$. We denote e_0 this red half-edge and our aim is to extend it to a fusion path. If the horizontal (or $\{B, G\}$) component of $S^*(\tau)$ containing i is a path,

then one can perform a horizontal cylinder shear on this path until e_0 becomes adjacent to a green half-edge. This $\{R, G\}$ -path (made of two half-edges) is a fusion path. Let us assume now that i belongs to a $\{B, G\}$ -cycle of $S^*(\tau)$ that we denote h_1 . By the Jordan Curve theorem, h_1 separates the sphere into two distinct connected components. If the connected component that does not contain e_0 is a terminal component, i.e. a leaf of $\Gamma(\tau)$, then it contains a red half-edge e_1 . Indeed, by assumption we have that $k^{(v)} > 0$ for every leaf $v \in \Gamma(\tau)$. In that case, one can apply a power of a horizontal cylinder shear on h_1 so that e_0 and e_1 become adjacent to a common green edge of h_1 . This terminates the construction in that situation. Now, if the connected component is not a leaf of $\Gamma(\tau)$ then there exists a red edge e_1 in that connected component connecting h_1 to a distinct $\{B, G\}$ -component h_2 . By performing a horizontal cylinder shear on h_1 one can make e_0 and e_1 adjacent to a common green edge. We can continue this construction starting from e_1 instead of e_0 and obtain a fusion path.

To handle the case of $i \neq \tau_R(i)$, we cut this red edge into two half-edges e_0 and e'_0 (the constructed square-tiled surface remains planar). The previous situation allows to build two fusion paths containing these two half-edges in some square-tiled surface τ' . Because τ' was obtained by performing only horizontal cylinder shears, e_0 and e'_0 belong to the same face of τ' and we can glue them back and obtain a fusion path in some square-tiled surface τ'' . It is easy to see that τ'' is obtained from τ . \square

We are now ready to prove Lemma 4.2.12.

Proof of Lemma 4.2.12. Let $S(\tau)$ be a planar square-tiled surface with profile (μ, k) with $\mu_1 \leq 1$. Assume that τ is not path-like. Then there is a red edge in $S^*(\tau)$ whose endpoints belong to two distinct $\{B, G\}$ -components. By Lemma 4.2.15, there exists a square-tiled surface τ' obtained from τ by a sequence of $\{B, G\}$ cylinder shears so that this red edge is part of a fusion path that ends in leaves of $\Gamma(\tau)$. By Lemma 4.2.14 performing a single vertical cylinder shear on this fusion path results in a square-tiled surface τ'' whose weighted stable graph $\Gamma(\tau'')$. One can continue this process until it reaches a path-like square-tiled surface.

Let us now count how many cylinder shears are performed. Each step involves some power of horizontal cylinder shears and a single vertical cylinder shear. We first derive a bound on the number of steps performed. Each step reduces the weighted stable graph by contracting a path between two leaves. The number of steps is hence bounded by the number of leaves, which is $O(k)$ (because each leaf of $\Gamma(\tau)$ corresponds to a component containing a half-edge). In particular, this bounds the number of vertical cylinder shears by $O(k)$. Let us now bound the number of powers of horizontal cylinder shears. At each step, this number is bounded by the length of the fusion path used at that step. But since the fusion

paths from different steps pass through different edges of the weighted stable graph their total number is at most the number of edges in $\Gamma(\tau)$ which also is $O(k)$. \square

4.2.3 Reconfiguring path-like configurations

We now show that any two path-like configurations can be reconfigured efficiently into one another. Namely we prove the following.

Lemma 4.2.16

Let (μ, k) be a spherical profile different from $([1^{\mu_1}, 2^{\mu_2}, 3], k)$ and with $\mu_1 \leq 1$. Then any pair of path-like square-tiled surfaces τ and τ' in $\text{ST}_{\text{quad}}(\mu)$ can be connected with a sequence of $O(k)$ vertical cylinder shears and powers of horizontal cylinder shears.

Throughout this subsection, all the square-tiled surfaces we consider will be spherical and path-like. Using Theorem 4.2.9, we work directly on decorated trees on which we will define a reconfiguration operation called *glue and cut* (see Subsubsection 4.2.3.1).

To show that two decorated trees are equivalent, a common technique in reconfiguration consists in showing that all decorated trees are equivalent to a specific one. The sketch of the proof is as follows. First, we show that all decorated trees are equivalent to a *nice* decorated tree: a decorated tree (T, L) with at most one leaf not in L and when it exists, this leaf is attached to a vertex of maximal degree (see Subsubsection 4.2.3.2). Second, we show that all nice decorated trees are equivalent up to a sequence of glue and cut operations to a nice decorated tree (T, L) where T is a caterpillar (see Subsubsection 4.2.3.3). Finally, we show that all nice decorated caterpillars are equivalent by reordering the vertices of degree at least three (see Subsubsection 4.2.3.3).


4.2.3.1 Glue and cut operation

A glue and cut operation on a decorated tree (T, L) can be decomposed in two steps. First, we delete two leaves of L and replace them by an edge between their parents. Then we delete an edge uv in the newly formed cycle and replace it by two pending leaves attached to u and v (see Figure 4.21).

Let $\mu = [1^{\mu_1}, 2^{\mu_2}, 3^{\mu_3} \dots]$ with no restriction on μ_1 , and let $k = \sum_i \mu_i(i-2) + 4$. The following lemma allows us to work directly on the decorated trees with glue and cut operations by decomposing glue and cut operations into cylinder shears.


Lemma 4.2.17

Let τ_1 and τ_2 be two spherical path-like configurations with profile $(\mu, k-2)$ such that the corresponding decorated trees (T_1, L_1) and (T_2, L_2) differ by one glue and



(a) Up to a power of a horizontal shear, the two half-edges impacted by the gluing are incident to a common green edge.

(b) The configuration obtained after the first vertical shear.



(c) The configuration obtained after the powers of horizontal shears on the two $\{G, B\}$ -paths P_1 and P_2

(d) Final configuration, obtained by cutting the desired edge by performing a vertical shear.

Figure 4.21: Glue and cut operation. The outer face is marked by a square node. The edges and leaves impacted by the glue and cut operation are thickened.

cut operation. Then τ_1 can be connected to τ_2 with a sequence of length $O(1)$ composed of vertical shears and powers of horizontal shears.

Proof. The reconfiguration sequence is illustrated by an example on Figure 4.21.

Recall that the leaves of L_i correspond to the red half-edges in τ_i . Let a and b denote the half-edges corresponding to the leaves removed by the glue and cut operation. Perform in τ_1 a power of the horizontal shear on the unique $\{G, B\}$ -path until the half-edges a and b are at distance 3 in the dual triangulation and separated by a green edge e (see Subfigure 4.21 (a)). Performing a vertical shear on (a, e, b) breaks the $\{G, B\}$ -path into two $\{G, B\}$ -paths P_1 and P_2 (see Subfigure 4.21 (b)).

Let f denote the edge of T_1 removed by the glue and cut operation. By performing a power of a horizontal shear on each of the $\{G, B\}$ -path P_i , we move the edge dual to f such that it connects the endpoints of P_i and is adjacent to a red half-edge h_i for each i (see Subfigure 4.21 (c)). We can then perform a vertical

shear on (h_1, f, h_2) , which merges P_1 and P_2 into a single $\{G, B\}$ -path, thereby obtaining a path-like configuration again (see Subfigure 4.21 (d)). This last shear replaces f by two pending leaves in the decorated tree, as desired. \square

4.2.3.2 Moving the conical singularities of degree π

In this subsection, we will present a new operation on decorated trees called *1-face transfer*, that given a decorated tree (T, L) , modifies L but not T . Recall that the leaves of T not in L correspond to conical singularities of angle π in the square-tiled surface and to faces of degree 3 in the tricolored dual cubic graph. We will describe how the glue and cut operation can be realized by combining the preceding operations and use it to obtain a nice decorated tree.

Let (T, L) be a decorated tree of profile $(\mu, k - 2)$ containing the two following subtrees: a leaf $u \in T \setminus L$, and v, w two leaves of L adjacent to a common vertex x and consecutive in the cyclic ordering of the vertices around x . Let L' be the set of leaves of T obtained by replacing v by u in L . We say that (T, L) and (T, L') differ by a *1-face transfer*.


Lemma 4.2.18

Let (T, L) and (T, L') be two decorated trees that differ by a 1-face transfer. Let τ and τ' be two path-like configurations that map to (T, L) and (T, L') respectively. The configurations τ and τ' differ by a sequence of operations composed of one vertical shear and $O(1)$ powers of horizontal shears.


Proof. The reconfiguration sequence is illustrated on the tricolored cubic dual graphs on an example drawn on Figure 4.22.

Let u be the leaf of $T \setminus L$ and v, w the two leaves of L impacted by the 1-face transfer. Let x be the common neighbor of v and w . Perform in $S^*(\tau)$ a power of the horizontal shear on the unique $\{G, B\}$ -path until the half-edge corresponding to w is adjacent to the face of degree 3 dual to u (see Subfigure 4.22 (a)). Perform a vertical shear on the $\{R, G\}$ -path composed of the half-edges dual to v and w and passing by the face of degree 3 dual to u . This separates the $\{G, B\}$ -path into two $\{G, B\}$ -paths P_1 and P_2 , adds v to L and replaces u and w by an edge e in the decorated tree (see Subfigures 4.22 (b) and 4.22 (c)). Perform on P_1 and P_2 a power of a horizontal shear, until the edge dual to e connects the extremities of both paths (see Subfigure 4.22 (d)). Perform a vertical shear on the edge dual to e and the incident green half-edges. This replaces e by two leaves u and w of L' and yields the desired decorated tree (T', L') . \square


Recall the following definition: A decorated tree (T, L) of profile $(\mu, k - 2)$ is *nice* if $\mu_1 = 0$ or if $\mu_1 = 1$ and the only leaf of T that does not belong to L is adjacent to a vertex of maximal degree.




(a) Up to a power of a horizontal shear, the half-edge corresponding to w is adjacent to the 1-face in the initial configuration




(b) After a vertical shear, the 1-face is moved and two half-edges are merged into an edge e



(c) Same configuration as Subfigure 4.22 (b)



(d) After a power of a horizontal shear, the edge e connect the extremities of the two $\{G, B\}$ -paths P_1, P_2



(e) After a vertical shear on e and the incident red half-edges, e is replaced by two leaves of L'

Figure 4.22: 1-Face Transfer on a decorated tree

The following technical lemma connects decorated trees to nice ones, but only in a subset of spherical profiles.

Lemma 4.2.19

Let (T, L) be a decorated tree of profile $(\mu, k - 2)$, such that $\sum \mu_i(i - 2) = k - 4$, with $\mu_1 \leq 1$ and $(\mu, k) \neq ([1, 2^{\mu_2}, 3], 4)$. There exists a sequence composed $O(1)$

1-face transfers and glue and cut operations that leads to a nice decorated tree of profile $(\mu, k - 2)$.

Proof. If $\mu_1 = 0$, then (T, L) is already nice. If $\mu_1 = 1$, then the assumption $(\mu, k) \neq ([1, 2^{\mu_2}, 3], 4)$ is equivalent to $k \geq 5$. The reconfiguration sequence is illustrated by an example on Figure 4.23.

There exists a vertex u such that three connected components of $T' - u$ intersect L . Let v, w, x be three leaves of L in distinct connected components of $T - u$, such that the connected components containing v and w appear consecutively around u (see Subfigure 4.23 (a)). Perform the glue and cut operation that replaces v and w by an edge and cuts one of the edges incident to u (see Subfigure 4.23 (a)). One of the two leaves y, z created by this operation is adjacent to u , say y (see Subfigure 4.23 (b)). Perform the glue and cut operation that replaces x and z by an edge and cuts the edge incident to u and consecutive to y (see Subfigure 4.23 (b)). In the decorated tree (T_f, L_f) obtained after this operation (see Subfigure 4.23 (d)), u is adjacent to two consecutive leaves of L_f and a 1-face transfer on t and these leaves results in a nice decorated tree. \square




Figure 4.23: Obtaining a nice decorated tree when $\mu_1 + k \geq 6$

4.2.3.3 Equivalence of nice decorated trees

Let $\mu = [1^{\mu_1}, 2^{\mu_2}, 3^{\mu_3}, \dots]$ and $k = \sum_i \mu_i(i - 2) + 4$, such that $\mu_1 \leq 1$. To prove Theorem 4.0.10, we only need to prove that all nice decorated trees with the profile $(\mu, k - 2)$ are equivalent up to a sequence of glue and cut operations. We will show that this is in fact true even in the strata $(\mu, k) = ([1, 2^{\mu_2}, 3], 4)$.

We will proceed in two steps: first proving that any nice decorated tree is equivalent up to a sequence of glue and cut operations to a nice decorated tree

(T, L) where T is a caterpillar, and then showing that all such decorated trees are equivalent up to a sequence of glue and cut operations.




Figure 4.24: Extending the longest path of a decorated tree. The dashed edge represents the gluing of leaves of L while the edge that is cut is crossed.

Lemma 4.2.20

Let (T, L) be a nice decorated tree of profile $(\mu, k - 2)$. There exists a sequence of at most k glue and cut operations and that connects (T, L) to a nice decorated tree (T', L') of with identical profile, where T' is a caterpillar.

Proof. The assumption $\mu_1 \leq 1$, guarantees that there is at most one leaf of T that does not belong to L . Let P be a maximal path containing the leaf that does not belong to L , if it exists. Let a and b be the two leaves at the extremities of P (see Figure 4.24), without loss of generality, we assume that $a \in L$. We prove that if T is not a caterpillar, then we can increase the length of P by performing one glue and cut operation.

Let c be a leaf at distance at least 2 from P . Let Q be the path from a to c and uv be the last edge of Q belonging to P . Perform the glue and cut operation that replaces the leaves a and c by an edge e and that removes the edge uv . It results in a longer path going from b to v via P , then to e via Q and finally to u via P . □

We call *spine* of a caterpillar its internal vertices, in the order they appear on the path going from one extremity to the other. Let (T^*, L^*) be the nice decorated tree of profile $(\mu, k - 2)$ such that T^* is the caterpillar whose spine (u_1, \dots, u_n) is a sequence of non-increasing degree, in which all leaves adjacent to u_2, \dots, u_n belong to L .

Lemma 4.2.21

Let (T, L) be a nice decorated trees with profile $(\mu, k - 2)$, such that T is a caterpillar. There exists a sequence of $O(k)$ glue and cut operations leading from (T, L) to (T^*, L^*) .

Proof of Lemma 4.2.21. Let (T, L) be a nice decorated tree of profile $(\mu, k - 2)$ with T a caterpillar, (u_1, \dots, u_n) its spine. If $k + \mu_1 = 4$, there is only one nice decorated tree of profile $(\mu, k - 2)$, so we are done. We now assume that $k + \mu_1 \geq 5$. We first prove a claim that allows us to perform some permutations of the spine. We will use it to sort the degree of the vertices on the spine and obtain (T^*, L^*) .

Claim 4.2.22

Assume that u_1 is adjacent to at least two leaves of L . Let u_i be a vertex of the spine. By performing at most three glue and cut operations, one can move u_i to obtain a caterpillar with spine $(u_i, u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n)$.


Proof. The procedure is illustrated by Figure 4.25. Denote a and b two leaves of L adjacent to u_1 . Note that u_n is adjacent to some leaf c of L : by assumption (T, L) is nice, so if $\deg(u_n) = 2 < \deg(u_1)$, u_n is adjacent to one leaf leaves of L ; and if $\deg(u_n) \geq 3$, then u_n is also adjacent to at least one leaf of L (see Subfigure 4.25 (a)).

Perform the glue and cut operation that replaces the leaves a and c by an edge and that cuts the edge $u_{i-1}u_i$ to replace it by a leaf d adjacent to u_{i-1} and a leaf e adjacent to u_i (see Subfigure 4.25 (a)). It produces a caterpillar with spine $(u_i, \dots, u_n, u_1 \dots u_{i-1})$ (see Subfigure 4.25 (b)). If $i = n$, we are done. Otherwise, perform a glue and cut operation that replaces the leaves b and e by an edge and that cuts the edge $u_i u_{i+1}$ to replace it by a leaf f adjacent to u_i and a leaf g adjacent to u_{i+1} (see Subfigure 4.25 (b)). It produces a tree that is not a caterpillar anymore, in which u_n and u_i are adjacent to u_1 (see Subfigure 4.25 (c)). Finally, perform the glue and cut operation that replaces the leaves c and g by an edge and cuts the edge $u_1 u_n$ (see Subfigure 4.25 (c)). It produces a caterpillar with the desired spine (see Subfigure 4.25 (d)).


□

We first ensure that the one of the extremities of the spin is of degree at least three. If that is not the case, each extremity is adjacent to a leaf of L , say a and b respectively. Since $k + \mu_1 \geq 5$, there is a internal vertex u that is adjacent to a leaf of L . Perform the glue and cut operation that replaces a and b by an edge and cuts an edge incident to u . It produces a caterpillar in which u is adjacent to two leaves of L and is placed at one extremity of the spin.


Let v_1, \dots, v_p be all the vertices of degree at least three in T , ordered such that for all i , $\deg(v_i) \leq \deg(v_{i+1})$ and for all $i < p$, the leaves adjacent to v_i all belong to L .



(a) Original caterpillar with spine (u_1, \dots, u_n) and u_1 adjacent to at least two leaves of L



(b) After one glue and cut operation, the spine becomes $(u_i, u_{i+1} \dots u_n, u_1, \dots, u_{i-1})$



(c) After one more glue and cut operation, the decorated is not a caterpillar anymore



(d) Final caterpillar with spine $(u_i, u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n)$

Figure 4.25: Moving u_i to the front of the caterpillar. Dashed edges represent the gluing of leaves of L , while the edge that is cut is crossed.

We can apply Claim 4.2.22 successively to each of the v_i . This results is the caterpillar (T^*, L^*) . \square

4.2.3.4 Proof of Lemma 4.2.16 and Theorem 4.0.10

Theorem 4.0.10 from the introduction follows directly from Lemma 4.2.12 that was proved in Subsection 4.2.2 and Lemma 4.2.16 that we prove now.

Proof of Lemma 4.2.16. Let τ be a path-like square-tiled spherical surface with profile (μ, k) and different from $([1^{\mu_1}, 2^{\mu_2}, 3], k)$. Let (T, L) be its associated decorated tree.

Let (T^*, L^*) be the decorated tree of profile $(\mu, k - 2)$ defined in the proof of Lemma 4.2.21. By Theorem 4.2.9, there is a path-like configuration τ^* with profile (μ, k) that maps to (T^*, L^*) .

If $\mu_1 = 1$, then by Lemma 4.2.19, (T, L) is connected to a nice decorated tree (T', L') using a sequence operations composed of $O(1)$ 1-face transfers and glue and cut operations. This is the only place where we need $(\mu, k) \neq ([1, 2^{\mu_2}, 3], 4)$.

By Theorem 4.2.9, let τ' be a path-like configuration with profile (μ, k) that maps to (T', L') . This reconfiguration between (T, L) and (T', L') translates into a reconfiguration sequence between τ and τ' composed of $O(1)$ powers of cylinder shears (by Lemmas 4.2.17 and 4.2.18).

If $\mu_1 = 0$, then $(T', L') := (T, L)$ is already nice.

By Lemmas 4.2.20 and 4.2.21, (T', L') can be connected to (T^*, L^*) via $O(k)$ glue and cut operations. Hence, τ and τ^* are equivalent up to $O(k)$ powers of cylinder shears by Lemma 4.2.17. This is true for all path-like square-tiled surfaces in $ST_{quad}(\mu, k)$, which concludes the proof of Lemma 4.2.16. \square

Combined with Lemma 4.2.12, this proves that any spherical square-tiled surfaces of profile (μ, k) are connected up to $O(k)$ powers of cylinder shears, if (μ, k) is different from $([1^{\mu_1}, 2^{\mu_2}, 3], k)$ and has $\mu_1 \leq 1$, thereby concluding the proof of Theorem 4.0.10.

4.2.4 Hyperelliptic square-tiled surfaces

4.2.4.1 Hyperelliptic components

Following [KZ03] we define the subsets $ST_{Ab}^{hyp}([2^{\mu_2}, (2g)^2])$ and $ST_{Ab}^{hyp}([2^{\mu_2}, 4g - 2])$ of respectively $ST_{Ab}([2^{\mu_2}, (2g)^2])$ and $ST_{Ab}([2^{\mu_2}, 4g - 2])$. These subsets correspond to the square-tiled surfaces that belong to the so-called hyperelliptic components of the moduli space of Abelian differentials. By Proposition 4.0.8 proven at the end of Subsubsection 4.2.1.3, these subsets are invariant under cylinder shears. However, we propose in this subsection an alternative combinatorial proof.

Let $\tau = (\tau_R, \tau_G, \tau_B) \in (S_n)^3$ be a square-tiled surface in some stratum $ST(\mu, k)$. An *automorphism* of τ is a permutation $\alpha \in S_n$ that normalizes simultaneously

τ_R, τ_G and τ_B that is : $\alpha\tau_i\alpha^{-1} = \tau_i$ for any color $i \in \{R, G, B\}$. We denote $\text{Aut}(\tau)$ the automorphism group of S .

Let H be a subgroup of $\text{Aut}(S)$. We define the quotient square-tiled surface S/H whose set of squares are the orbits of H on $[n]$ and $(\tau'_R, \tau'_G, \tau'_B)$ are the induced action on the orbits. Note that when taking the quotient, one needs to choose a labelling of the H -orbits in $[n]$ with $[n']$. A canonical choice consists in ordering these orbits according to the minimal element.

A square-tiled surface S of genus $g \geq 1$ is *hyperelliptic* if it admits an automorphism α which is of order 2 and such that the quotient $S/\langle\alpha\rangle$ has genus 0.

Let $g \geq 2$ and $\mu_2 \geq 0$. We define $\text{ST}_{Ab}^{hyp}([2^{\mu_2}, 4g - 2])$ to be the subset of $\text{ST}_{Ab}([2^{\mu_2}, 4g - 2])$ that are hyperelliptic and $\text{ST}_{Ab}^{hyp}([2^{\mu_2}, (2g)^2])$ to be the subset of $\text{ST}_{Ab}([2^{\mu_2}, (2g)^2])$ that are hyperelliptic and such that the hyperelliptic involution exchanges the two singularities of degree $2g$.

It is a standard result that a translation or half-translation surface of genus $g \geq 2$ admits at most one hyperelliptic involution. Furthermore, if it exists it is central in the automorphism group (i.e. all other automorphisms commute with the hyperelliptic involution). These two facts follow from the result that for a Riemann surface X , the induced action on homology gives rise to an injective map $\text{Aut}(X) \rightarrow \text{Sp}(H_1(X; \mathbb{Z}))$ and the image of the hyperelliptic involution is $-Id$, see [FM11, Theorem 6.8].

We warn the reader that for the profile $\mu = [2^{\mu_2}, (2g)^2]$ the condition that the hyperelliptic involution exchanges the two singularities of degree $2g$ is important. More precisely, there exists hyperelliptic square-tiled surfaces in $\text{ST}_{Ab}([2^{\mu_2}, (2g)^2])$ whose quotient belongs to a spherical square-tiled surface in some strata $\text{ST}_{quad}([1^{\mu'_1}, 2^{\mu'_2}, (g)^2], k)$, see Figure 4.26.

Lemma 4.2.23


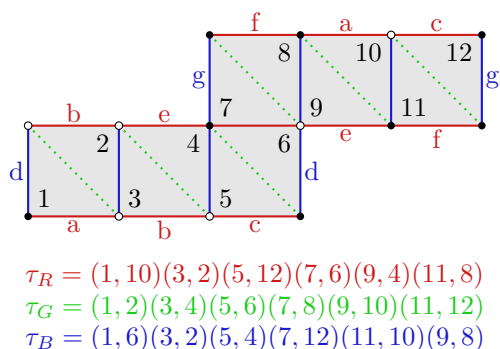
For $g \geq 2$ and $\mu_2 \geq 0$, the quotient by the hyperelliptic involution induces the bijections

$$\text{ST}_{Ab}^{hyp}([2^{\mu_2}, 4g - 2]) \rightarrow \bigcup_{\substack{k' + \mu'_1 = 2g + 1 \\ \mu'_1 + 2\mu'_2 = \mu_2}} \text{ST}_{quad}([1^{\mu'_1}, 2^{\mu'_2}, 2g - 1], k'). \quad (4.6)$$

and

$$\text{ST}_{Ab}^{hyp}([2^{\mu_2}, (2g)^2]) \rightarrow \bigcup_{\substack{k' + \mu'_1 = 2g + 2 \\ \mu'_1 + 2\mu'_2 = \mu_2}} \text{ST}_{quad}([1^{\mu'_1}, 2^{\mu'_2}, 2g], k'). \quad (4.7)$$

Proof. In the first case, the hyperelliptic involution preserves the singularity of degree $4g - 2$. The image of this singularity in the quotient by the hyperelliptic



(a) A surface in $ST_{Ab}([6^2])$ with hyperelliptic involution $\alpha = (1, 4)(2, 3)(5, 6)(7, 12)(8, 11)(9, 10)$.

(b) The quotient of Subfigure 4.26 (a) in $ST_{quad}([3^2], 6)$.

Figure 4.26: A hyperelliptic square-tiled surface in $ST_{Ab}([6^2])$ not in $ST_{Ab}^{hyp}([6^2])$: the quotient belongs to $ST_{quad}([3^2], 6)$.

involution is a singularity of degree $2g - 1$. By Euler's characteristic computation, one obtains that $k' + \mu'_1 = 2g + 1$.

In the second case, by definition, the hyperelliptic involution exchanges the two singularities of degree $2g$. The image of this pair of singularities in the quotient by the hyperelliptic involution is a singularity of degree $2g$. Again, by Euler's characteristic computation, one obtains that $k' + \mu'_1 = 2g + 2$.

This proves that the quotient by the hyperelliptic involution indeed belongs to the right hand sides of (4.6) and (4.7) respectively.

We now briefly explain why these maps are bijections. This follows from the standard construction that to any quadratic square-tiled surface, one can associate a unique Abelian square-tiled surface and an involution such that the quotient gives back the quadratic differential. \square

The following lemma is a consequence of a result of Kontsevich and Zorich [KZ03, Section 2.1]. We will provide a direct combinatorial proof using the half cylinder shears of the next subsection.

Lemma 4.2.24

Let $\mu = [2^{\mu_2}, 4g - 2]$ or $\mu = [2^{\mu_2}, (2g)^2]$ for some $g \geq 2$ and $\mu_2 \geq 0$. Then $ST_{Ab}^{hyp}(\mu)$ is preserved by cylinder shears.

Remark 4.2.25. Let us mention that for quadratic square-tiled surfaces, there also exist hyperelliptic connected components, see [Lan08]. They correspond exactly to quadratic square-tiled surfaces in genus $g \geq 1$ that are hyperelliptic and whose quotient belongs to $ST_{quad}(1^{\mu_1}, 2^{\mu_2}, a, b)$ where $a, b \geq 2$. That is, instead of having a single singularity $d \geq 3$ as in Lemma 4.2.23 on the sphere, there are two.

However, we were not able to apply the techniques we develop in this chapter in order to show the connectedness of $\text{ST}_{quad}^{hyp}(\mu)$.

4.2.4.2 Half cylinder shears

At first glance, Lemma 4.2.23 and Conjecture 4.0.7 seem in contradiction. Namely, from (4.6), the quotient by the hyperelliptic involution maps the subset of square-tiled surfaces $\text{ST}_{Ab}^{hyp}([2^{\mu_2}, 4g - 2])$ to a disjoint union of $\text{ST}_{quad}([1^{\mu'_1}, 2^{\mu'_2}, 2g - 1], k')$. And Conjecture 4.0.7 asserts that the left side of (4.6) is connected under cylinder shears, while the right side is obviously not by Lemma 4.2.2. The reason is that if τ is a square-tiled surface in $\text{ST}_{Ab}^{hyp}([2^{\mu_2}, 4g - 2])$ and τ' its quotient by the hyperelliptic involution then a cylinder shear in τ does not necessarily correspond to a cylinder shear in τ' . One sometimes obtains what we call a *half cylinder shear* that modifies the profile and which we study now.

