Communauté
UNIVERSITÉ Grenoble Alpes

THÈSE

Pour obtenir le grade de

DOCTEUR DE LA COMMUNAUTE UNIVERSITE GRENOBLE ALPES

Spécialité : **Informatique**

Arrêté ministériel : 25 mai 2016

Présentée par

Vincent DESPRÉ

Thèse dirigée par **Francis LAZARUS** et
codirigée par **András SEBŐ**

préparée au sein du Gipsa-lab et G-Scop Laboratoire des sciences pour la conception, l'optimisation et la production
dans **l'École Doctorale MSTII**

Topologie et algorithmes sur les cartes combinatoires

Thèse soutenue publiquement le **18 Octobre 2016**,
devant le jury composé de :

Olivier DEVILLERS
Directeur de recherche, INRIA Délégation Nancy Grand Est, Président
Xavier GOAOC
Professeur, Université Paris-Est Marne La Vallée, Rapporteur
Stefan FELSNER
Professeur, Université technique de Berlin, Allemagne, Rapporteur
Stéphan THOMASSÉ
Professeur, Ecole Normale Supérieure de Lyon, Membre
Nicolas BONICHON
Maître de conférences, Université de Bordeaux, Membre
Sergio CABELLO
Professeur, Université de Ljubljana, Slovénie, Membre
Francis LAZARUS
CNRS Délégation Alpes, Directeur de thèse
András SEBŐ CNRS Délégation Alpes, Co-directeur de thèse

Imre BÁRÁNY Invité

Abstract

In this thesis, we focus on the topological properties of surfaces, i.e. those that are preserved by continuous deformations. Intuitively, it can be understood as the properties that describe the general shape of surfaces. We describe surfaces as combinatorial maps. They have the double advantage of being well defined mathematical objects and of being straightforwardly transformed into data-structures.

We study three distinct problems. Firstly, we give algorithms to compute geometric intersection numbers of curves on surfaces. We obtain a quadratic algorithm to compute the minimal number of self-intersections in a homotopy class, a quartic one to construct a minimal representative and a quasi-linear one to decide if a homotopy class contains a simple curve. Secondly, we give counter-examples to a conjecture of Mohar and Thomassen about the existence of splitting cycles in triangulations. Finally, we use the recent work of Gonçalves and Lévèque about toroidal Schnyder woods to describe a bijection between toroidal triangulations and toroidal unicellular maps analogous to the well known bijection of Poulalhon and Schaeffer for planar triangulations.

Many different points of view are involved in this thesis. We thus propose a large preliminary chapter where we provide connections between the different viewpoints.

Résumé

Dans cette thèse, nous nous intéressons aux propriétés topologiques des surfaces, i.e. celles qui sont préservées par des déformations continues. Intuitivement, ces propriétés peuvent être imaginées comme étant celles qui décrivent la forme générale des surfaces. Nous utilisons des cartes combinatoires pour décrire les surfaces. Elles ont le double avantage d'être des objets mathématiques naturels et de pouvoir être transformées en structure de données.

Nous étudions trois problèmes différents. Premièrement, nous donnons des algorithmes de calcul du nombre géométrique d'intersection de courbes dessinées sur des surfaces. Nous avons obtenu: un algorithme quadratique de calcul du nombre minimal d'auto-intersections dans une classe d'homotopie, un algorithme quartique pour construire un représentant minimal et un algorithme quasi-linéaire pour déterminer si une classe d'homotopie contient bien une courbe simple. Nous donnons ensuite des contre-exemples à la conjecture de Mohar et Thomassen traitant de l'existence de cycles de partage dans les triangulations. Finalement, nous utilisons les travaux récents de Lévèque et Gonçalves à propos des bois toriques de Schnyder pour construire une bijection entre les triangulations du tore et certaines cartes unicellulaires analogues à la célèbre bijection de Poulalhon et Schaeffer pour les triangulations planaires.

Plusieurs points de vue sont utilisés au cours de cette thèse. Nous proposons donc un important chapitre préliminaire où nous insistons sur les connections entre ces différents points de vue.

Contents

Acknowledgments

A PhD appeared to be an adventure significantly deeper and intense that was I imagined. Fortunately, I was surrounded by a lot of people in both my professional and personal life. The first person I have to thank is clearly Francis. He was the perfect guide, giving me the advice I was needed and enough freedom to understand what is the real life of a researcher. In addition, by now, I consider him as a friend. I also thank András for having made this whole thing possible and for giving me hindsight on my work. It is difficult to imagine a better environment for a PhD.

I also have to thank all the 7 other people of the defense committee. All of them accepted to spend time on my work and I received a lot of very useful remarks. I particularly thank the two reviewers for whom it had required a lot of time during the summer. I also want to thank all the people that come to the defense including a lot of people from the gipsa-lab and the g-scop.

I am also thankful to the people I collaborate with. In particular, Daniel Gonçalves and Benjamin Lévêque who welcome me very nicely in Montpelier and Luca Castelli Aleardi. I had a lot of interesting discussions about my work some of them leading to drastic improvement and so I have to thank: Éric Colin de Verdière, Mark Ellingham, Thom Sulanke, Bojan Mohar, Patrice Ossona de Mendez, Victor Chepoi, Frédéric Meunier and Gilles Schaeffer. I am grateful to the people that invite me to seminars like Vincent Limouzy, Petru Valicov and Kolja Knauer.

J'ai vécu les derniers moments de ma thèse à Bordeaux et je voudrais aussi remercier tous les gens qui m'ont très chaleureusement accueilli. Je remercie en particulier Marthe, Claire et Patxi pour leur bonne humeur, Samuel, Orel et Pierre pour leur nombreux conseils et leur patience et Frédéric pour m'avoir, entre autres, fait répéter ma soutenance. J'adresse un remerciement particulier à Nicolas qui est l'une des personnes les plus agréables que j'ai rencontré et qui a grandement simplifié mon installation à Bordeaux à tous points de vue.

Quand j'ai décidé de reprendre les études, j'ai eu la chance de trouver recevoir un accueil très chaleureux au département informatique de l'ENS Lyon. Je voudrais remercier Guillaume Hanrot pour m'avoir fait confiance et pour avoir tout mis en

oeuvre pour faciliter l'organisation de ma reprise d'études. Toute l'équipe enseignante à d'ailleurs fait preuve de beaucoup de compréhension vis-à-vis de ma situation assez particulière et j'ai pu bénéficier de soutient de la part de: Nicolas Brisebarre, Bruno Salvy, Jean-Michel Muller, Laurent Théry, Damien Pous, Bruno Gaujal and Pascal Koiran. Je dois aussi parler plus particulièrement de Stéphan Thomassé qui a, en plus, été un excellent guide qui m'a permis de choisir une spécialité qui s'est avéré me convenir parfaitement.

Avant cela j'ai été initié aux mathématiques par une série d'incroyable professeurs, une liste exhaustive serait trop longue alors je vais tenter de me focaliser sur ceux qui m'ont le plus marqué. Je pense à Claude Danthony et Albert Fathi de l'ENS Lyon mais aussi à Franz Ridde qui fut sans l'ombre d'un doute la personne qui a le plus participer à mon éducation scientifique et qui reste mon modèle d'enseignant. Je remercie aussi les collègues qui m'ont initié aux spécificités de l'enseignement de l'informatique, Michel Burlet et Philippe Morat.

Ensuite, sur le plan personnel j'ai eu la chance d'être très entouré aussi bien par mes nouveaux camarades grenoblois Rémi, Lucas et Arnaud mon fidèle compagnon de chambre que par mes vieux amis Filou (je détiens encore plusieurs de tes caleçons!), Arnaud (un jour j'irai à new York avec toi...) et Bordo (on est de plus en plus nombreux pour les déménagements, ça devient trop facile non?). J'ai aussi bien profiter de la compagnie d'Onyx pendant la rédaction même si je n'ai pu garder aucun de ses textes (en partie parce qu'elle les avaient rédigés en s'asseyant sur le clavier!). J'ai, par ailleurs, été soutenu par ma famille qui a toujours été là pour moi, merci à Pascale, Fredo, Patrick, Chantal, Emeline ($\times 2$), Jean-Philippe, Céline et Mémé. Je remercie tout particulièrement mon frère d'avoir fait l'effort de me doubler dans les études pour que je puisse profiter de son expérience à mon tour et mes parents pour avoir toujours été la parfaite ancre à chaque fois que j'en ai eu besoin. Je termine par remercier Julie pour avoir rendu toute cette aventure bien plus agréable et confortable mais aussi pour avoir été la victime enthousiaste de toutes sortes de répétitions.

Introduction en Français

0.1 Sujet

La topologie est l'étude des propriétés des formes qui sont préservées par des déformations continues. Il n'y a pas de notion de longueur, d'aire ou de courbure qui appartiennent au domaine de la géométrie. Les propriétés géométriques sont préservées uniquement par des transformations rigides telles que les rotations ou les translations. Ainsi, dans un sens, la topologie manipule uniquement les propriétés les plus robustes. On attribue généralement les premiers travaux de nature topologique à Euler (1707-1783) qui a découvert la célèbre formule qui porte depuis son nom: $V - E + F = 2$. Ici V, E et F sont respectivement le nombre de sommets, d'arêtes et de faces d'un polyèdre donné. On peut se familiariser avec cette formule en vérifiant sa véracité dans le cas des solides platoniciens de la figure 1. Le fait que cette formule ne s'applique pas seulement aux polyèdres convexes mais à n'importe quelle subdivision de la sphère ne sera compris que bien plus tard quand la topologie sera devenue une discipline en tant que telle.

Figure 1: Les 5 polyhèdres platoniques: le tétrahèdre, le cube, l'octahèdre, le dodécahèdre et l'icosahèdre. http://whistleralley.com/polyhedra/platonic.htm

Il est possible de définir une notion de dimension qui caractérise la topologie d'un espace. Typiquement, un polyèdre a une dimension 3 tandis que son enveloppe a

une dimension égale à 2. Les espaces topologiques de dimension 2 correspondent à la notion intuitive de surface. Dans cette thèse, nous cherchons à développer des outils très généraux pour l'analyse des propriétés topologiques des surfaces. Pour illustrer la différence de points de vue entre le topologue et le géomètre dans le cas des surfaces, considérons la figure 2. Le topologue voit deux objets identiques qu'il appelle sphères car ces deux surfaces peuvent être déformées continûment en sphères. Au contraire un spécialiste de géométrie voit deux formes différentes. Évidemment, en ajoutant plus de structures, on devient capable de caractériser les formes plus précisément. Par exemple, un être humain ne verra pas seulement deux formes différentes mais un lapin et un dauphin, ajoutant une notion de sémantique. Clairement, la tâche de classification devient de plus en plus ardue quand on ajoute de la structure. Il est donc avisé de pratiquer une classification progressive en commençant par considérer les aspects purement topologiques puis la géométrie et enfin la sémantique. En pratique, la topologie apparaît comme la première étape dans de nombreux processus de classification.

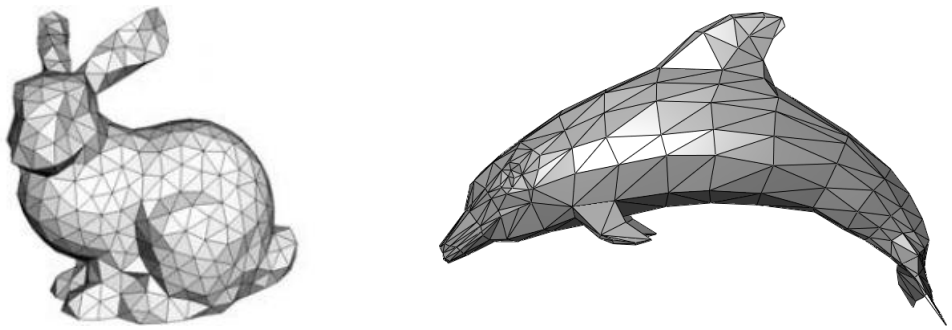

Figure 2: Deux sphères.

Plusieurs communautés s'intéressent aux problèmes topologiques sur les surfaces et les considèrent avec des points de vue différents. En topologie générale, on définit une surface comme une 2-variété. L'idée est de prendre le plan comme modèle de surface et ensuite de définir une surface comme un espace topologique qui ressemble localement au modèle en chacun de ses points. Ce point de vue prend sa source dans le célèbre Habilitationsvortrag de Riemann (1854). En parallèle une approche combinatoire s'est développée qui considère les espaces obtenus par assemblages de pièces élémentaires, généralement des simplexes (les points, les arêtes et les triangles dans le cas des surfaces). Le modèle des cartes combinatoires décrites à l'aide de systèmes de rotation est issu de ce type de construction. On les voit apparaître dans les travaux de Heffter (1891) puis dans les travaux d'Edmonds (1960). Elles sont

devenues populaires grâce, entre autres, à Youngs (1968) et Tutte (1973). Un point de vue intermédiaire a été adopté par les théoriciens des graphes. Ils visualisent les graphes comme des objets de dimension 1 qui peuvent être dessinés sur des surfaces. Cette idée a amené naturellement la notion de plongement cellulaire de graphes. Cette approche a été popularisée par Robertson et Seymour dans leur célèbre preuve du théorème des mineurs des graphes (1983-2004).

Ces différents points de vue sont illustrés par la figure 3. (a) montre la représentation classique d'un tore comme une forme plongée dans l'espace. (b) définit aussi un tore en identifiant les côtés opposés d'un carré (Cette représentation est préférée par la suite car elle permet d'avoir une meilleure vue d'ensemble). (c) représente un tore obtenu en assemblant des triangles, ce genre de cartes combinatoires est classique et est appelé triangulations. Finalement (d) montre un graphe à trois sommets, plongé sur un tore. Il est naturel de vouloir dessiner un graphe sur une surface et cela permet de formuler facilement de nombreux problèmes de théorie des graphes. Le contrepartie de cette dernière approche est la grande liberté qu'elle laisse aux courbes qui représentent les arêtes du graphe, les rendant potentiellement très complexes. Dans cette thèse, les cartes combinatoires seront systématiquement utilisées pour décrire les algorithmes et leurs preuves mais tous les points de vue seront utilisés dans le but d'obtenir la formulation la plus naturelle possible de chaque problème abordé.

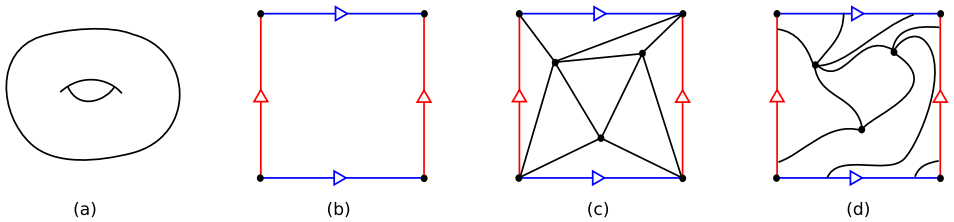

Figure 3: Des tores vu par différentes communautés.

Nous avons parlé des différences entre topologie et géométrie. Cependant, la géométrie peut parfois servir d'outil pour résoudre des problèmes purement topologiques. En effet, quand on s'intéresse à un problème topologique, il peut être utile d'ajouter une géométrie particulière à la surface considérée. Dans la mesure où la topologie est indépendante de la géométrie choisie, il est possible de définir n'importe quelle notion de distance ou d'angle qui a du sens vis-à-vis du problème étudié. Par exemple, considérons le problème suivant: "Une courbe donnée peut-elle être continûment déformée en une courbe simple (i.e. une courbe sans auto-intersection)?". Imaginons qu'on puisse trouver une géométrie particulière pour laquelle le processus de raccourcissement de courbe donnera toujours une courbe simple, s'il en existe une.

On pourrait alors trouver la solution à notre problème en étudiant ce qui est connu à propos des raccourcissements de courbes. Je voudrais insister sur le fait qu'il est difficile de construire la bonne géométrie à associer à la surface sous-jacente. Effectivement, si toutes les différentes surfaces topologiques peuvent être facilement obtenues comme objets discrets, ce n'est pas le cas pour toutes les géométries (par exemple, les métriques Riemanniennes sont problématiques).

0.2 Problèmes étudiés

Dans cette thèse nous nous intéressons à trois problèmes de natures assez différentes. Le premier propose des algorithmes pour un problème de topologie très classique. Ici, la difficulté est de proposer une discrétisation qui permette d'avoir accès aux résultats connus dans le cas continu. Ensuite, nous regardons en détail une conjecture à propos de la structure des cartes combinatoires. Finalement, la dernière étude concerne l'énumération des cartes combinatoires vérifiant certaines propriétés. Regardons ces trois problèmes plus en détail.

0.2.1 Nombre géométrique d'intersection de courbes

Nous souhaitons obtenir un algorithme qui prend en entrée une surface Σ et une courbe C sur Σ. La sortie recherchée est le nombre minimal d'auto-intersection qu'on peut obtenir pour C en l'autorisant à être continûment déformé. Regardons l'exemple de la figure 4. La partie gauche montre une entrée possible. C est dans une configuration assez complexe qui peut être très largement simplifiée comme le montre la partie droite de la figure. Dans ce cas la sortie devrait être 1.

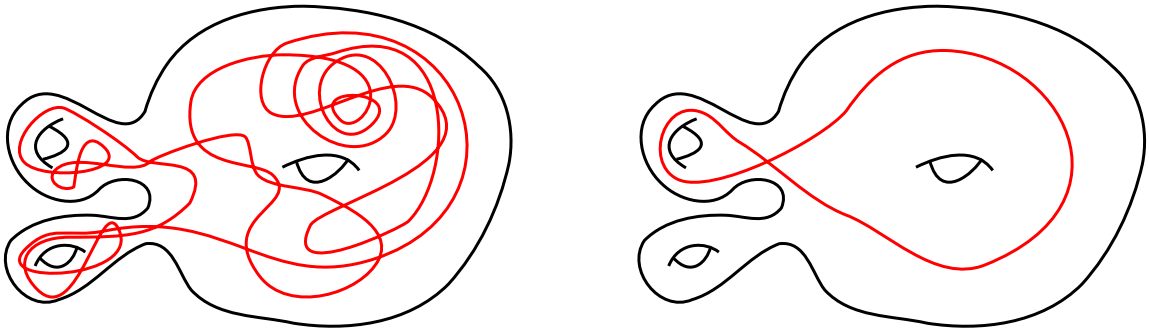

Figure 4: La même courbe (à déformation continue près) dans deux configurations.

Plusieurs variantes peuvent être imaginées. La sortie peut ne pas se contenter d'être le nombre de croisements minimum mais ajouter une configuration de C réalisant ce minimum. De façon assez surprenante, ce problème se révèle être significativement plus compliqué. Une autre variante est de juste vérifier si C peut être déformée en une courbe simple. Cela peut être intéressant car les techniques de chirurgie de surfaces (type copier-coller) sont souvent utilisées en topologie pour prouver des théorèmes. Pour pouvoir couper correctement une surface selon une courbe, il est nécessaire que celle-ci soit simple. Ce dernier problème apparaît dans les travaux de Poincaré. Finalement, les mêmes questions peuvent être posées pour les intersections de systèmes de courbes au lieu des auto-intersections d'une seule courbe.

Ces problèmes sont des problèmes naturels de topologie générale, ils ont donc été très étudiés depuis l'époque de Poincaré jusqu'à aujourd'hui. Plusieurs solutions intéressantes ont été proposées et certaines sont de nature algorithmique. D'autres idées viennent de la théorie des nœuds. Un nœud est un cercle plongé dans $\mathbb{R}^3$. La façon la plus simple de visualiser un nœud est de le projeter sur un plan et de considérer le diagramme obtenu (la figure 5 en donne des exemples). Ces diagrammes sont des courbes dessinées dans le plan avec une notion supplémentaire de dessus et dessous aux auto-intersections. La question la plus naturelle à propos d'un nœud est de savoir s'il est dénoué. L'une des approches classiques est d'appliquer une série de transformations particulières appelées mouvements de Redemeister dans le but de dénouer le diagramme. Un résultat très connu prouve que ces mouvements sont suffisants pour dénouer n'importe quel nœud qui peut l'être. Plus récemment, la nature polynomial [Lackenby, 2013] du nombre de mouvements nécessaires a été prouvée. Cette approche peut être utilisée pour résoudre notre problème en autorisant plus de mouvements puisque nous n'avons pas à respecter des notions de dessus et dessous.

Nous avons construit nos algorithmes en associant toutes ces idées dans un cadre discret. Nous avons obtenu des algorithmes simples dont les caractéristiques sont décrites dans les théorèmes 3.1.1, 3.1.2 et 3.1.3. La discrétisation du problème est la clé de nos résultats. Les détails sont décrits dans le chapitre correspondant mais on peut faire ici quelques remarques intéressantes. Nous avons utilisé des cartes combinatoires particulières appelées systèmes de quads qui ont des propriétés géométriques très utiles. Ils fournissent un compromis intéressant entre efficacité et précision (il est difficile d'obtenir des propriétés géométriques proches de la géométrie hyperboliques). Les systèmes de quads et leur géométrie sont décrits en détail dans les préliminaires.

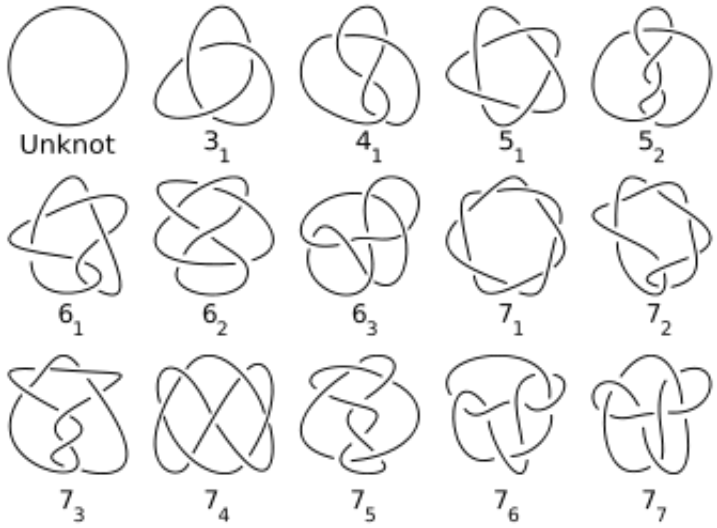

Figure 5: Classification des petits nœuds. Wikipedia

0.2.2 Cycle de partage

Le paramètre principal pour définir une surface topologique est son genre. Il correspond au nombre de poignées de la surface dans le cas orientable et le nombre de copies disjointes de plans projectifs troués dans le cas non-orientable. Nous nous concentrons sur le cas des surfaces orientables. Les courbes simples des surfaces ont elles deux caractéristiques importantes. Le première est la contractibilité. Une courbe est dite contractile si et seulement si elle peut être déformée continûment en un point. La seconde caractéristique est le caractère séparant. Une courbe simple est séparative si et seulement si son complément est déconnecté. Toutes les courbes simples contractiles sont séparatives. Les courbes contractile peuvent être caractérisées par le fait qu'elle séparent la surface en deux parties dont l'une est un disque. Cependant, les courbes non-contractile peuvent être séparatives ou non. Un cycle de partage est une courbe simple séparative et non-contractile. Une telle courbe sépare la surface en partie de genre au moins 1. Le genre est additif donc la somme des genres des deux parties séparées par un cycle de partage est égale au genre de la surface de départ. Ainsi seules les surfaces de genre au moins 2 peuvent avoir un cycle de partage. Dans le monde continu, toutes les surfaces de genre au moins 2 admettent un cycle de partage (voir figure 6).