Let τ be a square-tiled surface and c a $\{B, G\}$ -cycle that separates the square-tiled surface such that one of the connected components of the complement of c contains only hexagons, 1-faces and red half-edges. In more a combinatorial way, let $c = c_1 \sqcup c_2$ the decomposition into two $\tau_B \circ \tau_G$ -orbits. We ask that c_1 is stable under τ_R and that the faces bounded by c together with the red edges with ends in c_1 are only half-edges, 1-faces and hexagons. The *half cylinder shear* along c is obtained by changing τ_R along c_1 as follows

$$\tau'_R(i) = \begin{cases} \tau_G \circ \tau_B \circ \tau_R(i) & \text{if } i \in c_1 \\ \tau_R(i) & \text{otherwise.} \end{cases}$$

Similarly to the case of the cylinder shears, half cylinder shears induce a natural bijection on the vertices of $S^*(\tau)$, that we will not define here.




Figure 4.27: A (left) cylinder shear as the composition of two (left) half cylinder shears.

Lemma 4.2.26

Let τ be a square-tiled surface of genus $g \geq 2$ in either $\text{ST}_{Ab}^{hyp}([2^{\mu_2}, 4g - 2])$ or

$\text{ST}_{Ab}^{\text{hyp}}([2^{\mu_2}, (2g)^2])$. Let τ' be its quotient by the hyperelliptic involution. Let c be a $\{B, G\}$ -cycle in τ and c' its image in τ' .

If c' belongs to a horizontal cylinder bounded by a $\{B, G\}$ -path c'' (equivalently if the height h in the edge decoration (w, h) of the weighted stable graph $\Gamma(\tau')$ is a half integer) then the quotient of $S_{c, B, G}(\tau)$ by the hyperelliptic involution is isomorphic to $S_{c'', B, G}(\tau')$.

If c' belongs to a horizontal cylinder bounded by a $\{B, G\}$ -cycle c'' isolating a component of $S^*(\tau')$ made only of red half-edges, triangles and hexagons, then the quotient of $S_{c, B, G}(\tau)$ by the hyperelliptic involution is isomorphic to the half cylinder shear along c'' in τ' .

Proof. We only sketch the proof, which is straightforward.

Each cylinder of $S(\tau)$ is preserved by the hyperelliptic involution. The first case corresponds to the situation where the height of the cylinder containing the curve c has odd height. In that case, the circumference in the middle of the cylinder is mapped to a component of the critical graph that corresponds to a $\{B, G\}$ -path in $S^*(\tau)$. Performing a cylinder shear in c is the same as performing a cylinder shear along that $\{B, G\}$ -path.

In the second case, when the height is even, the circumference is mapped to a component of the critical graph made of red edges. In that case, the cylinder shear in c becomes a half cylinder shear in the quotient. \square

4.2.4.3 Path-like connection and 1-face edition

We continue to work on spherical square-tiled surface. We now describe how to circumvent the restriction $\mu_1 \leq 1$ on the profile when connecting to a path-like configuration as in Subsection 4.2.2 or when connecting two path-like configurations as in Subsection 4.2.3. The procedure uses the half cylinder shears introduced just before.

The analogue of Lemma 4.2.12 we prove in that context is the following.

Lemma 4.2.27

Let τ be a spherical square-tiled surface with profile $([1^{\mu_1}, 2^{\mu_2}, d], k)$ where $d \geq 3$. Then τ can be connected with a sequence of $O(k + \mu_1)$ powers of cylinder shears and $O(\mu_1)$ half cylinder shears to a path-like square-tiled surface.

Proof. Let τ be a spherical square-tiled surface with profile $([1^{\mu_1}, 2^{\mu_2}, d], k)$. The weighted stable tree $\Gamma(\tau)$ is a star whose center is the unique vertex v_0 whose decoration μ^{v_0} contains the element d of the profile. The other vertices, that necessarily are leaves, have decorations either equal to $([], 2)$, or $([1], 1)$ or $([2], 0)$. Using half horizontal cylinder shears, one can transform the leaves with decoration $([2], 0)$ into leaves with decorations $([], 2)$. If v_0 is not a leaf, then we are in the

setting of Lemma 4.2.15 and there exists a fusion path. Performing a vertical cylinder shear on that fusion path reduces the degree of the special vertex v_0 in $\Gamma(\tau)$.

Now it remains to treat the case when the degree of v_0 is one. In that case, $\Gamma(\tau)$ has a single other vertex v_1 with decoration either $(\square, 2)$, or $([1], 1)$ or $([2], 0)$. If v_0 is such that its decoration satisfies $k^{(v_0)} > 0$, then up to perform a half cylinder shear in v_1 , one can also use Lemma 4.2.15 to find a fusion path. We now further assume that $k(v_0) = 0$. Up to a half cylinder shear, one can assume that $\mu_1^{(v_1)} > 0$. It is easy to see that there exists a $\{R, G\}$ cycle c that “connects” a triangle in v_0 to a triangle in v_1 in the following sense: c crosses every $\{B, G\}$ components exactly twice and one of the connected component of the complement of c consists only of hexagons and two triangles. Performing a $\{R, G\}$ half cylinder shear in c results in a path like configuration (see Figure 4.28). \square




Figure 4.28: A $\{R, G\}$ -cycle in $S^*(\tau)$ as used in the proof of Lemma 4.2.27. The shaded region is made of two 1-faces (or triangles) and one 2-face (or hexagon).

We now describe an operation on decorated trees that we call *1-face edition*, that given a decorated tree (T, L) , adds or removes leaves of L , while not changing T much. We will show that this operation can be obtained as a sequence of half-shears and glue and cut operations. Finally, we will use it for two purposes: reducing the number of 1-faces once we have attained a path-like configuration to obtain $\mu_1 \leq 1$ and apply Theorem 4.0.10, and handling the strata $[1, 2^{\mu_2}, 3], 4)$ that is not covered by Theorem 4.0.10.

Let (T, L) and (T', L') be two decorated trees, and u and v be two leaves of $T \setminus L$, such that attaching a leaf w to v results in T' and $L' = L \cup \{u, w\}$. We say that (T, L) and (T', L') differ by a *1-face edition*. In particular, the profiles $(\mu, k - 2)$

and $(\mu', k' - 2)$ of (T, L) and (T', L') respectively are such that $\forall i > 2, \mu'_i = \mu_i, \mu'_2 = \mu_2 + 1, \mu'_1 = \mu_1 - 2,$ and $k' = k + 2.$

Lemma 4.2.28

Let (T, L) and (T', L') be two decorated trees that differ by a 1-face edition. Let τ and τ' be two path-like configurations that map to (T, L) and (T', L') respectively. The configurations τ and τ' differ by a sequence of operations composed of one half-shear, one vertical shear and $O(1)$ powers of horizontal shears.

Proof. The reconfiguration sequence is illustrated by an example on Figure 4.29.

Let u and v be the leaves of $T \setminus L$ impacted by the 1-face edition (see Subfigure 4.29 (a)). Perform in τ a power of the horizontal shear on the unique $\{G, B\}$ -path until the 1-faces dual to u and v share a green edge e (see Subfigure 4.29 (b)). Let e' be a red edge in one of the two 1-faces u and v . Perform a half-shear on these two 1-faces, this replaces e by two half-edges and thus cuts the $\{G, B\}$ -path in two $\{G, B\}$ -paths P_1 and P_2 . This also changes the profile by trading two 1-faces for two half-edges and a hexagon (see Subfigure 4.29 (c)). Let w be the vertex corresponding to this hexagon in the decorated tree. Perform a power of a horizontal shear on P_1 and P_2 until e' connects the extremities of both paths P_1 and P_2 (Subfigure 4.29 (d)). Perform a vertical shear on e' and the incident green half-edges to obtain again a single $\{G, B\}$ -path. This replaces the edge dual to e' by two leaves of L' (see (see Subfigure 4.29 (e)) and yields the desired decorated tree (T', L') . \square

4.2.4.4 Reducing to the spherical case: proof of Theorem 4.0.9


Lemma 4.2.29 (Generalization of Lemma 4.2.19 using half-shears))

Let (T, L) be a decorated tree of profile $(\mu, k - 2)$, such that $\sum \mu_i(i - 2) = k - 4.$ There exists a sequence composed of $O(\mu_1)$ 1-face editions, $O(1)$ 1-face transfers and glue and cut operations that leads to a nice decorated tree.


Let M be a maximum matching of the leaves in $T \setminus L$. For every pair $(u, v) \in M$, perform a deletion on u and v . This results in a decorated tree (T', L') of profile $(\mu', k' - 2)$, such that $\mu'_1 \leq 1$ and $\mu'_2 = \mu_2 + \lfloor \mu_1/2 \rfloor,$ for all $i > 2, \mu'_i = \mu_i,$ and $k' = k + 2\lfloor \mu_1/2 \rfloor.$

If μ_1 was even, then (T', L') is nice. However if μ_1 was odd, there is no guarantee that the remaining leaf t that does not belong to L' is adjacent to the vertex u of highest degree. In this case, we only need to place two leaves of L consecutively adjacent to u , before applying a 1-face transfer to obtain a nice decorated tree. If $\mu_1 + k \geq 6,$ then we conclude by Lemma 4.2.19.


The remaining case is $\mu_1 + k = 5,$ where T' has three leaves and $|L'| = 2.$ Moreover $\mu_3 = 1$ and $\mu_i = 0$ for $i \geq 4.$ Let u be the vertex of degree 3 and v




(a) Starting configuration.




(b) After a power of a horizontal shear the two 1-faces u and v are adjacent and separated by a green edge e .



(c) After a half-shear, e is replaced by two green half-edges. This impacts the profile.



(d) After a power of a horizontal shear, a red edge e' originally part of one of the two 1-faces is now adjacent to the extremities of both $\{G, B\}$ -paths.



(e) Final configuration, obtained by performing the vertical shear containing e' .

Figure 4.29: A 1-face edition on a decorated tree. The outer face is marked by a square node. The edges and leaves impacted by the 1-face edition are thickened.

and w be the leaves of L (see Subfigure 4.30 (a)). If u is not adjacent to both v and w , then perform the glue and cut operation that replaces v and w by an edge and that cuts one of the edges at distance 1 from u . This results in a decorated tree (T'', L'') with $L'' = \{x, y\}$, in which there is a vertex z of degree two that is adjacent to both u and a leaf x (see Subfigure 4.30 (b)). Thus we can perform a 1-face edition to remove x from T'' and y from L'' (see Subfigure 4.30 (c)). Another

1-face edition attaches a leaf s to t and adds s and y to L'' (see Subfigure 4.30 (d)). In this final decorated tree (T_f, L_f) , the only leaf not in L_f is z , which is adjacent to u , the vertex of highest degree. Hence, (T_f, L_f) is nice.




Figure 4.30: Obtaining a nice decorated tree when $\mu_1 + k = 5$

Proof of Theorem 4.0.9. Let us consider a square-tiled surface in either $ST_{Ab}^{hyp}([2^{\mu_2}, (2g)^2])$ or $ST_{Ab}^{hyp}([2^{\mu_2}, 4g - 2])$ for some $\mu_2 \geq 0$. By Lemma 4.2.23, its quotient by the hyperelliptic involution belongs to respectively either a stratum $ST_{quad}([1^{\mu'_1}, 2^{\mu'_2}, 2g - 1], k')$ or $ST_{quad}([1^{\mu'_1}, 2^{\mu'_2}, 2g], k')$ for some $\mu'_1, \mu'_2, k' \geq 0$. By Lemma 4.2.26, the cylinder shears in τ correspond to cylinder shears or half cylinder shears in the quotients. Hence, it is enough to connect $\bigcup_{\substack{k'+\mu'_1=2g+1 \\ \mu'_1+2\mu'_2=\mu_2}} ST_{quad}([1^{\mu'_1}, 2^{\mu'_2}, 2g - 1], k')$ or $\bigcup_{\substack{k'+\mu'_1=2g+2 \\ \mu'_1+2\mu'_2=\mu_2}} ST_{quad}([1^{\mu'_1}, 2^{\mu'_2}, 2g], k')$ using both cylinder shears and half cylinder shears.

By Theorem 4.0.10 we know that cylinder shears are enough to connect the spherical strata when $\mu'_1 = 0$ and $\mu'_1 = 1$ but the special case of $(\mu', k) = ([1, 2^{\mu'_2}, 3], 4)$ that we will treat separately. Let us first note that the parity of μ'_1 is fixed by the profile (μ, k) by Lemma 4.2.1. Hence, depending on (μ, k) only one of the quotient stratum with $\mu'_1 \leq 1$ is non-empty. What remains to do is to connect spherical square-tiled surface to these base cases and to treat $(\mu', k) = ([1^{\mu'_1}, 2^{\mu'_2}, 3], 5 - \mu'_1)$ with μ'_1 odd.

To do so, we first apply Lemma 4.2.27 that allows us to reach a path-like configuration $S^*(\tau)$.

Next, in the case of the non-exceptional strata, if there is more than one 1-face in $S^*(\tau)$, then Lemma 4.2.28 allows us to remove the extra 1-faces and connect to the base cases.

In the case of a path-like configuration whose profile is $(\mu', k) = ([1^{\mu'_1}, 2^{\mu'_2}, 3], 5 - \mu'_1)$ the only place in the proof Lemma 4.2.16 where this assumption is used is when applying Lemma 4.2.19. This step can be replaced by the usage of Lemma 4.2.29.

In all cases, we reduced the general case to the two known base cases with $\mu'_1 \leq 1$ in $O(k + \mu_1)$ powers of cylinder shears and half cylinder shears. This concludes the proof of the theorem. \square

Conclusion and perspectives

Equivalence via cylinder shears

The first open problem we consider is naturally Conjecture 4.0.7 in full generality.

Conjecture 4.0.7 (Delecroix and Legrand-Duchesne, generalizing Conjecture 4 in [BLd+22])

Let μ be an integer partition of n and k a non-negative integer such that

$$\sum_{i \geq 1} \mu_i(i - 2) = 4g - 4 + k.$$

Let S and S' be two square-tiled surfaces in $\text{ST}(\mu, k)$. Then S and S' are equivalent via cylinder shears if and only if they belong to the same connected component of the moduli space of quadratic differentials.

In this chapter, we confirmed Conjecture 4.0.7 in two special cases. First, in the Abelian hyperelliptic components (see Theorem 4.0.9). Second, on the sphere (see Theorem 4.0.10), in all strata when authorizing half-shears, but only in the strata such that $\mu_1 \leq 1$ and $(\mu, k) \neq ([1, 2^{\mu_2}, 3], 4)$ when restricting to regular shears. This suggests several milestones towards proving Conjecture 4.0.7, that isolate different points of failure of our proof. We present them ranked by increasing (supposed) difficulty:

Question 4.2.30

Let μ be an integer partition and k a positive integer such that $\sum \mu_i(i - 2) = k - 4$ and $\mu_1 \leq 1$. Is $\text{ST}(\mu, k)$ connected via cylinder shears only?

In other words, Question 4.2.30 asks whether Theorem 4.0.10 holds even in the spherical strata of the form $([1, 2^{\mu_2}, 3], 4)$. We hope to prove Question 4.2.30, as it suffices to prove that all corresponding decorated trees are equivalent to a nice

decorated tree, after what the rest of our proof of Theorem 4.0.10 also applies in these strata. This would mean improving Lemma 4.2.27 to avoid using half cylinder shears, that are slightly artificial in the non-hyperelliptic strata. Exhaustive computer search shows that this holds for $\mu_2 \leq 15$, using at most four powers of cylinder shears.

The second question is whether Theorem 4.0.9 can be generalized to quadratic hyperelliptic components:

Conjecture 4.2.31

Let $g \geq 1$ and μ be an integer partition such that $\sum \mu_i(i - 2) = 4g - 4$. If $\text{ST}(\mu)$ contains a quadratic hyperelliptic component $\text{ST}_{\text{quad}}^{\text{hyp}}(\mu)$, then any two square-tiled surfaces in $\text{ST}_{\text{quad}}^{\text{hyp}}(\mu)$ are connected by a sequence of at most $O(g)$ powers of cylinder shears.

Recall from Remark 4.2.25 and [Lan08] that the quadratic hyperelliptic square-tiled surfaces are those that have a quotient by an hyperelliptic involution that belongs to $\text{ST}_{\text{quad}}(1^{\mu_1}, 2^{\mu_2}, a, b)$ where $a, b \geq 2$. Lemma 4.2.27 connects hyperelliptic square-tiled surfaces to some path-like configuration, but is specific to the Abelian strata, that once quotiented have at most one singularity of degree more than two. If Lemma 4.2.27 holds for quadratic hyperelliptic components, then the rest of our proof of Theorem 4.0.9 implies Conjecture 4.2.31.

We conjecture the following which generalizes Theorem 4.0.10:

Conjecture 4.2.32

Let μ be an integer partition of n and k a non-negative integer. The set $\text{ST}(\mu, k)$ of spherical square-tiled surfaces, that is such that $\sum_i (i - 2)\mu_i = k - 4$, is connected by cylinder shears.

The following arguments suggest that we are far from proving Conjecture 4.2.32:

- If $k = 0$, there are no path-like configurations, which are the cornerstone of our approach.
- When $\mu_1 \geq 2$, we do not know how to reach a path-like configuration (if it exists) as soon as there is more than one singularity of degree more than two.
- When $\mu_1 \geq 2$, we do not know how to prove the equivalence of the path-like configurations (if they exist) without using half cylinder shears to reduce μ_1 .

Finally, our approach most probably collapses when trying to tackle Conjecture 4.0.7. Indeed, when $g \geq 1$ the Jordan Curve theorem does not hold. This causes two issues: the dual tricolored cubic graphs are not planar anymore and

our proof of Lemma 4.2.15 that finds the fusion paths relies crucially on Jordan's Curve theorem. Finally, when $g \geq 1$, it is possible the $S^*(\tau)$ contains neither 1-faces nor half-edges, which are at the core of our proofs. In this case, even the use of half cylinder shears does not change this fact.

Beyond Conjecture 4.0.7

Theorems 4.0.9 and 4.0.10 prove that in a subset of strata, any two square-tiled surfaces are equivalent up to $O(n)$ powers of cylinder shears, where n is the number of triangles of the surface. As we make extensive use of powers of horizontal shears, this gives reconfiguration sequences of length $O(n^2)$ when counting the shears one by one. By Proposition 4.2.11, there exist pairs of square-tiled surfaces in $ST(\mu, k)$ that are separated by $\Omega(n)$ cylinder shears, when (μ, k) is a spherical profile with $\mu_1 = 0$. Thus it would be interesting to close this gap by improving on either of these two bounds, even in the very restricted case of the sphere with $\mu_1 = 0$.

Finally, cylinder shears could be used to sample a random square-tiled surface in a fixed strata. One can easily define the *shear dynamics* that performs at each step a cylinder shear chosen at random, with probability proportional to its length. This raises the following question:

Question 4.2.33

In which connected component of the moduli space of quadratic differentials does the shear dynamics rapidly mix on the set of square-tiled surfaces?

A cylinder shear on the dual tricolored cubic graph $S^*(\tau)$ is a Kempe change on the 3-edge-coloring, followed by a slide of the incident edges to preserve the cyclic coloring of $S^*(\tau)$. Thus, cylinder shears are very similar to Kempe changes, and by extension, so are the WSK algorithm and the shear dynamics. In addition to this similarity, both the shear dynamics and the WSK dynamics are non-local and therefore hard to analyze with the state of the art methods. Hence, we hope that a better understanding of any of these two operations translates to the other.

Chapter 5

A structure theorem for locally finite quasi-transitive graphs avoiding a minor

This chapter presents results obtained with Louis Esperet and Ugo Giocanti [10], accepted but not yet published in JCTb. These results consists in a structure theorem on quasi-transitive graphs avoiding a minor and applications to their Hadwiger number and the decidability of the Domino Problem.

Introduction

In 1961, Wang initiated the study of some tilings of the plane, now known as Wang dominoes. The Wang tiling problem asks whether a given finite set of square tiles whose edges are colored can tile the plane, such that adjacent tiles share identically colored edges (see Figure 5.1)¹. Wang conjectured that the positive instances of the Wang tiling problem were exactly the sets of tiles for which there exists a periodic tiling, that is fixed by some translation. If true, this conjecture would imply that the Wang tiling problem is decidable, as the naive algorithm that tests all possible tilings of rectangles of increasing size would either find a pattern that can be repeated, or no tilings for some large enough rectangle.

In 1964, Berger disproved the conjecture of Wang by exhibiting a set of 20,426 tiles that can only tile the plane aperiodically. The size of such set of dominoes has been progressively reduced over the years, and recently Jeandel and Rao [JR21] gave a set of 11 tiles with this property (see Figure 5.2) and proved that no smaller set exists, thanks to an exhaustive computer search. More recently, the substitutive structure of the tiling of Jeandel and Rao was described

¹The tiles cannot be rotated or turned over.




Figure 5.1: A set of three Wang dominoes that can tile the plane periodically




Figure 5.2: Jeandel and Rao's aperiodic set of eleven Wang dominoes

by Labbé [Lab21b, Lab21a].

As an alternative proof of this result, Berger [Ber66] described how to encode a Turing machine in a set of tiles, thereby reducing the halting problem to the Wang tiling problem and laying the groundwork for what is now known as molecular programming.

We may wonder what happens when trying to tile other structures. For example, if we are trying to tile the infinite line with regular dominoes, by the pigeon hole principle, any tiling must use some tile twice (in fact an infinite number of times). Hence the pattern between the two occurrences of this tile can be repeated and there exists a periodic tiling. For this reason, the conjecture of Wang is true on the infinite line, and the Wang tiling problem is decidable in this case.

The original motivation of this chapter is to classify the spaces for which this problem is decidable. Graphs are a natural candidate to represent these other structures. Forbidding certain tiles to be adjacent comes down to describing a graph H whose vertices are the tiles and whose edges represent the forbidden adjacencies, then a tiling becomes a homomorphism from G to H . This shows that the domino problem can be viewed as a generalization of a coloring problem (see Subsection 1.1.3).

More generally, the Wang tiling problem can be extended to a problem on groups and their Cayley graphs. We will give a structure theorem reminiscent of Robertson and Seymour's Graph Minor Structure Theorem [RS03] and use it to

discuss this extension of the Wang tiling problem.

A structure theorem

A central result in modern graph theory is the Graph Minor Structure Theorem of Robertson and Seymour [RS03], later extended to infinite graphs by Diestel and Thomas [DT99]. This theorem states that any graph G avoiding a fixed minor has a tree-decomposition such that each piece of the decomposition, called a torso, is close to being embeddable on a surface of bounded genus. A natural question is the following: if the graph G has non-trivial symmetries, can we make these symmetries apparent in the tree-decomposition? In other words, do graphs avoiding a fixed minor have a tree-decomposition as above, but with the additional constraint that the decomposition is *canonical*, i.e., invariant under the action of the automorphism group of G ? With Louis Esperet and Ugo Giocanti [10], we answer this question positively for infinite, locally finite graphs G that are *quasi-transitive*, i.e., the vertex set of G has finitely many orbits under the action of the automorphism group of G . This additional restriction, which is a way of saying that the graph G is highly symmetric, has the advantage of making the structure theorem much cleaner: instead of being almost embeddable on a surface of bounded genus, each torso of the tree-decomposition is now simply finite or planar.

Theorem 5.0.1 (see Theorem 5.2.1)

Every locally finite quasi-transitive graph avoiding the countable clique as a minor has a canonical tree-decomposition whose torsos are finite or planar.

The tree-decomposition in Theorem 5.0.1 will be obtained by refining the tree-decomposition obtained in the following more detailed version of the result, which might be useful for applications.

Theorem 5.0.2 (see Theorem 5.2.3)

Every locally finite quasi-transitive graph G avoiding the countable clique as a minor has a canonical tree-decomposition with adhesion at most 3 in which each torso is a minor of G , and is planar or has bounded treewidth.

Interestingly, the proof does not use the original structure theorem of Robertson and Seymour [RS03] or its extension to infinite graphs by Diestel and Thomas [DT99]. Instead, we rely mainly on a series of results and tools introduced by Grohe [Gro16a] to study decompositions of finite 3-connected graphs into quasi-4-connected components, together with a result of Thomassen [Tho92] on locally finite quasi-4-connected graphs. The main technical contribution of our work consists in extending the results of Grohe to infinite, locally finite graphs and in addition, making sure that the decompositions we obtain are canonical (in a certain weak

sense). Our proof crucially relies on a recent result of Carmesin, Hamann, and Miraftab [CHM22], which shows that there exists a canonical tree-decomposition that distinguishes all tangles of a given order (in our case, of order 4).

Thomassen proved that if a locally finite quasi-transitive graph has only one end, then this end must be thick [Tho92, Proposition 5.6] (see below for a definition of a thick end). At some point in our proof, we also need to show the stronger result (see Proposition 5.2.7), of independent interest, that for any $k \geq 1$, a locally finite quasi-transitive graph cannot have only one end of degree k .

We now discuss some applications of Theorem 5.0.1.

Hadwiger number

As a consequence of Theorem 5.0.1, we obtain a result on the Hadwiger number of locally finite quasi-transitive graphs. The *Hadwiger number* of a graph G is the supremum of the sizes of all finite complete minors in G . We say that a graph G *attains its Hadwiger number* if the supremum above is attained, that is if it is either finite, or G contains an infinite clique minor. Thomassen [Tho92] proved that every locally finite quasi-transitive 4-connected graph attains its Hadwiger number, and suggested that the 4-connectedness assumption might be unnecessary. We prove that this is indeed the case.

Theorem 5.0.3 (see Theorem 5.3.1)

Every locally finite quasi-transitive graph attains its Hadwiger number.

We will indeed prove a stronger statement, namely that every locally finite quasi-transitive graph avoiding the countable clique as a minor also avoids a finite graph with crossing number 1 as a minor.

Accessibility in graphs

We now introduce the notion of accessibility in graphs considered by Thomassen and Woess [TW93]. To distinguish it from the related notion in groups (see below), we will call it vertex-accessibility in the remainder of this chapter. A *ray* in an infinite graph G is an infinite one-way path in G . Two rays of G are *equivalent* if there are infinitely many disjoint paths between them in G (note that this is indeed an equivalence relation). An *end* of G is an equivalence class of rays in G . When there is a finite set X of vertices of G , two distinct components C_1, C_2 of $G - X$, and two distinct ends ω_1, ω_2 of G such that for each $i = 1, 2$, all but finitely many vertices of all (equivalently any) rays of ω_i are in C_i , we say that X *separates* ω_1 and ω_2 . A graph G is *vertex-accessible* if there is an integer k such that for any two distinct ends ω_1, ω_2 in G , there is a set of at most k vertices that separates ω_1 and ω_2 .

It was proved by Dunwoody [Dun07] (see also [Ham18b, Ham18a] for a more combinatorial approach) that locally finite quasi-transitive planar graphs are vertex-accessible. Here we extend the result to locally finite quasi-transitive graphs excluding the countable clique K_∞ (and not necessarily K_5 and $K_{3,3}$) as a minor, and in particular to locally finite quasi-transitive graphs from any proper minor-closed family.

Theorem 5.0.4 (see Theorem 5.2 in [10])

Every locally finite quasi-transitive K_∞ -minor-free graph is vertex-accessible.

Accessibility in groups

The notion of vertex-accessibility introduced above is related to the notion of accessibility in groups. Given a finitely generated group Γ , and a finite set of generators S , the *Cayley graph* of Γ with respect to the set of generators S is the edge-labelled graph $\text{Cay}(\Gamma, S)$ whose vertex set is the set of elements of Γ and where for every two elements $g, h \in \Gamma$ we put an arc (g, h) labelled with $a \in S$ when $h = a \cdot g$. Cayley graphs have to be seen as highly symmetric graphs; in particular they are transitive: the right action of the group Γ onto itself can be easily seen to induce a transitive group action on $\text{Cay}(\Gamma, S)$. It is known that the number of ends of a Cayley graph of a finitely generated group does not depend of the choice of generators, so we can talk about the number of ends of a finitely generated group. A classical theorem of Stallings [Sta71] states that if a finitely generated group Γ has more than one end, it can be split as a non-trivial free product with finite amalgamation, or as an HNN-extension over a finite subgroup. If any group produced by the splitting still has more than one end we can keep splitting it using Stallings theorem. If the process eventually stops (with Γ being obtained from finitely many 0-ended or 1-ended groups using free products with amalgamation and HNN-extensions), then Γ is said to be *accessible*. Thomassen and Woess [TW93] proved that a finitely generated group is accessible if and only if at least one of its locally finite Cayley graphs is vertex-accessible, if and only if all of its locally finite Cayley graphs are vertex-accessible.

A finitely generated group is *minor-excluded* if at least one of its Cayley graphs avoids a finite minor. Similarly a finitely generated group is *K_∞ -minor-free* if one of its Cayley graphs avoids the countable clique as a minor, and *planar* if one of its Cayley graphs is planar. Note that planar groups are minor-excluded and Theorem 5.0.3 immediately implies that a finitely generated group is minor-excluded if and only if it is K_∞ -minor-free.

Droms [Dro06] proved that finitely generated planar groups are finitely presented, while Dunwoody [Dun85] proved that finitely presented groups are acces-

sible, which implies that finitely generated planar groups are accessible. Theorem 5.0.4 immediately implies the following, which extends this result to all minor-excluded finitely generated groups, and equivalently to all finitely generated K_∞ -minor-free groups.

Corollary 5.0.5

Every finitely generated K_∞ -minor-free group is accessible.

In fact, combining Theorem 5.0.1 with techniques introduced by Hamann in the planar case [Ham18b, Ham18a], we prove the following stronger result which also implies Corollary 5.0.5 using the result of Dunwoody [Dun85] that all finitely presented groups are accessible.

Theorem 5.0.6 (see Corollary 5.7 in [10])

Every finitely generated K_∞ -minor-free group is finitely presented.

The Domino Problem

We refer to [ABJ18] for a detailed introduction to the domino problem. In this chapter, colorings will not be necessarily proper, so a *coloring* of a graph G with colors from a set Σ is simply a map $V(G) \rightarrow \Sigma$. The *Domino Problem* for a finitely generated group Γ together with a finite generating set S is defined as follows. The input is a finite alphabet Σ and a finite set $\mathcal{F} = \{F_1, \dots, F_p\}$ of forbidden *patterns*, which are colorings with colors from Σ of the closed neighborhood of the neutral element 1_Γ in the Cayley graph $\text{Cay}(\Gamma, S)$, viewed as an edge-labelled subgraph of $\text{Cay}(\Gamma, S)$ (recall that each element of S corresponds to a different label). The problem then asks if there is a coloring of $\text{Cay}(\Gamma, S)$ with colors from Σ , such that for each $v \in \Gamma$, the coloring of the closed neighborhood of v in $\text{Cay}(\Gamma, S)$ (viewed as an edge-labelled subgraph of $\text{Cay}(\Gamma, S)$), is not isomorphic to any of the colorings F_1, \dots, F_p , where we consider isomorphisms preserving the edge-labels (equivalently, isomorphisms corresponding to the right multiplication by elements of Γ).

It turns out that the decidability of the Domino Problem for (Γ, S) is independent of the choice of the finite generating set S , hence we can talk of the decidability of the Domino Problem for a finitely generated group Γ . If we consider $\Gamma = (\mathbb{Z}^2, +)$, then the Domino Problem corresponds exactly to the well-known Wang tiling problem, as the forbidden pattern correspond to the pairs of tiles that cannot be adjacent, for each side. Thus the Domino Problem is undecidable in $\Gamma = (\mathbb{Z}^2, +)$. On the other hand, a similar reasoning as the example we gave on the infinite line yields a simple greedy procedure to solve the Domino Problem in free groups, which admit trees as Cayley graphs. More generally, the Domino Problem is decidable in virtually free groups, which can equivalently be

defined as finitely generated groups having a locally finite Cayley graph of bounded treewidth [ABJ18, Ant11]. A remarkable conjecture of Ballier and Stein [BS18] asserts that these groups are the only ones for which the Domino Problem is decidable.

Conjecture 5.0.7 Domino Problem conjecture [BS18]

A finitely generated group has a decidable Domino Problem if and only if it is virtually free.

Recall that virtually free groups are precisely the groups having a locally finite Cayley graph of bounded treewidth. Since having bounded treewidth is a property that is closed under taking minors, it is natural to ask whether Conjecture 5.0.7 holds for minor-excluded groups (or equivalently, using Theorem 5.0.3, to K_∞ -minor-free groups). Using Corollary 5.0.5, together with classical results on planar groups and recent results on fundamental groups of surfaces [ABM18], we prove that this is indeed the case.

Theorem 5.0.8 (see Theorem 5.3.2)

A finitely generated K_∞ -minor-free group has a decidable Domino Problem if and only if it is virtually free.

Overview of the proof of Theorems 5.0.1 and 5.0.2

Consider a locally finite quasi-transitive graph G that excludes the countable clique K_∞ as a minor. The graph G is said to be *quasi-4-connected* if it is 3-connected and for every set $S \subseteq V(G)$ of size 3 such that $G - S$ is not connected, $G - S$ has exactly two connected components and one of them consists of a single vertex. Thomassen proved that if G is quasi-4-connected, then G is planar or has finite treewidth [Tho92], which implies Theorem 5.0.2 in this case (with a trivial tree-decomposition consisting of a single node).

To deal with the more general case, the first step is to obtain a canonical tree-decomposition of G of adhesion at most 2 in which all torsos are minors of G that are 3-connected graphs, cycles, or complete graphs on at most 2 vertices. The existence of such a decomposition in the finite case is a well-known result of Tutte [Tut84] and was proved in the locally finite case in [DSS98b]. For our proof we need to go one step further. Grohe [Gro16a] proved that every finite graph G has a tree-decomposition of adhesion at most 3 whose torsos are minors of G and are complete graphs on at most 4 vertices or quasi-4-connected graphs. A crucial step for us would be to prove a version of this result in which the tree-decomposition would be canonical, and which would hold for locally finite graphs.

However, as observed by Grohe, even in the finite case the decomposition he obtains is not canonical in general. Our main technical contribution is to extend

the result of Grohe [Gro16a] mentioned in the previous paragraph to locally finite graphs, while making sure that most of the construction (except the very end) is canonical. For this, we proceed in two steps. First, we use a result of [CHM22] to find a canonical tree-decomposition of any 3-connected graph G that distinguishes all of its tangles of order 4. Using this result, we show that we can assume that the graph under consideration admits a unique tangle \mathcal{T} of order 4. We then follow the main arguments from [Gro16a] and show that G has a canonical tree-decomposition of adhesion 3 which is a star and whose torsos are all minors of G and finite, except for the torso H associated to the center of the star, which has the following property: there exists a matching $M \subseteq E(H)$ which is invariant under the action of the automorphism group of G and such that the graph $H' := H/M$ obtained after the contraction of the edges of M is quasi-transitive, locally finite, and quasi-4-connected. In particular, a result of Thomassen [Tho92] then implies that H' is planar or has bounded treewidth. We then prove that even if H itself is not necessarily quasi-4-connected, it is still planar or has bounded treewidth, which is enough to conclude the proof of Theorem 5.0.2. The final step to prove Theorem 5.0.1 consists in refining the tree-decomposition to make sure that torsos of bounded treewidth are replaced by torsos of finite size (moreover, this refinement has to be done in a canonical way).

Organization of this chapter

Section 5.1 is dedicated to the study of canonical tree-decompositions. It contains the main technical contribution of [10], a partial extension of results of Grohe [Gro16a] to infinite graphs. Section 5.2 contains the proof of Theorems 5.0.1 and 5.0.2. Section 5.3 is dedicated to the applications of Theorems 5.0.1 and 5.0.2 to the Hadwiger number and the Domino Problem. The applications to the accessibility of minor-free groups and the vertex-accessibility of locally finite quasi-transitive minor-free graphs are developed in [10] but are omitted here. We conclude with a number of open problems in Section 5.4.

5.1 Tree-decompositions and tangles

5.1.1 Separations and canonical tree-decompositions

As we will be heavily relying on results of Grohe [Gro16a], we use his notation for all objects related to separations and tangles. A *separation* in a graph $G = (V, E)$ is a triple (Y, S, Z) such that Y, S, Z are pairwise disjoint, $V = Y \cup S \cup Z$ and there is no edge between vertices of Y and Z . A separation (Y, S, Z) is *proper* if Y and Z are nonempty. In this case, S is a *separator* of G .

The separation (Y, S, Z) is said to be *tight* if there are some components C_Y, C_Z respectively of $G[Y], G[Z]$ such that $N_G(C_Y) = N_G(C_Z) = S$. The *order* of a separation (Y, S, Z) is $|S|$ and the *order* of a family \mathcal{N} of separations is the supremum of the orders of its separations. In what follows, we will always consider sets of separations of finite order. We will denote $\text{Sep}_k(G)$ (respectively $\text{Sep}_{<k}(G)$) the set of all separations of G of order k (respectively less than k).

The following lemma was originally stated in [TW93] for transitive graphs, but the same proof immediately implies that the result also holds for quasi-transitive graphs.

Lemma 5.1.1 Corollary 4.3 in [TW93]

Let G be a locally finite graph. Then for every $v \in V(G)$ and $k \geq 1$, there is only a finite number of tight separations (Y, S, Z) of order k in G such that $v \in S$. Moreover, for any group Γ acting quasi-transitively on G and any $k \geq 1$, there is only a finite number of Γ -orbits of tight separations of order at most k in G .

Canonical tree-decompositions

For a group Γ acting (by automorphisms) on a graph G , we say that a tree-decomposition (T, \mathcal{V}) of G is *canonical with respect to Γ* , or simply Γ -*canonical*, if Γ induces a group action on T such that for every $\gamma \in \Gamma$ and $t \in V(T)$, $V_t \cdot \gamma = V_{t \cdot \gamma}$. By definition of a group action on a graph, $t \mapsto t \cdot \gamma$ is an automorphism of T for any $\gamma \in \Gamma$. In particular, for every $\gamma \in \Gamma$, note that γ sends bags of (T, \mathcal{V}) to bags, and adhesion sets to adhesion sets. When (T, \mathcal{V}) is $\text{Aut}(G)$ -canonical, we simply say that it is *canonical*.

Remark 5.1.2. If (T, \mathcal{V}) is a Γ -canonical tree-decomposition of a graph G , then Γ acts both on G and T , so there are two different notions of a stabilizer of a node $t \in V(T)$: $\Gamma_t = \text{Stab}_\Gamma(t)$ (where we consider the action of Γ on T), and $\text{Stab}_\Gamma(V_t)$ (where we consider the action of Γ on G). Observe that for any $t \in V(T)$ we have $\Gamma_t \subseteq \text{Stab}_\Gamma(V_t)$. The reverse inclusion does not hold in general (when there are adjacent nodes $s, t \in V(T)$ with $V_s = V_t$, automorphisms of T exchanging s and t stabilize $V_s = V_t$ without stabilizing s or t). However, if $t \in V(T)$ is such that $\{t' \in V(T) : V_{t'} = V_t\} = \{t\}$, then $\Gamma_t = \text{Stab}_\Gamma(V_t)$. In particular, if (T, \mathcal{V}) has finite adhesion, then every bag V_t with $t \in V_\infty(T)$ appears only once in the decomposition, and thus for each such node $t \in V_\infty(T)$ we have $\Gamma_t = \text{Stab}_\Gamma(V_t)$. The property that the two notions of stabilizers coincide for infinite bags when the canonical tree-decomposition has finite adhesion will be used repeatedly in the remainder of this chapter.

Edge-separations and torsos

Consider a tree-decomposition (T, \mathcal{V}) of a graph G , with $\mathcal{V} = (V_t)_{t \in V(T)}$. Let A be an orientation of the edges of $E(T)$, i.e. a choice of either (t_1, t_2) or (t_2, t_1) for every edge $t_1 t_2$ of T . For an arbitrary pair $(t_1, t_2) \in A$, and for each $i \in \{1, 2\}$, let T_i denote the component of $T - \{t_1 t_2\}$ containing t_i . Then the *edge-separation* of G associated to (t_1, t_2) is (Y_1, S, Y_2) with $S := V_{t_1} \cap V_{t_2}$ and $Y_i := \bigcup_{s \in V(T_i)} V_s \setminus S$ for $i \in \{1, 2\}$.

Given a separation (Y, S, Z) and an automorphism γ of a graph G , let $(Y, S, Z) \cdot \gamma := (Y \cdot \gamma, S \cdot \gamma, Z \cdot \gamma)$. If $\Gamma \subseteq \text{Aut}(G)$ and \mathcal{N} is a family of separations of G , we say that \mathcal{N} is Γ -invariant if for every $(Y, S, Z) \in \mathcal{N}$ and $\gamma \in \Gamma$, we have $(Y, S, Z) \cdot \gamma \in \mathcal{N}$. Note that if (T, \mathcal{V}) is Γ -canonical, then the associated set of edge-separations is Γ -invariant.

The *torsos* of (T, \mathcal{V}) are the graphs with vertex set V_t and edge set $E(G[V_t])$ together with the edges xy such that x and y belong to a common adhesion set of (T, \mathcal{V}) . Note that this definition coincides with the general definition of torso $G[[V_t]]$ we gave in Subsection 1.1.7 when the edge-separations of (T, \mathcal{V}) are tight. To prevent any ambiguity between these definitions, we will always ensure that the tree-decompositions we work with have this property (i.e., all the associated edge-separations are tight).

Remark 5.1.3. If (T, \mathcal{V}) is a Γ -canonical tree-decomposition of a locally finite graph G whose edge-separations are tight, then by Lemma 5.1.1 the action of Γ on $E(T)$ must induce a finite number of orbits. In particular, Γ must also act quasi-transitively on $V(T)$.

Recall that the *treewidth* of a graph G is the infimum of the width of (T, \mathcal{V}) , among all tree-decompositions (T, \mathcal{V}) of G . Note that adding to a tree-decomposition of bounded width the restriction that it must be canonical can be very costly in the finite case: while it is well known that every cycle graph C_n on n vertices has treewidth 2, the example below shows that in any canonical tree-decomposition of C_n , some bag contains all the nodes of C_n .

Example 5.1.4

Let C_n be the cycle graph on n elements. Note that the additive group \mathbb{Z}_n acts transitively by rotation on C_n . We let a be a generator of \mathbb{Z}_n of order n . Let $(T, (V_t)_{t \in V(T)})$ be a \mathbb{Z}_n -canonical tree-decomposition of C_n . Without loss of generality we may assume that T is finite, by contracting every edge $tt' \in E(T)$ such that $V_t = V_{t'}$. We may also assume that no edge tt' of T is inverted by a , i.e. such that $(t, t') \cdot a = (t', t)$, as if it was the case we could subdivide the edge tt' (i.e. add a new vertex t^* between t and t') and let $V_{t^*} := V_t \cap V_{t'}$. If we let T' be the tree of the tree-decomposition obtained after performing such a subdivision, note that the

5. A structure theorem for locally finite quasi-transitive graphs avoiding a minor

obtained tree-decomposition is still \mathbb{Z}_n -canonical as a induces an automorphism of T' that stabilizes the vertex t^* and acts on $V(T)$ the same way that it did before the subdivision. After this operation none of the edges tt^* , t^*t' is inverted by a . It is an easy exercise to prove that if no edge of T is inverted by a , there exists a vertex $t \in V(T)$ stabilized by a , and hence by all the elements of \mathbb{Z}_n . Then as \mathbb{Z}_n acts transitively on G , we must have $V_t = V(C_n)$ for such a $t \in V(T)$.

It turns out that in the quasi-transitive locally-finite case, if a graph has bounded treewidth then adding the restriction that the tree-decomposition must be canonical is fairly inexpensive if we only care about having finite bags.

Theorem 5.1.5 Theorem 7.4 in [HLMR22], [MS83], [Woe89], [TW93]

Let G be a connected quasi-transitive locally finite graph. Then the following are equivalent:

- G has finite treewidth;
- all the ends of G are thin;
- there exists $k \geq 1$ such that every end of G has degree at most k ;
- there exists a canonical tree-decomposition of G with tight edge-separations and finite width.

Separations of order at most 3

If G is not connected, then the tree-decomposition (T, \mathcal{V}) where T is a star whose central bag is empty and where we put a bag for each connected component of G can easily be seen to be a canonical tree-decomposition with adhesion 0, as every automorphism of G acts on T by permuting some branches. If we start from a connected graph G , it is well-known that the block cut-tree of G is a canonical tree-decomposition $(T, (V_t)_{t \in V(T)})$ of G whose adhesion sets have size 1 and such that for each $t \in V(T)$, $G[V_t] = G[V_t]$ has either size at most 2 or is 2-connected. A similar result holds for separations of order 2 (this was proved by Tutte [Tut84] in the finite case, and generalized to infinite graphs in [DSS98b]).

Theorem 5.1.6 [DSS98b]

Every locally finite graph G has a canonical tree-decomposition of adhesion at most 2, whose torsos are minors of G and are complete graphs of order at most 2, cycles, or 3-connected graphs.

For separations of order 3, a similar result was obtained by Grohe for finite graphs [Gro16a].

Theorem 5.1.7 [Gro16a]

Every finite graph G has a tree-decomposition of adhesion at most 3 whose torsors are minors of G and are complete graphs on at most 4 vertices or quasi-4-connected graphs.

Our main technical contribution will be to extend Theorem 5.1.7 to locally finite graphs, while making sure that most of the construction (except the very end) is canonical. More precisely, we reproduce in Subsections 5.1.5 and 5.1.6 the main steps of the work of [Gro16a] and give the additional arguments to extend them to locally finite graphs. A consequence is that Theorem 5.1.7 extends to locally finite graphs. However in our case, the main difficulty we face is that the decomposition of Grohe is not canonical (although some parts of the construction are canonical, which will be crucial for our purposes). To give a rough idea of the bulk of the problem it is helpful to consider Example 5.1.8 below, which was introduced in [Gro16a] in the finite case.




Figure 5.3: Left: a finite section of the 3-connected infinite graph obtained by replacing in the infinite hexagonal planar grid each vertex by a triangle with three vertices of degree 3. Right: a finite section of the quasi-4-connected torso $G[V_{z_0}]$ of (T, \mathcal{V}) . Note that it does not depend of the choice of V_{z_0} .

Example 5.1.8

Consider the 3-connected infinite planar graph H obtained from the infinite hexagonal planar grid by replacing each vertex by a triangle with three vertices of degree 3 (see Figure 5.3 (left) for a finite part of this graph). We let M be the set of edges connecting pairs of triangles, or equivalently the set of edges that do not belong to any triangle (note that M is a perfect matching). The tree-decomposition (T, \mathcal{V}) of H obtained by extending the ideas in [Gro16a] to the infinite case has an infinite bag V_{z_0} obtained by selecting one endpoint of each edge of M (which is equivalent

to fixing an orientation of each of these edges). The tree T is a subdivision of a star with center z_0 , and its other bags are finite. While there are many different choices for V_{z_0} , none of them gives a canonical tree-decomposition. Indeed one can check more generally that no tree-decomposition of H satisfying the properties of Theorem 5.1.7 can be canonical. To see this, assume for the sake of contradiction that such a decomposition (T, \mathcal{V}) exists. Then one of its edge-separations should be proper of order 3. Note that the only such separations separate a subgraph of a triangle from the rest of the graph. Let (Y, S, Z) be such a separation, such that Z is finite. Then there exists an edge e from M with one endpoint in Z and the other in S . Note that there exists an automorphism $\gamma \in \text{Aut}(H)$ exchanging the two endpoints of e . In particular, as (T, \mathcal{V}) is canonical, both $(Y, S, Z) \cdot \gamma$ and (Y, S, Z) must be edge-separations of (T, \mathcal{V}) , which can be seen to be impossible.

This example illustrates the fact that in general it is impossible to obtain a canonical tree-decomposition having exactly the properties described in Theorem 5.1.7. However note that here, if we want a canonical tree-decomposition whose torsos are either planar or finite, it is sufficient to take the trivial tree-decomposition with a single bag containing all the vertices. This is exactly what our proof will do when applied to this graph. More precisely, we note that on this example, the set M of edges is invariant under the action of $\text{Aut}(H)$. Based on this observation, our solution to obtain a canonical tree-decomposition will be to start with the same decomposition as that of [Gro16a], but to keep the two endpoints of each edge of M instead of choosing only one of its endpoints as above.

Combining canonical tree-decompositions

Let (T, \mathcal{V}) and (T', \mathcal{V}') be tree-decompositions of two graphs G, G' , respectively. We say that (T, \mathcal{V}) and (T', \mathcal{V}') are *isomorphic* if there exists an isomorphism φ from G to G' , and an isomorphism ψ from T to T' such that for each $t \in V(T)$, we have: $V'_{\psi(t)} = \varphi(V_t)$.

Let G be a graph and let Γ be a group acting on G . Let (T, \mathcal{V}) , with $\mathcal{V} = (V_t)_{t \in V(T)}$, be a Γ -canonical tree-decomposition of G , and recall that for any $t \in V(T)$, $\Gamma_t = \text{Stab}_\Gamma(t)$ denotes the stabilizer of the node t in the action of Γ on the tree T . For each $t \in V(T)$, let (T_t, \mathcal{V}_t) be a Γ_t -canonical tree-decomposition of $G[V_t]$. Our goal will be to refine (T, \mathcal{V}) by combining it with the tree-decompositions $(T_t, \mathcal{V}_t)_{t \in V(T)}$. If we want the resulting refined tree-decomposition of G to be Γ -canonical, we need to impose a condition on the tree-decompositions $(T_t, \mathcal{V}_t)_{t \in V(T)}$ (namely that they are consistent with the action of Γ on G). This is captured by the following definition. For $g \in \Gamma$ and $t \in V(T)$, we define $(T_t, \mathcal{V}_t) \cdot g := (T_t, \mathcal{V}_t \cdot g)$ as the tree-decomposition of $G[V_t \cdot g] = G[V_{t \cdot g}]$ with underlying tree T_t and bags $\mathcal{V}_t \cdot g := (V_s \cdot g)_{s \in V(T_t)}$. Observe that $(T_t, \mathcal{V}_t) \cdot g$ is $\Gamma_{t \cdot g}$ -canonical. We say that the construction $t \mapsto (T_t, \mathcal{V}_t)$ is Γ -canonical if for

each $g \in \Gamma$ and $t \in V(T)$, the tree-decompositions $(T_t, \mathcal{V}_t) \cdot g$ and $(T_{t \cdot g}, \mathcal{V}_{t \cdot g})$ are isomorphic. We emphasize here that the first tree-decomposition is indexed by T_t , while the second is indexed by $T_{t \cdot g}$.

The *trivial tree-decomposition* of a graph G consists of a tree T with a single node, whose bag is $V(G)$. Note that the trivial tree-decomposition is canonical.

Lemma 5.1.9 (see Lemma 3.9 in [10] for a proof)

Assume that G is locally finite, Γ is a group acting on G , and (T, \mathcal{V}) is a Γ -canonical tree-decomposition of G with finitely bounded adhesion. Let $\{t_i, i \in I_\infty\}$ be a set of representatives of the Γ -orbits of $V_\infty(T)$, indexed by some set I_∞ . Assume that for every $i \in I_\infty$ there exists a Γ_{t_i} -canonical tree-decomposition $(T_{t_i}, \mathcal{V}_{t_i})$ of the torso $G[V_{t_i}]$ of finitely bounded adhesion. Then we can find some family $(T_t, \mathcal{V}_t)_{t \in V(T)}$ extending the family $(T_{t_i}, \mathcal{V}_{t_i})_{i \in I_\infty}$ such that each (T_t, \mathcal{V}_t) is a Γ_t -canonical tree-decomposition of $G[V_t]$, the construction $t \mapsto (T_t, \mathcal{V}_t)$ is Γ -canonical, and for each $t \in V(T) \setminus V_\infty(T)$, (T_t, \mathcal{V}_t) is the trivial tree-decomposition of $G[V_t]$.

Given two tree-decompositions $(T, \mathcal{V}), (T', \mathcal{V}')$ of a graph G , with $\mathcal{V} = (V_t)_{t \in V(T)}$ and $\mathcal{V}' = (V'_t)_{t \in V(T')}$, we say that (T', \mathcal{V}') *refines* (T, \mathcal{V}) with respect to some family $(T_t, \mathcal{V}_t)_{t \in V(T)}$ of tree-decompositions if for every $t \in V(T)$, T_t is a subtree of T' such that $V_t = \bigcup_{s \in V(T_t)} V'_s$ and the trees $(T_t)_{t \in V(T)}$ are pairwise vertex-disjoint, cover $V(T')$ and for every edge $uv \in E(T)$, there exist $u' \in V(T_u), v' \in V(T_v)$ such that $u'v' \in E(T')$.

We say that (T', \mathcal{V}') is a *subdivision* of (T, \mathcal{V}) if T' is obtained from T after considering a subset $E' \subseteq E(T)$ and doing the following for every edge $tt' \in E'$: we subdivide the edge tt' (by adding a new vertex t^* between t and t'), and we add a corresponding bag $V_{t^*} := V_t \cap V_{t'}$ in the tree-decomposition. Note that if (T', \mathcal{V}') is a subdivision of (T, \mathcal{V}) , the two tree-decompositions have the same edge-separations.

The following result from [CHM22] will allow us to construct canonical tree-decompositions inductively:

Proposition 5.1.10 Proposition 7.2 in [CHM22]

Assume that G is locally finite, Γ is a group acting on G and (T, \mathcal{V}) is a Γ -canonical tree-decomposition of G with finitely bounded adhesion, with $\mathcal{V} = (V_t)_{t \in V(T)}$. Assume that for every $t \in V(T)$, there exists a Γ_t -canonical tree-decomposition (T_t, \mathcal{V}_t) of the torso $G[V_t]$ of finitely bounded adhesion such that the edge-separations induced by (T_t, \mathcal{V}_t) in $G[V_t]$ are tight and pairwise distinct, and the construction $t \mapsto (T_t, \mathcal{V}_t)$ is Γ -canonical. Then there exists a Γ -canonical tree-decomposition (T', \mathcal{V}') of G that refines (T, \mathcal{V}) with respect to a family $(T'_t, \mathcal{V}'_t)_{t \in V(T)}$ such that for each $t \in V(T)$, (T'_t, \mathcal{V}'_t) is a Γ_t -canonical tree-decomposition of $G[V_t]$ which is a subdivision of (T_t, \mathcal{V}_t) , and such that every adhesion set of (T', \mathcal{V}') is either

5. A structure theorem for locally finite quasi-transitive graphs avoiding a minor

an adhesion set of (T, \mathcal{V}) or an adhesion set of some (T_t, \mathcal{V}_t) for some $t \in V(T)$. Moreover, the construction $t \mapsto (T'_t, \mathcal{V}'_t)$ is Γ -canonical.

Remark 5.1.11. In the original statement of [CHM22, Proposition 7.2], the fact that each tree-decomposition (T'_t, \mathcal{V}'_t) is Γ_t -canonical and that the construction $t \mapsto (T'_t, \mathcal{V}'_t)$ is also Γ -canonical is not stated explicitly, however the authors show it explicitly in the proof.

Hence putting Lemma 5.1.9 together with Proposition 5.1.10, we immediately get:

Corollary 5.1.12

Assume that G is locally finite, Γ is a group acting on G , and (T, \mathcal{V}) a Γ -canonical tree-decomposition of G of finitely bounded adhesion, with $\mathcal{V} = (V_t)_{t \in V(T)}$. Let $\{t_i : i \in I_\infty\}$ denote a set of representatives of the orbits $V_\infty(T)/\Gamma$ such that for each $i \in I_\infty$, there exists a Γ_{t_i} -canonical tree-decomposition $(T_{t_i}, \mathcal{V}_{t_i})$ of $G[V_{t_i}]$ with finitely bounded adhesion, such that the edge-separations induced by each $(T_{t_i}, \mathcal{V}_{t_i})$ in $G[V_{t_i}]$ are tight and pairwise distinct. Then there exists a Γ -canonical tree-decomposition of G that refines (T, \mathcal{V}) with respect to some family $(T'_t, \mathcal{V}'_t)_{t \in V(T)}$ of Γ_t -canonical tree-decompositions of $G[V_t]$ such that for each $i \in I_\infty$, $(T'_{t_i}, \mathcal{V}'_{t_i})$ is a subdivision of $(T_{t_i}, \mathcal{V}_{t_i})$, and for every $t \in V(T) \setminus V_\infty(T)$, (T'_t, \mathcal{V}'_t) is the trivial tree-decomposition of $G[V_t]$. Moreover, the construction $t \mapsto (T'_t, \mathcal{V}'_t)$ is Γ -canonical.

The main objects of study of this chapter are canonical tree-decompositions of quasi-transitive graphs. A crucial property is that the torsos or parts of the tree-decomposition are themselves quasi-transitive. This is proved in [HLMR22, Proposition 4.5] in the special case where Γ acts transitively on $E(T)$. We give in [10] a more general proof of this result, which is self-contained.

Lemma 5.1.13

Let $k \in \mathbb{N}$, let G be a locally finite graph, and let Γ be a group acting quasi-transitively on G . Let (T, \mathcal{V}) , with $\mathcal{V} = (V_t)_{t \in V(T)}$, be a Γ -canonical tree-decomposition of G whose edge-separations are tight and have order at most k . Then, for any $t \in V(T)$, the group $\Gamma_t := \text{Stab}_\Gamma(t)$ induces a quasi-transitive action on $G[V_t]$, and thus also on $G[V_t]$.