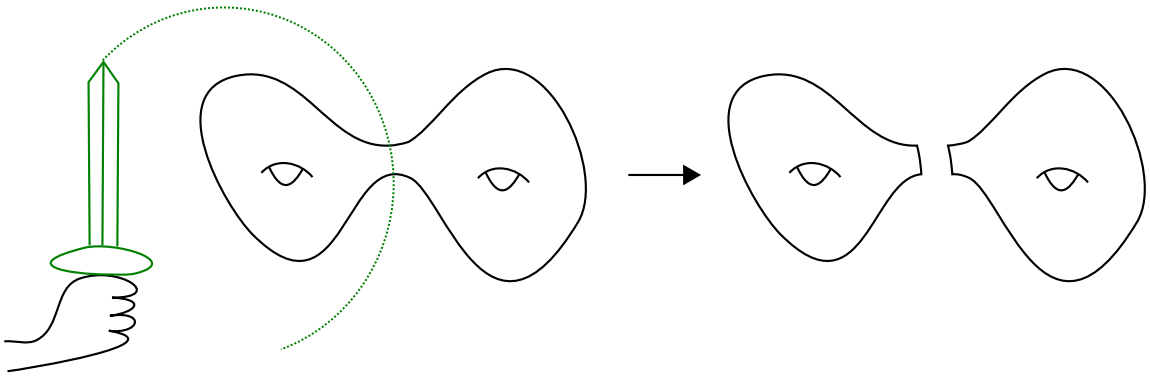

Figure 6: Un cycle de partage dans un cadre continu.

L'intérêt de trouver des cycles de partage est, entre autres, de pouvoir les utiliser pour découper la surface en plusieurs morceaux. Ce genre de découpes peut être utilisé pour faire des preuves par récurrence ou pour construire des algorithmes de type diviser pour régner. Il est donc naturel de rechercher des cycles simples. Une surface topologique a une infinité de cycles de partage. Effectivement, à partir d'un cycle comme celui de la figure 6 on peut obtenir d'autres cycles en appliquant les automorphismes de la surface. Dans le monde du discret, être un cycle simple est un propriété beaucoup plus difficile à obtenir. On veut obtenir des cycles qui sont des marches fermées simples dans le graphe sous-jacent. Il apparaît que certaines cartes combinatoires n'admettent pas de cycle de partage. En plus, il est difficile de savoir si une carte donnée à un cycle de partage car le problème est connu pour être NP-complet pour des cartes quelconques [Chambers et al., 2006].

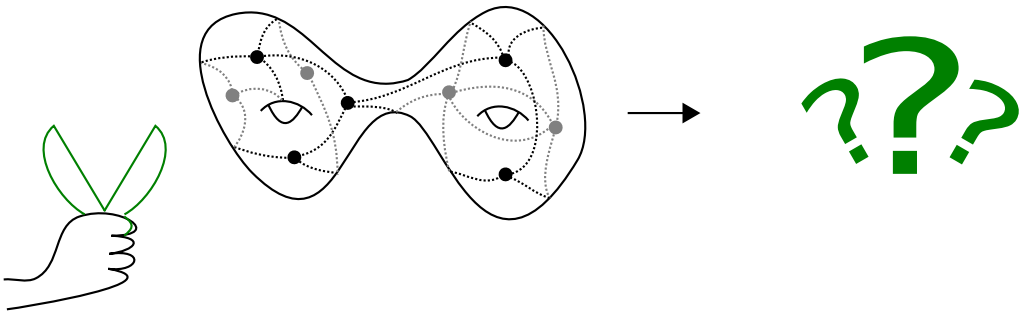

Figure 7: Un cycle de partage dans un cadre discret.

Il a été conjecturé par Barnette dans les années 80 que toutes les triangulations simples (telles que le graphe sous-jacent, qui n'ont pas d'arêtes boucles ni d'arêtes

multiples) doivent avoir un cycle de partage. Par la suite, Mohar et Thomassen ont ajouté que ces cartes devaient avoir des cycles de partage permettant de réaliser toutes les répartitions de genre possible. Cela nous a amené à introduire la notion de type d'un cycle de partage qui correspond au genre des deux parties obtenues en découpant la surface selon le cycle. Un cycle de partage est ainsi dit équilibré quand son type est équilibré. Nous avons trouvé des contre-exemples à cette dernière conjecture en utilisant des plongements de graphes complets. Par exemple, nous avons trouvé des plongements du graphe complet à 19 sommets sur la surface de genre 20, qui n'ont pas de cycle de partage équilibré (de type (10,10)).

0.2.3 Triangulations du tore

Considérer que les surfaces sont des cartes combinatoires, soulève immédiatement des questions. Par exemple, étant donnée une surface Σ et un nombre n, combien de cartes à n sommets sont homéomorphes à Σ? À quel point peuvent-elles être différentes? Nous nous intéressons ici à un type particulier de cartes, les triangulations. Pour les compter plus facilement, on considère en général des cartes enracinées. Une racine est une arête orientée ou alternativement un angle fait face. Un automorphisme de carte enracinée préserve la racine. La plupart des cartes enracinées n'ont aucun automorphisme [Richmond and Wormald, 1995] non trivial ce qui les rend plus pratique à compter.

De nombreux travaux existent avec de nombreux différents types de cartes, en particulier dans le cas planaire. Par exemple, le nombre de triangulations enracinées de la sphère avec $n+2$ sommets est égal à $\frac{2(4n-3)!}{n!(3n-1)!}$ et ce résultat est connu depuis 1962 [Tutte, 1962]. Il est assez rare d'obtenir une formule aussi simple. Le nombre de cartes à n sommets d'un type donné est plus souvent obtenu par une formule de récurrence sur n. On remarquera que ces formules sont généralement énoncées à l'aide de séries formelles.

L'histoire de l'énumération commence avec l'énumération d'ensembles faciles à compter de façon directe, comme typiquement l'ensemble des arbres enracinés. Ensuite pour compter un certain type d'objets, la première idée est d'établir une bijection entre ces objets et un ensemble bien connu. Pour ce qui nous concerne, considérons le cas des triangulations planaires enracinées. Elles sont en bijection avec un sous-ensemble des arbres enracinés. On ne donnera pas les détails de cette construction ici mais une illustration est donnée en figure 8. Cette bijection permet de compter le nombre de triangulations mais a de nombreuses autres conséquences. Celle qui est le plus en rapport avec le contenu global de cette thèse est la possibilité d'obtenir un tirage au sort uniforme des triangulations planaires enracinées à

n sommets en temps linéaire par rapport à n.

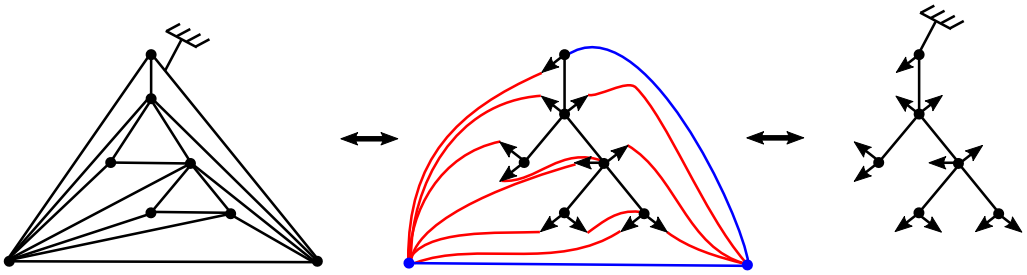

Figure 8: Une bijection pour les triangulations planaires enracinées.

Le problème qui nous intéresse dans cette thèse est la généralisation de cette bijection au cas des triangulations du tore. Premièrement, remarquons que les triangulations du tore ne peuvent pas naturellement être en bijection avec des arbres puisque ceux ci ne peuvent pas être plongés cellulairement sur une autre surface que le plan ou la sphère. La bonne généralisation des arbres en genre supérieur nécessite l'utilisation de cartes unicellulaires qui sont des cartes à une seule face. La bijection qu'on recherche doit être explicite et facilement calculable pour être intéressante. Dans le cas planaire, la construction de la bijection est connectée à la notion de bois de Schnyder. Un bois de Schnyder associé à une triangulation planaire permet une orientation ainsi qu'une coloration des arêtes du graphe sous-jacent (voir figure 9). Dans ce cas, les orientations et les couleurs sont très faciles à construire puisque que chaque couleur prend la forme d'un arbre couvrant orienté. Nous avons réussi à obtenir une généralisation satisfaisante pour le cas du tore.

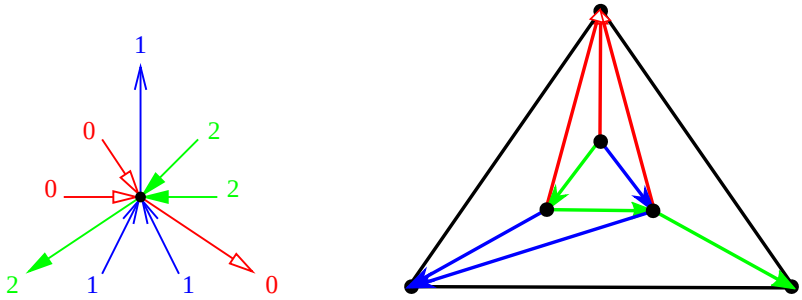

Figure 9: La règle local et un exemple de triangulation.

0.3 Organization

Dans le chapitre suivant nommé Préliminaires nous donnons toutes les définitions nécessaires aux études de nos trois problèmes. Comme nous l'avons dit au début de cette introduction, trois différents points de vue sont utilisés au cours de cette thèse. Nous décrivons en détail chacune de ces constructions ainsi que les connections qui les relient. Dans la mesure où nous souhaitons être lisible par des gens de chacune des communautés correspondantes, nous commençons par des rappels avant de rentrer dans les détails.

Ensuite nous développons les solutions que nous proposons pour les trois problèmes énoncés dans les chapitres 3, 4 et 5. Trois articles ont été rédigés qui correspondent à ces trois problèmes, deux d'entre eux sont en cours de relecture et celui sur le cycle de partage a été accepté dans la revue Graphs and Combinatorics. Chacune de ces parties a un contenu similaire à la version arXiv de l'article correspondant mais avec une formulation harmonisée et des remarques additionnelles. L'ordre dans lequel les chapitres sont présentés a été choisi en fonction des remarques suivantes. Premièrement, nous traitons le problème d'intersection géométrique de courbes car c'est un problème naturel et qu'il permet de se familiariser avec les détails de la traduction du continu au discret. Ensuite le problème du cycle de partage montre que les constructions discrètes peuvent aussi avoir des propriétés qui ne peuvent pas être facilement déduites du cas continu. En particulier, la définition de courbe simple est très différente dans ces deux premiers chapitres. Dans le chapitre 3 on donne un algorithme quasi-linéaire et donc très efficace pour tester la simplicité d'une classe d'homotopie alors que trouver un cycle simple de partage est difficile et nécessite d'utiliser des heuristiques pour être en pratique calculable. Dans cette seconde partie, nous nous sommes intéressés à la possibilité de tester notre conjecture sur des triangulations aléatoires. Le dernier chapitre nous rapproche de cette possibilité pour le tore et peut-être dans un avenir proche seront nous capables, en poursuivant cette étude, d'obtenir le tirage aléatoire uniforme pour les triangulations en genre quelconque.

Chapter 1

Introduction

1.1 Subject

Topology is the study of properties of shapes that are preserved by **continuous** deformations. There is no notion of length, area or curvature which belong to the world of geometry. Geometric properties are only preserved by rigid transformations such as rotations and translations. In some sense topology is thus concerned with more robust properties. We usually trace topology back to Euler (1707-1783) who discovered the famous formula, now bearing his name: V-E+F=2. Here V, E and F are respectively the number of vertices, edges and faces of a given polyhedron. This can be easily check on the platonic solids as illustrated in Figure 10. That this formula not only applies to convex polyhedron but to any subdivision of the sphere will be noticed much later when topology will become a discipline in its own.

Figure 10: The five platonic polyhedra: tetrahedron, cube, octahedron, dodecahedron and icosahedron. http://whistleralley.com/polyhedra/platonic.htm

It is possible to define a notion of dimension that characterizes the topology of a space. Typically a polyhedron has dimension 3 while its envelop has dimension 2.

Dimension 2 topological spaces capture the intuitive notion of surface. In this thesis, we aim at designing very general tools for the analysis of topological properties of surfaces. Again, to illustrate the difference of point of views between a topologist and a geometer, let us consider figure 11. A topologist sees two identical objects that he calls spheres because the two surfaces can be deformed continuously to a sphere. On the contrary a specialist of geometry would see two distinct shapes. Obviously, adding more structures allows you to distinguish shapes in a finer way. For instance, a human being will not only see distinct shapes but will see a rabbit and a dolphin thus adding some semantic. Clearly the task of classification becomes more complex as more structure is involved. It is thus a good idea to apply a progressive classification, first considering purely topological aspects, then refining with geometric considerations and finally semantic discriminators. In fact topology often appears as the first step in a classification process.

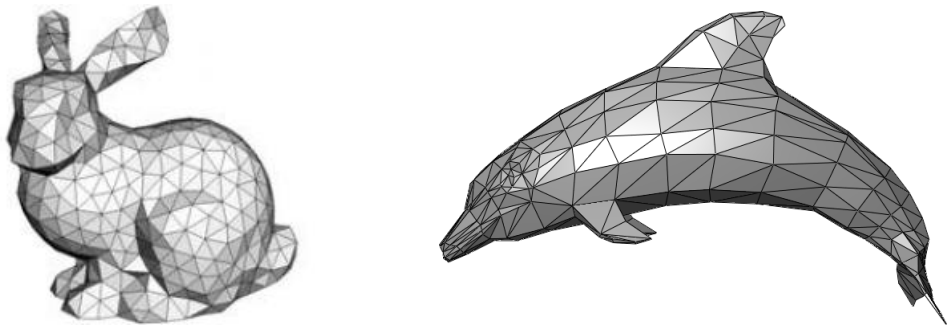

Figure 11: Two spheres.

Several communities are dealing with topological properties of surfaces with their own point of view. The classical topologist defines a surface as a 2-dimensional manifold. The idea is to consider the plane as the general model for a two dimensional object and then call surface any topological space that locally looks like that model. This point of view takes its source in the famous Habilitationsvortrag by Riemann (1854). In parallel to this approach a combinatorial viewpoint was developed. Here spaces are considered as assemblies of elementary pieces, usually simplexes (points, edges and triangles as far as surfaces are concerned). The model of combinatorial maps described by rotation schemes belongs to this trend. It first appeared in the work of Heffter (1891) and Edmonds (1960) and was further developed by people like Youngs (1968) and Tutte (1973). An intermediate viewpoint was adopted by many graph theorists. They envision a graph as a 1-dimensional structure that can be drawn on a topological surface. This naturally leads to the notion of cellular

embeddings of graphs. It was popularized by Robertson and Seymour in their proof of the graph minor theorem (1983-2004).

We illustrate those different points of view in figure 12. First (a) shows the classical representation of a torus as a shape embedded in $\mathbb{R}^3$. (b) also defines a torus by identifying the opposite sides of a square. This representation is used for the next cases because it allows to have a good view on the whole torus (nothing is hidden behind). (c) corresponds to a torus constructed from an assembly of triangles. Those kind of combinatorial maps are classical objects called triangulations. Finally (d) shows a graph on 3 vertices embedded on the torus. It is natural to try to draw a graph on a surface and it allows very natural description of problems in graph theory. The counterpart of that last representation is that it leaves a lot of freedom for the curves corresponding to edges, making them potentially very complicated. In this thesis, combinatorial maps will be used for the description of algorithms and their proofs but all the points of view will be used in order to try to give the most natural description possible of each problem considered.

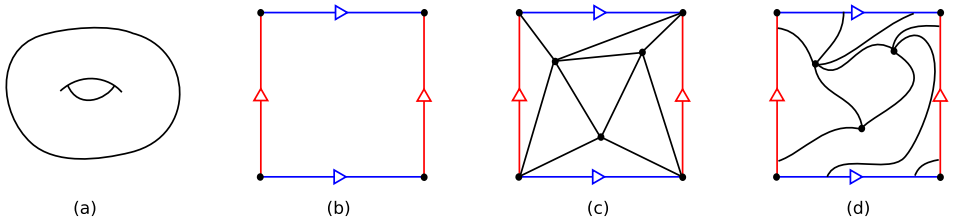

Figure 12: Three tori seen by three different communities.

We already discussed the difference between geometry and topology. However, it can be interesting to consider geometry even for purely topological questions. Indeed when looking at a topological problem it can be useful to add a specific geometry to the underlying surface. Since topological problems are independent of the geometry, it is possible to define any notion of distance or angle that makes sense with respect to the considered problem. For instance, let us look at a problem defined as "Can a given curve be deformed continuously to a simple curve (i.e. a curve without crossings)?". Imagine that we can find a geometry for which tightening a curve with respect to that geometry will always outputs a simple curve if it exists. By looking what is known about those tightened curves we can finally find the answer to the initial topological question. I want to emphasise that the difficulty here is to construct the interesting geometry from the initial surface. Indeed all topological surfaces can be obtained as discrete constructions but it is not the case for all possible geometries (consider for instance Riemannian metrics).

1.2 Studied Problems

In this thesis we study three particular problems. They have quite different flavors. The first one takes a classical problem of general topology and develops algorithms to deal with it. One difficulty is to correctly translate the continuous problem into an appropriate discrete setting. Then, we look in details at a conjecture about the structure of combinatorial maps. Combinatorial maps are of intrinsic interest as mathematical objects. In addition, it provides a very good data-structure to store a surface in a compact way. Thus, we will use them for all our algorithms. Finally the last study is about the enumeration of combinatorial maps with a fixed property. Let us look the corresponding problems in more details.

1.2.1 Geometric Intersection Number of Curves

We are looking for an algorithm that takes as input a surface Σ and a curve C on Σ. We want to output the minimal number of self-intersections of C by allowing it to be continuously deformed. Let us look at figure 13 for an example. The left part shows a possible entry. C has a quite complicated arrangement that can be strongly simplified as shown in the second part of the figure. In that case the output should be 1.

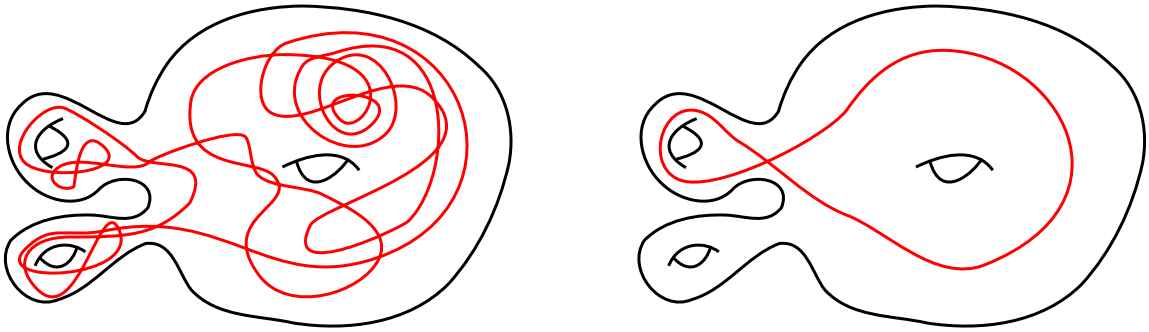

Figure 13: The same curve (up to continuous deformations) on a surface.

Many variants can be imagined. The output can be not only the number of intersections but also a configuration of C realizing the minimum. Last problem quite surprisingly appears to be significantly more expensive with our current algorithms. Another variant is to just ask if C can be deformed to a simple curve (i.e. without self-crossings). This can be interesting since surgery of surfaces by cut and paste operators are often used as an efficient tool for proving properties of topological

surfaces. In order to cut a surface properly we often use a simple curve to cut along it. This last problem first appeared in Poincaré's work. Finally, the same things can be asked for a system of curves instead of just one curve by considering crossings of curves instead of self-intersections.

Those problems are natural problems of general topology. They have thus been well studied. First the question has been looked in details starting from Poincaré up to now. A lot of interesting solutions are proposed and some of them are algorithmic. Other ideas come from knot theory. A knot is a circle embedded on $\mathbb{R}^3$. The easiest way to represent a knot is to project it on a plane and look at the obtained diagram (see figure 14 for examples). It corresponds to curves drawn in the plane with an added notion of top and bottom for the self-crossings of the curve. The most natural question about a knot is if it is unknoted. A classical approach is to perform particular transformations on the diagram called Redemeister moves in order to untie the diagram. A very well known result states that those particular moves are sufficient to untie any knot that can be and recently the number of moves necessary to unknot a knot as been shown to be polynomial [Lackenby, 2013]. This approach can be used to solve our problem by allowing more Redemeister-like moves since we do not have to respect rules about top and bottom at crossings.

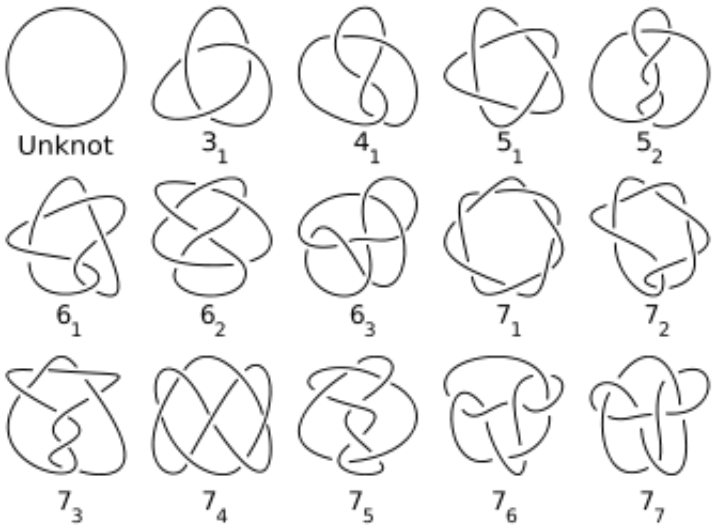

Figure 14: Classification of small knots. Wikipedia

We design algorithms by putting together all those ideas. It leads to simple algorithms whose precise characteristic are given by Theorems 3.1.1, 3.1.2 and 3.1.3.

It is not obvious to understand how the problem is discretized and this is extremely important in this work. This is described in full details in the corresponding section but here we can make an interesting remark. We use particular maps called system of quads that have very interesting geometric properties. It gives a good compromise between efficiency and a precise description of a geometry (often heavy in memory and expensive to manipulate). The system of quads and their geometry are described in details in the preliminaries.