Moreover, Lemma 5.1.13 still holds if we only require $E(T)/\Gamma$ to be finite (instead of requiring the edge-separations of (T, \mathcal{V}) to be tight).

5.1.2 Tangles

Tangles were introduced by Robertson and Seymour [RS91] and play a fundamental role in their proof of the Graph Minor Structure Theorem. We will consider here the equivalent definition used by Grohe [Gro16a]. A *tangle of order k* in G is a subset \mathcal{T} of $\text{Sep}_{<k}(G)$ such that

($\mathcal{T}1$) For all separations $(Y, S, Z) \in \text{Sep}_{<k}(G)$, either $(Y, S, Z) \in \mathcal{T}$ or $(Z, S, Y) \in \mathcal{T}$;

($\mathcal{T}2$) For all separations $(Y_1, S_1, Z_1), (Y_2, S_2, Z_2), (Y_3, S_3, Z_3) \in \mathcal{T}$, either $Z_1 \cap Z_2 \cap Z_3 \neq \emptyset$ or there exists an edge with an endpoint in each Z_i .

Note that ($\mathcal{T}2$) with $(Y_i, S_i, Z_i) = (Y, S, Z)$ for each $i \in [3]$ implies in particular that for every separation $(Y, S, Z) \in \mathcal{T}$, $Z \neq \emptyset$. A tangle in G will be called a G -tangle, for brevity. Intuitively, a G -tangle is a consistent orientation of the separations of G , pointing towards a highly connected region of G . We refer the reader to [RS91, Gro16a] for background on tangles. In general when G is finite there is a one-to-one correspondence between the G -tangles of order 1 and the connected components of G , between the G -tangles of order 2 and the *biconnected* components of G , and between the G -tangles of order 3 and the *triconnected* components of G (we omit the definitions as these notions will not be needed in the remainder of this chapter, and instead refer the interested reader to [RS91] and [Gro16b]).

In infinite graphs, tangles can also be seen as a notion generalizing the notion of ends: for each end ω of a graph G and every $k \geq 2$, we define the G -tangle \mathcal{T}_ω^k of order k induced by ω by:

$$\mathcal{T}_\omega^k := \{(Y, S, Z) \in \text{Sep}_{<k}(G) : \omega \text{ lives in a component of } Z\}.$$

The fact that \mathcal{T}_ω^k is indeed a tangle is a folklore result. One of the basic properties of tangles is that for any fixed model of a graph H in a graph G , any H -tangle of order k induces a G -tangle of order k . More precisely, if $\mathcal{M} = (M_v)_{v \in V(H)}$ is a model of H in G and (Y, S, Z) is a separation of order less than k in G , then its *projection* with respect to \mathcal{M} is the separation $\pi_{\mathcal{M}}(Y, S, Z) = (Y', S', Z')$ of H of order less than k defined by: $Y' := \{v \in V(H) : M_v \subseteq Y\}$, $S' := \{v \in V(H) : M_v \cap S \neq \emptyset\}$ and $Z' := \{v \in V(H) : M_v \subseteq Z\}$.

A proof of the following result can be found in [RS91, (6.1)] or in a more similar version in [Gro16a, Lemma 3.11]. Its proof extends to the locally finite case.

Lemma 5.1.14 Lemma 3.11 in [Gro16a]

Let G be a locally finite graph. Let $\mathcal{M} = (M_v)_{v \in V(H)}$ be a model of a graph H in G and \mathcal{T}' be an H -tangle of order k . Then the set

$$\mathcal{T} := \{(Y, S, Z) \in \text{Sep}_{<k}(G) : \pi_{\mathcal{M}}(Y, S, Z) \in \mathcal{T}'\}$$

is a G -tangle of order k , called the *lifting* of \mathcal{T}' in G with respect to \mathcal{M} .

Remark 5.1.15. Assume that \mathcal{M} is a faithful model of H in G with the property that for each $(Y', S', Z') \in \text{Sep}_{<_k}(H)$, there exists some $(Y, S, Z) \in \text{Sep}_{<_k}(G)$ such that $\pi_{\mathcal{M}}(Y, S, Z) = (Y', S', Z')$ and $S' = S$. Then the function that maps every tangle of order k in H to its lifting in G with respect to \mathcal{M} is injective. To see this, consider two distinct tangles $\mathcal{T}'_1 \neq \mathcal{T}'_2$ of order k in H . Then there exists some $(Y', S', Z') \in \mathcal{T}'_1$ such that $(Z', S', Y') \in \mathcal{T}'_2$. If we consider $(Y, S, Z) \in \text{Sep}_{<_k}(G)$ such that $\pi_{\mathcal{M}}(Y, S, Z) = (Y', S', Z')$, we have $(Y, S, Z) \in \mathcal{T}_1$ and $(Z, S, Y) \in \mathcal{T}_2$, where for each $i \in \{1, 2\}$, \mathcal{T}_i denotes the lifting of \mathcal{T}'_i with respect to \mathcal{M} . It then follows that $\mathcal{T}_1 \neq \mathcal{T}_2$, as desired. Note that if (T, \mathcal{V}) is a tree-decomposition with finitely bounded adhesion and $t \in V(T)$ is such that $G[[V_t]]$ is a faithful minor of G , then any faithful model \mathcal{M} of $G[[V_t]]$ has the property we just described.

If $\mathcal{M} = (M_v)_{v \in V(H)}$ is a model of H in G , and \mathcal{T} is a tangle of G , then $\mathcal{T}' := \{\pi_{\mathcal{M}}(Y, S, Z) : (Y, S, Z) \in \mathcal{T}\}$ is called the *projection* of \mathcal{T} . Note that \mathcal{T}' is not a tangle in general. Projecting is the converse operation of lifting in the sense that if \mathcal{M} is faithful, and \mathcal{T} and \mathcal{T}' are tangles of G and H , then \mathcal{T} is the lifting of \mathcal{T}' if and only if \mathcal{T}' is the projection of \mathcal{T} .

We define a partial order \prec over the set of separations of a graph G by letting for every two separations $(Y, S, Z), (Y', S', Z')$, $(Y, S, Z) \preceq (Y', S', Z')$ if and only if $S \cup Z \subsetneq S' \cup Z'$ or $(S \cup Z = S' \cup Z'$ and $S \subsetneq S')$. Intuitively, $(Y, S, Z) \preceq (Y', S', Z')$ means that (Y, S, Z) points towards a direction in a more accurate way than (Y', S', Z') does. Note that our definition of \preceq is the same as in [Gro16a, Subsection 3.2] and slightly differs from the more conventional one of [RS91, CHM22].

A partially ordered set $(X, <)$ is said to be *well-founded* if every strictly decreasing sequence of elements of X is finite. In particular, if $(X, <)$ is well-founded then for every $x \in X$, there exists $y \in X$ which is minimal with respect to $<$ and such that $y \leq x$. In the remainder of this chapter, whenever we consider a minimal separation or a well-founded family of separations, we always implicitly refer to the partial order \prec defined in the paragraph above.

We will distinguish two types of tangles in infinite graphs:

- the *region tangles*, defined as those which are well-founded (with respect to the order \prec), and
- the *evasive tangles*, which contain some infinite decreasing sequence of separations (with respect to the order \prec).

The tangles we consider in this work will always have order at most 4. Note that if G is 3-connected, an evasive tangle \mathcal{T} of order 4 is exactly a tangle \mathcal{T}_ω^4 induced by an end ω of degree 3. On the other hand, a region tangle is either a tangle of order 4 induced by some end ω of degree at least 4, or a tangle which is

not induced by an end. For example, one can check that both graphs in Figure 5.3 have a unique tangle of order 4 which is the tangle induced by their unique end (which is thick), and this tangle is a region tangle in both cases.

We say that a separation (Y, S, Z) *distinguishes* two tangles $\mathcal{T}, \mathcal{T}'$ if $(Y, S, Z) \in \mathcal{T}$ and $(Z, S, Y) \in \mathcal{T}'$, or vice versa. We say that (Y, S, Z) distinguishes \mathcal{T} and \mathcal{T}' *efficiently* if there is no separation of smaller order distinguishing \mathcal{T} and \mathcal{T}' . A tree-decomposition (T, \mathcal{V}) *distinguishes* a set of tangles \mathcal{A} if for every two distinct tangles $\mathcal{T}, \mathcal{T}' \in \mathcal{A}$ there exists an edge-separation of (T, \mathcal{V}) distinguishing \mathcal{T} and \mathcal{T}' . A separation is called *relevant* with respect to \mathcal{A} if it distinguishes at least two tangles of \mathcal{A} . A tree-decomposition is *nice* (with respect to \mathcal{A}) if all its edge-separations are relevant (with respect to \mathcal{A}).

We will need the following result, which extends earlier results of [RS91, DHL18], and which is a canonical version of one of the main results of the grid-minor series in the locally finite case. We will only use it with $k = 4$, but we nevertheless state the result in its most general form.

Theorem 5.1.16 Theorem 7.3 in [CHM22]

Let $k \geq 1$ and let G be a locally finite graph. Then there exists a canonical tree-decomposition (T, \mathcal{V}) of G that efficiently distinguishes the set \mathcal{A}_k of tangles of order at most k and that is nice with respect to \mathcal{A}_k .

Remark 5.1.17. The fact that (T, \mathcal{V}) is nice in Theorem 5.1.16 is not explicit in the original statement, however it directly follows from the proof. Moreover, the proof also ensures that the edge-separations of (T, \mathcal{V}) are pairwise distinct.

5.1.3 An example

We give here an example of a one-ended graph that excludes some minor and has infinitely many region tangles of order 4. We show how to distinguish them on this example with a canonical tree-decomposition. As the application of Theorem 5.1.16 allowing to distinguish all tangles of order 4 is the very first step of our proof of Theorems 5.2.1 and 5.2.3, this example may also be useful to have some intuition on it.

We consider the infinite graph G (a finite section of which is illustrated in Figure 5.4), which is obtained from the infinite triangular grid by adding in each triangular face $f = v_1v_2v_3$ three vertices $\{w_1, w_2, w_3\}$ inducing a $K_{3,3}$ with the vertices of the triangle, and another vertex z connected to each of the w_i 's. G has two types of tangles of order 4: one is the tangle \mathcal{T}_ω^4 induced by the unique end ω of G , and all the others are the tangles \mathcal{T}_f^4 pointing towards each face $f = v_1v_2v_3$ of the triangular grid; more precisely, \mathcal{T}_f^4 has the same set of separations as \mathcal{T}_ω^4 except for $(G - A, \{v_1, v_2, v_3\}, \{w_1, w_2, w_3, z\}) \in \mathcal{T}_f^4$, where $A := \{v_1, v_2, v_3, w_1, w_2, w_3, z\}$.




Figure 5.4: A useful example.

Note that with respect to our definition, all the tangles of order 4 of G are region tangles. We represented with red arrows the two separations of \mathcal{T}_f^4 that are minimal with respect to the order \prec but which are not minimal separations of \mathcal{T}_ω^k (for one fixed face f). The three red arrows crossing the red triangle correspond to the minimal separation $(G - A, \{v_1, v_2, v_3\}, \{w_1, w_2, w_3, z\})$ of \mathcal{T}_f^4 that points towards the triangular face $v_1v_2v_3$, while the three arrows directed away from z correspond to the minimal separation $(\{z\}, \{w_1, w_2, w_3\}, \{v_1, v_2, v_3\} \cup (G - A))$ of \mathcal{T}_f^4 . The tree-decomposition (T, \mathcal{V}) of Theorem 5.1.16 distinguishing all the tangles of order 4 is such that T is a star with center $t_0 \in V(T)$ such that $G[V_{t_0}]$ is the infinite planar triangular grid. Then T has one vertex t_f for each face $f = v_1v_2v_3$ of $G[V_{t_0}]$ and the bag V_{t_f} is finite and contains the 7 vertices $\{v_1, v_2, v_3, w_1, w_2, w_3, z\}$ associated to f . Note that such a tree-decomposition enjoys the properties of Theorems 5.2.1 and 5.2.3. However, this is not always the case and we need in general to decompose further some torsos of the tree-decomposition given by Theorem 5.1.16 in order to obtain such a decomposition.

5.1.4 Tangles of order 4: orthogonality and crossing-lemma

In this section we introduce some notions from [Gro16a] and briefly explain how to extend them in the locally finite case. The proofs being fairly technical, most of them will be either replaced by sketch of it or simply omitted (see [10] for the full version). Unless specified otherwise, we assume in the whole section that the graphs we consider are locally finite and 3-connected.

A separation $(Y, S, Z) \in \text{Sep}_{<4}(G)$ is said to be *degenerate* if

- (Y, S, Z) has order 3,
- $G[S]$ is an independent set, and
- $|Y| = 1$.

The following result from [Gro16a] immediately generalizes to locally finite graphs:

Lemma 5.1.18 Lemma 4.13 and Remark 4.14 in [Gro16a]

Let G be a locally finite 3-connected graph, and (Y, S, Z) be a proper separation of order 3. Then $G[[Z \cup S]]$ is a faithful minor of G if and only if (Y, S, Z) is non-degenerate.

We say that the edge-separations of a tree-decomposition (T, \mathcal{V}) are *non-degenerate* if for every $e \in E(T)$, none of the two edge-separations associated to e are degenerate.

Lemma 5.1.19

Let G be a locally finite 3-connected graph and let (T, \mathcal{V}) , with $\mathcal{V} = (V_t)_{t \in V(T)}$, be a tree-decomposition of G whose edge-separations have order 3 and are non-degenerate. Then $G[[V_t]]$ is a faithful minor of G for each $t \in V(T)$.

Proof. Let $t \in V(T)$, and t' be a neighbor of t in T . Let $(Y_{t'}, S_{t'}, Z_{t'})$ be the edge-separation of G associated to the (oriented) edge $(t', t) \in E(T)$, that is $S_{t'} = V_t \cap V_{t'}$, $V_{t'} \subseteq Y_{t'} \cup S_{t'}$, and $V_t \subseteq Z_{t'} \cup S_{t'}$. By Lemma 5.1.18, there is a faithful model $(M_v^{t'})_{v \in (Z_{t'} \cup S_{t'})}$ of $G[[Z_{t'} \cup S_{t'}]]$ in G . As the only edges of $G[[Z_{t'} \cup S_{t'}]]$ that are not edges of G must be between pairs of vertices of $S_{t'}$, we may assume that every $M_v^{t'}$ has size 1, except possibly when $v \in S_{t'}$, in which case the only vertices distinct from v that $M_v^{t'}$ can contain must lie in $Y_{t'}$, that is $M_v^{t'} \setminus \{v\} \subseteq M_v^{t'}$. For every $v \in V_t$, we let:

$$M_v := \bigcup_{\substack{t' \in V(T), \\ tt' \in E(T)}} M_v^{t'}.$$

We show that $(M_v)_{v \in V_t}$ is a faithful model of $G[[V_t]]$ in G . As (T, \mathcal{V}) is a tree-decomposition, for every two distinct neighbors t', t'' of t in T , $Y_{t'} \cap Y_{t''} = \emptyset$ so we must have $M_v^{t'} \cap M_v^{t''} = \{v\}$ and $M_v^{t'} \cap M_u^{t''} = \emptyset$ for each distinct vertices $u, v \in V_t$. As $(M_v^{t'})_{v \in (Z_{t'} \cup S_{t'})}$ is a model, we have $M_u^{t'} \cap M_v^{t'} = \emptyset$ for each $u \neq v \in V_t$. It follows that $M_u \cap M_v = \emptyset$ for each distinct $u, v \in V_t$. Now if $uv \in E(G[[V_t]])$ and $uv \notin E(G)$, there must exist some edge-separation $(Y_{t'}, S_{t'}, Z_{t'})$ such that $u, v \in S_{t'}$ and there exists a path from u to v in $G[S_{t'} \cup Y_{t'}]$. In particular, there must exist $u' \in M_u^{t'}$ and $v' \in M_v^{t'}$ such that $u'v' \in E(G)$. As $u' \in M_u$ and $v' \in M_v$, we proved that $(M_v)_{v \in V_t}$ is a faithful model of $G[[V_t]]$ in G . \square

For every tangle \mathcal{T} of a graph G , we denote by \mathcal{T}_{\min} its set of minimal separations (here and in the remainder, minimality of separations is always with respect to the partial order \preceq defined above). If \mathcal{T} has order 4, then we let \mathcal{T}_{nd} be its set of non-degenerate minimal separations.

Remark 5.1.20. Let G be locally finite, let \mathcal{T} be a G -tangle of order 4, and let (Y, S, Z) be a degenerate separation of G . Then $(Y, S, Z) \in \mathcal{T}$. This is a direct consequence of [Gro16a, Lemma 3.3], which states that if \mathcal{T} is a tangle of order k then for every separation (Y, S, Z) of order $k - 1$ such that $|Y \cup S| \leq \frac{3}{2}(k - 1)$ we have $(Y, S, Z) \in \mathcal{T}$.

For every tangle \mathcal{T} of order 4, we let:

$$X_{\mathcal{T}} := \bigcap_{\substack{(Y,S,Z) \in \mathcal{T}, \\ (Y,S,Z) \text{ is non-degenerate}}} (Z \cup S).$$

Note that if \mathcal{T} is an evasive tangle, then $X_{\mathcal{T}}$ is empty. In this case, and because G is 3-connected, there exists a unique end ω of degree 3 such that for any finite subset \S of \mathcal{T} , the end ω lies in

$$\bigcap_{\substack{(Y,S,Z) \in \S, \\ (Y,S,Z) \text{ is non-degenerate}}} (Z \cup S).$$

Remark 5.1.21. If $(Y, S, Z), (Y', S', Z') \in \mathcal{T}$ are such that $(Y', S', Z') \preceq (Y, S, Z)$ and (Y, S, Z) is non-degenerate, then it is easy to see that (Y', S', Z') is also non-degenerate (recall that G is 3-connected). Also if \mathcal{T} is a region tangle, for every $(Y, S, Z) \in \mathcal{T}$, there exists a separation $(Y', S', Z') \in \mathcal{T}_{\min}$ such that $(Y', S', Z') \preceq (Y, S, Z)$. These observations imply that if \mathcal{T} is a region tangle of order 4 and G is 3-connected and locally finite, then:

$$X_{\mathcal{T}} = \bigcap_{(Y,S,Z) \in \mathcal{T}_{\text{nd}}} (Z \cup S).$$

Two separations $(Y_1, S_1, Z_1), (Y_2, S_2, Z_2)$ are *orthogonal* if $(Y_1 \cup S_1) \cap (Y_2 \cup S_2) \subseteq S_1 \cap S_2$ (see Subfigure 5.5 (a)). A set \mathcal{N} of separations is said to be *orthogonal* if its separations are pairwise orthogonal. One can easily show that the set of minimal separations of a (region) tangle of order at most 3 is orthogonal. This does not hold for tangles of order 4, but Grohe [Gro16a] proved that for tangles of order 4, minimal separations can only cross in a restricted way. Two separations (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) are *crossing* if $Y_1 \cap Y_2 = S_1 \cap S_2 = \emptyset$ and there is an edge $s_1 s_2 \in E(G)$, with $S_1 \cap Y_2 = \{s_1\}$ and $S_2 \cap Y_1 = \{s_2\}$ (see Subfigure 5.5 (b)). In this case, we call $s_1 s_2$ the *crossedge* of (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) . We denote by $E_{\text{nd}}^{\times}(\mathcal{T})$ the set of crossedges of \mathcal{T}_{nd} . Lemma 4.16 from [Gro16a] generalizes to region tangles of order 4 of locally finite graphs:

Lemma 5.1.22 Lemma 4.16 and Corollary 4.20 in [Gro16a]

Let G be a locally finite 3-connected graph. Let \mathcal{T} be a region G -tangle of order 4. Then every two distinct minimal separations of \mathcal{T} are either crossing or orthogonal. Moreover, $E_{\text{nd}}^\times(\mathcal{T})$ forms a matching in G .




Figure 5.5: Interaction between minimal separations. The white zones represent empty sets while the grey represent potentially non-empty sets.

In [Gro16a], orthogonal sets of separations are presented as the nice case, as they allow to efficiently find quasi-4-connected regions. We show that, up to some additional assumptions, this observation still holds in the locally finite case. We recall that for a tangle \mathcal{T} of order 4, \mathcal{T}_{nd} denotes its set of minimal non-degenerate separations.

Lemma 5.1.23

Let G be a locally finite 3-connected graph. Let \mathcal{T} be a region G -tangle of order 4. Assume that \mathcal{T}_{nd} is orthogonal. Then $X_{\mathcal{T}} \neq \emptyset$ and the torso $G[[X_{\mathcal{T}}]]$ has size 3 or is a quasi-4-connected minor of G .

Proof. If every separation of order 3 in G is degenerate, then G is quasi-4-connected and all the separations of \mathcal{T}_{nd} are non-proper. It follows that $G = G[[X_{\mathcal{T}}]]$ and the desired properties hold.

Assume now that G has a proper non-degenerate separation (Y, S, Z) of order 3. As \mathcal{T} is a region tangle, there is a separation $(Y', S', Z') \in \mathcal{T}_{\text{min}}$ such that $(Y', S', Z') \preceq (Y, S, Z)$. As observed in Remark 5.1.21, $(Y', S', Z') \in \mathcal{T}_{\text{nd}}$. We claim that $S' \subseteq X_{\mathcal{T}}$ so $X_{\mathcal{T}} \neq \emptyset$: let $(Y_0, S_0, Z_0) \in \mathcal{T}_{\text{nd}} \setminus \{(Y', S', Z')\}$. As (Y_0, S_0, Z_0) and (Y', S', Z') are orthogonal, we must have: $S' \cap Y_0 = \emptyset$ so $S' \subseteq Z_0 \cup S_0$. As \mathcal{T} is a region tangle, the equality $X_{\mathcal{T}} = \bigcap_{(Y,S,Z) \in \mathcal{T}_{\text{nd}}} (Z \cup S)$ holds, and thus we proved that $S' \subseteq X_{\mathcal{T}}$, and so $X_{\mathcal{T}} \neq \emptyset$. Moreover, as G is 3-connected, the separations of \mathcal{T}_{min} have order 3 so $|X_{\mathcal{T}}| \geq |S'| \geq 3$.

We now assume that $|X_{\mathcal{T}}| \geq 4$ and show that $G[[X_{\mathcal{T}}]]$ is quasi-4-connected. Since G is 3-connected and $|X_{\mathcal{T}}| \geq 4$, $G[[X_{\mathcal{T}}]]$ is 3-connected (any proper separation of order at most 2 in $G[[X_{\mathcal{T}}]]$ would induce a proper separation of order at most 2 in G). If $|X_{\mathcal{T}}| = 4$, then $G[[X_{\mathcal{T}}]]$ is clearly also quasi-4-connected so we can assume that $|X_{\mathcal{T}}| \geq 5$. Suppose that $G[[X_{\mathcal{T}}]]$ is not 4-connected and let (Y_0, S_0, Z_0) be a proper separation of $G[[X_{\mathcal{T}}]]$ of order at most 3. We will prove that $|Y_0| = 1$ or $|Z_0| = 1$, which will immediately imply that $G[[X_{\mathcal{T}}]]$ is quasi-4-connected. We let Y_1 be the union of all connected components of $G - S_0$ that intersect Y_0 or have a neighbor in Y_0 . Let $S_1 := S_0$ and $Z_1 := V(G) \setminus (S_0 \cup Y_1)$. By definition of the torso $G[[X_{\mathcal{T}}]]$, (Y_1, S_1, Z_1) is a proper separation of order at most 3 in G , hence we must have $|S_0| = 3$ as G is 3-connected. Assume first that $(Y_1, S_1, Z_1) \in \mathcal{T}$. If (Y_1, S_1, Z_1) is non-degenerate, then $X_{\mathcal{T}} \cap Y_0 \subseteq X_{\mathcal{T}} \cap Y_1 = \emptyset$ by definition of $X_{\mathcal{T}}$. It follows that $Y_0 = \emptyset$, which contradicts the assumption that (Y_0, S_0, Z_0) is proper. If (Y_1, S_1, Z_1) is degenerate, then $|Y_0| \leq |Y_1| = 1$ so as (Y_0, S_0, Z_0) is proper we must have $|Y_0| = 1$. The case $(Z_1, S_1, Z_1) \in \mathcal{T}$ is symmetric. Hence we proved that every separation (Y, S, Z) of $G[[X_{\mathcal{T}}]]$ of order at most 3 satisfies $|Y| \leq 1$ or $|Z| \leq 1$ so we are done.

Finally the fact that $G[[X_{\mathcal{T}}]]$ is a minor of G easily follows from Lemma 5.1.19: we consider the tree-decomposition (T, \mathcal{V}) where T is a star with a central vertex z_0 and one edge $z_0 z_i$ for each $(Y_i, S_i, Z_i) \in \mathcal{T}_{\text{nd}}$. We let $V_{z_0} := X_{\mathcal{T}}$ and $V_{z_i} := Y_i \cup S_i$ for each $(Y_i, S_i, Z_i) \in \mathcal{T}_{\text{nd}}$. The fact that (T, \mathcal{V}) is a tree-decomposition follows from the orthogonality of \mathcal{T}_{nd} . Hence by Lemma 5.1.19, $G[[X_{\mathcal{T}}]]$ is a minor of G . \square

Whenever \mathcal{T}_{nd} is not orthogonal, Lemma 5.1.23 does not hold anymore and if we want to obtain a canonical tree-decomposition, we will need to consider a larger set, whose torso is not necessarily quasi-4-connected, but can be defined uniquely from the structural properties of \mathcal{T} , which will ensure that the resulting decomposition is canonical. For every region tangle \mathcal{T} of order 4 we let:

$$R_{\mathcal{T}} := \left(\bigcup_{(Y,S,Z) \in \mathcal{T}_{\text{nd}}} S \right) \cup \left(\bigcap_{(Y,S,Z) \in \mathcal{T}_{\text{nd}}} Z \right).$$

The set $R_{\mathcal{T}}$ corresponds to the set called $R^{(0)}$ in [Gro16a, Section 4.5]. Note that we always have $X_{\mathcal{T}} \subseteq R_{\mathcal{T}}$ and that equality holds when \mathcal{T}_{nd} is orthogonal. To illustrate the definition of $R_{\mathcal{T}}$, it is helpful to go back to the graph G on Figure 5.3 (left). Then G has a unique tangle \mathcal{T} of order 4, the set \mathcal{T}_{nd} is the set of separations (Y, S, Z) of order 3 where Y is a triangular face of G , $S = N(Y)$ and $Z = V(G) \setminus (Y \cup S)$. Hence on this example, the set of crossededges $E_{\text{nd}}^{\times}(G)$ is the set of edges joining two triangular faces, and $R_{\mathcal{T}} = V(G)$.

While the proof of the following result was originally written for finite graphs, it immediately generalizes to locally finite graphs.

Lemma 5.1.24 Lemma 4.32 in [Gro16a]

If G is a locally finite 3-connected graph and if \mathcal{T} is a region tangle of order 4 in G , then $G[[R_{\mathcal{T}}]]$ is a faithful minor of G .

For each $(Y, S, Z) \in \mathcal{T}_{\text{nd}}$, the fence $\text{fc}(S)$ of S in G is the union of

- the subset of vertices of S that are not the endpoint of some crossed edge of \mathcal{T} , and
- the subset of vertices s' such that ss' is a crossed edge of \mathcal{T} and $s \in S$.

In particular, as the crossed edges form a matching in G (Lemma 5.1.22), $|\text{fc}(S)| = |S| = 3$ for each $(Y, S, Z) \in \mathcal{T}_{\text{nd}}$. A consequence of Lemma 5.1.22 is the following:

Lemma 5.1.25

Let G be a locally finite 3-connected graph and \mathcal{T} be a region G -tangle of order 4. Then G has a tree-decomposition (T, \mathcal{V}) of adhesion 3 where $\mathcal{V} = (V_t)_{t \in V(T)}$ and T is a star with central vertex z_0 such that $V_{z_0} = R_{\mathcal{T}}$. If moreover, \mathcal{T} is the unique G -tangle of order 4, then (T, \mathcal{V}) is canonical and every bag except possibly V_{z_0} is finite.

Proof. We let:

$$V(T) := \{z_0\} \cup \{z_C : C \text{ connected component of } G - R_{\mathcal{T}}\}$$

where we choose the z_C 's to be pairwise distinct nodes. We let T be the star with vertex set $V(T)$ and central vertex z_0 , and we define $\mathcal{V} = (V_t)_{t \in V(T)}$ by setting $V_{z_0} := R_{\mathcal{T}}$, and for each connected component C of $G - R_{\mathcal{T}}$: $V_{z_C} := C \cup N(C)$. It is not hard to check that (T, \mathcal{V}) is a tree-decomposition of G . By [Gro16a, Lemma 4.31] (whose proof extends to the locally finite case), for each component C of $G - R_{\mathcal{T}}$ there exists a unique separator S such that $(Y, S, Z) \in \mathcal{T}_{\text{nd}}$ for some separation (Y, S, Z) , $N(C) = \text{fc}(S)$ and $C \subseteq Y$. This implies that (T, \mathcal{V}) has adhesion 3, so in particular its edge-separations are tight.

We now prove the second part of Lemma 5.1.25 and assume that \mathcal{T} is the unique tangle of order 4 of G . Then \mathcal{T} is $\text{Aut}(G)$ -invariant, and (T, \mathcal{V}) is clearly canonical. If some V_{z_C} is infinite for some $z_C \neq z_0$, then as G is locally finite, $G[V_{z_C}]$ has at least one infinite connected component, and hence there exists some end ω living in $G[V_{z_C}]$. In particular, ω induces some G -tangle \mathcal{T}_{ω} of order 4. We let $(Y, S, Z) \in \mathcal{T}_{\text{nd}}$ be the separation given by [Gro16a, Lemma 4.31] such that $N(C) = \text{fc}(S)$ and $C \subseteq Y$. Then (Y, S, Z) distinguishes \mathcal{T}_{ω} from \mathcal{T} , which contradicts the uniqueness of \mathcal{T} in G . \square

Note that in the non-orthogonal case, the tree-decomposition from Lemma 5.1.25 is not the same as the one from [Gro16a], as the torso $G[[R_{\mathcal{T}}]]$ associated to the

center of the star might not be quasi-4-connected. However, we will prove in Subsection 5.1.7 that $G[R_{\mathcal{T}}]$ still enjoys the same useful properties as a quasi-transitive quasi-4-connected graph, namely it is either planar or has bounded treewidth. We note that the crucial ingredient that allows us to obtain a canonical tree-decomposition in the second part of Lemma 5.1.25 (contrary to Grohe's decomposition) is the assumption that \mathcal{T} is the unique G -tangle of order 4. So one of the most important steps in the proof of our main results will be a reduction to the case where graphs have a single tangle of order 4.

5.1.5 Contracting a single crossed edge

In what follows, we let G be a locally finite 3-connected graph, and \mathcal{T} be a region tangle of order 4 in G . Recall that by Lemma 5.1.22 the set $E_{\text{nd}}^{\times}(\mathcal{T})$ of the crossed edges forms a matching. We will see that contracting a crossed edge results in a 3-connected graph G' that has a tangle \mathcal{T}' of order 4 induced by \mathcal{T} [Gro16a, Subsection 4.5]. More precisely, \mathcal{T}' contains as a subset the projection of \mathcal{T} with respect to the minor G' of G . We give here an overview of the lemmas stated in [Gro16a, Subsection 4.5] which all hold when G is locally finite instead of finite, by using the exact same proofs. The only additional property that we need in the locally finite case is that \mathcal{T}' is still a region tangle, which is proved in Lemma 5.1.29 below.

In the remainder of this subsection, we let (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) be two crossing separations of \mathcal{T}_{nd} with crossed edge s_1s_2 . *Contracting* s_1s_2 consists in deleting s_1 and s_2 and adding a new vertex s' whose neighborhood is equal to $N_G(s_1) \cup N_G(s_2) \setminus \{s_1, s_2\}$. We denote by G' the graph obtained after contracting s_1s_2 . The *projection* (referred to as *contraction* in [Gro16a]) of a set X of vertices of G is defined as

$$X^{\vee} := \begin{cases} X & \text{if } X \cap \{s_1, s_2\} = \emptyset \\ X \setminus \{s_1, s_2\} \cup \{s'\} & \text{if } X \cap \{s_1, s_2\} \neq \emptyset. \end{cases}$$

Given a set X' of vertices of G' , the *expansion* X'_{\wedge} of X' is defined as

$$X'_{\wedge} := \begin{cases} X' & \text{if } s' \notin X' \\ X' \setminus \{s'\} \cup \{s_1, s_2\} & \text{if } s' \in X'. \end{cases}$$

Observe that for all $X' \subseteq V(G')$, we have $(X'_{\wedge})^{\vee} = X'$ and for all $X \subseteq V(G)$, we have $X \subseteq (X^{\vee})_{\wedge}$ (where the inclusion might be strict). We also define for every $(Y, S, Z) \in \text{Sep}_{<4}(G)$:

$$(Y, S, Z)^{\vee} := \begin{cases} (Y^{\vee} \setminus \{s'\}, S^{\vee}, Z^{\vee} \setminus \{s'\}) & \text{if } S \cap \{s_1, s_2\} \neq \emptyset \\ (Y^{\vee}, S^{\vee}, Z^{\vee}) & \text{if } S \cap \{s_1, s_2\} = \emptyset. \end{cases}$$

Note that $(Y, S, Z)^\vee$ is exactly the projection $\pi_{\mathcal{M}}(Y, S, Z)$ with respect to the model $\mathcal{M} = (\{v\}_\wedge)_{v \in V(G')}$ of G' in G .

In the context of finite graphs, [Gro16a] proves the following lemmas that extend directly to the locally finite case:

Lemma 5.1.26 Corollary 4.24 in [Gro16a]

The graph G' resulting from the contraction of $s_1 s_2$ is 3-connected.

Lemma 5.1.27 Lemmas 4.26 and 4.27 in [Gro16a]

There exists a tangle \mathcal{T}' of order 4 in G' containing the projection of \mathcal{T} with respect to the model $\mathcal{M} = (\{v\}_\wedge)_{v \in V(G')}$.

Note that the projection of \mathcal{T} with respect to \mathcal{M} is exactly the set $\{(Y, S, Z)^\vee : (Y, S, Z) \in \mathcal{T}\}$. In the remainder of this subsection, we let \mathcal{T}' be the tangle given by Lemma 5.1.27. In [Gro16a], the author gives an explicit definition of \mathcal{T}' , but for the sake of clarity we only summarize here the properties of \mathcal{T}' that will be of interest for our purposes.

Note that the inclusion $\{(Y, S, Z)^\vee : (Y, S, Z) \in \mathcal{T}\} \subseteq \mathcal{T}'$ is strict in general, as some separations from $\text{Sep}_{<4}(G')$ might not be projections of separations from $\text{Sep}_{<4}(G)$. The next lemma intuitively states that every separation of \mathcal{T}' is close to an element from $\{(Y, S, Z)^\vee : (Y, S, Z) \in \mathcal{T}\} \subseteq \mathcal{T}'$.

Lemma 5.1.28 Definition of \mathcal{T}' and Lemmas 4.23 and 4.25 in [Gro16a]

For every separation $(Y', S', Z') \in \mathcal{T}'$ such that $s' \in S'$ and $G'[Z']$ is connected, there exists a separation $(Y, S, Z) \in \mathcal{T}$ such that $S^\vee = S'$ and $Z \setminus S_\vee = Z' \setminus \{s_1, s_2\} = Z'$.

As Lemma 5.1.28 is not exactly stated this way in [Gro16a], we briefly sketch how to obtain it. If $(Y', S', Z') \in \mathcal{T}'$ is such that $s' \in S'$, then by [Gro16a, Lemma 4.25] there exists a (unique) connected component C of $G \setminus S'_\vee$ such that the separations (Y'', S'', Z'') of \mathcal{T}' such that $S'' = S'$ are exactly the ones such that $C \subseteq Z''$, and for which every separation $(Y, S, Z) \in \mathcal{T}$ such that $S^\vee = S'$ satisfies $C \subseteq Z$. [Gro16a, Lemmas 4.23, 4.25] and the fact that \mathcal{T} is a tangle ensure the existence of a separation $(Y, S, Z) \in \mathcal{T}$ such that $S^\vee = S'$ and $C = Z \setminus \{s_1, s_2\} = Z \setminus S$. In particular by Lemma 5.1.27 the projection $(Y, S, Z)^\vee = (Y \setminus S', S', C)$ is in \mathcal{T}' so if we assume that $G'[Z']$ is connected, the choice of C imposes $C \subseteq Z'$, thus $Z' = C$. It implies that (Y, S, Z) satisfies the property described in Lemma 5.1.28.

Lemma 5.1.29

\mathcal{T}' is a region tangle.

Proof. Assume for the sake of contradiction that \mathcal{T}' contains an infinite strictly decreasing sequence of separations $(Y'_n, S'_n, Z'_n)_{n \in \mathbb{N}}$. By Lemma 5.1.26, G' is 3-connected, so the only possible non-proper separation (Y_n, S_n, Z_n) is for $n = 0$, thus we may assume that all the separations (Y_n, S_n, Z_n) are tight. By Lemma 5.1.1, there are finitely many integers n for which $s' \in S'_n$. Up to extracting an infinite subsequence, one can assume that for all n , either $s' \in Y'_n$ or $s' \in Z'_n$. If there exists N such that $s' \in Y'_N$ then for all $n \geq N$ we must have $s' \in Y_n$ by definition of \preceq . Up to extracting another infinite subsequence, we can assume that either $s' \in Y'_n$ for all n or $s' \in Z'_n$ for all n . As a result, and because \mathcal{T}' contains the projection of \mathcal{T} with respect to \mathcal{M} , $((Y'_n)_\wedge, S'_n, (Z'_n)_\wedge)_{n \in \mathbb{N}}$ is an infinite decreasing sequence of separations of order 3 in \mathcal{T} , contradicting the fact that \mathcal{T} is well-founded. \square

We conclude this subsection with the following result relating the degeneracy of minimal separations in G and G' . Its proof is the same as the proof of [Gro16a, Lemma 4.28], which directly translates to the locally finite case. To be more precise, we also need the additional assumption that \mathcal{T}' is a region tangle to make the proof work, which is given by Lemma 5.1.29.

Lemma 5.1.30 Lemma 4.28 and Corollary 4.29 in [Gro16a]
 Either G' is 4-connected and $\mathcal{T}'_{\min} = \{(\emptyset, \emptyset, V(G'))\}$, or

$$\mathcal{T}'_{\min} = \{(Y, S, Z)^\vee : (Y, S, Z) \in \mathcal{T}_{\min} \text{ and } S^\vee \text{ is a separator of } G'\}.$$

In the latter case, for all $(Y, S, Z) \in \mathcal{T}_{\min}$, (Y, S, Z) is non-degenerate if and only if $(Y, S, Z)^\vee$ is non-degenerate. Moreover, $E_{\text{nd}}^\times(\mathcal{T}') = E_{\text{nd}}^\times(\mathcal{T}) \setminus \{s_1 s_2\}$.

5.1.6 Contracting all the crossedges

In the previous subsection we studied the consequences of contracting a single crossedge in G . However, in our application we will need to contract *all* crossedges of $E_{\text{nd}}^\times(\mathcal{T})$ (which form a matching in G). We now study how this affects G .

Before going further, we will need to introduce some notation, extending the notation from [Gro16a] to the infinite case. For convenience we write $M := E_{\text{nd}}^\times(\mathcal{T})$ (and recall that M is a matching in G). For every subset $L \subseteq M$ of crossedges, we let $G^{\setminus L}$ be the graph obtained from G after contracting each edge $uv \in L$ into a new vertex $s_{u,v}$. Note that the order in which the edges are contracted is irrelevant in the definition of $G^{\setminus L}$.

We denote $\bar{L} = M \setminus L$. In this section we will also often use the notation $L - L'$ instead of $L \setminus L'$, to avoid any possible confusion when reading superscripts (for instance we will write $G^{\setminus L - L'}$ instead of $G^{\setminus L \setminus L'}$).

For every $L \subseteq M$, for every vertex $x \in V(G)$, we let

$$x^{\setminus L} := \begin{cases} x & \text{if } x \in X \setminus V(L) \\ s_{u,v} & \text{if } x \text{ is the endpoint of a crossedge } uv \in L. \end{cases}$$

For every subset $X \subseteq V(G)$ of vertices, we let $X^{\setminus L/} := \{x^{\setminus L/} : x \in X\}$ be the projection of X to $G^{\setminus L/}$.

Remark 5.1.31. Note that for every disjoint subsets $K, L \subset M$ and for all $X \in V(G)$, $X^{\setminus K \cup L/} = (X^{\setminus L/})^{\setminus K/} = (X^{\setminus K/})^{\setminus L/}$.

For every $X' \subseteq V(G^{\setminus M/})$ and $L \subseteq M$, let $X'_{/L\setminus}$ denote the maximal set $X \subseteq V(G^{\setminus \bar{L}/})$ such that $X^{\setminus L/} = X'$. In other words $X'_{/L\setminus}$ is the set of vertices obtained after “uncontracting” the edges of L in X . Note that with the notation introduced above we have $G = G^{\setminus \emptyset/}$. Given a separation (Y, S, Z) of G , we define

$$(Y, S, Z)^{\setminus L/} := (Y^{\setminus L/} \setminus S^{\setminus L/}, S^{\setminus L/}, Z^{\setminus L/} \setminus S^{\setminus L/}).$$

Note that when $L = \{s_1 s_2\}$ consists of a single edge, we recover the notions of the previous subsection; with our previous notation this gives: $x^{\setminus L/} = x^\vee$, $X^{\setminus L/} = X^\vee$ and $(Y, S, Z)^{\setminus L/} = (Y, S, Z)^\vee$.

For each finite subset of crossed edges $L \subseteq M$, and every enumeration (e_1, \dots, e_ℓ) of the edges of L , we let $\mathcal{T}^{\setminus (e_1, \dots, e_\ell)/}$ denote the tangle of $G^{\setminus L/}$ obtained after iteratively applying Lemma 5.1.27 to the graphs $G_0 := G, G_1, \dots, G_\ell$ with $G_i := G^{\setminus \{e_1, \dots, e_i\}/}$ for each $i \in [\ell]$.

Lemma 5.1.32 Lemma 4.30 (5) in [Gro16a]

For every enumeration (e_1, \dots, e_ℓ) of a finite set $L \subseteq M$ of crossed edges and every permutation σ of $[\ell]$, $\mathcal{T}^{\setminus (e_1, \dots, e_\ell)/} = \mathcal{T}^{\setminus (e_{\sigma(1)}, \dots, e_{\sigma(\ell)})/}$.

In the remainder of the subsection, for every finite subset $L \subseteq M$, we will denote with $\mathcal{T}^{\setminus L/}$ the unique tangle associated to any enumeration of L given by Lemma 5.1.32.

Intuitively when G is finite, one of the main properties of $\mathcal{T}^{\setminus L/}$ is that separations of $\mathcal{T}_{\text{nd}}^{\setminus L/}$ are in correspondence with separations of \mathcal{T}_{nd} , and that the only crossing pairs between elements of $\mathcal{T}^{\setminus L/}$ correspond to pairs which were already crossing in \mathcal{T} . Thus after each contraction, we reduce the number of crossed edges, hence when $L = M$, the family $\mathcal{T}_{\text{nd}}^{\setminus L/}$ must be orthogonal and we can apply results from previous subsections to the graph $G^{\setminus L/}$. We now show formally how to extend the relevant proofs of [Gro16a] to the locally finite case.

In [Gro16a], the author proved that if G is finite and 3-connected, for every $L \subseteq M$, the graph $G^{\setminus L/}$ is 3-connected and that $\mathcal{T}^{\setminus L/}$ is a tangle of order 4 induced by \mathcal{T} in $G^{\setminus L/}$. Using the results from the previous subsection, this immediately extends to $G^{\setminus L/}$ and $\mathcal{T}^{\setminus L/}$ when G is locally finite and L is finite, by induction on the size of L .

Theorem 5.1.33 Generalization of Lemma 4.30 in [Gro16a]

Let $L \subseteq M$ be a finite set of crossed edges. Then we have

5. *A structure theorem for locally finite quasi-transitive graphs avoiding a minor*

1. $G^{\setminus L/}$ is 3-connected
2. $\mathcal{T}^{\setminus L/}$ is a region tangle of order 4 of $G^{\setminus L/}$ such that

$$\mathcal{T}_{\min}^{\setminus L/} = \{(Y, S, Z)^{\setminus L/} : (Y, S, Z) \in \mathcal{T}_{\min} \text{ such that } S^{\setminus L/} \text{ is a separator of } G^{\setminus L/}\}$$
 or $\mathcal{T}_{\min}^{\setminus L/} = \{(\emptyset, \emptyset, V(G'))\}$ if $L = M$ is finite and $G^{\setminus L/}$ is 4-connected.
3. $E_{\text{nd}}^{\times}(\mathcal{T}^{\setminus L/}) = E_{\text{nd}}^{\times}(\mathcal{T}) \setminus L$
4. $\mathcal{T}^{\setminus L/}$ contains the projection $\{(Y, S, Z)^{\setminus L/} : (Y, S, Z) \in \mathcal{T}\}$ of \mathcal{T} with respect to the model $\mathcal{M} = (\{v\}_{/L\setminus})_{v \in V(G^{\setminus L/})}$.

We will now extend Theorem 5.1.33 to the case where $L \subseteq M$ is infinite. Given a set $X \subseteq V(G^{\setminus M/})$, we denote by $M(X) \subseteq M$ the subset of edges of G contracted to a vertex in X .

Lemma 5.1.34

The graph $G^{\setminus M/}$ is 3-connected.

Proof. Assume for the sake of contradiction that $G^{\setminus M/}$ has a separator S of order at most 2. Then the set $L := M(S)$ has size at most 2 and S is a separator of order 2 of $G^{\setminus L/}$. This contradicts Theorem 5.1.33. \square

We now let $L \subseteq M$ be any (not necessarily finite) subset of crossed edges and give a general definition of $\mathcal{T}^{\setminus L/}$ extending the previous one. For every $(Y', S', Z') \in \text{Sep}_{<4}(G^{\setminus L/})$, we let $L' := M(S')$. Note that L' is finite, and that $(Y'_{/L'\setminus}, S', Z'_{/L'\setminus})$ is a separation of order at most 3 in $G^{\setminus L'/}$. We define $\mathcal{T}^{\setminus L/}$ as the family of separations (Y', S', Z') of $G^{\setminus L/}$ such that $(Y'_{/L'\setminus}, S', Z'_{/L'\setminus}) \in \mathcal{T}^{\setminus L'/}$. Note that when L is finite, $(Y', S', Z') = (Y'_{/L\setminus}, S', Z'_{/L\setminus})^{\setminus L-L'}$ and thus iterative applications of Lemma 5.1.27 together with Lemma 5.1.32 imply that our definition of $\mathcal{T}^{\setminus L/}$ coincides with the one we gave above for finite subsets $L \subseteq M$.

Thanks to Remark 5.1.31, for all $L \subseteq M$, $\mathcal{T}^{\setminus M/} = (\mathcal{T}^{\setminus L/})^{\setminus \bar{L}/}$ and $G^{\setminus M/} = (G^{\setminus L/})^{\setminus \bar{L}/}$. We say that a set $X \subseteq V(G)$ hits the edges of L once if for all $e \in L$, $|X \cap e| = 1$.

Lemma 5.1.35

$\mathcal{T}^{\setminus M/}$ is a region tangle of order 4 in $G^{\setminus M/}$ such that $\{(Y, S, Z)^{\setminus M/} : (Y, S, Z) \in \mathcal{T}\} \subseteq \mathcal{T}^{\setminus M/}$.

Proof. We first prove that $\mathcal{T}^{\setminus M/}$ is a tangle of order 4. To prove (T1), let $(Y', S', Z') \in \text{Sep}_{<4}(G^{\setminus M/})$ and $L := M(S')$. Then L has size at most 3, so by Theorem 5.1.33, $\mathcal{T}^{\setminus L/}$ is a region tangle of order 4 of $G^{\setminus L/}$. As $(Y, S, Z) :=$

$(Y'_{/\bar{L}}, S', Z'_{/\bar{L}})$ is a separation of order 3 of $G^{\setminus L/}$ and $\mathcal{T}^{\setminus L/}$ is a tangle of order 4, either $(Y, S, Z) \in \mathcal{T}^{\setminus L/}$ or $(Z, S, Y) \in \mathcal{T}^{\setminus L/}$. By definition of $\mathcal{T}^{\setminus M/}$ we then have either $(Z', S', Y') \in \mathcal{T}^{\setminus M/}$ or $(Y', S', Z') \in \mathcal{T}^{\setminus M/}$, implying that $\mathcal{T}^{\setminus M/}$ satisfies **(T1)**.

To prove **(T2)**, let $(Y'_1, S'_1, Z'_1), (Y'_2, S'_2, Z'_2), (Y'_3, S'_3, Z'_3) \in \mathcal{T}^{\setminus M/}$. Let $L := M(S'_1 \cup S'_2 \cup S'_3)$. Once again L is finite with size at most 9 and for all $i \in \{1, 2, 3\}$, $(Y_i, S_i, Z_i) := ((Y'_i)_{/\bar{L}}, S'_i, (Z'_i)_{/\bar{L}})$ is a separation of order 3 of $G^{\setminus L/}$.

Claim 5.1.36

For every $i \in \{1, 2, 3\}$, $(Y_i, S_i, Z_i) \in \mathcal{T}^{\setminus L/}$.

Proof of the Claim: Assume that $i = 1$, the other cases being symmetric. We let $L_1 := M(S'_1)$. Then by definition of $\mathcal{T}^{\setminus M/}$, $(Y''_1, S''_1, Z''_1) := ((Y'_1)_{/\bar{L}_1}, S'_1, (Z'_1)_{/\bar{L}_1}) \in \mathcal{T}^{\setminus L_1/}$. Our goal is to show that $(Y_1, S_1, Z_1) = (Y''_1, S''_1, Z''_1)^{\setminus L-L_1/}$. As both L and L_1 are finite and $L_1 \subseteq L$, iterative applications of Lemmas 5.1.27 and Lemma 5.1.32 imply $\mathcal{T}^{\setminus L/}$ must contain the projection of $\mathcal{T}^{\setminus L_1/}$ with respect to the model $\mathcal{M} = (\{v\}_{/(\setminus L \setminus L_1) \setminus v} \in V(G)^{\setminus L/}$. Thus if we succeed to prove that

$$(Y_1, S_1, Z_1) = (Y''_1, S''_1, Z''_1)^{\setminus L-L_1/}, \quad (5.1)$$

we immediately obtain that $(Y_1, S_1, Z_1) \in \mathcal{T}^{\setminus L/}$, which concludes the claim.

To prove that (5.1) holds, note first that every edge of M is contracted in $G^{\setminus M/}$ so in particular it has its endpoints in exactly one of the three sets $(Y'_1)_{/M \setminus}$, $(S'_1)_{/M \setminus}$ and $(Z'_1)_{/M \setminus}$. In particular by definition of L_1 , the edges of L_1 are all disjoint from $(Y'_1)_{/M \setminus}$ and thus $(Y'_1)_{/M \setminus} = (Y'_1)_{/\bar{L}_1}$. This implies that $Y_1 = (Y'_1)_{/\bar{L}} = ((Y'_1)_{/\bar{L}_1})^{\setminus L-L_1/} = (Y''_1)^{\setminus L-L_1/}$. As S'_1 is disjoint from Y'_1 in $G^{\setminus M/}$, it is also disjoint from Y_1 in $G^{\setminus L/}$ so we have $Y_1 = (Y''_1)^{\setminus (L \setminus L_1)/} \setminus S''_1$. Symmetric arguments give $Z_1 = (Z''_1)^{\setminus L-L_1/} \setminus S''_1$, and as $S''_1 = S'_1 = S_1$, we get the desired equality. \diamond

By Theorem 5.1.33, $\mathcal{T}^{\setminus L/}$ is a region tangle of order 4 so Claim 5.1.36 implies that either $Z_1 \cap Z_2 \cap Z_3 \neq \emptyset$ or there exists an edge of $G^{\setminus L/}$ with both endpoints in $Z_1 \cup Z_2 \cup Z_3$. If $Z_1 \cap Z_2 \cap Z_3 \neq \emptyset$, then $Z'_1 \cap Z'_2 \cap Z'_3 = (Z_1 \cap Z_2 \cap Z_3)^{\setminus \bar{L}/} \neq \emptyset$. Otherwise, there is an edge $e \in E(G^{\setminus L/})$ that has an endpoint in each Z_i , in which case, either $Z'_1 \cap Z'_2 \cap Z'_3 \neq \emptyset$ if $e \in \bar{L}$, or $e^{\setminus \bar{L}/}$ is an edge of $G^{\setminus M/}$ which has an endpoint in each Z'_i . This proves **(T2)** and shows that $\mathcal{T}^{\setminus M/}$ is a tangle of order 4.

We now prove the inclusion $\{(Y, S, Z)^{\setminus M/} : (Y, S, Z) \in \mathcal{T}\} \subseteq \mathcal{T}^{\setminus M/}$. Let $(Y, S, Z) \in \mathcal{T}$ and $L := M(S)$. By Theorem 5.1.33 (4), $(Y, S, Z)^{\setminus L/} \in \mathcal{T}^{\setminus L/}$. Write $(Y', S', Z') = (Y, S, Z)^{\setminus M/}$ and note that $(Y'_{/\bar{L}}, S', Z'_{/\bar{L}}) = (Y, S, Z)^{\setminus L/}$, thus by definition of $\mathcal{T}^{\setminus M/}$, $(Y, S, Z)^{\setminus M/} \in \mathcal{T}^{\setminus M/}$.

We now prove that $\mathcal{T}^{\setminus M/}$ is a well-founded set. For the sake of contradiction, let $((Y'_n, S'_n, Z'_n))_{n \in \mathbb{N}}$ be an infinite decreasing sequence of separations of $\mathcal{T}^{\setminus M/}$. The contradiction will follow from the next claim, whose proof we will omit here

Claim 5.1.37 Claim 3.37 in [10]

There exists an infinite decreasing sequence $((Y_n'', S_n'', Z_n''))_{n \in \mathbb{N}}$ in $\mathcal{T}^{\setminus M/}$ such that for each $n \geq 0$, $(Y_n'', S_n'', Z_n'') = (Y_n, S_n, Z_n)^{\setminus M/}$ for some $(Y_n, S_n, Z_n) \in \mathcal{T}$.

It remains to show how to derive a contradiction from Claim 5.1.37. For this let $((Y_n'', S_n'', Z_n''))_{n \in \mathbb{N}}$ and $((Y_n, S_n, Z_n))_{n \in \mathbb{N}}$ be as in Claim 5.1.37. Again, up to considering a subsequence, we can assume that for all n , $S_n'' \subseteq Y_{n+1}''$ and $S_{n+1}'' \subseteq Z_n''$. As $S_n'' \cap S_{n+1}'' = \emptyset$, the separators S_n and S_{n+1} cannot contain two vertices of a common crossed edge of M . Thus, $(S_n'')_{/M \setminus} \subseteq Y_{n+1}$ and $(S_{n+1}'')_{/M \setminus} \subseteq Z_n$ and hence $(Y_{n+1}, S_{n+1}, Z_{n+1}) \prec (Y_n, S_n, Z_n)$. This proves that $((Y_n, S_n, Z_n))_{n \in \mathbb{N}}$ is an infinite decreasing sequence of separations of G with respect to \mathcal{T} , contradicting the fact that \mathcal{T} is a region tangle. \square

Lemma 5.1.38

Either $G^{\setminus M/}$ is 4-connected and $\mathcal{T}_{\min}^{\setminus M/} = \{(\emptyset, \emptyset, V(G)^{\setminus M/})\}$ or

$$\mathcal{T}_{\min}^{\setminus M/} = \{(Y, S, Z)^{\setminus M/} : (Y, S, Z) \in \mathcal{T}_{\min} \text{ such that } S^{\setminus M/} \text{ is a separator of } G^{\setminus M/}\}.$$

Finally, we have $E_{\text{nd}}^{\times}(G^{\setminus M/}) = \emptyset$.