1.2.2 Splitting Cycle

A key parameter of a surface is its genus. It can be seen as the number of handles of the surface in the orientable case and twice the number of crosscaps otherwise. We focus here on orientable surfaces. Simple curves on surfaces have two important characteristics. The first one is contractibility. A curve is contractible when it can be continuously deformed into a point. The second characteristic is separability. A simple curve is separating if its complement has two connected components. A contractible curve must be separating. Contractible simple curves can be characterized by the fact that they must bound a disk on one side. However, non-contractible ones can be separating or not. A splitting cycle is a simple closed curve that separates the surface and that is non-contractible. It thus cuts the surface into two pieces of genus at least 1. Note that the genus of the starting surface is the addition of the genus of the two parts obtained by a cut along a splitting cycle. So only surfaces of genus at least 2 may have a splitting cycle. In the continuous world every surface of genus at least 2 can be split that way [Chambers et al., 2006] (see figure 15).

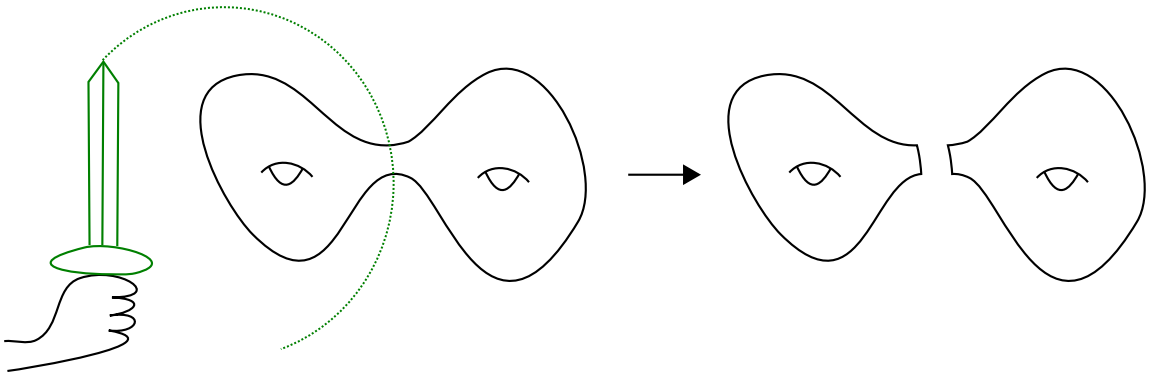

Figure 15: A splitting cycle in a continuous setting.

The interest of finding a splitting cycle is to use it to cut the surface into pieces of smaller genus. It could be used to make proofs by induction or to build divide

and conquer algorithms. It is thus natural to require a cycle in order to make a proper cutting. A topological surface has infinitely many such cycles. Indeed, It is easy to find some cycles on a topological surface like the one of figure 15 for instance. Then any self-homeomorphism of the surface will transform any splitting cycle another one. Self-homeomorphisms of surfaces can be seen as symmetries and local contractions or dilatations. They lead to infinitely many different copies to any cycle on the surface. In the discrete world things are totally different. The point is that being simple has a different meaning. To cut properly a combinatorial map it is reasonable to restrict to simple closed walks in the underlying graph. A splitting cycle in the discrete sense does not exist in all combinatorial maps. In addition, it is not easy to decide if a given map has one or not (see figure 16) since the problem is known to be NP-complete for general combinatorial maps.

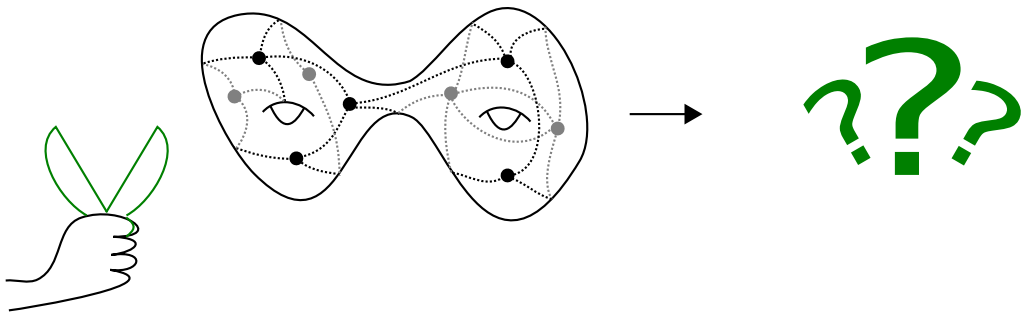

Figure 16: A splitting cycle in a discrete setting.

It was conjectured by Barnette in the eighties that maps constructed as assembly of triangles with a simple induced graph (in other words simplicial triangulations) must have a splitting cycle. It then was added by Mohar and Thomassen that it may be possible to find a splitting cycle for any partition of the genus of the simple triangulation. It has led us to introduce the **type** of a splitting cycle that is, in the orientable case, the genus of each connected components of the surface after the removal of the splitting cycle. A splitting cycle is said balanced if its type is balanced. We found counter-examples to the last conjecture using embeddings of complete graphs. For instance, we find that there exist embeddings of the complete graph on 19 vertices on the orientable surface of genus 20 without balanced splitting cycles.

1.2.3 Triangulations of the Torus

Considering surfaces as combinatorial maps rises some natural questions. For instance, given a surface Σ and a number n, how many maps with n vertices are homeomorphic to Σ? How "different" can they be? We focus here on a particular kind of maps: triangulations. These are the maps with triangle faces without loops or multiple edges. In order to be able to count them correctly it is generally easier to count the rooted maps. A root is an oriented edge of a map or equivalently an angle of a face. An automorphism of rooted maps is one that preserves the root. Most of rooted triangulations have a trivial automorphism group [Richmond and Wormald, 1995] making them easier to count than unrooted ones.

A lot of work has been made for a lot of different kinds of maps, in particular in the planar case. For instance, the number of distinct rooted triangulations of the sphere with $n+2$ vertices is known to be $\frac{2(4n-3)!}{n!(3n-1)!}$ since 1962 [Tutte, 1962]. It is quite unusual to find such a simple formula. The number of maps with n vertices of a given type is generally given by an inductive formula on n. Note that those formulae are generally obtained by manipulating formal series.

The story of enumeration starts with the enumeration of sets that can be easily counted. Typically the set of rooted embedded trees. Then when aiming at counting a special kind of objects the first idea became to find a bijection between those objects and a well known set. For our purpose let us consider the case of rooted planar triangulations on $n+2$ vertices. They are in bijection with rooted embedded trees of a special form. The details will not be given here but see figure 17 for an illustration. This bijection allows to count the triangulations but also has a lot of interesting consequences. The one that mainly makes sense with respect to the content of this thesis is that we obtain uniform sampling of triangulated maps of size n in linear time.

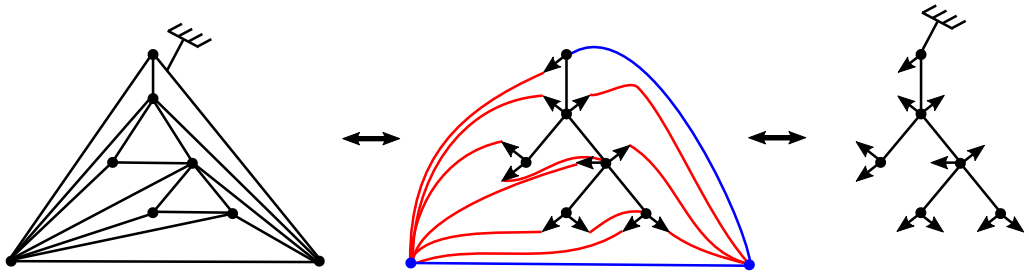

Figure 17: A bijection for rooted planar triangulations.

Our point in this thesis is the generalization of this bijection to toroidal triangulations. First we remark that triangulations of the torus cannot be put in bijection

with embedded trees because any map associated to a tree is planar (an embedded tree cannot encode the topology of the initial map). The natural generalization of trees here is unicellular toroidal maps meaning combinatorial maps with only one face and the topology of the torus. The bijection we are looking for has to be explicit and efficiently computable in order to be useful. The bijection in the planar case is connected to Schnyder woods. A Schnyder wood of a planar triangulations is a colored orientation of the edges of the underlying graph of the triangulation (see figure 18). In that case the orientation and coloration is really easy to construct since each color forms an oriented spanning tree. We achieve the generalization to the torus by keeping the main advantages of the planar bijection.

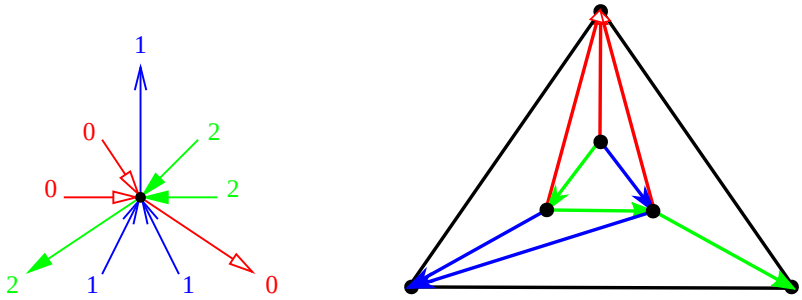

Figure 18: The local rule and an example.

1.3 Organization

In the next "preliminary" chapter we explain all the notions that are necessary to understand our solutions to the problems described above. As explained in the introduction three points of view are involved in this thesis. We describe in details the corresponding structures. We first recall some basic definitions. After those basics we give the precise definitions we use for the continuous point of view and for the discrete one.

Then we develop the Chapters 3, 4 and 5 corresponding to our solutions to the three announced problems. It has lead to three articles that are submitted at the time when this thesis was written. Each chapter is quite similar to the arXiv version of the corresponding article with homogenized formulation and additional personal remarks. The order in which the problems are presented has been chosen as follows. First the geometric intersection problem is very natural and highlight a particular translation of a continuous problem in a discrete setting. Then the splitting cycle

problem illustrate the fact that combinatorial maps have particular properties that cannot be easily guessed from continuous intuition. In particular, the concept of simplicity is not naturally defined in the same way. Chapter 3 gives a quasi-linear algorithm to test continuous simplicity while discrete simplicity is very difficult to evaluate and needs to use heuristic to be calculated in practice. We get interested in testing conjectures against random triangulations. The last problem brings us closer to uniform sampling of triangulations of the torus and gives hope to a generalization on general surfaces.

Chapter 2

Preliminaries

2.0 Basics

2.0.1 Topology

The definition of topology can be applied to define a lot of very different spaces. We first give the most general definition. Then we look at the particular case of $\mathbb{R}^n$ which is the most natural example and which is used as model to define particular topological spaces called manifolds.

Definition Giving a topology to a set S can be understood as giving a sense to the notion of neighborhood. A neighborhood of a point is a particular subset containing that point. In practice defining a topology consists in giving the list $\mathcal{O} \subset \mathcal{P}(S)$ of all the basic subsets that can be considered as neighborhoods. Those subsets are called **open subsets**. In order to have a topology some axioms are added. First $\mathcal{O}$ must contain $\emptyset$ and S. Then $\mathcal{O}$ has to be preserved by unions and finite intersections. A subset F is called a **closed subset** if and only if its complement is open. A function $f : (S, \mathcal{O}) \to (S', \mathcal{O}')$ is continuous if and only if for all $O' \in \mathcal{O}'$, $f^{-1}(O') \in \mathcal{O}$. An **homeomorphism** between $(S, \mathcal{O})$ and $(S', \mathcal{O}')$ is a bijective function $h : S \to S'$ such that f and f^{-1} are continuous. In other words, a homeomorphism preserves the topology and homeomorphic spaces can be considered as identical.

Particular Case of $\mathbb{R}^n$ There are various ways to define topologies on $\mathbb{R}^n$. We can define trivial topology like $\mathcal{O} = \{\emptyset, \mathbb{R}^n\}$ or $\mathcal{O} = \mathcal{P}(\mathbb{R}^n)$. This is not of great interest for our purpose. Let us look at the canonical way to define the topology of $\mathbb{R}^n$. It is done via the classical Euclidean norm on $\mathbb{R}^n$ ($||(a_1, \cdots, a_n)||_2 = \sqrt{a_1^2 + \cdots + a_n^2}$).

The subsets defined as $B(x,r) = \{y \in \mathbb{R}^n / ||y-x||_2 < r\}$ are called open balls. We want open balls to be open subsets. So the canonical topology of $\mathbb{R}^n$ is the one generated by the open balls. Elements of $\mathcal{O}$ are thus subsets obtained by union of finite intersections of open balls. In practice a subset O is open if and only if for all $x \in O$, there exists $l > 0$ such that $B(x,l) \subset O$.

Properties We describe here some classical topological notions of particular interest for our purpose. The first one is compactness. A topological set C (or a subset) is compact if and only if from all set of open sets whose union contains C, a finite subset can be extracted that covers C. In the particular case of $\mathbb{R}^n$ with its classical topology compact subsets are exactly those that are closed and bounded. The other notion that we will use is connectedness. A topological space S is connected if and only if there does not exist two disjoint open subsets covering S. It can be understood as S is connected if it cannot be decomposed into multiple parts with respect to the topology. A connected component of S is a maximal subset of S that is connected. Note that this can alternatively be defined as follows. S is said arc-connected if and only if every pair of points of S can be joined by a continuous path in S. Note that the two points of view are different in a general setting but will be equivalent for our particular topological spaces. We quite always ask for connectedness because our algorithms will be just run independently on each connected components otherwise.

Manifolds The idea of the definition of manifolds is to characterize the topological spaces that locally "look like" $\mathbb{R}^n$. Note that "looks like" in topology means there exists a homeomorphism and we thus obtain the following definition. A **n-manifold** is a non-empty (completely separable) topological space M such that for all $x \in M$ there exists an open neighborhood U_x of x and a homeomorphism $\varphi_x : U_x \to \mathbb{R}^n$ (note that completely separable is a technical assumption ensuring that M is not "too big"). For instance the sphere $S^n = \{x \in \mathbb{R}^{n+1} \text{ s.t. } ||x||_2 = 1\}$ is a n-manifold. Manifolds with boundaries can be defined in the same way but instead of considering the unique model $\mathbb{R}^n$ we also allow points of M to have neighborhoods homeomorphic to the closed half-plane $\{(x_1, \cdots, x_n) \in \mathbb{R}^n / x_1 \geq 0\}$. In this thesis, we will focus on topological issues of 2-manifolds. However, it is sometimes useful to consider the geometric point of view (mainly in Chapter 3). In order to state properly some interesting results we give the following definitions. That point is not critical in this thesis so we do not go in full details and we assume some familiarity with differential calculus.

Additional Structures on Manifolds The canonical topology of $\mathbb{R}^n$ comes together with an inner product and so a notion of distance. In addition, $\mathbb{R}^n$ has a structure of vector space. Let us define a structure that allows to carry over those properties to manifolds. The main difficulty at first sight is that to each point $x \in M$, the corresponding homeomorphism φ_x suggests a local behavior on U_x and there is no reason that those behaviors should be coherent. To minimize those problems we define the notion of atlas. An **atlas** of M is a collection of pairs (U_x, φ_x) for a particular set of elements $x \in M$ such that the union of the U_x covers M. Such an atlas is sufficient to define a given structure on M. However, it remains that the φ_x have to be coherent on the intersection of the corresponding U_x. Thus M is said differentiable if and only if it admits an atlas all of which **transition maps** of the form $\varphi_{xy} = \varphi_x \circ \varphi_y^{-1}|_{\varphi(U_x \cup U_y)}$ are C^1. When M has that structure one can define the tangent space T_xM of M at x. It is a quotient of the set of all curves of M hitting x and differentiable at x. Two curves are said to be equivalent in that setting if their derivative are the same after going through φ_x. This is an equivalence relation and allows to define T_xM. T_xM appears to be a n dimensional vector space so one can define inner products on it: $g_x : (T_xM, T_xM) \to \mathbb{R}$. For it to be interesting we need some regularity on those inner products. Let us consider $X, Y : M \to T_xM$ two smooth vector fields. If for all such X and Y, $x \in M \to g_x(X(x), Y(x))$ is smooth then one say that g is a **Riemannian metric** and that (M, g) is a **Riemannian manifold**.

Another important characteristic that can be easily defined using the same tools is the orientability of a manifold. First let us look what we call an orientation of $\mathbb{R}^n$. An orientation of a vector space is a partition of the set of all its ordered basis. We first arbitrary choose an ordered basis B_0 and set it as a direct basis. Then another ordered basis B is direct if the determinant of the transition matrix from B to B_0 is positive and indirect otherwise. Now let go to manifolds. As we see homeomorphism can carry those notions of orientations. Once again a coherent atlas will define an orientation on the manifold. In practice an atlas allows to construct an orientation if and only if its transition maps are C^1 and have positive Jacobian determinant.

Immersions and Embeddings In order to represent a n-manifold in an intuitive way it is natural to think of representing it drawn in some $\mathbb{R}^m$ with m bigger than n. For instance, a circle is a 1-manifold that is not possible to represent on the line ($\mathbb{R}^1$). However it is natural to represent a circle in the plane. The example we are particularly interested in is embedding 2-manifolds in $\mathbb{R}^3$. In that direction, we define two ways to put a manifold into a bigger one. First an **immersion** of a differentiable manifold N into another one M is a differentiable function $f : N \to M$ whose

derivative is everywhere injective. Then in the same setting we call f an embedding if it is a diffeomorphism onto its image. Note that an embedding can be defined in a purely topological context by replacing diffeomorphism by homeomorphism. The important point is the difference between the two notions. An immersion can create self-crossings while an embedding is not allowed to. With that vocabulary, a regular curve is just the circle immersed into a bigger manifold.

2.0.2 Algebra

We fix here all the vocabulary that is useful to describe the algebraic elements of this thesis.

Groups We recall that a **group** $\mathcal{G} = (G, \cdot)$ is a set G and a binary relation $\cdot$ on G such that:

- there is an element $1_G \in G$ such that $\forall g \in G$, $1_G \cdot g = g \cdot 1_G = g$.
- $\forall g_1, g_2, g_3 \in G$, $g_1 \cdot (g_2 \cdot g_3) = (g_1 \cdot g_2) \cdot g_3$.
- $\forall g \in G$, $\exists g^{-1} \in G$ such that $g \cdot g^{-1} = g^{-1} \cdot g = 1_G$.

A function $\varphi : (G, \cdot) \to (G', \times)$ is a homomorphism if it respects the group structures associated to G and G'. It means that $\forall g_1, g_2 \in G$, $\varphi(g_1 \cdot g_2) = \varphi(g_1) \times \varphi(g_2)$. We often drop the relation and say that G is a group as an abuse of notation. A subgroup H of a group G is a subset of G containing 1_G that is a group with the relation it inherits from G. The left cosets of H in G are the sets of the form $g \cdot H$. Those sets form a partition of G. It is natural to try to build the quotient of G by H as the group formed by the left classes of H with an induced relation. It means that $(g_1 \cdot H) \cdot (g_2 \cdot H)$ should be equal to $(g_1 \cdot g_2) \cdot H$. In order to obtain that natural construction it is necessary that H verify the following relation: $\forall g \in G$, $g \cdot H \cdot g^{-1} = H$. Such an H is called a normal subgroup. We will often speak of group quotients without explicitly checking the normality of the subgroup but we have to keep in mind that it is required. Remark that the elements $g' \cdot g \cdot g'^{-1}$ are called the conjugates of g.

A group is abelian (or commutative) if for all $g_1, g_2 \in G$, $g_1 \cdot g_2 = g_2 \cdot g_1$. Alternatively the condition can be written $g_1 \cdot g_2 \cdot g_1^{-1} \cdot g_2^{-1} = 1_G$. The elements $g_1 \cdot g_2 \cdot g_1^{-1} \cdot g_2^{-1}$ are called the commutators of G. If G is not abelian we can obtain a maximal abelian group by quotienting it by the group generated by its commutators. We then obtain the abelianization of G.

Free Groups A **free group** $\mathcal{F} = (G, \cdot)$ is defined by giving a set of elements A called the **generators** of G. We then construct an alphabet from A as $\mathcal{A} = A \cup A^{-1}$ (where A^{-1} is the set of abstract inverses of elements of A). A word on $\mathcal{A}$ is a finite sequence of elements of $\mathcal{A}$. It is called a **reduced word** if it does not contain two consecutive inverse elements. Then G is the set of all reduced words of elements of $\mathcal{A}$ and $\cdot$ is the concatenation operation (with possible reductions). Note that the element 1_G is the empty word. For instance the free group with only one generator is isomorphic to $(\mathbb{Z}, +)$. This notion is extremely useful since all groups can be expressed as a quotient of a free group by a subgroup. In that idea, a group $\mathcal{G}$ can be defined as a couple $\langle A; R\rangle$ where R is a set of elements of the free group generated by A that are set to 1. Then we say that G is represented by **generators and relations**. It is the quotient of the free group generated by A by the smallest normal subgroup that contains R. For instance $G = \langle a; a \cdot a\rangle$ is isomorphic to $\mathbb{Z}/2\mathbb{Z}$. The relation $a \cdot a = 1_G$ can be interpreted as "a is of order 2".

Group Actions An **action** φ of a group G on a set S is an application from $G \times S$ to S such that: $\forall g_1,\, g_2 \in G$ and $s \in S$, $\varphi(g_1, \varphi(g_2, s)) = \varphi(g_1 \cdot g_2, s)$. φ is often noted as a binary operator $\cdot$ making the last equation written as $g_1 \cdot (g_2 \cdot s) = (g_1 \cdot g_2) \cdot s$. The **orbit** of s under the action φ is the set of all $s' \in S$ such that there exists a $g \in G$ such that $s' = g \cdot s$. Two orbits $\mathcal{O}(s_1)$ and $\mathcal{O}(s_2)$ can only be equal or disjoint making the orbits a partition of S.

Lattices This part is about set orderings. This is a wide subject of study on its own and we just give some useful definitions. Let S be a set. A binary relation $\leq$ is a partial order on S if it verifies the following rules: (reflexivity) $\forall x \in S, x \leq x$, (anti-symmetry) $\forall x, y \in S, x \leq y$ and $y \leq x \Rightarrow x = y$ and (transitivity) $\forall x, y, z \in S, x \leq y$ and $y \leq z \Rightarrow x \leq z$. A minimal element m of $(S, \leq)$ is an element such that $\forall x \in S, m \leq x$. A similar definition holds for the maximum M. Let us consider x and y in S. Let $m(x, y)$ be the set of elements z in S such that $z \leq x$ and $z \leq y$. If $m(x, y)$ is not empty and admits a unique maximum we say that x and y admit a **meet**. By reversing minimum and maximum we define the **join** of a pair of elements. S is a **lattice** if any pair of elements of S admits a meet and a join. The first property of lattices that we will use is that a lattice always have a unique minimal element. Indeed if m_1 and m_2 are minimal elements then any element of $m(m_1, m_2)$ is smaller than m_1 and m_2 and is different from at least one of them leading to a contradiction. A lattice is said distributive if and only if the two operations that associate the meet and the join to a pair of elements are distributive on each other. Lattice are often represented by their **Hasse diagram**. The Hasse diagram of an order relation

is a diagram representing all the element of the set such that if $x < y$ then x is represented under y. An edge is added between x and y if $x < y$ and there exists no z such that $x < z < y$. In the case of a lattice we obtain a diagram that looks like the one of Figure 19.

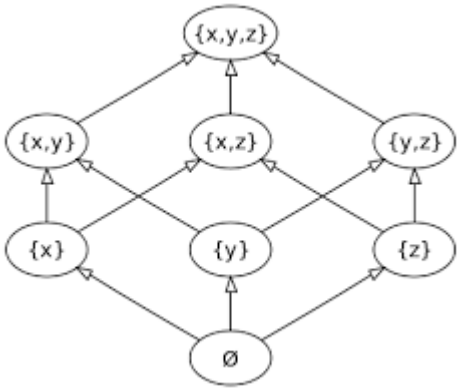

Figure 19: The Hasse diagram of a classical lattice. Wikipedia

2.1 Surfaces: a Continuous Point of View

2.1.1 Dimension 2 Manifolds

Dimension 2 corresponds to the notion of surface. In practice, we call **surfaces** 2-manifolds that are connected and compact. A surface can have boundaries or not. In that setting surfaces can be efficiently classified:

Theorem 2.1.1 (Classification of surfaces)**.** *Any compact connected 2-manifold without boundary is homeomorphic to one of the following surfaces:*

- *The sphere.*
- *A connected sum of g tori $T\# \cdots \#T$ (meaning g tori glued together).*
- *A connected sum of g projective planes $P\# \cdots \#P$.*

In other words a surface without boundary is defined by its orientability and an integer g called its **genus**. By Whitney's theorem it is known that every 2-manifold can be immersed in $\mathbb{R}^3$ (see figure 20). In fact, all orientable surfaces can be embedded in $\mathbb{R}^3$ but no non-orientable surface (without boundary) can. Note that

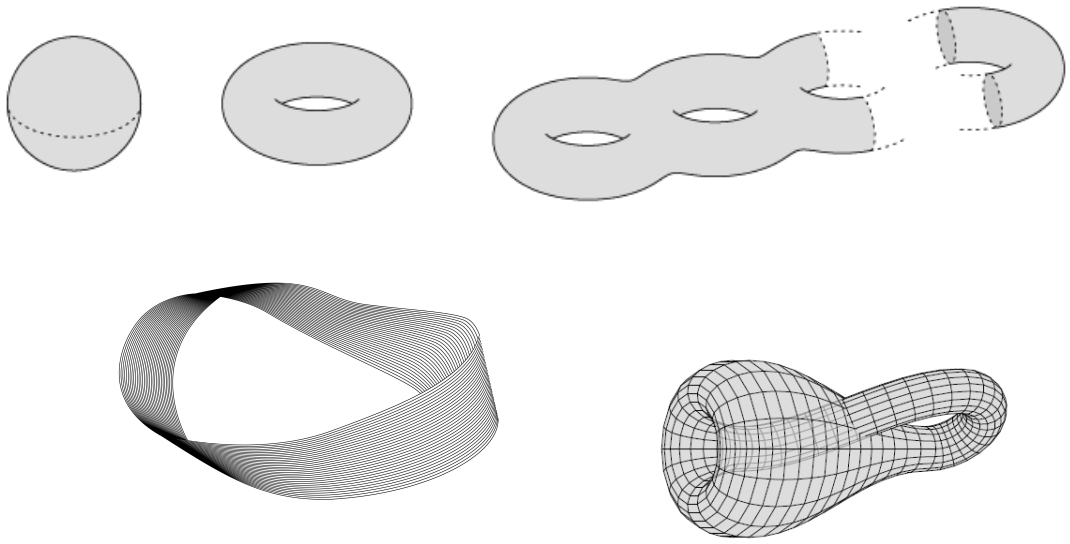

Figure 20: Orientable surfaces, a Möbius strip and a klein bottle

a surface is non-orientable if and only if it contains a Möbius strip. For surfaces with boundary it works the same way but with an additional parameter b corresponding to the number of boundaries of the surface. This means that there is only one surface (up to homeomorphism) with fixed orientation, genus and number of boundaries.

2.1.2 Homotopy of Curves

A curve on a surface Σ is a continuous function $\gamma : [0,1] \to \Sigma$. It is a closed curve if $\gamma(0) = \gamma(1)$. In this case, it is possible to consider $\mathbb{R}/\mathbb{Z}$ or S^1 instead of $[0,1]$. In topology, we usually consider as equivalent two curves that can be continuously deformed from one to the other. Let us look at it in more details.

Free Homotopy Let $\gamma_1, \gamma_2 : [0,1] \to \Sigma$ be two curves on a surface Σ. We say that $H : [0,1] \times [0,1] \to \Sigma$ is a free homotopy between γ_1 and γ_2 if and only if H has the following properties: H is continuous, $H(0,\cdot) = \gamma_1(\cdot)$ and $H(1,\cdot) = \gamma_2(\cdot)$. If such an H exists we say that γ_1 and γ_2 are freely homotopic. It defines an equivalence relation between curves of Σ. A curve that is freely homotopic to a simple point is called a **contractible curve**. A surface where all curves are contractible is said **simply connected**. The sphere and the disk are the only two compact simply connected surfaces.

Fixed Point Homotopy In order to set an algebraic structure on the homotopy classes of Σ it is required to have a binary relation to define a group structure. This is correctly achieved by concatenation (noted $\circ$) after enforcing curves to have a common initial point. So let us fix a point $x_0 \in \Sigma$. We note C_{x_0} the set of all curves $\gamma(\cdot)$ such that $\gamma(0) = \gamma(1) = x_0$. a homotopy H is a free homotopy that verify: $\forall t \in [0,1]$, $H(t,0) = H(t,1) = x_0$. This defines an equivalence relation $\sim$ on C_{x_0}. We then define the first homotopy group or fundamental group of Σ as $\pi_1(\Sigma) = (C_{x_0}/\sim, \circ)$. The group $\pi_1(\Sigma)$ is independent of the choice of x_0 up to a (non-canonical) isomorphism. Two curves in C_{x_0} are freely homotopic if and only if their homotopy classes in $\pi_1(\Sigma)$ are conjugated. Homotopy of paths can be described analogously by fixing the two extreme points of the paths instead of just one point. However, there is no group structure.

Fundamental Groups of Orientable Surfaces Let $\Sigma_{g,b}$ be the surface of genus g with b boundaries. The most natural way to describe fundamental groups is to define them by generators and relations. We recall that a presentation of a group is of the following form: $G = \langle e_1, \cdots, e_n; R_1, \cdots, R_k \rangle$ where $e_1, \cdots, e_n$ are elements of G and $R_1, \cdots, R_k$ relations. If $b > 0$ then $\pi_1(\Sigma)$ is a free group generated by $2g + b - 1$ elements. If $b = 0$ then $\pi_1(\Sigma)$ is canonically represented by $\langle e_1, \cdots, e_{2g}; e_1 e_2 e_1^{-1} e_2^{-1} \cdots e_{2g-1} e_{2g} e_{2g-1}^{-1} e_{2g}^{-1} \rangle$. For instance $\pi_1(\Sigma_{1,0}) = \langle a, b; aba^{-1}b^{-1} \rangle$. Since $aba^{-1}b^{-1} = 1$ is equivalent to $ab = ba$, the fundamental group of a torus is abelian. This is a very strong property but it is not the case for surfaces of genus at least 2.

Configurations It is possible to have a more constrained notion of equivalence of curves than free homotopy. Let us define an **ambient isotopy** as a continuous map $I : [0,1] \times M \to M$ such that $I(0,\cdot)$ is the identity map and for all $s \in [0,1]$, $I(s,\cdot)$ is a homeomorphism of M. Then we say that γ_1 and γ_2 are equivalent with respect to I if $I(1, \gamma_1(\cdot)) = \gamma_2(\cdot)$. The equivalence classes under that relation are called **configurations**. A given homotopy class is the union of many different configurations. For instance a cycle is not in the same configuration as an "eight" even if the two are contractible and thus homotopic. In fact, isotopy do not allow to change the "arrangement" of a curve meaning that it cannot gain or lose crossings for instance.

2.1.3 Geometric Structures

Even for studying purely topological properties it may be useful to add a geometry on the considered surface. We thus define the minimum structures that we need.

Geodesic A **geodesic** on a surface is a curve that locally minimizes the distance. It is the kind of curves we imagine when we think of tight curves. In more details, a curve $\gamma : I \to \Sigma$ is a geodesic if and only if for all $x \in I$ there is a neighborhood U_x of x such that $\forall x_1, x_2 \in U_x$, $d(\gamma(x_1), \gamma(x_2)) = |x_2 - x_1|$ (where d is the distance function of Σ). Alternatively, in the Riemannian case we can imagine a geodesic as a curve with "constant" tangent vector. It is not absolutely straightforward since the tangent vectors lies on different tangent spaces. It is easily defined if Σ is embedded in $\mathbb{R}^n$. In that case, we look at the tangent vector of $\gamma(x)$ by considering γ as a curve in $\mathbb{R}^n$. Its acceleration can be decomposed into two parts, one in the tangent space of Σ at x and one that is normal to it. Then γ is a geodesic if this acceleration have a constant null tangent component. Note that a geodesic only locally minimizes the distance (see figure 21). In this figure we consider a torus represented in the plane with the corresponding geometry then the two green lines are geodesics. Only the first one realize the shortest path from A to B. So the second one does not minimize the distance between A and B but do it in a neighborhood of each point (such as the red zone on the right part of the figure).

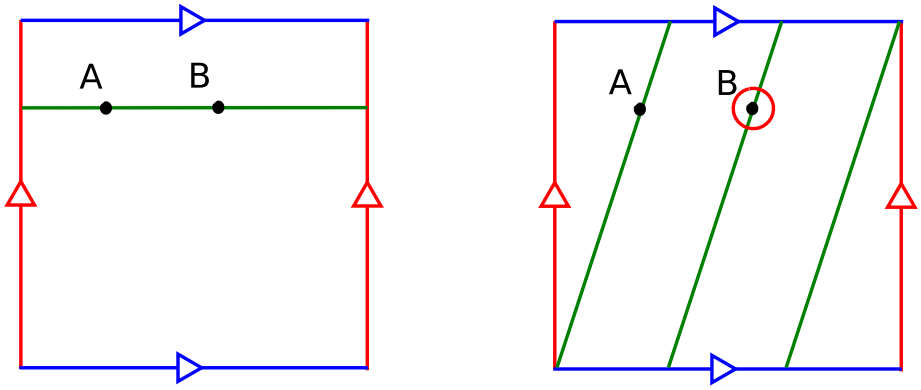

Figure 21: Two distinct geodesics between A and B

Curvature Along with the definition of distances on Σ, the notion of curvature appears. It describes how the surface looks like locally. Although its definition is quite natural for curves (the inverse of the curvature radius) it is less intuitive for

surfaces and some variants exists. Let us define the (Gaussian) **curvature** of Σ at a point x_0. We denote it by $\kappa(x_0)$. There are many equivalent ways to define the curvature, we choose the easiest one with respect to our setting. To do that we denote by $A(r)$ the area of the set of elements of Σ at distance shorter than r from x_0. The curvature is given by the following formula:

$$\kappa(x_0) = lim_{r\to 0}\ 12\frac{\pi r^2 - A(r)}{\pi r^4}$$

Let us look how the neighborhood of x_0 looks like with respect to its curvature. If x_0 has a neighborhood with the geometry of the plane then $A(r) = \pi r^2$ and so $\kappa(x_0) = 0$. On a sphere this area is shorter than πr^2 and so $\kappa(x_0) > 0$. A negative curvature can be found on the central point of a saddle (see figure 22).

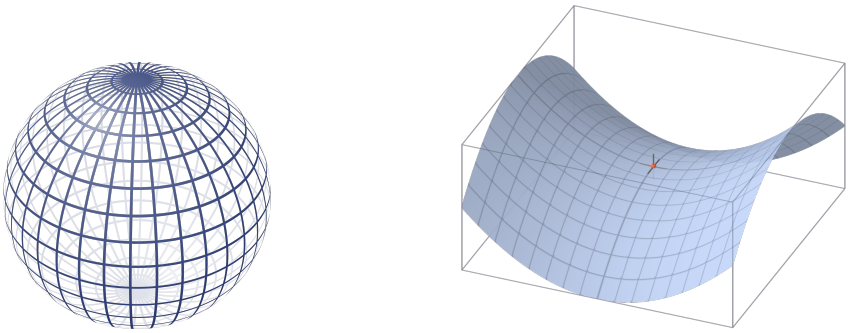

Figure 22: Curvature.

The curvature of the surface is related to its topology via the Gauss-Bonnet theorem:

Theorem 2.1.2 (Gauss-Bonnet)**.** *Let Σ be a surface without boundary, then*

$$\int_\Sigma \kappa dA = 2\pi \cdot \chi(\Sigma)$$

where $\chi(\Sigma)$ is the Euler characteristic of Σ.

The surfaces with constant curvature metrics are interesting. By the Gauss-Bonnet theorem we see that the corresponding constant have the same sign as $\chi(\Sigma)$. It leads to positive curvature for a genus 0 surface (realized by the two dimensional sphere S^2 with the induced metric of $\mathbb{R}^3$). A torus may have a flat geometry (the one

of a Euclidean square with identified opposite sides). All surfaces of genus at least 2 can only have constant negative curvature. In that case we say that the surface has a **hyperbolic metric**. It is difficult to illustrate Hyperbolic surfaces since Hilbert's theorem states that no hyperbolic surface can be immersed isometrically into $\mathbb{R}^3$ (with a C^2 immersion).

Hyperbolic Metrics Surfaces with constant negative curvature (generally fixed at -1) have very interesting properties and are intensely studied. We give the most important point for our purpose: the unicity of geodesics. Indeed given a path P from A to B in Σ there is only one geodesic curve homotopic to P that goes from A to B. The properties of hyperbolic surfaces are significantly different from euclidean ones. For instance consider the case of hyperbolic triangles (triangles whose three sides are geodesics). The sum of the angles of a triangle T is strictly less than π. We say that the complement to π of this sum is the **angle defect** d of T. The area of T is proportional to d, so that $A(T) = d$ with normalized curvature. So the maximal area of a hyperbolic triangle is realized by **ideal triangles** with all angles null and is π.

2.1.4 Covering Spaces

Intuitively a covering of a surface is another surface that wraps around the first one possibly infinitely many times. For instance it is possible to wrap an infinite band around a cylinder. More precisely a surface C is a **covering space** of Σ if there is a surjective map $\rho : C \to M$ such that for every $x \in \Sigma$ there is a neighborhood U_x of x such that $\rho^{-1}(U_x)$ is an union of disjoint open sets U_i in C each of which induces a homeomorphism onto U_x. The U_i are called the **covering sheets** of U_x. Figure 23 illustrates the notion in the case of the cylinder and the infinite band. The first part of the figure shows the idea of the wrapping of the band. Then let us assume that the red circle is U_x. That part of the cylinder has infinitely many different pre-images that we can see on the band. Let γ be the green curve on the cylinder. We call the green line on the band a **lift** of γ.

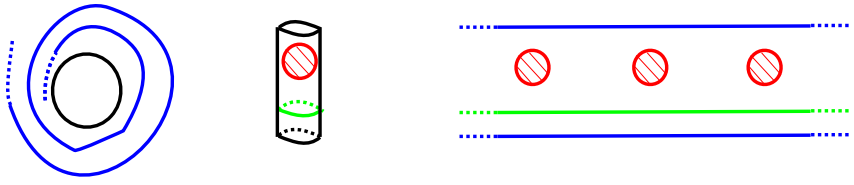

Figure 23: A cylinder covered by an infinite band.

Universal Cover If C is **simply connected** it is called a **universal cover** of Σ. A surface has a unique universal cover C^∞ up to homeomorphism. C^∞ covers all coverings of Σ so that we have the following diagram for any covering C of Σ: $C^\infty \to C \to \Sigma$. If Σ has no boundary and is different from the sphere then its universal cover is homeomorphic to the plane. If it has boundary then its universal cover is an infinite thick tree. It is extremely interesting to consider the **automorphisms** of the universal cover. Here an automorphism g is defined as follows: it is a homeomorphism of C^∞ such that $\rho \circ g = \rho$ where ρ is the covering map associated to C^∞. Those automorphisms form a group $\mathcal{G} = (G, \circ)$ such that $\mathcal{G} \simeq \pi_1(\Sigma)$. In addition, $C^\infty/\mathcal{G} \simeq \Sigma$ (note that it is not true for any covering). It follows that every intermediate covering of Σ is the quotient of the universal cover by a subgroup of the automorphisms group of the surface. It also has interesting properties for lifts of curves.

Lemma 2.1.3. *Let γ be a curve of Σ. The lifts of γ in C^∞ are closed if and only if γ is contractible on Σ. In addition, $\pi(\Sigma)$ acts on the set of lifts of γ via the automorphisms of C^∞.*

Note that the point of view can be reversed. The surfaces without boundary can be constructed directly from a quotient of the plane by particular discrete groups of homeomorphisms of the plane. The corresponding constructions appear naturally with the support of geometry. Let us look at the useful models.

Geometric Models for Universal Covers The canonical metric on $\mathbb{R}^2$ has constant null curvature. It is a good model for the universal cover of the characteristic 0 surface, the torus. Let us consider $\mathcal{G}$ the group of the translations by integral vectors (a, b). The classical representation of the torus as $[0, 1]^2$ with opposite sides identified, correspond to the quotient of the plane by $\mathcal{G}$ (see figure 24). The two corresponding identification correspond to the translations of vector $(1, 0)$ and $(0, 1)$ that generates $\mathcal{G}$. That construction is a natural way to define a torus with a constant null curvature metric.

The constant curvature surfaces of higher genera cannot be obtained in a similar way using a quotient of the Euclidean plane by a group of Euclidean isometries. Instead we can use the hyperbolic model. A model of a surface homeomorphic to the plane with a constant negative curvature is Poincaré disk, which is the open unit disk with a particular metric. We do not detail the construction of that metric but focus on the properties that are interesting for us. The main question is: What do the maximal geodesics look like? The answer is that those geodesics are the intersections between the unit disk and circles that are orthogonal to the unit circle.

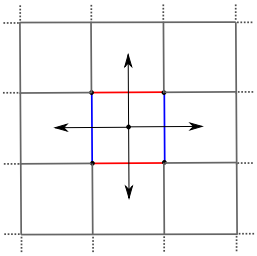

Figure 24: The universal cover of the torus.

Considering Euclidean lines as circles with infinite radius, diameters also appear as maximal geodesics. We can make an analogous construction as the one of the null curvature torus for hyperbolic surfaces by quotienting the Poincaré disk by groups of hyperbolic translations. The polygon obtained as a fundamental domain of the quotient is no longer a square like in the torus case but has $2g$ sides, see figure 25 for an illustration on the genus 2 surface.

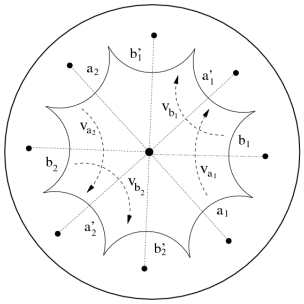

Figure 25: The universal cover of the hyperbolic double torus and a fundamental domain with the corresponding translations.

Hyperbolic Translations In order to facilitate the study of the state of the art of Chapter 3, we give more details about hyperbolic translations. We are looking for isometries of the Poincaré disk. The good candidates are the **Möbius transformations**. They are the transformations of $\mathbb{C}$ of the form $z \to \frac{az+b}{cz+d}$ with $a, b, c, d \in \mathbb{C}$. The isometries of the Poincaré disk are the Möbius transformations that send the unit disk into itself. Note that those isometries form a group isomorphic to $PSL_2(\mathbb{R})$. Möbius transformations are characterized by their 2 fixed points (possibly at infinity). A **hyperbolic translation** is an isometry of the Poincaré disk whose two fixed points are on the unit circle. Its axis is the unique hyperbolic line between those two points.

2.2 Discrete Surfaces

2.2.1 Polygons

Definition We first present the very natural construction of surfaces as assemblies of polygons. A polygon is a plane figure bounded by an alternate sequence of points and segments. Alternatively it is possible to think of circular polygons whose boundary is a circle. The model is pointless here since we want to define topological surfaces but becomes significant for geometrical constructions. In full generality a polygon may have one or two sides as in figure 26. The corresponding polygons are called monogons and bigons.

Figure 26: A monogon and a bigon.

All polygons are homeomorphic to disks. The points are named **vertices** and the lines **edges**. It is possible to describe a surface by gluing together pairs of oriented edges of a set of polygons. To identify two oriented edges $\vec{e}$ and $\vec{e'}$ we associate parametrizations of the corresponding lines. Note that in a discrete setting we just need to know the correspondence of edges and the fact that the interior of each polygon is an open disk to understand the topology of the surface. The edges that are glued are interior edges of the final surface and the other remain boundary edges. To obtain a 2-manifold, all the edges must appear in at most one identified pair. However, a vertex can appear in any number of associations. So a vertex of the final surface corresponds to an arbitrary set of vertices of the initial polygons.

Torus Case We can construct a torus in many ways, we focus on two different constructions (see figure 27). A 4-gon with boundary $(v_1, e_1, v_2, e_2, v_3, e_3, v_4, e_4)$ with the following identification $\overrightarrow{v_1v_2} \equiv \overrightarrow{v_4v_3}$ and $\overrightarrow{v_2v_3} \equiv \overrightarrow{v_1v_4}$. The first identification gives $v_1 \equiv v_4$ and $v_2 \equiv v_3$. The other $v_2 \equiv v_1$ and $v_3 \equiv v_4$ leading to a unique vertex on the final torus. We can also construct a torus starting from a hexagon with the identification depicted in figure 27. In that case we obtain two vertices. So the "6-gon torus" has one additional vertex and one additional edge compared to the 4-gon version. In fact, it satisfies to a topological invariant called the Euler characteristic.

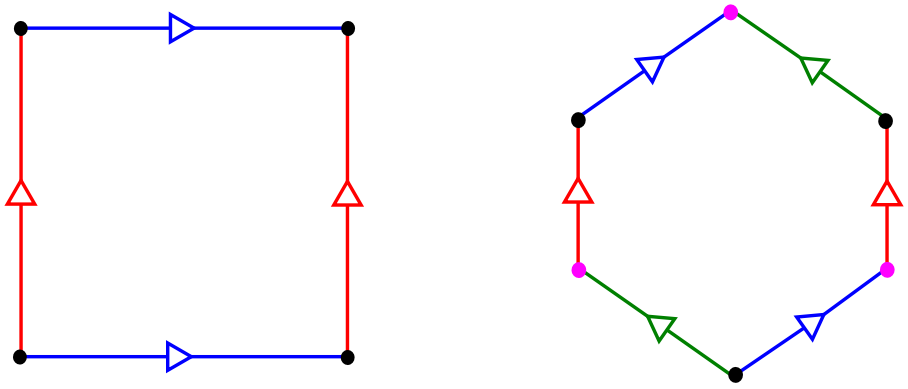

Figure 27: Two different tori

Theorem 2.2.1 (Euler characteristic)**.** *Let $\Sigma_{g,b}$ be a surface of genus g with b boundaries. Let V and E be the numbers of vertices and edges after identifications. Let F be the number of faces (or the number of polygons involved in the construction of $\Sigma_{g,b}$). Then, we have:*

$$\chi(\Sigma_{g,b}) = V - E + F = 2 - 2g - b$$

In the previous examples we obtain $1 - 2 + 1 = 0$ for the 4-gon, $2 - 3 + 1 = 0$ for the 6-gon and it corresponds to $\chi(\Sigma_{1,0}) = 2 - 2 * 1 - 0 = 0$. The proof of the theorem is made by applying transformations to any surface in order to transform it into a canonical one. For instance on our example, if we contract the green edge of the 6-gon we obtain exactly the 4-gon representation. In general, the sketch is the following (assuming for simplicity that the surface has no boundary). First contract the edges of a spanning tree. Edge contraction can be shown to maintain the relation. Then all the edges that bounds two different faces are deleted. At this point, we have constructed a representation of the initial surface as a sole polygon with identifications. Then It remains to show that we can always go from any one polygon construction to a canonical one. Note that this can be made algorithmically [Lazarus et al., 2001] and is often useful as a subroutine for topological algorithms.