Proof. Assume that $G^{\setminus M/}$ is not 4-connected. We first prove the direct inclusion. Let $(Y', S', Z') \in \mathcal{T}_{\min}^{\setminus M/}$, $L := M(S')$ and $(Y_0, S_0, Z_0) := (Y'_{/\bar{L}}, S', Z'_{/\bar{L}})$. Then by definition of $\mathcal{T}^{\setminus M/}$, $(Y_0, S_0, Z_0) \in \mathcal{T}^{\setminus L/}$. We prove that (Y_0, S_0, Z_0) is a minimal element of $\mathcal{T}^{\setminus L/}$. Assume for a contradiction that there exists $(Y_1, S_1, Z_1) \prec (Y_0, S_0, Z_0)$ in $\mathcal{T}^{\setminus L/}$. Then $(Y_1, S_1, Z_1)^{\setminus \bar{L}/} \preceq (Y_0, S_0, Z_0)^{\setminus \bar{L}/} = (Y', S', Z')$ and $(Y_1, S_1, Z_1)^{\setminus \bar{L}/} \neq (Y', S', Z')$ because $\bar{L} \cap S' = \emptyset$. Moreover, by Theorem 5.1.33 (2), we may assume that (Y_1', S_1', Z_1') is minimal and that $(Y_1, S_1, Z_1) = (Y_1', S_1', Z_1')^{\setminus L/}$ for some $(Y_1', S_1', Z_1') \in \mathcal{T}$. Thus as $(Y_1, S_1, Z_1)^{\setminus \bar{L}/} = ((Y_1', S_1', Z_1')^{\setminus L/})^{\setminus \bar{L}/} = (Y_1', S_1', Z_1')^{\setminus M/}$, we must have $(Y_1, S_1, Z_1) \in \mathcal{T}^{\setminus M/}$, contradicting the minimality of (Y_0, S_0, Z_0) in $\mathcal{T}^{\setminus L/}$. Hence $(Y', S', Z') = (Y_0, S_0, Z_0)^{\setminus \bar{L}/}$ for some $(Y_0, S_0, Z_0) \in \mathcal{T}_{\min}^{\setminus L/}$. Again we can apply Theorem 5.1.33 and write $(Y_0, S_0, Z_0) = (Y, S, Z)^{\setminus L/}$ for some $(Y, S, Z) \in \mathcal{T}_{\min}$. Thus we have:

$$(Y', S', Z') = ((Y, S, Z)^{\setminus L/})^{\setminus \bar{L}/} = (Y, S, Z)^{\setminus M/},$$

so we are done with the direct inclusion.

Conversely, let $(Y_1, S_1, Z_1) \in \mathcal{T}_{\min}$ such that $S_1^{\setminus M/}$ is a separator of $G^{\setminus M/}$. Note that by previous inclusion and because $\mathcal{T}^{\setminus M/}$ is a region tangle, it is enough to prove that for any $(Y_2, S_2, Z_2) \in \mathcal{T}$ such that $(Y_2, S_2, Z_2)^{\setminus M/} \preceq (Y_1, S_1, Z_1)^{\setminus M/}$, we have $(Y_2, S_2, Z_2)^{\setminus M/} = (Y_1, S_1, Z_1)^{\setminus M/}$. Let $(Y_2, S_2, Z_2) \in \mathcal{T}$ such that $(Y_2, S_2, Z_2)^{\setminus M/} \preceq (Y_1, S_1, Z_1)^{\setminus M/}$, $L := M(S_1 \cup S_2)$ and $\bar{L} := M \setminus L$. For $i \in \{1, 2\}$, the edges in \bar{L} are either contained in Y_i or in Z_i . Note that for each i and $L \subseteq M$:

$$S_i^{\setminus L/} \cup (Z_i^{\setminus L/} \setminus S_i^{\setminus L/}) = (Z_i \cup S_i)^{\setminus L/}$$

and

$$S_i^{\setminus L/} \cup (Y_i^{\setminus L/} \setminus S_i^{\setminus L/}) = (Y_i \cup S_i)^{\setminus L/}.$$

Thus as $(Y_1, S_1, Z_1)^{\setminus M/} \preceq (Y_2, S_2, Z_2)^{\setminus M/}$, we have $(S_1 \cup Z_1)^{\setminus M/} \subseteq (S_2 \cup Z_2)^{\setminus M/}$. By previous remark that no edge of \bar{L} is contained in Z_i , this implies that $(S_1 \cup Z_1)^{\setminus L/} \subseteq (S_2 \cup Z_2)^{\setminus L/}$. Likewise $(S_1 \cup Y_1)^{\setminus L/} \subseteq (S_2 \cup Y_2)^{\setminus L/}$. As a result, $(Y_1, S_1, Z_1)^{\setminus L/} \preceq (Y_2, S_2, Z_2)^{\setminus L/}$. By Theorem 5.1.33 (2), $(Y_1, S_1, Z_1)^{\setminus L/} \in \mathcal{T}_{\min}^{\setminus L/}$, thus we have $(Y_1, S_1, Z_1)^{\setminus L/} = (Y_2, S_2, Z_2)^{\setminus L/}$ and $(Y_1, S_1, Z_1)^{\setminus M/} = (Y_2, S_2, Z_2)^{\setminus M/}$, showing that $(Y_1, S_1, Z_1)^{\setminus M/} \in \mathcal{T}_{\min}^{\setminus M/}$.

We now prove that $E_{\text{nd}}^{\times}(G^{\setminus M/}) = \emptyset$. Assume that there are two crossing non-degenerate minimal 3-separations (Y_1'', S_1'', Z_1'') and (Y_2'', S_2'', Z_2'') in $\mathcal{T}^{\setminus M/}$, let $L = M(S_1'' \cup S_2'')$. For $i \in \{1, 2\}$, all crossed edges of $G^{\setminus L/}$ lie in Y_i'' or in Z_i'' , hence $(Y_i', S_i', Z_i') = ((Y_i'')_{/\bar{L}}, S_i'', (Z_i'')_{/\bar{L}})$ is the only 3-separation of $\text{Sep}_{<4}(G^{\setminus L/})$ such that $(Y_i', S_i', Z_i')^{\setminus L/} = (Y_i'', S_i'', Z_i'')$. Since $(Y_i'', S_i'', Z_i'') \in \mathcal{T}_{\min}^{\setminus M/}$, from what we just proved, we must have $(Y_i', S_i', Z_i') \in \mathcal{T}_{\min}^{\setminus M/}$. Note that the separations (Y_i', S_i', Z_i') are non-degenerate in $G^{\setminus L/}$. Furthermore, as $L = M(S_1'' \cup S_2'')$, (Y_1', S_1', Z_1') and (Y_2', S_2', Z_2') must be also crossing in $G^{\setminus L/}$, but this contradicts $E_{\text{nd}}^{\times}(T^{\setminus L/}) = E_{\text{nd}}^{\times}(\mathcal{T}) \setminus L$ (third item of Theorem 5.1.33). \square

For each $L \subseteq M$, we let $R^{\setminus L/} := R_{\mathcal{T}}^{\setminus L/}$. Note that for each $L \subseteq M$, $(R_{\mathcal{T}^{\setminus L/}})_{/\bar{L}} = R_{\mathcal{T}}$. Thus together with Lemma 5.1.23, this immediately gives the following, which is the locally finite extension of one of the main results from [Gro16a]:

Theorem 5.1.39

Let G be a locally finite 3-connected graph, and let \mathcal{T} be a region tangle of order 4 in G . Let $M := E_{\text{nd}}^{\times}(\mathcal{T})$ be the set of crossed edges between non-degenerate minimal separations of \mathcal{T} . Then the graph $G^{\setminus M/}[\![R^{\setminus M/}]\!]$ is a quasi-4-connected minor of G .

In order to obtain a proof of Theorem 5.1.7 in the locally finite case, one can either reuse the arguments from [Gro16a, Section 5], or equivalently adapt our proof from Subsection 5.2.3.

5.1.7 Planarity after uncontracting crossed edges

The difficulty is that in general, nice properties of $G^{\setminus M/}[\![R^{\setminus M/}]\!]$ are not satisfied anymore by $G[\![R_{\mathcal{T}}]\!]$. To circumvent this and find a quasi-4-connected region in G , it is proved in [Gro16a] that for every subset X' of $R_{\mathcal{T}}$ obtained by deleting one endpoint of each edge of M , the graph $G[\![X']\!]$ is isomorphic to $G^{\setminus M/}[\![R^{\setminus M/}]\!]$. However we cannot choose such a subset X' canonically in general, as illustrated in Example 5.1.8. Despite the fact that uncontracting the edges of M does not preserve the quasi-4-connectivity of torsos, we now prove that at least planarity is preserved by this operation.

Proposition 5.1.40

If $G^{\setminus M}/\llbracket R^{\setminus M} \rrbracket$ is planar, then so is $G\llbracket R_{\mathcal{T}} \rrbracket$.

This is obtained by combining the following two lemmas:

Lemma 5.1.41

For every subset $L \subseteq M$, we denote $\bar{L} := M \setminus L$. Assume that for every finite subset $L \subseteq M$, $G^{\setminus \bar{L}}/\llbracket R_{\mathcal{T}}^{\setminus \bar{L}} \rrbracket$ is planar. Then $G\llbracket R_{\mathcal{T}} \rrbracket$ is also planar.

Proof. Assume for the sake of contradiction that G enjoys the properties described above but that $G\llbracket R_{\mathcal{T}} \rrbracket$ is not planar. Then by Wagner's theorem [Wag37], $G\llbracket R_{\mathcal{T}} \rrbracket$ admits F as a minor, for some $F \in \{K_5, K_{3,3}\}$. Note that we can find a model $(V_v)_{v \in V(F)}$ of F such that each set V_v is finite. Then $X := \bigcup_{v \in V(F)} V_v$ is a finite subset of $V(G)$ and as G is locally finite, there are only finitely many edges in $M(X^{\setminus M/})$ (recall that for each subset $X' \subseteq V(G^{\setminus M/})$, $M(X')$ is the set of crossed edges of M that contract to a vertex in X'). We let $L := M(X^{\setminus M/})$ denote this finite set of edges and note that the sets V_v are also subsets of $V(G^{\setminus L/})$. It follows that $(V_v)_{v \in V(F)}$ is also a model of F in $G^{\setminus \bar{L}}/\llbracket R^{\setminus \bar{L}} \rrbracket$, a contradiction. \square

Lemma 5.1.42 Planar contraction of a single crossed edge

Let G be locally finite and 3-connected, and \mathcal{T} be a region tangle of order 4 in G . Let (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) be two minimal non-degenerate crossing separations of \mathcal{T} . Let $s_1 s_2$ be the corresponding crossed edge and G' be the graph obtained from G after contracting $s_1 s_2$. Let $(Y'_i, S'_i, Z'_i) := (Y_i, S_i, Z_i)^\vee$ be the projection of (Y_i, S_i, Z_i) to G' for each $i \in \{1, 2\}$. Let $R := R_{\mathcal{T}} \subseteq V(G)$ and $R' := R^\vee$. If $G'\llbracket R' \rrbracket$ is planar, then so is $G\llbracket R \rrbracket$.

Proof. We let $H := G\llbracket R \rrbracket$ and $H' := G'\llbracket R' \rrbracket$ and for $i \in \{1, 2\}$, we write $S_i = \{s_i, t_i, r_i\}$ such that $s_1 s_2$ is the crossed edge between (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) . Since (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) are crossing, the edge $s_1 s_2$ belongs to $E(G[R]) \subseteq E(H)$. Note that in particular we have $6 = |S_1 \cup S_2| \leq |R|$.

Claim 5.1.43

The neighborhood of s_1 in H is:

$$N_H(s_1) = \{s_2\} \cup (\text{fc}(S_2) \setminus \{s_1\}),$$

and $\text{fc}(S_2)$ is a triangle in H .

Proof of the Claim: Note that by definition of the torso, the projection $(Y_i \cap R, S_i \cap R, Z_i \cap R)$ of (Y_i, S_i, Z_i) to H is a separation of H . Hence the only possible neighbors of s_1 in H must lie in $(R \cap Z_1 \cap Y_2) \cup \{t_2, r_2\}$ (see Subfigure 5.5 (b)).

Note that as G is 3-connected, H must also be 3-connected: this comes from the fact that $|V(H)| \geq 6$ and from the observation that any separator $S \subseteq R = V(H)$

of H is also a separator of G . Thus in particular every vertex of H has degree at least 3.

Then, as $|\{s_2\} \cup (\text{fc}(S_2) \setminus \{s_1\})| = 3$, it is enough to prove that $N_H(s_1) \subseteq \{s_2\} \cup (\text{fc}(S_2) \setminus \{s_1\})$ as equality will be immediately implied as $d_H(s_1) \geq 3$. For this we let $t \in N_H(s_1) \setminus \{s_2\}$. We distinguish two cases:

- If $t \in S_2$, then without loss of generality let $t = t_2$. First note that if t is not an endpoint of some crossed edge then $t \in S_2 \cap \text{fc}(S_2)$ and there is nothing to prove. Thus we assume that there exists a crossed edge t_2s_3 incident to t_2 for some s_3 and we prove that this case implies a contradiction, which will imply the desired inclusion. As $t_2 \neq s_2$ and $E_{\text{nd}}^{\times}(\mathcal{T})$ is a matching, we have $s_3 \neq s_1$ and there exists $(Y_3, S_3, Z_3) \in \mathcal{T}_{\text{nd}}$ that crosses (Y_2, S_2, Z_2) via the crossed edge t_2s_3 . As (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) cross, we have $s_1 \in Y_2$. As (Y_2, S_2, Z_2) and (Y_3, S_3, Z_3) cross, we have $t_2 \in Y_3$ and $S_3 \setminus \{s_3\} \subseteq Z_2$. As we assumed that $s_1t_2 \in E(H)$, $s_1 \neq s_3$ and as $t_2 \in Y_3$, we must have $s_1 \in Y_3 \cup (S_3 \setminus \{s_3\})$. This implies a contradiction as $Y_2 \cap Y_3 = \emptyset$ and $Y_2 \cap (S_3 \setminus \{s_3\}) = \emptyset$.
- If $t \notin S_2$, then we must have: $t \in R \cap Y_2 \cap Z_1$ and by definition of R , as $t \notin \bigcap_{(Y,S,Z) \in \mathcal{T}_{\text{nd}}} Z$, this means that $t \in S_3$ for some $(Y_3, S_3, Z_3) \in \mathcal{T}_{\text{nd}} \setminus \{(Y_1, S_1, Z_1), (Y_2, S_2, Z_2)\}$. By Lemma 5.1.22, (Y_2, S_2, Z_2) and (Y_3, S_3, Z_3) are either orthogonal or crossing. If we were in the former case, then we should have $S_3 \cap Y_2 = \emptyset$, which is impossible as $t \in S_3 \cap Y_2$. Hence (Y_2, S_2, Z_2) and (Y_3, S_3, Z_3) are crossing, and if s_3 denotes the endpoint of the crossed edge between S_2 and S_3 , as $S_3 \cap Y_2 = \{s_3\}$ we must have $t = s_3$ so we are done as $s_3 \in \text{fc}(S_2) \setminus \{s_1\}$.

The fact that $\text{fc}(S_2)$ forms a clique follows from [Gro16a, Lemma 4.33]. \diamond

Note that by symmetry, Claim 5.1.43 also implies that we have $N_H(s_2) = \{s_1\} \cup (\text{fc}(S_1) \setminus \{s_2\})$ and that $\text{fc}(S_1)$ is a clique.

Recall that by Wagner's theorem [Wag37], a graph is planar if and only if it is K_5 and $K_{3,3}$ -minor-free. Hence, it is enough to prove that if H contains K_5 or $K_{3,3}$ as a minor, then so does H' . We write $\text{fc}(S_i) = \{s_{3-i}, u_i, v_i\}$ for $i \in \{1, 2\}$, and we recall that the vertex of H' resulting from the contraction of s_1 and s_2 is denoted by s' .

Claim 5.1.44

If H contains a K_5 -minor, then so does H' .

Proof of the Claim: Let (V_1, \dots, V_5) be a model of K_5 in H . Let V'_1, \dots, V'_5 be the projection of the sets V_i to H' . If s_1 and s_2 are in the same set V_i , then (V'_1, \dots, V'_5) is also a model of K_5 in H' , so we can assume that the vertices s_1 and s_2 belong to distinct sets V_i , say $s_1 \in V_1$ and $s_2 \in V_2$. As by Claim 5.1.43, s_1 has degree 3




Figure 5.6: Left: The graph G when S_2 is incident to exactly 2 crossededges. Here t_2 is part of no crossededge and S_2 and S_3 are crossing via the crossededge r_2s_3 . Hence, $\text{fc}(S_2) = \{s_1, t_2, s_3\}$.

Right: The graph G' obtained after contracting the crossededge s_1s_2 . The dashed edges are edges that appear in H and H' respectively. The situation is identical when S_2 is incident to 3 crossededges, but harder to illustrate in 2 dimensions.

in H , we have $V_1 \neq \{s_1\}$, so V_1 contains one neighbor of s_1 distinct of s_2 , say u_2 . Since u_2 and v_2 are adjacent in H , the edge u_2s' in H' has an endpoint in $V_1' \setminus \{s'\}$ and an endpoint in V_2' . Moreover as $u_2v_2 \in E(H')$, the set $V_1' \setminus \{s'\}$ is connected in H' . Thus $(V_1' \setminus \{s'\}, V_2', V_3', V_4', V_5')$ is a model of K_5 in H' , as desired. \diamond

Claim 5.1.45

If H contains a $K_{3,3}$ -minor, then H' contains a $K_{3,3}$ -minor or a K_5 -minor.

Proof of the Claim: Let (V_1, \dots, V_6) be a model of $K_{3,3}$ in H , such that V_i is adjacent to V_j if i and j have different parities. Let V_1', \dots, V_6' be their projection to H' . If s_1 and s_2 are in the same set V_i , then (V_1', \dots, V_6') is also a model of $K_{3,3}$ in H' , so we can assume that the vertices s_1 and s_2 belong to distinct sets V_i , say $s_1 \in V_1$ and $s_2 \notin V_1$.

If $u_2 \in V_1$, then the edges $s'u_2$ and u_2v_2 in H' ensure that $V_1' \setminus \{s'\}$ remains connected and that $(V_1' \setminus \{s'\}, V_2', \dots, V_6')$ is a model of $K_{3,3}$ in H' . Thus we can assume that $u_2 \notin V_1$ and similarly $v_2 \notin V_1$. Since V_1 is connected, we must have $V_1 = \{s_1\}$.

As s_1 has degree three and V_1 is adjacent to V_2, V_4 and V_6 , this implies that s_2, u_2 and v_2 must belong to different sets V_{2i} , say $s_2 \in V_2, u_2 \in V_4$ and $v_2 \in V_6$. By applying the same reasoning as for s_2 , we obtain $V_2 = \{s_2\}, u_1 \in V_3$ and $v_1 \in V_5$. But then $(\{s'\}, V_3', V_4', V_5', V_6')$ is a model of K_5 in H' . \diamond

This concludes the proof of Lemma 5.1.42. \square

Proof of Proposition 5.1.40. Assume that $G^{\setminus M}/[[R^{\setminus M}]]$ is planar and for every $L \subseteq M$, set $\bar{L} := M \setminus L$. Then, using Lemma 5.1.42, we can easily prove by induction

on $|L| \in \mathbb{N}$ that for any *finite* set $L \subseteq M$, $G^{\setminus \bar{L}} / \llbracket R_{\mathcal{T}}^{\setminus \bar{L}} \rrbracket$ is planar. In order to be able to use induction, we also need to observe that for the contraction of a single crossed edge, the equality $R_{\mathcal{T}'} = R_{\mathcal{T}}^{\vee}$ holds. This is proved in [Gro16a, Section 4.5] and can be deduced from item (2) of Theorem 5.1.33. We thus conclude by Lemma 5.1.41 that $G \llbracket R_{\mathcal{T}} \rrbracket$ is also planar. \square

5.2 The structure of quasi-transitive graphs avoiding a minor

5.2.1 Main results

Our main result in this section is the following more precise version of Theorem 5.0.1.

Theorem 5.2.1

Let G be a locally finite graph excluding K_{∞} as a minor and let Γ be a group with a quasi-transitive action on G . Then there is an integer k such that G admits a Γ -canonical tree-decomposition (T, \mathcal{V}) , with $\mathcal{V} = (V_t)_{t \in V(T)}$, whose torsos $G \llbracket V_t \rrbracket$ either have size at most k or are Γ_t -quasi-transitive 3-connected planar minors of G . Moreover, the edge-separations of (T, \mathcal{V}) are tight.

Remark 5.2.2. A natural question is whether we can bound the maximum size k of a finite bag in Theorem 5.2.1 by a function of the forbidden minor, when G excludes some finite minor instead of the countable clique K_{∞} . By taking the free product of the cyclic groups \mathbb{Z}_k and \mathbb{Z} , with their natural sets of generators, we obtain a 4-regular Cayley graph consisting of cycles of length k arranged in a tree-like way. This graph has no K_4 minor, but in any canonical tree-decomposition, each cycle C_k has to be entirely contained in a bag, and thus there is no bound on the size of a bag as a function of the forbidden minor in Theorem 5.2.1. We can replace \mathbb{Z}_k in this construction by the toroidal grid $\mathbb{Z}_k \times \mathbb{Z}_k$, and obtain a Cayley graph with no K_8 -minor, such that the bags in any (non-necessarily canonical) tree-decomposition of finitely bounded adhesion are arbitrarily large.

We will also prove the following version of Theorem 5.0.2 at the same time.

Theorem 5.2.3

Let G be a locally finite graph excluding K_{∞} as a minor and let Γ be a group with a quasi-transitive action on G . Then there is an integer k such that G admits a Γ -canonical tree-decomposition (T, \mathcal{V}) , with $\mathcal{V} = (V_t)_{t \in V(T)}$, of adhesion at most 3, and whose torsos $G \llbracket V_t \rrbracket$ are Γ_t -quasi-transitive minors of G which are either planar or have treewidth at most k . The edge-separations of (T, \mathcal{V}) are all non-degenerate.

Remark 5.2.4. If we carefully consider the proof of Theorem 5.2.3, we can check that if G has only one end (as in the example of Figure 5.4), then the tree-decomposition we obtain has adhesion 3 and consists of a star with one infinite bag associated to its central vertex z_0 , and finite bags on its branches. In particular, $G[[V_{z_0}]]$ cannot have bounded treewidth, as otherwise it would have more than one end, hence it must be planar. Thus, Theorem 5.2.3 implies that every one-ended locally finite quasi-transitive graph that excludes a minor can be obtained from a one-ended quasi-transitive planar graph by attaching some finite graphs on it along separators of order at most 3.

5.2.2 Tools

Our proof of Theorems 5.2.1 and 5.2.3 mainly consists in an application of Theorem 5.1.39 together with the following result of Thomassen:

Theorem 5.2.5 Theorem 4.1 in [Tho92]

Let G be a locally finite, quasi-transitive, quasi-4-connected graph G . If G has a thick end, then G is either planar or admits the countable clique K_∞ as a minor.

A direct consequence of Theorem 5.2.5 is the following, which will be our base case in what follows.

Corollary 5.2.6

Let G be a quasi-transitive, quasi-4-connected, locally finite graph which excludes the countable clique K_∞ as a minor. Then G is planar or has finite treewidth.

Proof. Assume that G is non-planar. As G is K_∞ -minor-free, by Theorem 5.2.5, all its ends have finite degree. Then by Theorem 5.1.5, G has finite treewidth. \square

Thomassen proved that if a quasi-transitive graph has only one end, then this end must be thick [Tho92, Proposition 5.6]. We prove the following generalization, which might be of independent interest.

Proposition 5.2.7

Let $k \geq 1$ be an integer, and let G be a locally finite quasi-transitive graph. Then G cannot have exactly one end of degree exactly k .

Proof. Assume without loss of generality that G is connected, since otherwise each component of G is also quasi-transitive locally finite, and we can restrict ourselves to a single component containing an end of degree exactly k . Let Γ be a group acting quasi-transitively on G . Assume that G has an end ω of degree exactly k for some integer $k \geq 1$. As explained in [TW93, Section 4], there exists an infinite sequence of sets S_0, S_1, \dots of size k such that for each $i \geq 0$, S_{i+1} belongs

to the component G_i of $G_{i-1} - S_i$ where ω lives (where we set $G_0 := G$), and such that there exist k vertex-disjoint paths $P_{1,i}, \dots, P_{k,i}$ from the k vertices of S_i to the k vertices of S_{i+1} . By concatenating these paths, we obtain k vertex-disjoint rays in G living in ω . As G is connected and locally finite, note that up to extracting a subsequence of $(S_i)_{i \geq 0}$, we may assume that the k paths $P_{1,i}, \dots, P_{k,i}$ are in the same component of $G - (S_i \cup S_{i+1})$. Hence if we set $(Y_i, S_i, Z_i) := (G - (G_i \cup S_i), S_i, G_i)$ for each $i \geq 1$, (Y_i, S_i, Z_i) is a tight separation such that for each $i \geq 1$, $(Y_{i+1}, S_{i+1}, Z_{i+1}) \prec (Y_i, S_i, Z_i)$. Hence by Lemma 5.1.1, as there are only finitely many Γ -orbits of tight separations of size k , there exist $i < j$ and $g \in \Gamma$ such that $(Y_i, S_i, Z_i) \cdot g = (Y_j, S_j, Z_j)$. Assume without loss of generality that $(i, j) = (0, 1)$. Note that by definition of \prec , the action of g preserves the order \prec , i.e. for each $(Y, S, Z) \prec (Y', S', Z')$, we must have $(Y, S, Z) \cdot g \prec (Y', S', Z') \cdot g$. We now consider the sequence of separations $(Y'_i, S'_i, Z'_i)_{i \geq 0}$ defined for each $i \geq 0$ by: $(Y'_i, S'_i, Z'_i) := (Y_0, S_0, Z_0) \cdot g^i$. Then the sequence $(Y'_i, S'_i, Z'_i)_{i \geq 0}$ is strictly decreasing according to \prec . Recall that there exist k vertex-disjoint paths from S_0 to S_1 that extend to k disjoint rays belonging to ω . Then for each $i \geq 0$, there exist k vertex-disjoint paths from S'_i to S'_{i+1} such that their concatenations consists in k vertex-disjoint rays that belong to some end ω' of degree exactly k (the fact that the end has degree at most k follows from the fact that all the sets S'_i are separators of size k in G). If $\omega' \neq \omega$ then we are done, so we assume that $\omega' = \omega$. Now, observe that the sequence $(Y''_i, S''_i, Z''_i)_{i \geq 0}$ defined for each $i \geq 0$ by $(Y''_i, S''_i, Z''_i) := (Y_0, S_0, Z_0) \cdot g^{-i}$ also satisfies that for each $i \geq 0$, there exists k vertex-disjoint paths $P''_{j,i} := P_{j,0} \cdot g^{-i}$ for $j \in [k]$ from S''_{i+1} to S''_i . If we consider the k vertex-disjoint rays obtained from the concatenation of the paths $P''_{j,k}$, these rays must belong to the same end ω'' as for each i , the paths $P''_{j,k}$ are in the same component of $G - (S''_i \cup S''_{i+1})$. The end ω'' must have degree exactly k as each (Y''_i, S''_i, Z''_i) is a separation of order k . Moreover the sequence $(Y''_i, S''_i, Z''_i)_{i \geq 0}$ is strictly increasing according to \prec , hence ω and ω'' cannot live in the same component of $G - S_0$. Thus we found an end ω'' distinct from ω of degree k . \square

Proposition 5.2.7 and its proof are reminiscent of Halin's classification of the different types of action an automorphism of a quasi-transitive locally finite graph G can have on the ends of G [Hal73, Theorem 9]. However it is not clear for us whether Proposition 5.2.7 can be seen as an immediate corollary of Halin's work.

5.2.3 Proof of Theorems 5.2.1 and 5.2.3

Let G be a locally-finite quasi-transitive graph excluding K_∞ as a minor and let Γ be a group inducing a quasi-transitive action on G . Let (T, \mathcal{V}) , with $\mathcal{V} = (V_t)_{t \in V(T)}$, be a Γ -canonical tree-decomposition of adhesion at most 2 obtained by applying Theorem 5.1.6 to G . By Lemma 5.1.13, for each $t \in V_t$, Γ_t acts quasi-transitively

on $G_t := G[V_t]$. Moreover, as G_t is a minor of G , it must also exclude K_∞ as a minor.