Classical Constructions We just see that any surface constructed by polygons can be transformed into a surface constructed from only one polygon. Reciprocally, every surface can be constructed that way. For every orientable surface without boundary we thus construct a canonical fundamental polygon that corresponds to the surface by using a $2g$-gon where g is the genus of the surface. Then the identification can be made such that we can "read" the canonical relation of the fun-

damental group of the surface around the polygon (see figure 28 where we read $aba^{-1}b^{-1}cdc^{-1}d^{-1}$). Surfaces that have boundaries or that are non-orientable can have canonical representation in the same idea.

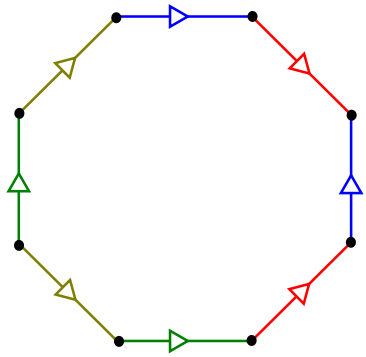

Figure 28: Canonical polygon

Another special case occurs when all the polygons are triangles. In that case we say that we have a **general triangulation**. It is important to remark that there are many different possible definitions for triangulations. This one is the most general one but is not the one used in this thesis.

2.2.2 Embedded Graphs

It is very natural to try to draw graphs on surfaces. In the previous construction using polygons we see that graphs appear naturally as the identified edges of the polygons. In that subsection we start from a continuous surface and we define what drawing a graph on it means.

Graph Let V be a set of elements called **vertices**. A graph G is a couple (V, E) where E is a collection of elements of $V \times V$. The elements of E are called the **edges** of G. Since E is a collection, a given pair can appears many times. In that case we say that G has **multiple edges**. E can also have pairs (v, v) that we call **loop edges**. If G has no multiple nor loop edges then it is a **simple graph**. A lot of practical problems can be expressed in a graph setting. Therefore graphs are intensively studied.

Embedding Let Σ be a surface. An **embedding** of G on Σ is a drawing of the graph on the surface. It is natural to consider that vertices have to be represented

as points and edges as curves between those points. In details, an edge $e = (v_0, v_1)$ is embedded as a curve $\gamma_e : [0, 1] \to \Sigma$ such that $\gamma_e(0) = v_0$ and $\gamma_e(1) = v_1$. We also require all $\gamma_e|_{]0,1[}$ to be disjoint and to not intersect embedded vertices. A graph that can be embedded on the sphere (or the plane) is called a **planar graph**. The image of G corresponding to the embedding will often be considered as G on Σ (which is a language abuse). Note that this definition becomes quite difficult to handle if one wants to use it without avoiding any details.

Faces The faces of an embedding are defined as the connected components of $\Sigma \backslash G$. To a face can be attached a cycle in the graph called its **facial walk**. Reciprocally it is possible to describe an embedding by giving the paths in G corresponding to the faces. If each face corresponds to an open disk then we have a description of the surface with polygons. In that case we say that the embedding is **cellular** and the Euler characteristic is correctly defined by the formula $V - E + F$. In full generality faces can be other open surfaces. Those surfaces can be exactly defined by their orientability and two numbers g and b corresponding to their genus and number of boundaries as usual. That kind of definition of surfaces is not very interesting at first sight but it must be considered since it can be obtained after performing transformations on a graph that was initially cellularly embedded (such as edge deletion for instance).

2.2.3 Combinatorial Maps

Combinatorial maps give a good way to speak of graph embedding in practice. Graphs embeddings remain interesting since they are useful to design geometric algorithms. However, when the geometry of the defined surface is not fixed there is no point to use something else than combinatorial maps. For more details on combinatorial maps see [Mohar and Thomassen, 2001; Gross and Tucker, 1987].

Definition We want to have a concise description of topological surfaces. So a combinatorial map M is defined as a base graph G with an added information that describes the local behavior of the embedding around the vertices of G. For a vertex $v \in G$ it is well done by a circular ordering of all the edges incident to v. The set of all that orderings are called the **rotation scheme** of M. This description leads to efficient data-structures that are described below. Those data structure will be used for the input of all the algorithms described in this thesis.

Half-Edges Here we store combinatorial maps as a set of **half-edges**. Each half-edge corresponds to an oriented edge on the map. There are two half-edges per edge. A half-edge needs to have the information of which half-edge is the other half-edge of its support edge. We call σ_0 the correspondence (see figure 29). This edge also needs to know the next half-edge in the rotation scheme around the origin vertex of the half-edge. This correspondence is denoted by σ_1. σ_0 and σ_1 can be seen as permutations on the half-edges and are sufficient to construct the whole map. The vertices are the orbits of the action of $< \sigma_1 >$ on the set of all half-edges (recall: $< \sigma_1 >$ is the group generated by σ_1 with the natural relation $\circ$, since σ_1 is define as a function on the set of half-edges the action is clearly defined). The faces are the orbits of the action of $< \sigma_0 \circ \sigma_1 >$. Note that the edges are also defined that way using $< \sigma_0 >$. Here all the orbits have size 2 since σ_0 is an involution.

Flag Representation The same idea can be used in a more general way. We define a **flag** as a triple (v, e, f) where v is a vertex, e an edge and f a face all incident to each other. Then we define on the set of flags three involutions α_i for $i = 1,\ 2$ or 3. Each α_i acts on a triple by changing its cell of dimension i (see figure 29). As in the previous case we can define all entities as orbits. Vertices are given by $< \alpha_1, \alpha_2 >$, edges by $< \alpha_0, \alpha_2 >$ and faces by $< \alpha_0, \alpha_1 >$. An edge is always a set of four flags but vertices and faces may have any cardinality. Remark that by globally identify the pairs of flags linked by α_2 we obtain the half-edge representation. However only the flag representation allows to store non orientable maps.

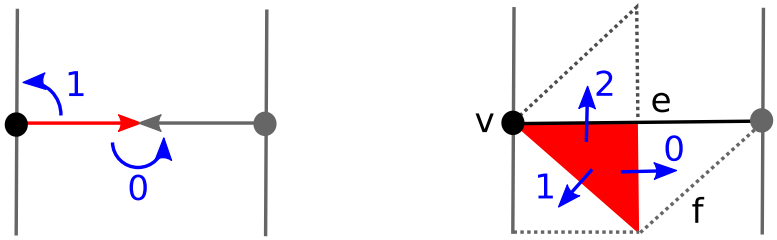

Figure 29: Half-edge and flag representations

Duality There is a natural notion of duality for combinatorial maps. In particular, the flag representation shows that the vertices and the faces have symmetric roles. Thus the dual of a map M noted M^* is given by taking the set of the faces of M as the vertices of M^*. Then two vertices of M^* are linked by an edge for each edge of M that separates the two corresponding faces in M. Finally the rotation scheme

of a vertex of M^* is exactly equivalent to the facial walk of the corresponding face in M. The key point is that the faces of M^* corresponds to the vertices of M. So the bidual of M is M. Note that if M has a flag representation $(\alpha_0, \alpha_1, \alpha_2)$ then M^* admits $(\alpha_2, \alpha_1, \alpha_0)$ as flag representation.

Tree-Cotree Decomposition By considering both the map M and its dual it is possible to give a nice partition of the edges. We begin by taking a spanning tree T of the primal map. It uses one less edges than the number of vertices of the primal graph. We remove the edges of T from the dual graph. The tree that we removed is not separating, so the remaining dual graph is connected. It is thus possible to take a spanning tree C of the remaining dual. We call X the set of the remaining edges (X can be empty). The triple (T, C, X) is called a **tree-cotree decomposition** of M. Let us look the cardinality of X. Let v, e and f be the number of vertices, edges and faces of M. T contains $v-1$ edges and C, $f-1$. Then, $e = v+f-2+2g = (v-1)+(f-1)+2g$. We conclude that X has $2g$ edges. In particular, X is empty if and only if M has 0 genus (see figure 30).

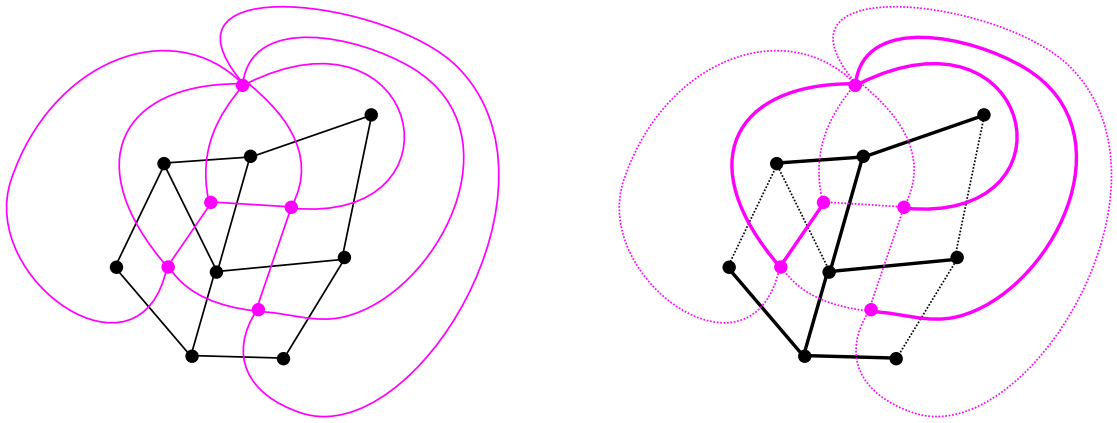

Figure 30: A tree-cotree decomposition of a plane map.

2.2.4 Simplicial Complexes

We quickly define simplicial complexes. It is a particular discrete construction that is useful in a lot of situations. In particular it is very easily extended to higher dimension and it allows to give a simple definition for homology.

Algebraic Description Let us consider a non-empty set V of elements called vertices. Then a **simplicial complex** is a set of non-empty subsets of V containing all the singletons and with the following hereditary property: if C is in S then all non-empty subsets of C are in S. We call a n-cell of S a $C \in S$ that contains $n+1$ elements of V. The dimension of a simplicial complex is the maximum n such that S has a n-cell. This structure can be interpreted as a topological object, we describe the case of 2-complexes.

Topology A simplicial complex of dimension 2 has three kinds of elements different from the empty set. Singletons are considered as vertices. 1-cells are the edges of the complex and the 2-cells its faces. It can easily be interpreted as a surface defined by polygons. Here the polygons are triangles since faces contain exactly 3 vertices. By the definition of simplicial complexes the boundary of the faces are edges and vertices that belong to S. Just as with polygons, not all simplicial complexes of dimension 2 correspond to a 2-manifold. For instance, it is necessary that at most 2 faces share a given edge.

Simplicial Triangulations As all faces of a surface defined as a simplicial complex are triangles it is natural to think of it as triangulations. However, all the general triangulations cannot be described that way. Simplicial complexes do not allow more than one k-cell defined by a given subset of vertices. So a simplicial triangulation cannot have multiple edges. It cannot have loop edges neither. Finally every general triangulation whose induced embedded graph is simple is simplicial except for the triangle on the sphere. In that particular case there are two triangles with the same vertices.

2.2.5 Homology of Surfaces

Homology is an alternative to homotopy. It gives another way to describe the equivalence of two curves. For convenience we describe homology on simplicial complexes. The same definition works in the other settings we described above with slight modifications.

k-Chains A path in a simplicial complex has to be composed of 1-cells since it has dimension 1. The walks in the 1-skeleton of the complex (the subcomplex obtain from the initial one by only keeping cells of dimension less or equal to 1) correctly fill that intuition but we are looking for a slightly more algebraic setting. Generally, we would like to define an abelian group structure on the set of k-cells. The case $k = 1$

corresponding to the ideas of paths. The k-cells must be considered as oriented so if e is an edge with a given orientation then $-e$ is the opposite edge. Such a group needs to be able to associate something to $c+c'$ for any k-cells c and c'. We consider all formal sums of k-cells with $\mathbb{Z}$ coefficients that we call k-chains. Let C_k be the set of k-chains. We define the boundary operator $\partial_k : C_k \to C_{k-1}$ as a homomorphism that associates to each oriented k-cell the alternate sum of its boundary $(k-1)$-cells. With those definitions, it can be understood that a 1-chain is a generalized path.

First Homology Group Let us consider $Ker(\partial_1)$. It is composed of 1-chains with null boundary. Paths without boundary are just closed paths we thus call the elements of $Ker(\partial_1)$ **cycles**. More generally, a cycle is a formal sum of closed paths. Then $Im(\partial_2)$ is the set of cycles that bound a union of faces. Typically a cycle that bounds a face has to be equivalent to the trivial cycle. The first homology group is defined as $H_1 = Ker(\partial_1)/Im(\partial_2)$. It is the set of cycles quotiented by the subset of the cycles generated by the face cycles. H_i can be defined in the same way. H_1 is related to the fundamental group quite directly. The first homology group is the abelianization of the first homotopy group. For instance $\pi_1(S_1) \simeq H_1(S_1)$ since $\pi_1(S_1)$ is abelian but $\pi_1(S_2) \neq H_1(S_2)$ where S_i is the orientable surface of genus g. In general, we have $H_1(S_g) \simeq \mathbb{Z}^{2g}$.

2.3 Constructions on Discrete Surfaces

2.3.1 Curves

Edge-Paths The natural way to obtain curves on a combinatorial map is to make a walk on the edges of the map. It leads to a discrete analogue for curves, let us describe it in details. An **edge-path** is defined as a sequence of oriented edges such that each edge end coincides with the next edge origin. If we consider embedded graphs then we can also directly interpret edge paths as curves on the underlying surface Σ. Since faces are open disks it is possible to show that edge paths faithfully represent the curves on a surface. More formally, we have:

Proposition 2.3.1. *Let Σ be a surface, γ a closed curve on Σ and M a combinatorial map homeomorphic to Σ. There is an edge path P in M such that the curve corresponding to P in Σ is freely homotopic to γ.*

The same kind of results holds for paths on Σ.

Immersions When an edge path uses twice the same edge its direct image has infinitely double points. In that case we say that the curve is not in general position. In the continuous setting it is avoided by adding a little perturbation to the curve. In the discrete world, we manage to replace perturbations by adding a structure to the edge paths. We thus obtain general position realizations of our edge paths. More precisely, we call each edge of an edge path an **occurrence**. For each edge of the initial map used by the path we add a transverse ordering on the multiple occurrences going through it. Then a continuous realization of the total structure consists in attaching thin bands replacing the edges of the map and small disks for the vertices. The transverse orderings of occurrences allow to construct a continuous curve in general position (see figure 31).

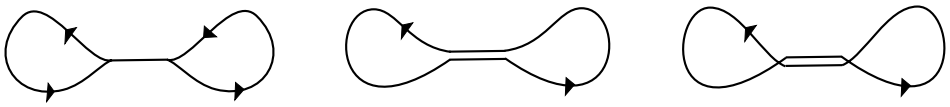

Figure 31: An ambiguous representation and its two possible immersions.

Combinatorial Crossings For continuous curves a crossings between two curves or a self-crossing of a curve is naturally defined as a double point. In our framework, we want to give a combinatorial definition of a crossing that corresponds to the continuous idea and that is not ambiguous. With our definition of immersion as edge-paths with additional transverse order, it can be done quite naturally. The idea is that a continuous realization of an immersion should be the following. We thicken the graph of the combinatorial map by replacing vertices by small disk and edges by small cylinders (see figure 32). Now it is possible to put the occurrences inside the cylinders in an order corresponding to the local transverse ordering. It implies that no double point can appear in the cylinders. In the disk we decide to just add a segment to join the occurrences that are adjacent in the edge path. Then, for each pair of occurrences linked in the same vertex there is a double point if and only if the four border points of the occurrences alternate (as in figure 32 on the right). This can be described purely combinatorially since the order in which the occurrences appears around a vertex is given by the rotation scheme of the vertex and the transverse orderings.

Cross-Metric Surfaces Alternatively, a dual approach is possible. We start with a combinatorial map but instead of considering curves as edge paths we consider curves in general position with respect to the map. It means that we only consider

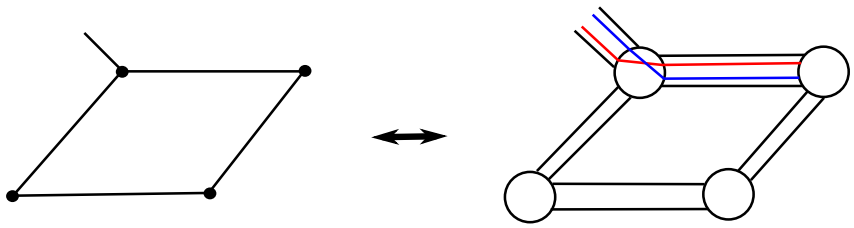

Figure 32: A graph and its corresponding thicken construction.

curves that intersects the edges of the map transversely and do not intersect the vertices. In practice, to store such a curve it is needed to store the intersections of the curve with the map and the relative order of the crossings along a fixed edge. Then the length of a curve is given by its number of intersections with the map that is its length as a walk in the dual map.

2.3.2 Diagrams and Coverings

Diagrams and coverings are maps along with a projection onto another map. Coverings naturally appear as continuous constructions and are easily discretized. Diagrams are less constrained constructions that are interesting for our purpose. Let us give the detailed definitions.

Coverings The translation from the continuous notion is quite straightforward. A combinatorial map C is a covering of a combinatorial map M if and only if there exists a covering function $\delta : C \to M$. Such a function has to send vertices to vertices, edges to edges and faces to faces. It also has to be locally bijective meaning that each rotation around a vertex and each face walk are preserved by δ. Note that we need to allow infinite maps for C because we often want to work in the universal cover. It is quite strange since we claimed that combinatorial maps are used in order to make direct implementations. In practice this will be handled by constructing just a finite part of the universal cover.

Diagrams The definition is very similar to the one of coverings. A map D is a diagram on M if there is a diagram function: $\partial : D \to M$. The only difference between a diagram map and a covering one is that no bijection is required. It implies for instance that the neighborhood of a vertex v in D can be sent many times on the neighborhood of $\partial(v)$. The question is: why is it interesting? The main answer for our purpose is the following property illustrated by figure 33. A diagram is said

reduced if it does not send two adjacent faces to the same face of M with opposite orientations.

Proposition 2.3.2 ([Van Kampen, 1933]). *A cycle γ in M is contractible if and only if there is a reduced disk diagram Δ_γ whose boundary cycle is sent to γ by the diagram map.*

Figure 33: A contractible cycle on the torus and the corresponding diagram.

It substantiates the natural idea that a contractible cycle morally bounds a disk. Unfortunately the disk is not necessarily embedded and the diagram is the best we can have in the general case. In the same way we can characterize pairs of homotopic cycles by trying to make them bound a cylinder. Once again the diagrams give a good translation of the intuition.

Proposition 2.3.3 ([Lyndon, 1966]). *Two cycles cycle γ_1 and γ_2 in Σ are freely homotopic if and only if there is a reduced annular diagram $\Delta_{\gamma_1,\gamma_2}$ whose boundary cycles are sent to γ_1 and γ_2 by the diagram function.*

2.3.3 Discrete Geometry

Construction We want to put a particular geometry on a combinatorial map. First the length of the curves is defined as the length of the corresponding edge-path (eventually with associated weights). For triangulated maps a natural idea is to set a fixed geometry for the triangles and then glue them as smoothly as possible. Models for triangles can be Euclidean models or hyperbolic ones. Again we can let some freedom on the choice of the angles or ask for equilateral Euclidean triangles or ideal hyperbolic ones. In the general case, many construction are possible, we decide to use Euclidean triangles. We assume that each face subdivided into flat triangles as in figure 34. Now everything is flat except the neighborhood of each vertex and the neighborhood of the virtual vertex that we have added in each face. We need to define the angles of the corners of the faces. Since no natural construction appears let us assume that we set an arbitrary θ_a value to every angle a. Since it is purely

combinatorial and in order to simplify formulae we normalize the angles such that a full tour is set to 1 (a direct approach should have lead to 2π). So for each flat triangle the sum of its three angles is $\frac{1}{2}$.

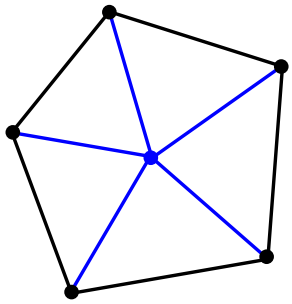

Figure 34: The geometry of the faces.

Discrete Curvature We are now able to set a notion of curvature. Let us consider the formula we used to define the curvature. Let r be a positive number and v a vertex of the underlying combinatorial map M. In each triangle corner at v with angle θ_a the part of area that is added to $A(r)$ is $\pi r^2\theta_a$ by the normalization convention. The area around v is the sum on the corners around it. Let $N(v)$ be the set of corners around v. The curvature should be $\kappa(v) = 12\frac{\pi r^2 - \pi r^2(\sum_{a\in N(v)}\theta_a)}{\pi r^4} = \frac{12}{r^4}(1 - \sum_{a\in N(v)}\theta_a)$. The factor $\frac{12}{r^4}$ shows that the vertices are singularities in the continuous corresponding model. After another normalization we define $\kappa(v) = 1 - \sum_{a\in N(v)}\theta_a$. If the sum is less than 1 we obtain positive curvature like on a sphere and if it is more than 1 then we obtain a negative curvature. In the case where the surface has boundaries we set at a flat angle $(\frac{1}{2})$ each boundary angles of each vertex. Now let us look at the face vertices. The sum of the angles of a triangle is $\frac{1}{2}$ so if v_f is such a vertex its curvature should be: $\kappa(f) = 1 - \sum_{a\in N(v_f)}\theta_a = 1 - \sum_{a\in N(v_f)}(\frac{1-\theta_{a_1}-\theta_{a_2}}{2}) = 1 - \frac{d(f)}{2} + \sum_{a\in f}\theta_a$ where $d(f)$ is the degree of the face f. This definition works well since it leads to the following theorem.

Theorem 2.3.4 (Discrete Gauss-Bonnet)**.** *Let M be a combinatorial map. In the setting described above the following formula holds:*

$$\chi(M) = \sum_{v\in V(M)}\kappa(v) + \sum_{f\in F(M)}\kappa(f)$$

Proof. We first make the calculation for surfaces without boundary.
$\sum_{v\in V(M)}\kappa(v) + \sum_{f\in F(M)}\kappa(f)$

$= \sum_{v \in V(M)} (1 - \sum_{a \in N(v)} \theta_a) + \sum_{f \in F(M)} (1 - \frac{d(f)}{2} + \sum_{a \in f} \theta_a)$
$= |V(M)| - \sum_a \theta_a + |F(M)| - \frac{2*|E(M)|}{2} + \sum_a \theta_a$
$= |V(M)| - |E(M)| + |F(M)| = \chi(M)$
Consider the case of surfaces with boundary. Let B be the number of boundary angles (or equivalently the total length of the boundaries).
$\sum_{v \in V(M)} \kappa(v) + \sum_{f \in F(M)} \kappa(f)$
$= \sum_{v \in V(M)} (1 - \sum_{a \in N(v)} \theta_a) - \frac{1}{2}B + \sum_{f \in F(M)} (1 - \frac{d(f)}{2} + \sum_{a \in f} \theta_a)$
$= |V(M)| - \sum_a \theta_a + |F(M)| - \frac{1}{2}B - \frac{2*|E(M)| - B}{2} + \sum_a \theta_a$
$= |V(M)| - |E(M)| + |F(M)| = \chi(M)$ □

2.3.4 An Interesting Example: the Test of Homotopy

We now have set all the general constructions we need. Let us look at how it behaves in practice. The problem we are interested is the following: For two given curves, how to decide if they are homotopic or not? This is extremely classical, let us look at what can be done about this. With the continuous definition to show that two curves are homotopic it is necessary to give a (free) homotopy between them. This is not easily translated into a discrete setting. The following construction appears to be the good one.

Combinatorial Homotopies An **elementary homotopy** on a combinatorial curve c consists in adding or removing a **spur**(two consecutive identical edges with opposite orientations), or replacing in c a possibly empty part of a facial walk by its complementary part. For a free elementary homotopy we can also apply a circular shift to the indices of the closed curve c. The equivalence relation generated by (free) homotopies is called combinatorial (free) homotopy.

An **elementary move** of a combinatorial immersion consists either in an elementary homotopy or in an **adjacent transposition**, i.e. in exchanging the left-to-right order of two occurrences in a same arc where one occurrence is next to the right of the other. We further require before performing an elementary homotopy that the immersion is in **good position**. This means, if the elementary homotopy applies to a nonempty part u of a facial walk of some face, that each arc occurrence in u should be the rightmost element of its arc, i.e. the most interior to the face. When removing a spur, we just require that the two arc occurrences to be removed are adjacent in left-to-right order. See Figure 35. The elementary homotopy does not modify the order of the remaining arc occurrences. When inserting a spur we make the inserted arc occurrences adjacent and when inserting part of a facial walk we

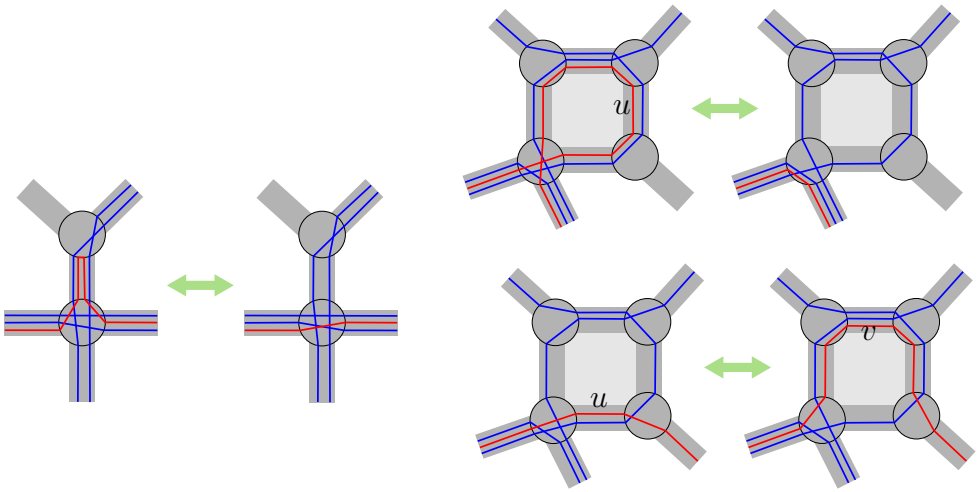

Figure 35: Left, the removal of a spur in good position. Right, two elementary homotopies replacing a subpath u of a facial walk by the complementary subpath v. In the upper right, v is empty. The immersions being in good positions, crossings may only appear or disappear by pairs.

insert each arc occurrence after the rightmost element of its arc. Remark that by a sequence of adjacent transpositions we can always enforce an immersion to be in good position.

Algebraic Solution However such a direct approach is not very efficient. In an algorithmic point of view the most natural idea is to consider curves as words in a given basis (if we are looking for free homotopy it may be needed to add a spur to the curves to make them reach the base point). First we look at the elementary problem of deciding if a given word is equivalent to the empty word. It is sufficient to solve that problem since to know if w and w' are words that correspond to homotopic curves it is sufficient to be able to decide if $w \cdot w'^{-1}$ is the empty word. It first can be frightening since this problem is known to be undecidable in the general case. Fortunately it is not the case for surfaces. An important point is to be able to obtain a decomposition of the input curves in a given basis. This can be done the following way. Starting from a combinatorial map it is always possible to contract the edges of a spanning tree to obtain a map with one vertex. Then edges that bound two different faces can be deleted to finally obtain a set of $2g$ loops generating the fundamental group of the underlying surface. Such a map is called a **reduced map**. By this construction, a curve on the initial map can be traduced into a word on the

generators corresponding to the $2g$ loops.

Universal Cover and Dehn's Algorithm The universal cover gives a nice characterization of contractible cycles. A cycle is contractible if and only if its pre-image on the universal cover is a union of closed curves. After going to a bouquet of circles Dehn has been able to show the following lemma.

Lemma 2.3.5 ([Dehn, 1912]). *Let Σ be a surface. Consider a presentation of $\pi_1(\Sigma)$ with $2g$ generators and 1 relation $r = 1$. Let w be a reduced word in the generators of $\pi_1(\Sigma)$ and their inverses that is equivalent to the empty word. Then w contains more than half a copy of r as a subword.*

This leads to a gready algorithm that runs in time $O(f(g)(n+l))$ where g is the genus of the map, n is the complexity of the initial map and l is the length of the curves. This can be strongly improved by considering geometric matters. We then describe the construction of the systems of quads that are combinatorial maps with extremely interesting geometric properties.

2.3.5 System of Quads

Definition Let us first describe what we call a system of quads. Let Σ be the surface of genus g. We want to start from a reduced map with $2g$ loops. Alternatively we can start with a $4g$-gon whose sides are pairwise identified leading to a genus g surface. It is equivalent to give a representation of the $\pi(\Sigma)$ and choose the identification such that we read the relation around the polygon. For instance if $g = 2$, the following different words can be read in clockwise order around an octogon to obtain Σ: $abcda^{-1}b^{-1}c^{-1}d^{-1}$, $aba^{-1}b^{-1}cdc^{-1}d^{-1}$ etc. All the different representations are allowed, systems of quads do not require to be based on a canonical one. Then we add a vertex in the sole face of the map and link it to all the copies of the initial vertex in the polygonal representation. Finally, we remove the initial edges. The resulting graph is called the **radial graph** of the initial reduced map. Clearly all faces are quadrangles (see figure 36). Just to get a better understanding of what happen let us check Euler characteristic. Since the polygon has $4g$ vertices the system of quads has $4g$ edges. Then we have $2g$ faces because all faces have four sides. It gives $2 - 4g + 2g = 2 - 2g = \chi(\Sigma)$.

Transformation into a System of Quads A first thing to remark is that it is easy to transform any combinatorial map into a system of quads. We just apply the procedure described in the previous subsection (contract a spanning tree and

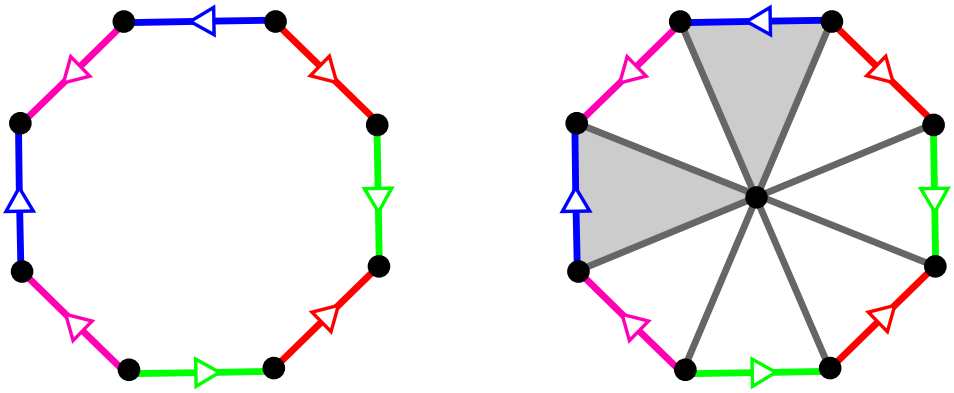

Figure 36: A system of quads of genus 2

remove edges) to obtain a reduced map. Then the system of quads is obtained by constructing the corresponding radial graph. This procedure is easily done in linear time with respect to the size of the initial map. Here an important point has to be developed: what happen to curves in that process? The operations needed to obtain the system of quads do not allow to preserve the edge paths. What we want for a curve on the initial map is to be transformed along that process in order to have a final edge path homotopic to the initial one. While contracting edges it is easy to maintain any given curve without increasing its size. The removing part is more tricky. Let us fix the final reduced map and consider the corresponding polygonal representation. If the curve go through an edge that it removed then this edge can be represented as a cord of the polygon. it is necessary to replace that edge by a path on the boundary of the polygon. Since we can choose the shorter one it may multiply the length of the curve by g. Better can be obtained if the curve is not updated directly but when the second vertex is added. At this moment the curve can be stored as a edge path of edges of the polygon and cords. Any pair of points of the polygon can be joined by a length 2 path in the radial graph. It leads to a length that can be multiplied by at most 2. This is a key point in the study of the complexity of algorithm involving systems of quads.

Geometry of Systems of Quads We use the construction described below. We need to make a choice for the angles at the corners of the quads. We thus decide to set the angles such that the face vertices have flat geometry. It is done by putting $\frac{1}{4}$ for all the angles. Indeed in that case $\kappa(f) = 1 - \frac{4}{2} + 4 * \frac{1}{4} = 1 - 2 + 1 = 0$. With that setting the faces of our systems of quads are flat squares. We will be mainly interested by surfaces of genus at least 2. In that case we have a strictly negative

curvature at every vertex. Adding that geometry to the systems of quads allows to give simple proofs to some structural theorems. We are interested in the shapes of geodesics in such a surface. In general the interest of geodesics in hyperbolic metrics is their uniqueness. Here, we hope to have a geometry that is not far from a discrete hyperbolic metric. Ideally we would like to have a nice way to start from any curve and then tighten it to the corresponding unique geodesic. Some easy modifications will tighten a curve (see figure 37). The first curve has a spur that can be just removed. The second one has a **bracket** that is a turn sequence $1, 2 \cdots 2, 1$ or $-1, -2 \cdots -2, -1$. The **turn sequence** of a path is the number of corners it leaves on its left at each of its vertex. We define the edge-paths without spurs nor brackets as **combinatorial geodesics**. The following proposition is a first justification of this definition.

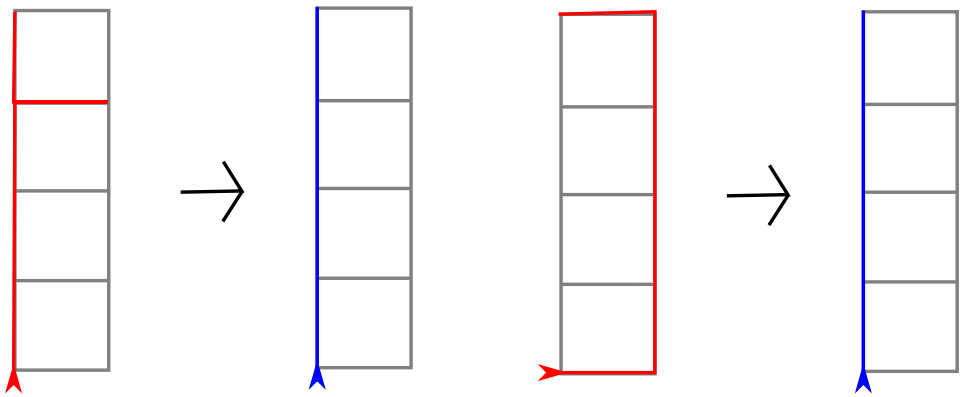

Figure 37: Spur and bracket removing.

Theorem 2.3.6. *A non-trivial contractible cycle in a system of quads of genus at least 2 has either a spur or four brackets.*

Proof. Let C be a non-trivial contractible cycle of a combinatorial map M. Let us consider the disk diagram Δ_C given by proposition 2.3.2. The boundary of Δ_C is labeled by C. Gauss-Bonnet gives us that $\sum_{v \in V(\Delta_C)} \kappa(v) + \sum_{f \in F(\Delta_C)} \kappa(f) = 1$ since the Euler characteristic of a disk is 1 (null genus and 1 boundary). Because of the assignments of the angles it becomes $\sum_{v \in V(\Delta_C)} \kappa(v) = 1$. Let us consider a vertex v in the interior of Δ_C. $\kappa(v) = 1 - \sum_{a \in N(v)} = 1 - \frac{1}{4} d(v)$. $d(v)$ is at least the degree of the vertex of M it is projected to by the diagram map. It implies that $d(v)$ is at least $2g$ and so at least 8. So $\kappa(v) < 0$. Assume now that v is on the boundary of Δ_C then $\kappa(v) = 1 - \sum_{a \in N(v)} -\frac{1}{2} B_v$ where B_v is the number of boundary angles around

v. As soon as B_v is at least 2 then $\kappa(v) \leq 0$. In fact Δ_C has the structure of a tree of open non-empty disk with cycle boundary (see figure 38). Without loss of generality we can assume that Δ_C is a disk with a cycle boundary and non-empty interior (a connected component of the dual of Δ_C in practice do that role, if there is no then Δ_C is flat and C has a spur). Now we label each vertex v on the boundary of Δ_C with the number of angle n_v it has inside Δ_C. If $n_v = 1$ then $\kappa(v) = \frac{1}{4}$, if $n_v = 2$ then $\kappa(v) = 0$ else $\kappa(v) \leq -\frac{1}{4}$. Let us consider the sequence of the n_v with respect to their order on the boundary. If we drop the 2 we have a cyclic word with at least 4 more 1 than other numbers in order to have a sum of curvature of 1. It implies that there are 4 pairs of consecutive 1. Adding the 2 back we obtain sequences of the form $1, 2 \cdots 2, 1$ that will project on the 4 announced brackets (signs are positive because we do not define the orientation of Δ_C but it can be positive or negative after the projection). □

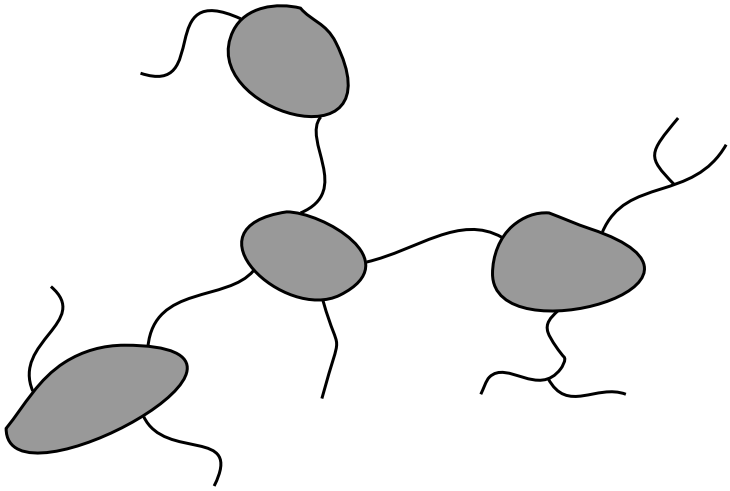

Figure 38: The general shape of Δ_C.

By looking the proof of the previous proposition we see that if Δ_C has an interior vertex then we must have more vertices corresponding to an angle $\frac{1}{4}$ to have the correct curvature. It proves the following.

Theorem 2.3.7. *A non-trivial contractible cycle in a system of quads of genus at least 2 that labels the boundary of a disk diagram with at least one interior vertex has either a spur or* ***five*** *brackets.*

Consequences of the Previous Results Those results will be useful in chapter 3 to describe the exact behavior of combinatorial geodesics in the systems of quads. As announced it is also very useful to design an optimal algorithm for the problem of testing homotopy. That work has been first imagined by Lazarus and Rivaud [Lazarus and Rivaud, 2012] and then its formulation has been significantly simplified by Erickson and Whittelsey [Erickson and Whittelsey, 2013]. The idea is the following. We want to construct a canonical curve homotopic to each curve we want to compare. A unique geodesic would be the perfect canonical curve for that purpose. Unfortunately, using the tightening of Figure 37 is not sufficient to obtain a unique element. It can be proved that adding a move that pushes the curve to the right is sufficient (see figure 39). It leads to a linear algorithm to test the homotopy of two curves. In details it runs in time $O(n+l)$ where n is the size of the map and l the length of the curves. The n comes from the reduction from the initial map to a system of quads and the l to the tightening procedure.

Theorem 2.3.8 ([Lazarus and Rivaud, 2012; Erickson and Whittelsey, 2013])**.** *The canonical form of a combinatorial closed curve of length ℓ on a system of quads can be computed in $O(\ell)$ time. It is the unique homotopic curve that contains no spurs nor brackets and whose turning sequence contains no -1's and contians at least one turn that is not -2.*

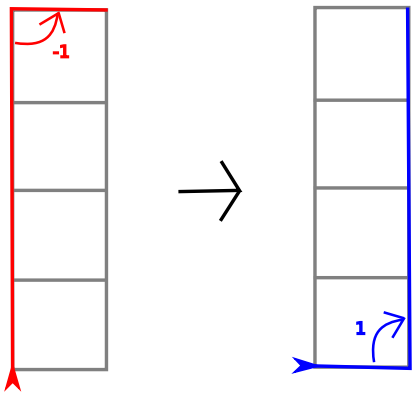

Figure 39: Pushing to the right.

2.4 Structure of Combinatorial Maps

2.4.1 Planar Maps

The result of Chapter 5 is a generalization to the torus of a result about plane triangulations. It is important for us to try to have the best possible formulation for our result because we aim at extend our generalization to higher genus. We first describe the classical constructions for the planar case.

Rooted Planar Maps A planar map is a combinatorial map of genus 0. Unless specified we will only consider maps maps without boundary. As described in Section 2.2.3 we store a map in a flag data-structure. We will consider rooted maps where the root can be any flag of the map. The main reason for considering rooted maps is that it is very simple to check if two rooted maps are isomorphic. Indeed, it is sufficient to fix a labeling of the flags of one of the maps and then, starting from the root of the other map, we can report the labeling of the first map to the second one. If it can be done consistently with respect to the three involutions α_0, α_1 and α_2 then the two maps are isomorphic (we implicitly consider connected maps). An unrooted map without symmetry i.e. whose group of automorphisms is trivial, corresponds to k rooted maps where k is the number of flags. When the map has symmetries, its unrooted version corresponds to a smaller number of rooted maps. For this reason, it is not obvious to obtain a uniform sampling of unrooted maps from a uniform sampling of rooted maps. However, since most of the maps have trivial automorphism group, it is generally considered as a good approximation.

Leaf-Rooted Embedded Binary Trees We begin with a very specific kind of maps that can be easily counted. A tree is binary if every vertex has degree degree 3 or 1. As usual, vertices of degree 1 are called **leaves** and the other vertices are called **nodes**. A map whose underlying graph is a tree has only one face. Let us denote by $\mathcal{B}$ the set of binary trees rooted at a leaf. Remark that an element in $\mathcal{B}$ can either be a single edge between a leaf and the root or a node connected to the root and to two binary trees. Let B_n be the number of binary trees with n non-root leaves (and one more leaf for the root). The above remark leads to the following recurrence relation:

$$B_1 = 1 \quad \text{and} \quad B_n = \sum_{n_1,n_2\in\mathbb{N}^*/n_1\leq n_2, n_1+n_2=n} B_{n_1}B_{n_2} \quad \text{for} \quad n \geq 2$$

Let $B(x)$ be the formal power series $\sum_{n>0} B_n \cdot x^n$. The recurrence relation we easily get $B(x) = x + B(x)^2$. In general, this kind of functional relation is the best we can hope but in this case, it leads to the following counting formula.

Theorem 2.4.1. *Let* $C_n = \frac{1}{n+1}\binom{2n}{n}$ *be the* n^{th} *Catalan number. With the above notation we have:* $B_{n+1} = C_n$.

Triangulated Polygons The Catalan numbers also appear in the counting of triangulated polygons. It is interesting for us to see how the triangulated polygons can be counted via a bijection with binary trees. The bijection is illustrated in Figure 40. The idea is that the dual of a triangulated polygon is a binary tree. Since triangulated polygons are maps with boundaries the notion of dual of Section 2.2.3 has to be adapted. To every triangulated polygon we associate a binary tree whose nodes are the triangles of the polygon and whose leaves are the boundary edges of the polygon. To node are connected by an edge if the corresponding triangles are adjacent. A leaf is connected to the triangle node that it bounds. We define a rooted polygon as a polygon with one boundary edge marked. The rooted triangulated polygons with n boundary edges are in bijection with the rooted binary trees with $n-1$ non-root leaves. It follows that there are $C_{n-2} = \frac{1}{n-1}\binom{2n-4}{n-2}$ distinct rooted triangulations of a polygon with n edges.

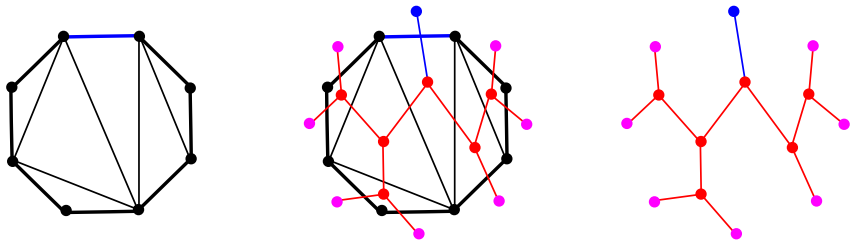

Figure 40: The bijection between triangulated polygons and binary trees.