We let t_1, \dots, t_m be representatives of the orbits of $V(T)/\Gamma$. For each finite torso G_{t_i} of (T, \mathcal{V}) , we define $(\tilde{T}_{t_i}, \tilde{\mathcal{V}}_{t_i})$ as the trivial tree-decomposition of G_{t_i} (in which the tree \tilde{T}_{t_i} contains a single node). For each infinite, 3-connected torso G_{t_i} of (T, \mathcal{V}) , we let $(\tilde{T}_{t_i}, \tilde{\mathcal{V}}_{t_i})$ be a Γ_{t_i} -canonical tree-decomposition of G_{t_i} obtained by applying Theorem 5.1.16 to G_{t_i} , i.e. $(\tilde{T}_{t_i}, \tilde{\mathcal{V}}_{t_i})$ distinguishes efficiently all the tangles of G_{t_i} of order 4. By Remark 5.1.17, the edge-separations of $(\tilde{T}_{t_i}, \tilde{\mathcal{V}}_{t_i})$ in G_{t_i} are all distinct. By Remark 5.1.20, the edge-separations of $(\tilde{T}_{t_i}, \tilde{\mathcal{V}}_{t_i})$ in G_{t_i} are non-degenerate. Hence by Lemma 5.1.19, the torsos of $(\tilde{T}_{t_i}, \tilde{\mathcal{V}}_{t_i})$ are minors of G_{t_i} . We now use Corollary 5.1.12 and find a refinement (T_1, \mathcal{V}_1) of (T, \mathcal{V}) with respect to some family $(T_t, \mathcal{V}_t)_{t \in V(T)}$ of Γ_t -canonical tree-decompositions of G_t such that the construction $t \mapsto (T_t, \mathcal{V}_t)_{t \in V(T)}$ is Γ -canonical and such that for each $i \in I$, $(T_{t_i}, \mathcal{V}_{t_i})$ is a subdivision of $(\tilde{T}_{t_i}, \tilde{\mathcal{V}}_{t_i})$. Since the construction $t \mapsto (T_t, \mathcal{V}_t)_{t \in V(T)}$ is Γ -canonical, for each $t \in V(T_1)$ the decomposition (T_t, \mathcal{V}_t) is Γ_t -canonical and efficiently distinguishes the tangles of order 4 of G_t (by a slight abuse of notation, we keep denoting by G_t the torso of the tree-decomposition (T_1, \mathcal{V}_1) associated to the node $t \in V(T_1)$). Note that by construction, the adhesion sets of (T_1, \mathcal{V}_1) have size at most 3 and all the edge-separations are tight. Moreover, the torsos of each tree-decomposition (T_t, \mathcal{V}_t) are minors of G_t for each $t \in V(T)$, and as the torsos of (T, \mathcal{V}) are minors of G , we also have that the torsos of (T_1, \mathcal{V}_1) are minors of G . In particular, they also exclude K_∞ as a minor. Moreover, by Lemma 5.1.13, for each $t \in V(T_1)$, Γ_t acts quasi-transitively on G_t . By Lemma 5.1.1, since all edge-separations of (T_1, \mathcal{V}_1) are tight and have order at most 3, the graph G_t is locally finite for each $t \in V(T_1)$.

Claim 5.2.8

For each $t \in V(T_1)$ such that G_t is infinite, G_t is 3-connected and has a unique tangle \mathcal{T}_t of order 4. Moreover \mathcal{T}_t is a Γ_t -invariant region tangle and every end of G_t has degree at least 4.

Proof of the Claim: Consider a node $t \in V(T_1)$ such that G_t is infinite. As all torsos are cycles, subgraphs of complete graphs of size at most 3, or 3-connected, G_t itself is 3-connected. Since G_t is connected and infinite, it contains some end ω . Let $\mathcal{T}_t := \{(Y, S, Z), |S| \leq 3 \text{ and } \omega \text{ lives in } Z\}$ be defined in G_t . Note that \mathcal{T}_t is a tangle of order 4 in G_t . As G_t is a minor of G , by Lemma 5.1.14 every tangle \mathcal{T}' of order 4 in G_t induces a tangle \mathcal{T} of order 4 in G , and by Remark 5.1.15 this mapping is injective. Moreover, note that if (Y, S, Z) is an edge-separation of (T_1, \mathcal{V}_1) such that $V_t \subseteq Z \cup S$, then if \mathcal{M} is any faithful model of G_t in G , the projection $(Y', S', Z') := \pi_{\mathcal{M}}(Y, S, Z)$ is such that $Y' = \emptyset$. Thus $(Y', S', Z') \in \mathcal{T}'$, hence $(Y, S, Z) \in \mathcal{T}$. This means that every edge-separation of (T_1, \mathcal{V}_1) is oriented

toward t by \mathcal{T} . Hence if G_t admits two distinct tangles $\mathcal{T}'_1, \mathcal{T}'_2$ of order 4, the two associated tangles $\mathcal{T}_1, \mathcal{T}_2$ given by Lemma 5.1.14 must be distinct and not distinguished by (T_1, \mathcal{V}_1) , a contradiction. This proves the existence and uniqueness of a tangle \mathcal{T}_t of order 4 in G_t .

Note that as Γ_t acts on G_t and \mathcal{T}_t is the unique tangle of order 4 in G_t , the tangle \mathcal{T}_t is Γ_t -invariant (as a family of separations).

We can also observe that if the end ω in G_t has degree at most 3, then by Proposition 5.2.7, G_t has another end ω' of degree at most 3 and the construction of \mathcal{T}_t using the end ω' instead of ω yields a different tangle of order 4, which contradicts the uniqueness of \mathcal{T}_t . So every end of G_t has degree at least 4.

It remains to prove that \mathcal{T}_t is a region tangle. If not we can find an infinite decreasing sequence of separations of order 3 in G_t , and this sequence defines an end of degree 3 in G_t , which contradicts the fact that every end of G_t has degree at least 4. \diamond

We will need to decompose further the infinite torsos of the tree-decomposition (T_1, \mathcal{V}_1) . Let $t \in V(T_1)$ be such that G_t is infinite, and let \mathcal{T}_t be the region tangle of order 4 in G_t given by Claim 5.2.8. We let $M_t := E_{\text{nd}}^\times(\mathcal{T}_t)$ denote the set of crossed edges of \mathcal{T}_t , (T'_t, \mathcal{V}'_t) be the Γ_t -canonical tree-decomposition of G_t given by Lemma 5.1.25, and $z_0 \in V(T'_t)$ be the center of the star T'_t . By Lemmas 5.1.13 and 5.1.24, the graph $H := G_t \llbracket V'_{z_0} \rrbracket$ is a Γ_t -quasi-transitive faithful minor of G_t , thus it must also exclude K_∞ as a minor.

Now we observe that Γ_t induces a quasi-transitive group action on $H^{\setminus M_t/}$: for each $w \in V(H^{\setminus M_t/})$ and every $\gamma \in \Gamma_t$, we set:

$$w \cdot \gamma := \begin{cases} s_{u \cdot \gamma, v \cdot \gamma} & \text{if } w = s_{u, v}, \text{ for some } \{u, v\} \in M_t, \text{ and} \\ w \cdot \gamma & \text{otherwise,} \end{cases}$$

where we recall that the notation $s_{u, v}$, for $\{u, v\} \in M_t$, is introduced at the beginning of Subsection 5.1.6. As M_t is Γ_t -invariant, we easily see that the mapping γ defines a bijection over $V(H^{\setminus M_t/})$. We let the reader check that it gives a graph isomorphism of $H^{\setminus M_t/}$. Note that the number of Γ_t -orbits of $V(H^{\setminus M_t/})$ is at most the number of Γ_t -orbits of $V(H)$, hence it must be finite.

As $H^{\setminus M_t/}$ is a minor of H , it also excludes the countable clique K_∞ as a minor. It follows from Theorem 5.1.39 that $H^{\setminus M_t/}$ is quasi-4-connected. Hence, by Corollary 5.2.6, $H^{\setminus M_t/}$ either has finite treewidth or it is planar. It is not hard to observe that the treewidth of H is at most twice the treewidth of $H^{\setminus M_t/}$ so in particular if we are in the first case, H has also bounded treewidth. In the second case, Proposition 5.1.40 implies that H is also planar. In both cases, we obtain that (T'_t, \mathcal{V}'_t) is a Γ_t -canonical tree-decomposition of G_t with non-degenerate edge-separations, adhesion 3 and where each torso is a minor of G_t and has either

bounded treewidth or is planar. Eventually we can use Proposition 5.1.10 together with Lemma 5.1.9 as we did before to find a tree-decomposition (T^*, \mathcal{V}^*) of G with the properties of Theorem 5.2.3.

We now explain how to derive Theorem 5.2.1: every torso $G[[V_t]]$ of (T^*, \mathcal{V}^*) which is neither finite nor planar must have bounded treewidth, hence by Theorem 5.1.5 it must admit a Γ_t -canonical tree-decomposition where each torso has bounded width. Exactly as before we can apply Corollary 5.1.12 to find a refinement of (T^*, \mathcal{V}^*) with the properties of Theorem 5.2.1. \square

5.3 Applications

5.3.1 The Hadwiger number of quasi-transitive graphs

We say that a graph H is *singly-crossing* if H can be embedded in the plane with a single edge-crossing. It was observed by Paul Seymour that Theorem 5.2.3 bears striking similarities with a structure theorem of Robertson and Seymour [RS93] related to the exclusion of a singly-crossing graph as a minor. Their theorem states that if H is singly-crossing, then there is a constant k_H such that any graph excluding H as a minor has a tree-decomposition with adhesion at most 3 in which all torsos are planar or have treewidth at most k_H . On the other hand, for any integer k there is a finite singly-crossing graph H_k such that any graph with a tree-decomposition with adhesion at most 3 in which all torsos are planar or have treewidth at most k must exclude H_k as a minor (this can be seen by taking H_k to be a 4-connected triangulation of a sufficiently large grid, and adding an edge between two non-adjacent vertices lying on incident faces). Using this observation, the following strengthening of Theorem 5.0.3 is now an immediate consequence of Theorem 5.2.3.

Theorem 5.3.1

For every locally finite quasi-transitive graph G avoiding the countable clique K_∞ as a minor, there is a finite singly-crossing graph H such that G is H -minor-free. In particular there is an integer k such that G is K_k -minor-free.

Note that in this application we have not used explicitly the property that the underlying tree-decomposition was canonical, but it is used implicitly in the sense that this is what guarantees that the treewidth of the torsos is uniformly bounded in Theorem 5.2.3.

5.3.2 The Domino Problem in minor-excluded groups

In this section we prove the following version of Theorem 5.0.8.

Theorem 5.3.2

Let Γ be a finitely generated group excluding the countable clique K_∞ as a minor. Then the Domino Problem is undecidable on Γ if and only if

- Γ is one-ended, or
- Γ has an infinite number of ends and has a one-ended planar subgroup which is finitely generated.

In particular, these situations correspond exactly to the cases where Γ is not virtually free.

We will need the following results on one-ended planar groups. One-ended planar groups are exactly infinite *planar discontinuous groups*, i.e. groups that act properly discontinuously on simplicial complexes homeomorphic to \mathbb{R}^2 where every nontrivial automorphism fixes neither a face nor maps an edge (u, v) to its inverse (v, u) (see [ZVC80, Section 4] for a complete survey on these groups).

Theorem 5.3.3 [BN46, Fox52]

Every planar discontinuous group contains the fundamental group of a closed orientable surface as a subgroup of finite index.

It follows that every one-ended planar group contains the fundamental group of a closed orientable surface of genus $g \geq 1$ as a subgroup (if $g = 0$ the fundamental group is trivial and cannot be a subgroup of finite index). We will use Theorem 5.3.3 in combination with the following result of Aubrun, Barbieri and Moutot [ABM18] on the Domino Problem in surface groups (note that the case $g = 1$ corresponds to the undecidability of the domino problem in \mathbb{Z}^2).

Theorem 5.3.4 [ABM18]

For any $g \geq 1$, the Domino Problem in the fundamental group of the closed orientable surface of genus g is undecidable.

For groups with infinitely many ends we will use the existence of finite graphs of groups (which follows from the accessibility of the groups we consider). The following result of Babai [Bab77] will be crucial.

Theorem 5.3.5 [Bab77]

If Γ' is a finitely generated subgroup of Γ , then there exists some finite generating set S' of Γ' such that $\text{Cay}(\Gamma', S')$ is a minor of $\text{Cay}(\Gamma, S)$.

Finally we will need the following result of Thomassen, which is tightly related to Theorem 5.2.5, and which will allow us to reduce the K_∞ -minor-free case to the planar case.

Theorem 5.3.6 Theorem 5.7 in [Tho92]

Let G be a locally finite, transitive, non-planar, one-ended graph. Then G contains the countable clique K_∞ as a minor.

We are now ready to prove Theorem 5.3.2.

Proof. Assume that Γ has a Cayley graph $G = \text{Cay}(\Gamma, S)$ excluding K_∞ as a minor. If Γ has 0 end, then it is finite and the Domino Problem is decidable. If Γ has 2 ends then Γ is virtually \mathbb{Z} and the Domino Problem is also decidable. If Γ is one-ended, then since G is transitive and excludes the countable clique K_∞ as a minor, it follows from Theorem 5.3.6 that G is (one-ended) planar. By Theorem 5.3.3, Γ contains the fundamental group Γ' of a closed surface of genus $g \geq 1$ as a subgroup. By Theorem 5.3.4, the domino problem is undecidable for Γ' . It was proved that for every finitely generated subgroup Γ' of a finitely generated group Γ , if the Domino Problem is undecidable for Γ' then it is also undecidable for Γ [ABJ18, Proposition 9.3.30]. This implies that if Γ is one-ended, the Domino Problem is undecidable for Γ .

Assume now that Γ has an infinite number of ends. By Corollary 5.0.5, Γ is accessible and by Theorem 5, if Γ is not virtually free then one of the (finitely presented) vertex-groups Γ_u (which is a subgroup of Γ) in its associated finite graph of groups is one-ended. By Theorem 5.3.5, Γ_u must also have a locally finite Cayley graph that excludes K_∞ as a minor. By the paragraph above, the Domino Problem is undecidable for Γ_u , and since Γ_u is a subgroup of Γ , the Domino Problem is also undecidable for Γ . \square

5.4 Open problems

The first problem we consider is related to Conjecture 5.0.7, which states that a group has a decidable Domino Problem if and only if it is virtually free (or equivalently, it has bounded treewidth). Using the planar case of Theorem 5.3.2, this conjecture would be a direct consequence of a positive answer to the following problem.

Problem 5.4.1

Is it true that every finitely generated group which is not virtually free has a finitely generated one-ended planar subgroup?

It was pointed out to us by Agelos Georgakopoulos (personal communication) that the lamplighter group $L = \mathbb{Z}_2 \wr \mathbb{Z}$ provides a negative answer to Problem 5.4.1. This follows from a characterization of all subgroups of L by Grigorchuk and Kravchenko [GK14], which implies that every finitely generated subgroup of L is

either finite or isomorphic to another lamplighter group (and the latter cannot be planar).

A related problem on the graph theory side is the following.

Problem 5.4.2

Is it true that every locally finite quasi-transitive graph of unbounded treewidth contains a quasi-transitive planar graph of unbounded treewidth as a subgraph?

It could be the case that Cayley graphs of lamplighter groups also provide a negative answer to Problem 5.4.2, but this is not clear (the assumption of being quasi-transitive and subgraph containment are both much weaker than being a Cayley graph and subgroup containment).

A crucial component of our result is the strategy of Grohe [Gro16a] to obtain a canonical tree-decomposition up to the contraction of a matching (see also Theorem 5.1.39). At one point we tried to figure out whether the following was true (in the end it turned out that we did not need to answer these questions, but we believe they might be of independent interest).

Problem 5.4.3

Let G be a locally finite quasi-transitive graph. Is there a proper coloring of G with a finite number of colors such that the colored graph G itself is quasi-transitive (where automorphisms have to preserve the colors of the vertices)?

Problem 5.4.4

Let G be a locally finite quasi-transitive graph. Is there an orientation of the edges of G such that the oriented graph G itself is quasi-transitive (where automorphisms have to preserve the orientation of the edges)?

Note that Problem 5.4.4 has a positive answer for Cayley graphs whose generating set does not contain any element of order 2. An example showing that Problem 5.4.3 has a negative answer was recently constructed by Hamann and Möller. It was then observed by Abrishami, Esperet and Giocanti, and independently by Norin and Przytycki, that a variant of this example could be used to provide a negative answer to Problem 5.4.4 as well.

Finally, we mentioned in the introduction the following natural question: do graphs avoiding a fixed minor have a tree-decomposition in the spirit of the Graph Minor Structure Theorem of Robertson and Seymour [RS03], but with the additional constraint that the decomposition is canonical? We gave a positive answer to this question for quasi-transitive graphs but the general case is still open (although Example 5.1.4 and Remark 5.2.2 give natural limitations to the properties of such a canonical tree-decomposition).

5.4.1 Related work

Independently of our work, Carmesin and Kurkofka [CK23] recently worked on decompositions of 3-connected graphs with an approach that differs from Grohe's approach. They obtained a canonical decomposition into basic pieces consisting in quasi-4-connected graphs, wheels or thickenings of $K_{3,m}$ for $m \geq 0$. It is possible that their approach could also imply some of the applications we describe. However it is not clear to us whether this work directly implies the existence of canonical tree-decompositions with the properties described in Theorems 5.0.1 and 5.0.2, as they consider *mixed separations*, i.e. separations containing both vertices and edges, while we focus on vertex-separations. Vertex-separations yield tree-decompositions, while mixed separations do not yield tree-decompositions in a traditional sense.

Conclusion and Perspectives

This thesis presents results mostly centered around combinatorial reconfiguration and structural graph theory. Chapter 1 contains all the prerequisites to the understanding of this thesis. In the next two chapters, we focus on graph recoloring with Kempe changes. In Chapter 2, we survey the recent results on Kempe reconfiguration, while Chapter 3 presents personal results. In Chapter 4, we study some reconfiguration problems arising from low-dimensional topology: Reidemeister moves on knot diagrams and reconfiguration of square-tiled surfaces with shearing moves. Finally, in Chapter 5, we move away from reconfiguration and give a structure theorem on locally finite quasi-transitive graphs avoiding a minor, with applications to the decidability of the domino problem. We present here a selection of questions that these results raise.

Recoloring with Kempe changes

We point out two main questions on the Kempe recolorability.

Obstacles to Kempe recolorability

What are the obstacles to being k -recolorable? As of today, we only know of two arguments to prove that a graph is not k -recolorable. The first and most general one consists in exhibiting a frozen coloring, while the second is a topological argument of Fisk, Mohar and Salas. However, this topological argument is specific to 4-colorings of planar or toroidal graphs and does not seem to generalize to a higher number of colors or to higher genus. Thus we conjecture that frozen colorings are the only “general” reason behind non-Kempe equivalence:

Question 1 (see Question 2.1.25)

Is there another obstacle to k -recolorability, beside frozen colorings and the topological argument of Fisk, Mohar and Salas?

To tackle this question, we suggest studying two graph classes that cannot admit frozen k -colorings. Triangle-free planar graphs are known to be 3-colorable by

Grötzsch’s theorem but cannot admit frozen 3-colorings by an elementary counting argument. This motivates the following conjecture:

Conjecture 2 (Bonamy, Legrand-Duchesne and [SS22], see Conjecture 2.1.28)
All triangle-free planar graphs are 3-recolorable.

As a first step towards this conjecture, we show that the 3-recolorability of planar graphs of girth 5 implies that of triangle-free planar graphs.

The second graph class we would like to draw attention to are odd-hole-free graphs. A conjecture of Reed places the chromatic number of any graph closer to trivial lower bound (clique number) than to the trivial upper bound (maximum degree plus one). We show in Chapter 2 that all the k -colorings of any odd-hole-free graph G are Kempe equivalent to some $\lceil \frac{\omega(G)+\Delta(G)+1}{2} \rceil$ -coloring, thereby proving Reed’s conjecture on odd-hole-free graphs and proving that they admit no frozen coloring above this threshold. However, the following question remains open:

Conjecture 3 (Bonamy, Kaiser and Legrand-Duchesne, see Conjecture 2.1.42)
All odd-hole free graphs of maximum degree Δ and clique number ω are k -recolorable for $k \geq \lceil \frac{\Delta+1+\omega}{2} \rceil$.

Another motivation to solve Conjecture 3 is that the advances on Reed’s conjecture all used the probabilistic method. So far the vast majority of coloring theorems for which a recoloring version holds can be proved by induction or by directly using Kempe changes. However, we do not know of any k -recolorable graph class whose only proof of k -colorability uses the probabilistic method.

Recoloring at and below the degeneracy threshold

Perhaps the most elementary yet central result in Kempe recoloring is the lemma of Las Vergnas and Meyniel [VM81] that states that d -degenerate graphs are k -recolorable for $k > d$. There are two trade-offs for the simplicity of this proof. The first one is that the corresponding bound on the number of Kempe changes is exponential in the number of vertices. The second one is that going below the degeneracy threshold in more specific graph classes often proves difficult and is achieved thanks to very different techniques.

Addressing the first issue, Bonamy, Bousquet, Feghali and Johnson conjectured the following, which is one of the most inspiring conjectures in the field. Note that it is reminiscent of Cereceda’s recoloring conjecture:

Conjecture 4 ([BBFJ19], see Conjecture 2.2.1)
Let G be an n -vertex d -degenerate graph. Any two k -colorings of G are Kempe equivalent up to $O(n^2)$ Kempe changes for $k \geq d + 1$.

Although we proved polynomial bounds on the recoloring diameter in several related graph classes (see Chapter 3), we do not even know if a polynomial version of Conjecture 2.2.1 holds. Part of the reason is that the Kempe recoloring distance is very poorly understood: as of today, we do not know of any explicit construction of colorings at Kempe recoloring distance superlinear in the number of vertices, even though we know colorings at superpolynomial distance exist assuming that $\text{PSPACE} \neq \text{NP}$.

As an example of the second issue, we advertise this recoloring version of the Linear Hadwiger conjecture:

Conjecture 5 ([2], see Conjecture 2.1.13)

There exists a constant c such that any K_t -minor-free graph is ct -recolorable.

With Marthe Bonamy, Marc Heinrich and Jonathan Narboni [2] we proved that if such c exists, then $c \geq 3/2$, thereby proving that a recoloring version of Hadwiger’s conjecture cannot hold. In the 1980s, [Kos82, Kos84] and [Tho84] proved independently that a graph with no K_t -minor has degeneracy $O(t\sqrt{\log t})$, with current best hidden factor in [KP20]. Hence, the lemma of Las Vergnas and Meyniel [VM81] shows that K_t -minor-free graphs are $O(t\sqrt{\log t})$ -(re)colorable. We do not know of any better bound for Conjecture 2.1.13, unlike the Linear Hadwiger’s conjecture for which the best bound as of today is $O(t \log(\log(t)))$, below the degeneracy threshold [DP21].

Reconfiguration of square-tiled surfaces

In Chapter 4, we give a partial answer to the following conjecture:

Conjecture 6 (Delecroix and Legrand-Duchesne, see Conjecture 4.0.7)

Let μ be an integer partition of n and k a non-negative integer such that

$$\sum_{i \geq 1} \mu_i(i - 2) = 4g - 4 + k.$$

Let S and S' be two square-tiled surfaces in $\text{ST}(\mu, k)$. Then S and S' are equivalent via cylinder shears if and only if they belong to the same connected component of the moduli space of quadratic differentials.

We addressed the case of the Abelian hyperelliptic components and a subset of spherical strata, however this conjecture remains widely open in full generality.

Last but not least, we would like to draw attention to Monte Carlo Markov Chain algorithms that use non-local updates. This non-local character grants the Markov Chain more flexibility in cases where local changes are inefficient, but in

return the ergodicity and mixing time of such Markov chains are much harder to analyze with the state of the art methods. Both cylinder shears and Kempe changes are examples of such operations, that present surprising similarities. We hope that a better understanding of any of these two operations translates to the other. In this direction, we wonder whether the WSK algorithm can efficiently sample random colorings in some meaningful graph classes where the Glauber dynamics mixes torpidly and we ask the following question:

Question 7 (Delecroix and Legrand-Duchesne, see Question 4.2.33)

In which connected component of the moduli space of quadratic differentials does the shear dynamics rapidly mix on the set of square-tiled surfaces?

References

- [AB26] J. W. Alexander and G. B. Briggs. On Types of Knotted Curves. *Annals of Mathematics*, 28(1/4):562–586, 1926.
- [ABJ18] Nathalie Aubrun, Sebastián Barbieri, and Emmanuel Jeandel. About the domino problem for subshifts on groups. In *Sequences, Groups, and Number Theory*, Trends in Mathematics, pages 331–389. Birkhäuser, Cham, 2018.
- [ABM18] Nathalie Aubrun, Sebastián Barbieri, and Etienne Moutot. The domino problem is undecidable on surface groups. *arXiv preprint arXiv:1811.08420*, 2018.
- [AH89] Kenneth I. Appel and Wolfgang Haken. *Every Planar Map Is Four Colorable*. American Mathematical Soc., 1989.
- [Aks74] Valerii. A. Aksionov. On continuation of 3-colouring of planar graphs. *Diskret Analiz*, pages 3–19, 1974.
- [Ald94] David Aldous. Triangulating the circle, at random. *The American Mathematical Monthly*, 101(3):223–233, 1994.
- [AM24] Jayadev S. Athreya and Howard Masur. *Translation Surfaces (to Appear)*. Providence, RI: American Mathematical Society (AMS), 2024.
- [Ant11] Yago Antolin. On Cayley graphs of virtually free groups. *Groups – Complexity – Cryptology*, 3(2), January 2011.
- [Asg12] Nabiha Asghar. Grötzsch’s Theorem. Master’s thesis, University of Waterloo, 2012.
- [Bab77] László Babai. Some applications of graph contractions. *Journal of Graph Theory*, 1(2):125–130, 1977.

-
- [BB18] Marthe Bonamy and Nicolas Bousquet. Recoloring graphs via tree decompositions. *European Journal of Combinatorics*, 69:200–213, March 2018.
- [BBFJ19] Marthe Bonamy, Nicolas Bousquet, Carl Feghali, and Matthew Johnson. On a conjecture of Mohar concerning Kempe equivalence of regular graphs. *Journal of Combinatorial Theory, Series B*, 135:179–199, March 2019.
- [BC09] Paul Bonsma and Luis Cereceda. Finding Paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoretical Computer Science*, 410(50):5215–5226, November 2009.
- [BD97] Russ Bubley and Martin Dyer. Path coupling: A technique for proving rapid mixing in Markov chains. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 223–231. IEEE, 1997.
- [BDGJ99] Russ Bubley, Martin Dyer, Catherine Greenhill, and Mark Jerrum. On Approximately Counting Colorings of Small Degree Graphs. *SIAM Journal on Computing*, 29(2):387–400, January 1999.
- [BDK⁺23] Marthe Bonamy, Oscar Defrain, Tereza Klimošová, Aurélie Lagoutte, and Jonathan Narboni. On Vizing’s edge colouring question. *Journal of Combinatorial Theory, Series B*, 159:126–139, 2023.
- [Ber66] Robert L. Berger. The undecidability of the domino problem. *Memoirs of the American Mathematical Society*, 1966.
- [BH22] Nicolas Bousquet and Marc Heinrich. A polynomial version of Cereceda’s conjecture. *Journal of Combinatorial Theory, Series B*, 155:1–16, July 2022.
- [BHI⁺20] Marthe Bonamy, Marc Heinrich, Takehiro Ito, Yusuke Kobayashi, Haruka Mizuta, Moritz Mühlenthaler, Akira Suzuki, and Kunihiro Wasa. Diameter of colorings under Kempe changes. *Theoretical Computer Science*, 838:45–57, October 2020.
- [BKT10] Thomas Böhme, Alexandr Kostochka, and Andrew Thomason. Hadwiger numbers and over-dominating colourings. *Discrete Mathematics*, 310(20):2662–2665, October 2010.

- [BLd⁺22] Maike Buchin, Anna Lubiw, Arnaud de Mesmay, Saul Schleimer, and Florestan Brunck. Computation and Reconfiguration in Low-Dimensional Topological Spaces (Dagstuhl Seminar 22062). Technical report, 2022.
- [BN46] Svend Bundgaard and Jakob Nielsen. Forenklede beviser for nogle sætninger i fladetopologien. *Matematisk Tidsskrift. B*, pages 1–16, 1946.
- [Bod96] Hans L. Bodlaender. A linear time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on computing*, 25(6):1305–1317, 1996.
- [BP16] Nicolas Bousquet and Guillem Perarnau. Fast recoloring of sparse graphs. *European Journal of Combinatorics*, 52:1–11, February 2016.
- [BPP22] Marthe Bonamy, Thomas Perrett, and Luke Postle. Colouring graphs with sparse neighbourhoods: Bounds and applications. *Journal of Combinatorial Theory, Series B*, 155:278–317, July 2022.
- [Bro87] R. Leonard Brooks. On Colouring the Nodes of a Network. In *Classic Papers in Combinatorics*, Modern Birkhäuser Classics, pages 118–121. Birkhäuser, Boston, MA, 1987.
- [BS18] Alexis Ballier and Maya Stein. The domino problem on groups of polynomial growth. *Groups, Geometry, and Dynamics*, 12(1):93–105, 2018.
- [Bud17] Thomas Budzinski. On the mixing time of the flip walk on triangulations of the sphere. *Comptes Rendus. Mathématique*, 355(4):464–471, 2017.
- [CESB⁺23] Barry Cimring, Remy El Sabeh, Marc Bacvanski, Stephanie Maaz, Izzat El Hajj, Naomi Nishimura, Amer E. Mouawad, and Alexandre Cooper. Efficient algorithms to solve atom reconfiguration problems. I. Redistribution-reconfiguration algorithm. *Physical Review A*, 108(2):023107, August 2023.
- [CHM22] Johannes Carmesin, Matthias Hamann, and Babak Miraftab. Canonical trees of tree-decompositions. *Journal of Combinatorial Theory, Series B*, 152:1–26, 2022.
- [CK23] Johannes Carmesin and Jan Kurkofka. Canonical decompositions of 3-connected graphs. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1887–1920. IEEE, 2023.