Bijection for Planar Maps We have seen that trees can be counted easily thanks to a recursive formula. In order to count planar maps we shall define a bijection with a particular kind of trees. Given a rooted combinatorial map M the next procedure provides the desired tree. We perform a depth-first search from the root of M considering the rightmost non-visited edge at each step (see Figure 41). We construct the depth-first search spanning tree and we cut the edges that are not used by the spanning tree. Then this tree is decorated by oriented half-edges that represent the non-used edges. A non-tree edge gives rise to an outgoing edge from its most recently visited vertex and to ingoing edge on the other side. This procedure is called a depth-first search opening and was introduced by Schaeffer. We can recover M from the decorated tree. For this we remark that the half-edges form

a well parenthesized word when read from the root. For more details on both this constructions and its various implications see [Schaeffer, 1998]. For instance, this bijection provides a nice and simple explanation for Tutte's formula:

Theorem 2.4.2. *[Tutte, 1963] The number of rooted planar maps with n edges is:*

$$\frac{2 \cdot 3^n (2n)!}{n!(n+2)!} = \frac{2 \cdot 3^n}{n+2} C_n$$

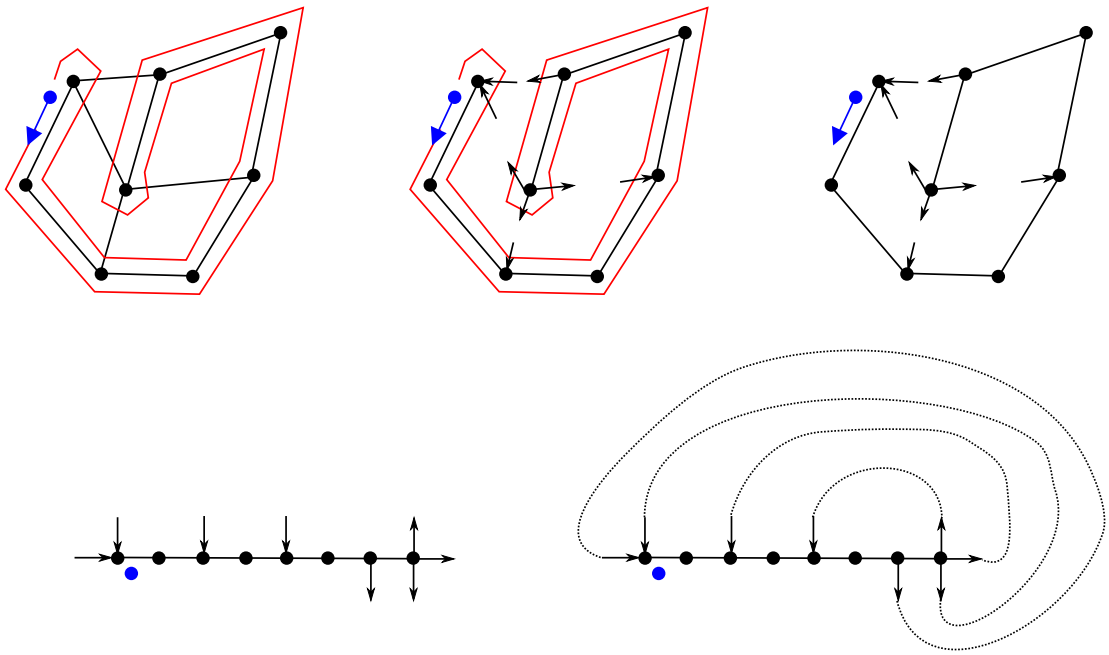

Figure 41: The depth-first search opening procedure.

2.4.2 Schnyder Woods

3-Orientation Let G be a graph. An orientation O of G is a choice of orientation for each of its edges. O is a 3-orientation if and only if each vertex has outdegree 3. Let M be a planar simple triangulation with underlying graph G. We fix a face of M as the outer face that we call the **special face**. We denote by v_0, v_1 and v_2 the 3 vertices of the special face. Let n, e and f be respectively the number of vertices, edges and faces of M. We have $2e = 3f$ since M is a triangulation and $n - e + f = 2$ by Euler formula. It follows $e = 3n - 6$. If O was a 3-orientation of G then we

would have $e = 3n$ which is a contradiction. In order to define 3-orientations for triangulations, we decide to forget the 3 vertices and the 3 edges of the special face. We obtain $n' = n - 3$ interior vertices and $e' = e - 3 = 3n - 9 = 3(n - 3) = 3n'$ interior edges. We call a 3-orientation of M any orientation of its internal edges such that all its internal vertices have outdegree 3.

Theorem 2.4.3. *[Kampen, 1976] Every plane triangulation admits a 3-orientation.*

Proof. All plane triangulations can be reduced to a tetrahedron by performing a sequence of edge-contractions [Rademacher and Steinitz, 1934]. Since the tetrahedron trivially has a 3-orientation, it is sufficient to observe that edge-contractions preserve the existence of 3-orientations. □

Since $e' = 3n'$, we remark that all the outgoing edges exhaust the internal edges. In particular, all the internal edges incident to the v_i's are ingoing.

Definition of Schnyder Woods A Schnyder wood [Schnyder, 1989] on a plane simple triangulation M is a 3-orientation of M with a labeling of the edges with 3 distinct labels 0, 1, 2 following the rule depicted on Figure 42. In clockwise order around a vertex we must see: an outgoing red edge, any number of incoming green edges (can be 0), an outgoing blue edge, any number of incoming red edges, an outgoing green edge and any number of incoming blue edges. We also require that all the edges incident to v_i receive color i for $i \in \{0, 1, 2\}$. The three colors play symmetric roles in the definition and we consider the Schnyder woods up to cyclic permutation of the colors. We denote by T_i the set of edges with color i.

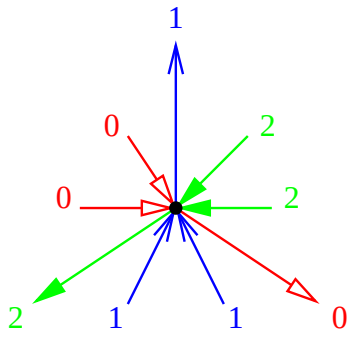

Figure 42: The local rule of Schnyder woods.

Proposition 2.4.4. *T_i contains no oriented cycle.*

Proof. By way of contradiction, assume that T_i contains an oriented cycle C. We consider the restriction M' of M to the interior of C included. Let n_1 be the length of C and let n_2 be the number of interior vertices of M'. Since the coloration forms a Schnyder wood, each vertex on C has one outgoing edge inside M' and one outside. Hence every vertex in C has outdegree 2 and every interior vertex has outdegree 3. We deduce that the number of edges of M' is $e = 2n_1 + 3n_2$. Let f be the number of triangles of M', $3f$ counts twice the interior edges and once those of C so that $3f = 2e - n_1$. We obtain from the Euler relation $1 = n_1 + n_2 - 2n_1 - 3n_2 + \frac{2}{3}(2n_1 + 3n_2) - \frac{1}{3}n_1 = 0$ which is a contradiction. □

Corollary 2.4.5. *Each T_i is an oriented spanning tree rooted at v_i. In particular, T_i is connected as an undirected graph.*

That last corollary motivates the name of Schnyder woods and is illustrated in Figure 43. In the plane, Schnyder woods are entirely determined by their underlying 3-orientation:

Theorem 2.4.6. *[de Fraysseix and Ossona de Mendez, 2001] Each 3-orientation of a plane simple triangulation admits a unique coloring leading to a Schnyder woods.*

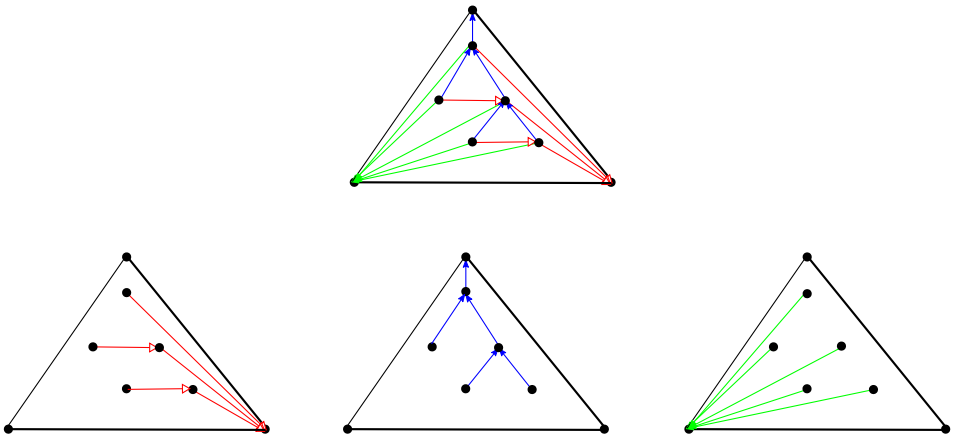

Figure 43: The three monochromatic trees.

Applications of Colorings Since orientations and colorings are equivalent, one may wonder if the colorings are useful. We present a well known application that is naturally described using the colorings. We first give a property of the colorings.

Proposition 2.4.7. *Let v be an internal vertex. There is a unique monochromatic oriented path from v to v_i for $i \in \{1, 2, 3\}$. In addition, the three paths do not cross.*

We fix an internal vertex v. The three paths of the proposition define three regions of the triangulation that we denote by R_i (see Figure 44). R_0 is the region bounded by the blue and green paths and the edge between v_1 and v_2. We associate a triple (x, y, z) to v where $x = |R_0|$ (the cardinality of R_0 is the number of faces in the corresponding region), $y = |R_1|$ and $z = |R_2|$. In addition, we set v_0 to $(f-1, 0, 0)$, v_1 to $(0, f-1, 0)$ and v_2 to $(0, 0, f-1)$ where f is the total number of faces of the triangulation (the -1 accounts for the outer face).

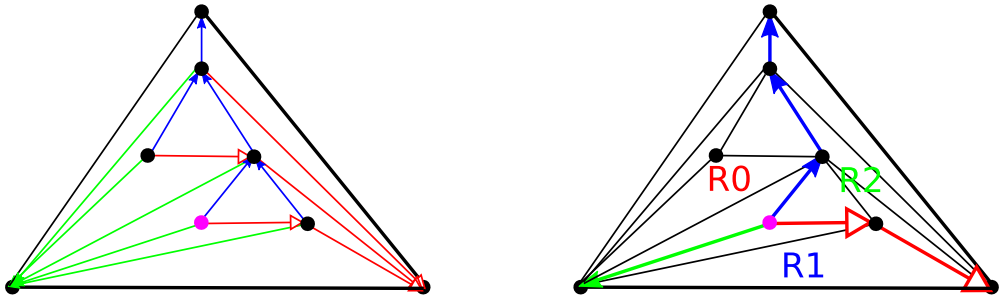

Figure 44: The three parts R_i.

Theorem 2.4.8. *[Schnyder, 1990] With the notation described above, using the triple as coordinates in $\mathbb{R}^3$ lead to a correct embedding of the triangulation in the plane $\{(x, y, z) \mid x + y + z = f - 1\}$. In addition, the orthogonal projection on the plane $\{z = 0\}$ leads to a correct embedding of the triangulation on a $(f-1) \times (f-1)$ grid.*

We can also count the number of Schnyder woods (with a root fixed at v_0) of plane simple triangulations by considering the colors [Bonichon, 2002]. Let us use the same construction for the dual than the one we used for polygons by considering the triangle $v_0v_1v_2$ as a boundary (see Figure 46). Then we cut edges dual to the red edges of the tree T_0. We obtain a tree-cotree decomposition $(T_0, C_0, X = \emptyset)$ of triangulations with one boundary. We associate a particular tree to every (rooted) Schnyder wood. The number of Schnyder woods over plane simple triangulations with $n + 3$ vertices can be obtained by counting the trees that can be obtained by this association and is $C_{n+2}C_n - C_{n+1}^2$ ([Bonichon, 2002, Theorem 3.3.1, p.51]).

Lattice Structure Let M be a plane simple triangulation and let $\mathcal{SW}_M$ be the set of all Schnyder woods on M. Consider two distinct elements sw_1 and sw_2 of

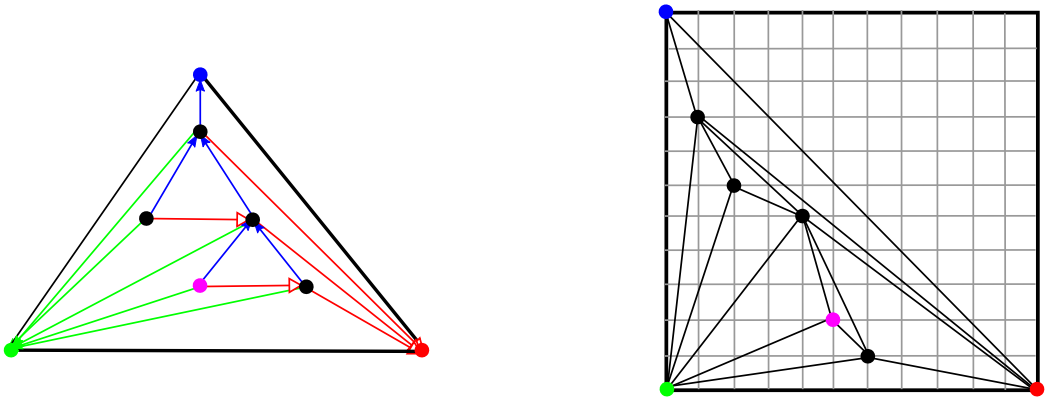

Figure 45: The grid embedding of our example.

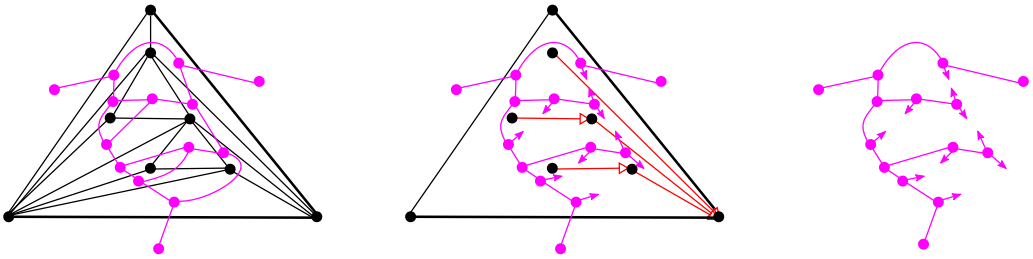

Figure 46: The grid embedding of our example.

$\mathcal{SW}_M$. By definition their difference is the set of edges with opposite orientation in sw_1 and sw_2. We write $sw_1 \leq sw_2$ when their difference is the boundary of a unique triangle (not necessarily a face triangle) that is oriented counterclockwise in sw_1 and clockwise in sw_2. We consider the partial order on $\mathcal{SW}_M$ generated by those relations.

Theorem 2.4.9. *[Propp, 1993; Ossona de Mendez, 1994; Felsner, 2004]* $(\mathcal{SW}_M, \leq)$ *has a structure of distributive lattice.*

The minimal element of the lattice can be chosen as a canonical representative. Moreover the canonical element can be computed in linear time by swapping triangles. The minimal element is the unique 3-orientation of M with no clockwise oriented cycle. This characterization comes from the following lemma.

Lemma 2.4.10. *If a 3-orientation of a simple triangulation has a clockwise oriented cycle then it has a clockwise oriented triangle (not necessarily a face triangle).*

Proof. We fix a Schnyder wood coloring of the 3-orientation (see Theorem 2.4.6). Let C be a clockwise oriented cycle such that the interior map of C contains no clockwise oriented cycle different from C. The proof is similar as the one of Proposition 2.4.4. We use the same notations. We recall that n_1 denotes the length of the oriented cycle C and n_2 the number of interior vertices of M'. The main change is that we do not know the outdegree of the vertices of C in M'. Let k be the sum of the outdegrees of the vertices of C minus n_1 (we remove the n_1 edges of C). We obtain $1 = n_1 + n_2 - n_1 - k - 3n_2 + \frac{2}{3}(n_1 + k + 3n_2) - \frac{1}{3}n_1 = \frac{1}{3}n_1 - \frac{1}{3}k$ by Euler formula. It implies $k = n_1 - 3$. If C is not a triangle then $k \geq 1$. There exists an edge e incident to a vertex v of C and oriented outward v. Let i be the color of e. It is the first edge of the oriented path linking v to v_i. This path contains a subpath P that link v to another vertex v' of C (see Figure 47). C can be decomposed into two oriented paths P_1 from v to v' and P_2 from v' to v. $P \cdot P_2$ is a clockwise cycle that contradicts the choice on C. So C is a triangle. □

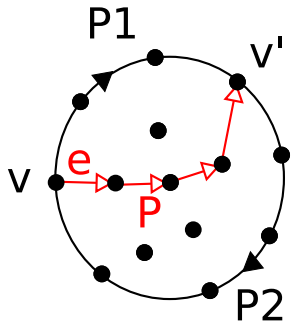

Figure 47: A shortcut.

2.4.3 Generalizations to Higher Genus

Unicellular Maps We saw that the plane maps are in bijection with particular trees. The trees can be seen as the skeletons of the maps. The question is: can we define the skeleton of a map of non-null genus? The good generalization of trees happens to be **unicellular maps**. A unicellular map is a combinatorial map with a unique face. Unicellular maps of genus 0 are exactly the combinatorial maps whose underlying graph is a tree. Unicellular maps are well studied. For a detailed and purely combinatorial study we advise to look at [Chapuy, 2009]. Let M be a combinatorial map of genus g. If (T, C, X) is a tree-cotree decomposition of M then the restriction of M to $T \cup X$ is a unicellular map. Thus, when looking for the

skeleton of M we will look at something corresponding to a tree with $2g$ additional edges.

Toroidal Schnyder Woods The definition of Schnyder woods for toroidal triangulations is the following. Given a toroidal triangulation M (here a triangulation is a general triangulation without contractible loops or homotopic double edges), a **(toroidal) Schnyder wood** of M is an orientation and coloring of the edges of M with the colors 0, 1, 2, where each vertex satisfies the Schnyder property (see Figure 48 for an example). Note that in the torus case we have $e = 3n$, so we can ask for the Schnyder property everywhere and we no longer need to specify a triangle as in the planar case. We consider that a Schnyder wood and its underlying orientation are the same object since one can easily recover a coloring of the edges in a greedy way (by choosing the color of an edge arbitrarily and then satisfying the Schnyder property at every vertex).

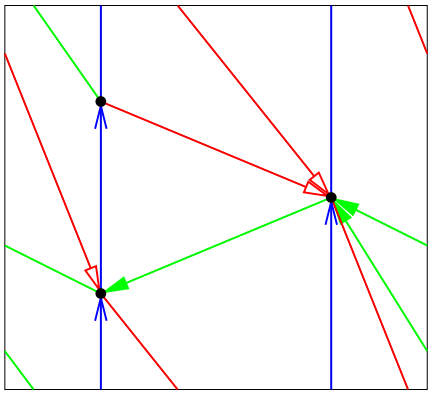

Figure 48: A Schnyder wood in a toroidal triangulation (opposite sides are identified in order to form a torus).

This introduction of non-planar Schnyder woods is due to [Gonçalves and Lévêque, 2013]. Their study is pushed forward in [Gonçalves et al., 2015]. All the results stated in this section and in the next one comes from those 2 articles. An actualize and very detailed summary can be found in [Lévêque, 2016]. Note that they provide generalization of Schnyder woods to maps of any genus. However, the situation becomes quite different for triangulations of genus 2 or more since we have $e > 3n$.

Properties of Toroidal Schnyder Woods The situation is quite different from the planar case. In a Schnyder wood of a toroidal triangulation, each vertex has exactly one outgoing arc in each color so that there are monochromatic cycles contrarily

to the planar case. Moreover the graph T_i induced by the color i is not necessarily connected. However, Proposition 2.4.4 does not disappear but can be weakened as follows.

Proposition 2.4.11. *T_i contains no oriented **null-homologous** cycle.*

By a recursive argument based on edge contraction we can show that:

Theorem 2.4.12. *Every toroidal triangulation admits a 3-orientation.*

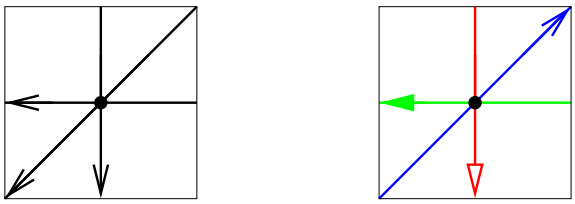

Figure 49: Two different orientations of a toroidal triangulation. Only the one on the right corresponds to a Schnyder wood.

In the planar case, any orientation corresponds to a Schnyder wood (Theorem 2.4.6). This is not true for toroidal triangulations since there exists 3-orientations that do not correspond to a Schnyder wood (see Figure 49). However, it is still true that every toroidal triangulation admits a Schnyder wood. This again shown by contracting edges until we reach a one vertex triangulation. It leaves us with the following question: how to characterize and compute a canonical Schnyder wood? It is the subject of the next section. Just before we go deeper into this, we quickly justify the use of homology over homotopy.

Homology vs. Homotopy The two notions are equivalent for simple curves on the torus. However, it is interesting to try to make an informed choice in order to facilitate the generalization to higher genus. First, the operation we use to travel in the lattice of the Schnyder woods of a plane map is swapping triangles. If we want to find a lattice structure based upon the same kind of operation then it should be possible to go from any Schnyder wood in the lattice to any other one by applying a sequence of triangle swaps. It implies that two Schnyder woods in the same lattice will differ in a set of edges that is a null-homologous cycle (that is not necessarily a simple cycle). It implies that a splitting cycle in genus at least two should be considered like contractible ones. It thus justify the use of homology. Let us give a formal definition of the notion of homology that we will use on 3-orientations. An

orientation of a graph can naturally be seen as a 1-chain. Let O_1 and O_2 be two 3-orientations of the same graph. We consider the **difference** d between O_1 and O_2 defined as $d = \frac{1}{2}(O_1 - O_2)$. The half factor is correct because the 1-chain $O_1 - O_2$ only have coefficients that are 0, 2 or -2. We say that O_1 and O_2 are homologous if and only if d is null-homologous cycle (not necessarily simple or connected). Remark that we use $\mathbb{Z}$-homology here.

2.4.4 Canonical Toroidal Schnyder woods

Crossing Schnyder Woods A Schnyder wood of a toroidal triangulation is said to be a **crossing Schnyder Woods**, if for each pair (i, j) of different colors, there exists a monochromatic cycle of color i intersecting a monochromatic cycle of color j.

Theorem 2.4.13. *Every toroidal triangulation admits a crossing Schnyder wood.*

Figure 50 depicts two different Schnyder woods of the same graph where just the one on the left is crossing. On the right of the figure, the blue cycle crosses both green and red cycles but none of the red cycles cut the green cycle and we say that the Schnyder wood is **half-crossing**. Note that the Schnyder wood on the right is obtained from the one on the left by flipping a clockwise triangle into a counterclockwise triangle.

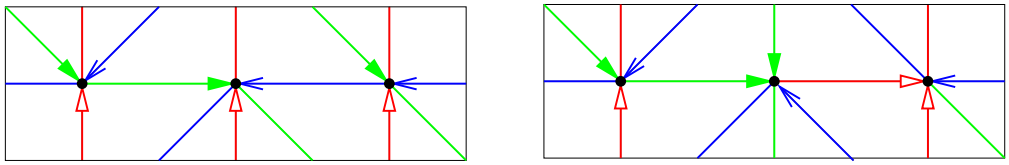

Crossing Half-crossing

Figure 50: A crossing and a half-crossing Schnyder wood.

Consider a toroidal triangulation G given with a crossing Schnyder wood. Let D_0 be the corresponding 3-orientation of G and let $O(G)$ be the set of all the orientations of G that are homologous to D_0.

Lemma 2.4.14. *The crossing Schnyder woods of G are pairwise homologous.*

It implies that $O(G)$ contains all the crossing Schnyder woods of G. Thus the definition of $O(G)$ does not depend on the particular choice of D_0 and thus it is uniquely defined. Thanks to the next theorem, we call the elements of $O(G)$ the **homologous-to-crossing Schnyder woods**, or **HTC Schnyder woods** for short.