-
- [CLRS22] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT press, 2022.
- [CR15] Daniel W. Cranston and Landon Rabern. Brooks' Theorem and Beyond. *Journal of Graph Theory*, 80(3):199–225, 2015.
- [CRST06] Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The Strong Perfect Graph Theorem. *Annals of Mathematics*, 164(1):51–229, 2006.
- [CS12] Maria Chudnovsky and Paul Seymour. Packing seagulls. *Combinatorica*, 32(3):251–282, April 2012.
- [CS20] Alessandra Caraceni and Alexandre Stauffer. Polynomial mixing time of edge flips on quadrangulations. *Probability Theory and Related Fields*, 176(1):35–76, 2020.
- [Cul97] Joseph Culberson. Sokoban is PSPACE-complete. 1997.
- [Dan84] Jerome Dancis. Triangulated n -manifolds are determined by their $[n/2]+1$ -skeletons. *Topology and its Applications*, 18(1):17–26, 1984.
- [DD89] Warren Dicks and Martin J. Dunwoody. *Groups Acting on Graphs*, volume 17. Cambridge University Press, 1989.
- [DF13] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, London, 2013.
- [DGZ⁺20] Vincent Delecroix, Élise Goujard, Peter Zograf, Anton Zorich, and Engel. Contribution of one-cylinder square-tiled surfaces to Masur–Veech volumes. In *Some Aspects of the Theory of Dynamical Systems: A Tribute to Jean-Christophe Yoccoz. Volume I*, pages 223–274. Paris: Société Mathématique de France (SMF), 2020.
- [DGZZ20] Vincent Delecroix, Elise Goujard, Peter Zograf, and Anton Zorich. Enumeration of meanders and Masur–Veech volumes. In *Forum of Mathematics, Pi*, volume 8, page e4. Cambridge University Press, 2020.
- [DGZZ21] Vincent Delecroix, Élise Goujard, Peter Zograf, and Anton Zorich. Masur–Veech volumes, frequencies of simple closed geodesics, and intersection numbers of moduli spaces of curves. *Duke Mathematical Journal*, 170(12):2633–2718, 2021.

- [DHL18] Reinhard Diestel, Fabian Hundertmark, and Sahar Lemanczyk. Profiles of separations: In graphs, matroids, and beyond. *Combinatorica. An International Journal on Combinatorics and the Theory of Computing*, 39, November 2018.
- [Die17] Reinhard Diestel. *Graph Theory: 5th Edition*. Springer Graduate Texts in Mathematics. Springer-Verlag, 2017.
- [Dir52] Gabriel Andrew Dirac. Some theorems on abstract graphs. *Proceedings of the London Mathematical Society*, 3(1):69–81, 1952.
- [DJM93] Reinhard Diestel, Heinz A Jung, and Rögnvaldur G Möller. On vertex transitive graphs of infinite degree. *Archiv der Mathematik*, 60:591–600, 1993.
- [DKT21] Zdeněk Dvořák, Daniel Král', and Robin Thomas. Three-coloring triangle-free graphs on surfaces V. Coloring planar graphs with distant anomalies. *Journal of Combinatorial Theory, Series B*, 150:244–269, September 2021.
- [DL15] Zdeněk Dvořák and Bernard Lidický. 3-coloring triangle-free planar graphs with a precolored 8-cycle. *Journal of Graph Theory*, 80(2):98–111, October 2015.
- [DP17] Michelle Delcourt and Luke Postle. On the List Coloring Version of Reed's Conjecture. *Electronic Notes in Discrete Mathematics*, 61:343–349, August 2017.
- [DP19] Valentina Disarlo and Hugo Parlier. The geometry of flip graphs and mapping class groups. *Transactions of the American Mathematical Society*, 372(6):3809–3844, 2019.
- [DP21] Michelle Delcourt and Luke Postle. Reducing linear Hadwiger's conjecture to coloring small graphs. *arXiv preprint arXiv:2108.01633*, 2021.
- [DPP18] Michelle Delcourt, Guillem Perarnau, and Luke Postle. Rapid mixing of Glauber dynamics for colorings below Vigoda's 11/6 threshold. *arXiv preprint arXiv:1804.04025*, 2018.
- [Dro06] Carl Droms. Infinite-ended groups with planar Cayley graphs. *Journal of Group Theory*, 9(4):487–496, 2006.

-
- [dRST21] Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. The unbearable hardness of unknotting. *Advances in Mathematics*, 381:107648, April 2021.
- [DS93] Persi Diaconis and Laurent Saloff-Coste. Comparison Theorems for Reversible Markov Chains. *The Annals of Applied Probability*, 3(3):696–730, August 1993.
- [DSS98b] Carl Droms, Brigitte Servatius, and Herman Servatius. The structure of locally finite two-connected graphs. *Electronic Journal of Combinatorics*, 2, January 1998.
- [DT99] Reinhard Diestel and Robin Thomas. Excluding a countable clique. *Journal of Combinatorial Theory, Series B*, 76(1):41–67, 1999.
- [Dun85] Martin J. Dunwoody. The accessibility of finitely presented groups. *Inventiones Mathematicae*, 81:449–458, 1985.
- [Dun93] Martin J. Dunwoody. An inaccessible group. *Geometric group theory*, 1:75–78, 1993.
- [Dun07] Martin J. Dunwoody. Planar graphs and covers. *arXiv preprint arXiv:0708.0920*, 2007.
- [EF20] Eduard Eiben and Carl Feghali. Toward Cereceda’s conjecture for planar graphs. *Journal of Graph Theory*, 94(2):267–277, 2020.
- [EF22] David Eppstein and Daniel Frishberg. Improved mixing for the convex polygon triangulation flip walk. *arXiv preprint arXiv:2207.09972*, 2022.
- [Erd59] Paul Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, 11:34–38, 1959.
- [ESBD⁺23] Remy El Sabeh, Jessica Bohm, Zhiqian Ding, Stephanie Maaz, Naomi Nishimura, Izzat El Hajj, Amer E. Mouawad, and Alexandre Cooper. Efficient algorithms to solve atom reconfiguration problems. II. Assignment-rerouting-ordering algorithm. *Physical Review A*, 108(2):023108, August 2023.
- [Feg19] Carl Feghali. Paths between colourings of sparse graphs. *European Journal of Combinatorics*, 75:169–171, January 2019.
- [Feg23] Carl Feghali. Kempe equivalence of 4-critical planar graphs. *Journal of Graph Theory*, 103(1):139–147, 2023.

- [FF98] Rudolf Fritsch and Gerda Fritsch. *Four-Color Theorem*. Springer, 1998.
- [FG06] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [Fis73] Steve Fisk. Combinatorial structure on triangulations. I. The structure of four colorings. *Advances in Mathematics*, 11(3):326–338, December 1973.
- [Fis77a] Steve Fisk. Geometric coloring theory. *Advances in Mathematics*, 24(3):298–340, June 1977.
- [FJP17] Carl Feghali, Matthew Johnson, and Daniël Paulusma. Kempe equivalence of colourings of cubic graphs. *European Journal of Combinatorics*, 59:1–10, January 2017.
- [FM11] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups (Pms-49)*, volume 41. Princeton university press, 2011.
- [Fox52] Ralph H. Fox. On Fenchel’s Conjecture about F-Groups. *Matematisk Tidsskrift. B*, pages 61–65, 1952.
- [Fre44] Hans Freudenthal. Über die enden diskreter räume und gruppen. *Commentarii Mathematici Helvetici*, 17(1):1–38, 1944.
- [FV07] Alan Frieze and Eric Vigoda. A survey on the use of Markov chains to randomly sample colourings. *Oxford Lecture Series in Mathematics and its Applications*, 34:53, 2007.
- [GK14] Rostislav Grigorchuk and Rostyslav Kravchenko. On the lattice of subgroups of the lamplighter group. *International Journal of Algebra and Computation*, 24(6):837–877, 2014.
- [GPT] Mario Gutiérrez, Hugo Parlier, and Paul Turner. Quadratis. <https://www.unige.ch/~turnerpa/Quadratis.html>.
- [GPT22] Mario Gutiérrez, Hugo Parlier, and Paul Turner. Quadratis puzzles. *arXiv preprint arXiv:2208.01367*, 2022.
- [Grö59] Herbert Grötzsch. Zur theorie der diskreten gebilde. vii. ein dreifarbensatz für dreikreisfreie netze auf der kugel. *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg, Math-Nat.*, 8:109–120, 1959.

-
- [Gro16a] Martin Grohe. Quasi-4-connected components. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPICs*, pages 8:1–8:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [Gro16b] Martin Grohe. Tangles and connectivity in graphs. In *Language and Automata Theory and Applications*, pages 24–41. Springer, 2016.
- [Gyá75] András Gyárfás. On ramsey covering-numbers. *Infinite and Finite Sets*, 2:801–816, 1975.
- [Hak61] Wolfgang Haken. Theorie der normalflächen. *Acta Mathematica*, 105(3):245–375, 1961.
- [Hal65] Rudolf Halin. Über die Maximalzahl fremder unendlicher Wege in Graphen. *Mathematische Nachrichten*, 30(1-2):63–85, 1965.
- [Hal73] Rudolf Halin. Automorphisms and endomorphisms of infinite locally finite graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 39(1):251–283, September 1973.
- [Ham18a] Matthias Hamann. Accessibility in transitive graphs. *Combinatorica. An International Journal on Combinatorics and the Theory of Computing*, 38(4):847–859, 2018.
- [Ham18b] Matthias Hamann. Planar transitive graphs. *Electronic Journal of Combinatorics*, 25(4):research paper p4.8, 18, 2018.
- [HdK21] Eoin Hurley, Rémi de Joannis de Verclos, and Ross J. Kang. An improved procedure for colouring graphs of bounded local density. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 135–148. Society for Industrial and Applied Mathematics, January 2021.
- [HLMR22] Matthias Hamann, Florian Lehner, Babak Miraftab, and Tim Rühmann. A Stallings type theorem for quasi-transitive graphs. *Journal of Combinatorial Theory, Series B*, 157:40–69, 2022.
- [HLP99] Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *Journal of the ACM*, 46(2):185–211, March 1999.
- [Hop43] Heinz Hopf. Enden offener Räume und unendliche diskontinuierliche Gruppen. *Commentarii Mathematici Helvetici*, 16(1):81–100, 1943.

- [Jer95] Mark Jerrum. A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Structures & Algorithms*, 7(2):157–165, September 1995.
- [Jør94] Leif K. Jørgensen. Contractions to k_8 . *Journal of Graph Theory*, 18(5):431–448, 1994.
- [JR21] Emmanuel Jeandel and Michael Rao. An aperiodic set of 11 Wang tiles. *Advances in Combinatorics*, January 2021.
- [JVV86] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, January 1986.
- [Kal94] Gil Kalai. Some Aspects of the Combinatorial Theory of Convex Polytopes. In *Polytopes: Abstract, Convex and Computational*, pages 205–229. Springer Netherlands, Dordrecht, 1994.
- [Kar72] Richard M. Karp. Reducibility among Combinatorial Problems. In *Complexity of Computer Computations: Proceedings of a Symposium on the Complexity of Computer Computations, Held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and Sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*, The IBM Research Symposia Series, pages 85–103. Springer US, Boston, MA, 1972.
- [Kem79] Alfred B. Kempe. On the Geographical Problem of the Four Colours. *American Journal of Mathematics*, 2(3):193–200, 1879.
- [KL12] Louis H. Kauffman and Sofia Lambropoulou. Hard unknots and collapsing tangles. *Introductory lectures on knot theory, Ser. Knots Everything*, 46:187–247, 2012.
- [KMP23] Tom Kelly, Alp Müyesser, and Alexey Pokrovskiy. Optimal spread for spanning subgraphs of Dirac hypergraphs. *arXiv preprint arXiv:2308.08535*, 2023.
- [Kos82] Alexandr V. Kostochka. The minimum Hadwiger number for graphs with a given mean degree of vertices. *Metody Diskret. Analiz.*, (38):37–58, 1982.
- [Kos84] Alexandr V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica. An International*

-
- Journal on Combinatorics and the Theory of Computing*, 4(4):307–316, 1984.
- [KP20] Tom Kelly and Luke Postle. A local epsilon version of Reed’s conjecture. *Journal of Combinatorial Theory, Series B*, 141:181–222, 2020.
- [KPS73] Abraham Karrass, Alfred Pietrowski, and Donald Solitar. Finite and infinite cyclic extensions of free groups. *Journal of the Australian Mathematical Society*, 16(4):458–466, 1973.
- [KR16] Andrew D. King and Bruce A. Reed. A Short Proof That χ Can be Bounded ε Away from $\Delta + 1$ toward ω . *Journal of Graph Theory*, 81(1):30–34, 2016.
- [Kri21] Matthias Kriesell. A note on uniquely 10-colorable graphs. *Journal of Graph Theory*, 98(1):24–26, 2021.
- [KT21] Dale Koenig and Anastasiia Tsvietkova. NP-hard problems naturally arising in knot theory. *Transactions of the American Mathematical Society, Series B*, 8(15):420–441, May 2021.
- [Kup14] Greg Kuperberg. Knottedness is in NP, modulo GRH. *Advances in Mathematics*, 256:493–506, May 2014.
- [KZ03] Maxim Kontsevich and Anton Zorich. Connected components of the moduli spaces of Abelian differentials with prescribed singularities. *Inventiones Mathematicae*, 153(3):631–678, September 2003.
- [Lab21a] Sébastien Labbé. Markov partitions for toral \mathbb{Z}^2 -rotations featuring Jeandel–Rao Wang shift and model sets. *Annales Henri Lebesgue*, 4:283–324, 2021.
- [Lab21b] Sébastien Labbé. Substitutive structure of Jeandel–Rao aperiodic tilings. *Discrete & Computational Geometry*, 65(3):800–855, 2021.
- [Lac15] Marc Lackenby. A polynomial upper bound on Reidemeister moves. *Annals of Mathematics*, 182(2):491–564, 2015.
- [Lac17] Marc Lackenby. Elementary knot theory. In *Lectures on Geometry*, Clay Lect. Notes, pages 29–64. Oxford Univ. Press, Oxford, 2017.
- [Lac21a] Marc Lackenby. The efficient certification of knottedness and Thurston norm. *Advances in Mathematics*, 387:107796, August 2021.

- [Lac21b] Marc Lackenby. Marc Lackenby announces a new unknot recognition algorithm that runs in quasi-polynomial time | Mathematical Institute. <https://www.maths.ox.ac.uk/node/38304>, 2021.
- [Lan08] Erwan Lanneau. Connected components of the strata of the moduli spaces of quadratic differentials. In *Annales Scientifiques de l'École Normale Supérieure*, volume 41, pages 1–56, 2008.
- [LP17] David A. Levin and Yuval Peres. *Markov Chains and Mixing Times*, volume 107. American Mathematical Soc., 2017.
- [ŁV05] Tomasz Łuczak and Eric Vigoda. Torpid mixing of the Wang–Swendsen–Kotecký algorithm for sampling colorings. *Journal of Discrete Algorithms*, 3(1):92–100, March 2005.
- [Mey75] Henry Meyniel. Les 5-colorations d'un graphe planaire forment une classe de commutation unique. *Journal of Combinatorial Theory, Series B*, 24:251–257, 1975.
- [Moh85] Bojan Mohar. Akempic triangulations with 4 odd vertices. *Discrete Mathematics*, 54(1):23–29, 1985.
- [Moh07] Bojan Mohar. Kempe Equivalence of Colorings. In *Graph Theory in Paris: Proceedings of a Conference in Memory of Claude Berge*, Trends in Mathematics, pages 287–297. Birkhäuser, Basel, 2007.
- [MR99] Meena Mahajan and Venkatesh Raman. Parameterizing above Guaranteed Values: MaxSat and MaxCut. *Journal of Algorithms*, 31(2):335–354, May 1999.
- [MRS97] Michael Molloy, Bruce A. Reed, and William Steiger. On the mixing rate of the triangulation walk. 1997.
- [MS83] David E. Muller and Paul E. Schupp. Groups, the Theory of ends, and context-free languages. *Journal of Computer and System Sciences*, 26(3):295–310, 1983.
- [MS09] Bojan Mohar and Jesús Salas. A new Kempe invariant and the (non)-ergodicity of the Wang–Swendsen–Kotecký algorithm. *Journal of Physics A: Mathematical and Theoretical*, 42(22), 2009.
- [MS10] Bojan Mohar and Jesús Salas. On the non-ergodicity of the Swendsen–Wang–Kotecký algorithm on the kagomé lattice. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(05):P05016, May 2010.

-
- [Nar23] Jonathan Narboni. Vizing’s edge-recoloring conjecture holds. *arXiv preprint arXiv:2302.12914*, 2023.
- [NZ23] Isabella Novik and Hailun Zheng. Reconstructing simplicial polytopes from their graphs and affine 2-stresses. *Israel Journal of Mathematics*, 255(2):891–910, June 2023.
- [Pad13] Arnau Padrol. Many neighborly polytopes and oriented matroids. *Discrete & Computational Geometry*, 50:865–902, 2013.
- [Pin73] Mark S. Pinsky. On the complexity of a concentrator. In *7th International Teletraffic Conference*, volume 4, pages 1–318. Citeseer, 1973.
- [PP24] Jinyoung Park and Huy T. Pham. A proof of the Kahn–Kalai conjecture. *Journal of the American Mathematical Society*, 37(1):235–243, 2024.
- [PSSS22] Huy T. Pham, Ashwin Sah, Mehtaab Sawhney, and Michael Simkin. A toolkit for robust thresholds. *arXiv preprint arXiv:2210.03064*, 2022.
- [PST03] Michael Plummer, Michael Stiebitz, and Bjarne Toft. On a special case of hadwiger’s conjecture. *Discussiones Mathematicae Graph Theory*, 23(2):333–363, 2003.
- [Ree93] Bruce A. Reed. Counterexamples to a conjecture of Las Vergnas and Meyniel. *Contemporary Mathematics*, 147:157–157, 1993.
- [Ree98] Bruce A. Reed. ω , Δ , and χ . *Journal of Graph Theory*, 27(4):177–212, 1998.
- [Rei27] Kurt Reidemeister. Elementare begründung der knotentheorie. In *Abhandlungen Aus Dem Mathematischen Seminar Der Universität Hamburg*, volume 5, pages 24–32. Springer, 1927.
- [RKDD14] Tomas Rokicki, Herbert Kociemba, Morley Davidson, and John De-
thridge. The Diameter of the Rubik’s Cube Group Is Twenty. *SIAM Review*, 56(4):645–670, January 2014.
- [RP77] G. Ravindra and K.R. Parthasarathy. Perfect product graphs. *Discrete Mathematics*, 20:177–186, 1977.
- [RS91] Neil Robertson and Paul Seymour. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, 1991.

- [RS93] Neil Robertson and Paul Seymour. Excluding a graph with one crossing. In *Graph Structure Theory*, pages 669–675. American Mathematical Society, 1993.
- [RS03] Neil Robertson and Paul Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.
- [RS04] Neil Robertson and Paul D Seymour. Graph minors. XX. Wagner’s conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.
- [RSST97] Neil Robertson, Daniel Sanders, Paul Seymour, and Robin Thomas. The four-colour theorem. *Journal of Combinatorial Theory, Series B*, 70(1):2–44, 1997.
- [RTL76] Donald J. Rose, R. Endre Tarjan, and George S. Lueker. Algorithmic Aspects of Vertex Elimination on Graphs. *SIAM Journal on Computing*, 5(2):266–283, June 1976.
- [Sav70] Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of computer and system sciences*, 4(2):177–192, 1970.
- [Ser80] Jean-Pierre Serre. *Trees*. Springer-Verlag, Berlin, 1980.
- [Sey16] Paul Seymour. Hadwiger’s conjecture. *Open problems in mathematics*, pages 417–437, 2016.
- [She82] Ido Shemer. Neighborly polytopes. *Israel Journal of Mathematics*, 43:291–314, 1982.
- [Son05] Zi-Xia Song. The extremal function for K_8 - minors. *Journal of Combinatorial Theory, Series B*, 95(2):300–317, November 2005.
- [SS97] Jesús Salas and Alan D. Sokal. Absence of phase transition for antiferromagnetic Potts models via the Dobrushin uniqueness theorem. *Journal of Statistical Physics*, 86(3):551–579, February 1997.
- [SS22] Jesús Salas and Alan D. Sokal. Ergodicity of the Wang–Swendsen–Kotecký algorithm on several classes of lattices on the torus. *Journal of Physics A: Mathematical and Theoretical*, 55(41):415004, 2022.

-
- [SSSZ20] Ashwin Sah, Mehtaab Sawhney, David Stoner, and Yufei Zhao. A reverse Sidorenko inequality. *Inventiones mathematicae*, 221(2):665–711, August 2020.
- [Sta71] John R. Stallings. *Group Theory and Three-Dimensional Manifolds*. Yale University Press, 1971.
- [STT88] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Journal of the American Mathematical Society*, 1(3):647–681, July 1988.
- [STT92] Daniel D. Sleator, Robert E. Trajan, and William P. Thurston. Short encodings of evolving structures. *SIAM Journal on Discrete Mathematics*, 5(3):428–450, 1992.
- [Sum81] David P. Sumner. Subtrees of a graph and chromatic number. *The theory and applications of graphs*, pages 557–576, 1981.
- [Tho84] Andrew Thomason. An extremal function for contractions of graphs. *Mathematical Proceedings of the Cambridge Philosophical Society*, 95(2):261–265, March 1984.
- [Tho92] Carsten Thomassen. The Hadwiger number of infinite vertex-transitive graphs. *Combinatorica. An International Journal on Combinatorics and the Theory of Computing*, 12:481–491, 1992.
- [Tho03] Carsten Thomassen. A short list color proof of Grötzsch’s theorem. *Journal of Combinatorial Theory, Series B*, 88(1):189–192, May 2003.
- [Tut69] William T. Tutte. Even and odd 4-colorings. In *Proof Techniques in Graph Theory*, pages 161–169. Academic Press New York, 1969.
- [Tut84] William T. Tutte. *Graph Theory*. Encyclopedia of Mathematics and Its Applications. Cambridge University Press, 1984.
- [TW93] Carsten Thomassen and Wolfgang Woess. Vertex-transitive graphs and accessibility. *Journal of Combinatorial Theory, Series B*, 58(2):248–268, 1993.
- [vdH13] Jan van den Heuvel. The Complexity of Change. *arXiv preprint arXiv:1312.2816*, December 2013.
- [Vig00] Eric Vigoda. Improved bounds for sampling colorings. *Journal of Mathematical Physics*, 41(3):1555–1569, 2000.

- [Viz64] Vadim G. Vizing. On an estimate of the chromatic class of a p -graph. *Discret Analiz*, 3:25–30, 1964.
- [Viz65] Vadim G. Vizing. The chromatic class of a multigraph. *Cybernetics*, 1(3):32–41, May 1965.
- [VM81] Michel Las Vergnas and Henri Meyniel. Kempe classes and the Hadwiger Conjecture. *Journal of Combinatorial Theory, Series B*, 31(1):95–104, August 1981.
- [Wag37] Klaus Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114(1):570–590, 1937.
- [Wan61] Hao Wang. Proving theorems by pattern recognition—II. *Bell system technical journal*, 40(1):1–41, 1961.
- [Wei16] Vera Weil. A note on Reed’s Conjecture about ω , Δ , and χ with respect to vertices of high degree. *arXiv preprint arXiv:1611.02063*, 2016.
- [Woe89] Wolfgang Woess. Graphs and groups with tree-like properties. *Journal of Combinatorial Theory, Series B*, 47(3):361–371, 1989.
- [Wol89] Ulli Wolff. Collective Monte Carlo Updating for Spin Systems. *Physical Review Letters*, 62(4):361–364, January 1989.
- [Zor02] Anton Zorich. Square tiled surfaces and Teichmüller volumes of the moduli spaces of abelian differentials. *Rigidity in Dynamics and Geometry: Contributions from the Programme Ergodic Theory, Geometric Rigidity and Number Theory, Isaac Newton Institute for the Mathematical Sciences Cambridge, United Kingdom, 5 January–7 July 2000*, pages 459–471, 2002.
- [Zor06] Anton Zorich. Flat surfaces. *arXiv preprint arXiv:math/0609392*, 2006.
- [Zor08] Anton Zorich. Explicit Jenkins-Strebel representatives of all strata of Abelian and quadratic differentials. *Journal of Modern Dynamics*, 2(1):139–185, 2008.
- [ZVC80] Heiner Zieschang, Elmar Vogt, and Hans-Dieter Coldewey. *Surfaces and Planar Discontinuous Groups. Revised and Expanded Transl. from the German by J. Stillwell*, volume 835 of *Lect. Notes Math.* Springer, Cham, 1980.

Publications

Articles in Peer-reviewed Conferences

- [1] Marthe Bonamy, Vincent Delecroix, and Clément Legrand-Duchesne. Kempe Changes in Bounded Treewidth Graphs. In *Extended Abstracts EuroComb 2021*, Trends in Mathematics, pages 714–719, Cham, 2021. Springer International Publishing.
- [2] Marthe Bonamy, Marc Heinrich, Clément Legrand-Duchesne, and Jonathan Narboni. On a recolouring version of Hadwiger’s conjecture. *European Conference on Combinatorics, Graph Theory and Applications*, pages 142–148, August 2023.
- [3] Guilherme C. M. Gomes, Clément Legrand-Duchesne, Reem Mahmoud, Amer E. Mouawad, Yoshio Okamoto, Vinicius F. dos Santos, and Tom C. van der Zanden. Minimum Separator Reconfiguration. In *18th International Symposium on Parameterized and Exact Computation (IPEC 2023)*, volume 285 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:12. Schloss-Dagstuhl - Leibniz Zentrum für Informatik, 2023.
- [4] Louis Esperet, Ugo Giocanti, and Clément Legrand-Duchesne. The structure of quasi-transitive graphs avoiding a minor with applications to the domino problem. *European Conference on Combinatorics, Graph Theory and Applications*, pages 425–432, August 2023.
- [5] Clément Legrand-Duchesne, Ashutosh Rai, and Martin Tancer. Parameterized Complexity of Untangling Knots. In *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 88:1–88:17. Schloss-Dagstuhl - Leibniz Zentrum für Informatik, 2022.

Articles in Peer-reviewed Journals

- [6] Marthe Bonamy, Vincent Delecroix, and Clément Legrand–Duchesne. Kempe changes in degenerate graphs. *European Journal of Combinatorics*, page 103802, November 2023.
- [7] Marthe Bonamy, Marc Heinrich, Clément Legrand-Duchesne, and Jonathan Narboni. On a recolouring version of Hadwiger’s conjecture. *Journal of Combinatorial Theory, Series B*, 164:364–370, January 2024.
- [8] Quentin Deschamps, Carl Feghali, František Kardoš, Clément Legrand-Duchesne, and Théo Pierron. Strengthening a Theorem of Meyniel. *SIAM Journal on Discrete Mathematics*, 37(2):604–611, June 2023.
- [9] Clément Legrand-Duchesne, Ashutosh Rai, and Martin Tancer. Parameterized Complexity of Untangling Knots. *SIAM Journal on Computing*, 53(2):431–479, April 2024.

Preprints

- [10] Louis Esperet, Ugo Giocanti, and Clément Legrand-Duchesne. The structure of quasi-transitive graphs avoiding a minor with applications to the domino problem, April 2024.
- [11] Guilherme C. M. Gomes, Clément Legrand-Duchesne, Reem Mahmoud, Amer E. Mouawad, Yoshio Okamoto, Vinicius F. dos Santos, and Tom C. van der Zanden. Minimum Separator Reconfiguration, July 2023.