Theorem 2.4.15. *Every orientation of $O(G)$ admits a coloring that respects the Schnyder property.*

Examples Figure 51 gives an example of a HTC Schnyder wood that is not crossing and a Schnyder woods that is not HTC. The example on the left is obtained from the crossing Schnyder wood of Figure 50 by flipping two triangles (one to obtain the half-crossing Schnyder wood of Figure 50 and then another one flipped from counterclockwise to clockwise). Thus it is HTC since the difference with a crossing Schnyder wood is a 0-homologous oriented subgraph. The example on the right of Figure 51 is obtained from the crossing Schnyder wood of Figure 50 by reversing the three vertical red monochromatic cycles. The union of these three cycles is not a 0-homologous oriented subgraph, thus the resulting orientation is not HTC.

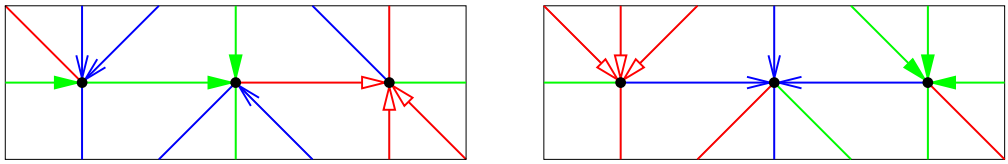

Non-crossing and HTC Not HTC

Figure 51: Non-crossing Schnyder woods.

Properties of Monochromatic Cycles Here we state some properties of monochromatic cycles. Since we do not want to consider the orientation of the monochromatic cycles we use $\mathbb{Z}/2\mathbb{Z}$-homology in this paragraph. In particular a cycle C is $\mathbb{Z}/2\mathbb{Z}$-homologous to C' if one of the following is true: C and C' are $\mathbb{Z}$-homologous or C and $-C'$ are $\mathbb{Z}$-homologous.

Lemma 2.4.16. *Monochromatic cycles of the same color are $\mathbb{Z}/2\mathbb{Z}$-homologous.*

Proof. Let C and C' be two monochromatic cycles of the same color. By Proposition 2.4.11, none of C or C' is null-homologous. For a contradiction, we assume that C and C' are not homologous. Since they are drawn on the torus they must intersect. Hence they must have a common edge (there is only one outgoing edge of the correct color on the intersection vertex). Since every vertex has a unique outgoing edge of each color, C and C' must actually be equal. We have thus reached a contradiction. □

Lemma 2.4.17. *Let C and C' be two monochromatic cycles of different colors that intersect. Then C and C' are not $\mathbb{Z}/2\mathbb{Z}$-homologous.*

Proof. By way of contradiction, we assume that C and C' are homologous. Let v be a common vertex of C and C'. Let $\tilde{v}$ be any lift of v in the universal cover of the torus. Let $\tilde{C}$ and $\tilde{C'}$ be the lifts of C and C' starting at $\tilde{v}$. Since C and C' are homologous, $\tilde{C}$ and $\tilde{C'}$ have the same endpoints (note that as we consider homology in $\mathbb{Z}/2\mathbb{Z}$, the orientation of C' has to be chosen consistently with the one of C). It implies that there is a disk in the universal cover whose boundary consists in a simple subpath of C concatenated with a simple subpath of C'. As in the proof of Lemma 2.4.10, the boundary of the disk must have exactly 3 less outedges than its length. Since monochromatic cycles have an outedge on each of their sides, the vertices of the boundary of the disk all have an outedge directed inside the disk except maybe for the two vertices of the concatenation. This leads to a contradiction. □

The first remark of the paragraph and the last lemma gives:

Lemma 2.4.18. *Let C and C' be two monochromatic cycles of different colors of a crossing Schnyder wood. Then C and C' are not $\mathbb{Z}/2\mathbb{Z}$-homologous.*

HTC Lattice It is proved in [Gonçalves et al., 2015] that on any oriented surface the set of orientations of a given map having the same homology carries a structure of distributive lattice. Thus, in particular, the set of HTC Schnyder wood carries a structure of distributive lattice. We choose an arbitrary face f_0 of G. This choice of a particular face corresponds to the choice of the outer face in the planar case and defines clockwise and counterclockwise for the oriented null-homologous (simple) cycle of G. Indeed such a cycle is separating. It thus define two connected components on the surface, one not containing f_0. We say that the cycle is oriented counterclockwise if this connected component is on the left of the cycle. We define the order we want for the lattice as follows. Let O_1 and O_2 be two 3-orientations of G. Let $d = \frac{1}{2}(O_1 - O_2)$ be their difference. Then we say that $O_1 \leq_{f_0} O_2$ if and only if there exists a set of faces oriented counterclockwise (different from f_0) such that $d = \sum_{F \in \mathcal{F}} \partial(F)$. In the same idea, that a directed subgraph is oriented clockwise if the corresponding 1-chain can be written as a sum of faces oriented clockwise.

Theorem 2.4.19. *$(O(G), \leq_{f_0})$ is a distributive lattice.*

Structure of Lattices of 3-Orientations [Gonçalves et al., 2015] provides different formulations and many proofs for the same theorems. One of their formulations is especially relevant for us. It relies on a certain function γ defined as follows. Consider a particular orientation of G. Let C be a cycle that is given with an arbitrary

direction (not necessarily induced by the orientation of G). Then $\gamma(C)$ is defined by:

$$\gamma(C) \quad = \quad |\{\text{edges leaving } C \text{ on its right}\}| \quad - \quad |\{\text{edges leaving } C \text{ on its left}\}|$$

Leaving here means that the edges are outgoing when seen from their vertex in C. We can now give an alternative characterization of $O(G)$ using the function γ. We remark that by the Schnyder property, it is clear that in a toroidal Schnyder woods, a monochromatic cycle C always satisfies $\gamma(C) = 0$ (each vertex of C has one outgoing edge on each side). Suppose that the orientation of G corresponds to a crossing Schnyder wood and let C_1, C_2 be two monochromatic cycles of different colors. By the preceding remark we have $\gamma(C_1) = \gamma(C_2) = 0$. By Lemma 2.4.18, the two cycles C_1, C_2 are non-homologous. Since they are also non-null-homologous they form a basis for the homology. When swapping a null-homologous oriented subgraph, the value of γ on a given cycle does not change. Thus any HTC Schnyder wood also satisfies $\gamma(C_1) = \gamma(C_2) = 0$. We can show that:

Proposition 2.4.20. *If a 3-orientation of a toroidal triangulation satisfies γ equals 0 for two cycles forming a basis for the homology, then γ equals 0 for any non-null-homologous cycle.*

It implies that any HTC Schnyder woods satisfies γ equals 0 for any non-null-homologous cycle. We call this property the γ_0 **property**. In general, it is proved that γ can be used to characterize a partition of the 3-orientations of a given triangulation into sets that carry a structure of lattice for our order relation. This theorem is far from trivial, see [Lévêque, 2016] for a detailed proof. Note that computing γ is the easiest way to check that our examples of 3-orientation actually have the claimed properties.

Theorem 2.4.21. *Let G be a triangulation and C_1 and C_2 two simple cycles in G that form a basis of the homology of G. The couple $\mathcal{C}_\gamma = (\gamma(C_1), \gamma(C_2))$ of a 3-orientation O characterizes its lattice as follows:*

- *If $\mathcal{C}_\gamma = (0, 0)$, then O is a HTC Schnyder woods.*
- *If $\mathcal{C}_\gamma = (a, b)$ with a and b being multiples of 3 and a or b different from 0, then O is non-HTC Schnyder woods.*
- *If $\mathcal{C}_\gamma = (a, b)$ and 3 does not divide a or b, then O does not correspond to a Schnyder woods.*

Canonical Schnyder Woods Since $(O(G), \leq_{f_0})$ is a distributive lattice, it has a unique minimal element. This minimal element will be our canonical Schnyder woods. The following lemma gives a characterization of our canonical structure. The proof is similar to the one of Lemma 2.4.10.

Lemma 2.4.22. *The minimal element of $(O(G), \leq_{f_0})$ is the only HTC Schnyder wood that contains no clockwise (non-empty) 0-homologous oriented subgraph with respect to f_0.*

The crossing Schnyder wood of Figure 52 is the minimal HTC Schnyder wood for the choice of f_0 corresponding to the shaded face.

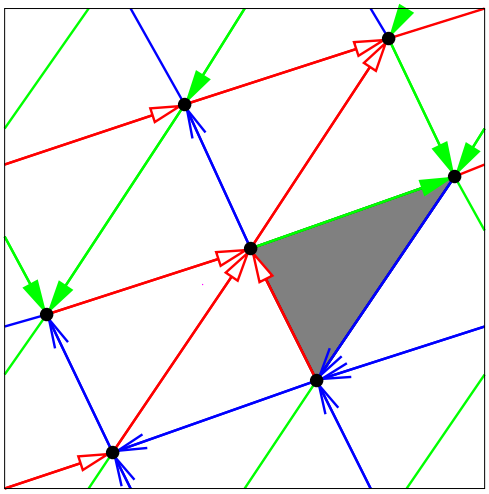

Figure 52: The minimal HTC Schnyder wood of K_7 with respect to the shaded face.

The two HTC Schnyder woods of Figure 50 are not minimal (for any choice of special face f_0) since they contain several triangles that are oriented clockwise. On the contrary, the HTC Schnyder wood of Figure 51 is minimal with respect to its only face oriented clockwise. These examples shows that the minimal HTC Schnyder wood is not always crossing.

An Important Lemma We define the dual orientation D^* of an orientation D of G as an orientation of the edges of the dual map G^* of G satisfying the following rule: the dual e^* of an edge e goes from the face on the left of e to the face on the right of e. The following lemma gives the key property of HTC Schnyder woods that we need in Chapter 5:

Lemma 2.4.23. *If D is an orientation corresponding to a HTC Schnyder wood, then the dual orientation D^* contains no oriented non-null-homologous cycle.*

Proof. We first prove the property for a crossing Schnyder wood and then show that it is stable by reversing a 0-homologous oriented subgraph. Thus it is true for all HTC Schnyder woods.

Consider a crossing Schnyder wood of G whose existence is ensured by Theorem 2.4.13. Let D_0 be the corresponding orientation. For $i \in \{0, 1, 2\}$, let C_i be a monochromatic cycle of color i, it is not null-homologous by Proposition 2.4.11. By Lemma 2.4.18, the cycles C_i are pairwise non-homologous. Thus for $i \in \{0, 1, 2\}$, the two cycles C_{i-1} and C_{i+1} form a basis B_i for the $\mathbb{Q}$-homology. It means that the matrix of a B_i in a $\mathbb{Z}$-basis has non-null determinant but it can be anything in $\mathbb{Z}$. In practice it is the algebraic number of intersections between the two cycle of B_i. By the Schnyder property, cycle C_{i-1} is crossing C_i (maybe several time) from left to right. Thus the homology of any closed curve can be expressed in at least one of the basis B_i with only non-negative coefficients (in any case the third cycle that is not in B_i has two negative coefficients, then some basic algebra gives the result).

Suppose now by contradiction that D_0^* contains an oriented non-null-homologous cycle C^*. If this cycle is homologous to some C_i then it crosses then C_i crosses C^* at least one time from left to right, a contradiction. Let i in $\{0, 1, 2\}$, such that C^* is homologous to $\lambda_{i-1}C_{i-1} + \lambda_{i+1}C_{i+1}$ with λ_{i-1} and λ_{i+1} in $\mathbb{Q}$ with $\lambda_{i-1} > 0$ and $\lambda_{i+1} > 0$. Then C_{i+1} is crossing C^* at least once from left to right, contradicting the fact that C^* is an oriented cycle of D_0^*. So D_0^* contains no oriented non-null-homologous cycle.

Consider now a HTC Schnyder wood of G and let D be the corresponding orientation. Since D and D_0 are both element of $O(G)$ they are homologous to each other. Let T be the 0-homologous oriented subgraph of D such that $T = D \setminus D_0$. Thus D_0 is obtained from D by reversing the edges of T.

Suppose by contradiction that D^* contains an oriented non-null-homologous cycle C^*. The oriented subgraph T is 0-homologous thus it intersects C^* exactly the same number of time from right to left than from left to right. Since C^* is oriented T cannot intersect it from left to right. So T does not intersect C^* at all. Thus reversing T to go from D to D_0 does not affect C^*. Thus C^* is an oriented non-null-homologous cycle of D_0^*, a contradiction. □

For the non-HTC Schnyder woods of Figure 51, one can see that there is a horizontal oriented non-null-homologous cycle in the dual, so it does not satisfy the conclusion of Lemma 2.4.23. Note that this property is not a characterization of being

HTC. Figure 53 is a Schnyder wood that is not HTC but satisfies the conclusion of Lemma 2.4.23.

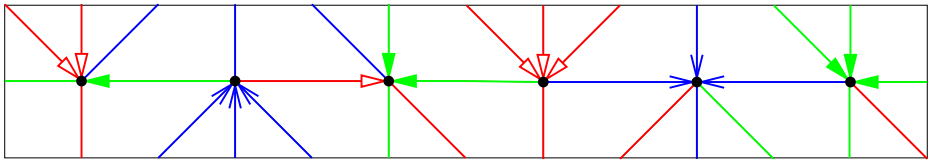

Figure 53: A Schnyder wood that is not HTC but contains no oriented non-null-homologous cycle in the dual.

Chapter 3

Computing the Geometric Intersection Number of Curves

3.1 Introduction

Let S be a surface and let $\alpha, \beta : \mathbb{R}/\mathbb{Z} \to S$ be two freely homotopic curves. Assuming the curves in generic position, their number of intersections is

$$|\alpha \cap \beta| = |\{(t, t') \mid t, t' \in \mathbb{R}/\mathbb{Z} \text{ and } \alpha(t) = \beta(t')\}|.$$

Their **geometric intersection number** only depends on their free homotopy classes and is defined as

$$i(\alpha, \beta) = \min_{\alpha' \sim \alpha, \beta' \sim \beta} |\alpha' \cap \beta'|$$

Likewise, the number of self-intersections of α is given by

$$\frac{1}{2}|\{(t, t') \mid t \neq t' \in \mathbb{R}/\mathbb{Z} \text{ and } \alpha(t) = \alpha(t')\}|,$$

and its minimum over all the curves freely homotopic to α is its **geometric self-intersection number** $i(\alpha)$. Note the one half factor that comes from the identification of (t, t') with (t', t).

The geometric intersection number is an important parameter that allows to stratify the set of homotopy classes of curves on a surface. The surface is usually endowed with a hyperbolic metric, implying that each homotopy class is identified by its unique geodesic representative. Extending a former result by Mirzakhani [Mirzakhani, 2008], Sapir [Sapir, 2016; Mirzakhani, 2016] has recently provided tight asymptotics for the number of closed geodesics with bounded length

and bounded geometric intersection number. Chas and Lalley [Chas and Lalley, 2012] also proved that the distribution of the geometric intersection number with respect to the length of geodesics approaches the Gaussian distribution as the length grows to infinity. Other more experimental results were obtained with the help of a computer to show the existence of length-equivalent homotopy classes with distinct geometric intersection numbers [Chas, 2014]. Hence, for both theoretical and practical reasons various aspects of the computation of geometric intersection numbers have been studied in the past including the algorithmic ones. Nonetheless, all the previous approaches rely on rather complex mathematical arguments and to our knowledge no exact complexity analysis has yet appeared. In this chapter, we make our own the words of Dehn who noted that the metric on words (on some basis of the fundamental group of the surface) can advantageously replace the hyperbolic metric [de La Harpe, 2010]. We propose a combinatorial framework that leads to simple algorithms of low complexity to compute the geometric intersection number of curves or to test if this number is zero. Our approach is based on the computation of canonical forms as recently introduced in the purpose of testing whether two curves are homotopic [Lazarus and Rivaud, 2012; Erickson and Whittelsey, 2013]. Canonical forms are instances of combinatorial geodesics who share nice properties with the geodesics of a hyperbolic surface. On such surfaces each homotopy class contains a unique geodesic that moreover minimizes the number of self-intersections. Although a combinatorial geodesic is generally not unique in its homotopy class, it must stay at distance one from its canonical representative and a careful analysis of its structure leads to the first result of the chapter.

Theorem 3.1.1. *Given two curves represented by closed walks of length at most ℓ on a combinatorial surface of complexity n we can compute the geometric intersection number of each curve or of the two curves in $O(n + \ell^2)$ time.*

As usual the complexity of a combinatorial surface stands for its total number of vertices, edges and faces. A key point in our algorithm is the ability to compute the primitive root of a canonical curve in linear time. We next provide an algorithm to compute an actual curve immersion that minimizes the number of self-interactions in its homotopy class.

Theorem 3.1.2. *Let c be a closed walk of length ℓ in canonical form. We can compute a combinatorial immersion with $i(c)$ crossings in $O(\ell^4)$ time.*

We also propose a nearly optimal algorithm that answers an old problem studied by Poincaré [Poincaré, 1904, §4]: decide if the geometric intersection number of a curve is null, that is if the curve is homotopic to a simple curve.

Theorem 3.1.3. *Given a curve represented by a closed walk of length ℓ on a combinatorial surface of complexity n we can decide if the curve is homotopic to a simple curve in $O(n + \ell \log^2 \ell)$ time. In the affirmative we can construct an embedding of c in the same amount of time.*

In the next section we review some of the previous relevant works. We introduce our combinatorial framework in Section 3.3 and discuss the structure of combinatorial geodesics in Section 3.4. Section 4.3 presents our general simple strategy to compute the geometric intersection number. The proof of Theorem 3.1.1 is given in the next three Sections where the case of non-primitive curves is also treated. The computation of a minimally crossing immersion is presented in Section 3.9 and 3.10 with the proof of Theorem 3.1.2. We finally propose a simple algorithm to detect and embed curves homotopic to simple curves (Theorem 3.1.3) in Section 3.11.

3.2 Historical Notes

In the fifth supplement to its **Analysis situs** Poincaré [Poincaré, 1904, §4] describes a method to decide whether a given closed curve γ on a surface is homotopic to a simple curve. For this, he considers the surface as the quotient $\mathbb{D}/\Gamma$ of the Poincaré disk $\mathbb{D}$ by a particular group Γ of hyperbolic translations. The endpoints of a lift of γ in the Poincaré disk are related by a hyperbolic translation whose axis is a hyperbolic line L representing the unique geodesic homotopic to γ (see Figure 54). He concludes that γ is homotopic to a simple curve if and only if all the transforms of L by Γ are pairwise disjoint or equal (in the example of the figure two distinct lines have a crossing so that γ is not homotopic to a simple curve). This method was turned into an algorithm by Reinhart [Reinhart, 1962] who worked out the explicit computations in the Poincaré disk using the usual representation of hyperbolic translations by two-by-two matrices. The entries of the matrices being algebraic the computation could indeed be performed accurately on a computer. The ability to recognize curves that are **primitive curves**, i.e. whose homotopy class cannot be expressed as a proper power of another class, happens to be crucial in this algorithm though computationally expensive.

Birman ans Series [Birman and Series, 1984] subsequently proposed an algorithm for the case of surfaces with nonempty boundary that avoids manipulating algebraic numbers. While their arguments applies to a hyperbolic structure, their algorithm is purely combinatorial. Intuitively, a surface with boundary deform retracts onto a fat graph (in fact a fat bouquet of circles) whose universal covering space embeds as a fat tree in the Poincaré disk. The successive lifts of a curve γ trace a bi-infinite

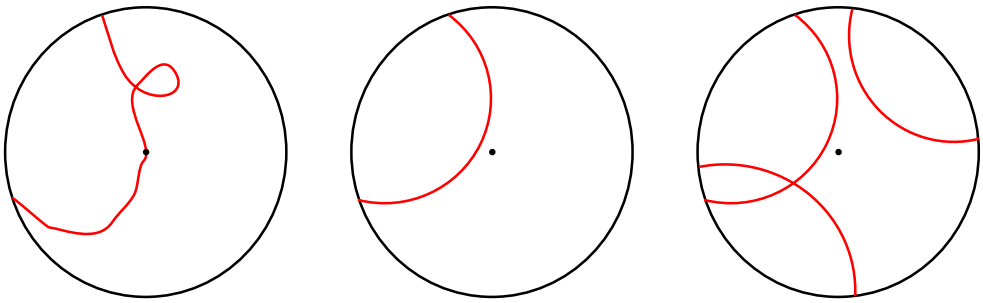

Figure 54: On the left γ, on the middle the corresponding line L and on the right some translations of L.

path in this tree. The limit points of this path belongs to the circle $\partial\mathbb{D}$ at infinity and coincide with the ideal endpoints of the axis of the hyperbolic transformation corresponding to any of the lifts of the curve in the path. The question as to whether two lifts give rise to intersecting axes can thus be reduced to test if the corresponding bi-infinite paths have separating limit points on $\partial\mathbb{D}$. As for Reinhart, Birman and Series assume that the homotopy class of γ is given by a word W on some given set of generators of the fundamental group of the surface. In turn, the above test on bi-infinite paths boils down to considering the cyclic permutations of W and W^{-1} in a cyclic lexicographic order and to check if this ordering is well-parenthesized with respect to some pairing of the words.

Cohen and Lustig [Cohen and Lustig, 1987] further observed that the approach of Birman and Series could be extended to count the geometric intersection number of one or two curves. In a second paper Lustig [Lustig, 1987] tackles the case where the curves are taken on a surface without boundary. He considers a closed surface with negative Euler characteristic as the quotient of the Poincaré disk by a group of transformations isomorphic to the fundamental group of the surface. The main contribution of the paper is to define a canonical representative for every free homotopy class α given as a word in some fixed system of loops generating the fundamental group. Lustig first notes that there is no known way of choosing a unique word for α. In particular, the shortest words are far from unique. He rather represents (a lift of) α by a path in the union $G \cup N \cup H$ of three tessellations of $\mathbb{D}$, where the edges of G are all the lifts of the generating loops, N is the dual tessellation, and the edges of H joins the vertices of G with the vertices of N. (Although the graphs of G, N and H are embedded their union is not as every edge of G crosses its dual edge in N.) Lustig gives a purely combinatorial characterization of canonical paths and argues that the method in his first paper [Cohen and Lustig, 1987] can be applied

to the canonical representative of α. Overall the two papers adds up to 60 pages with essential arguments from hyperbolic geometry and the complexity of the whole procedure remains to be determined.

Other approaches were developed without assuming any hyperbolic structure. Based on the notion of winding number, Chillingsworth [Chillingworth, 1969; Chillingworth, 1972] provides an algorithm to test whether a curve is homotopic (this time with fixed basepoint) to a simple curve on a surface with nonempty boundary. He also proposed an algorithm for determining when a given set of **simple closed curves** can be made disjoint by (free) homotopy [Chillingworth, 1971]. While the winding number relies on a differentiable structure, Zieschang [Zieschang, 1965; Zieschang, 1969] used the connection between the automorphisms of a topological surface and the automorphisms of its fundamental group in order to detect the homotopy classes of simple closed curves.

A related work by Hass and Scott [Hass and Scott, 1985] is concerned with curves which have excess intersection, i.e. that can be homotoped so as to reduce their number of intersections. Hass and Scott introduce various types of monogons and bigons that can be either embedded, singular or weak. A **singular monogon** of a curve $\gamma : \mathbb{R}/\mathbb{Z} \to S$ is a contractible subpath of γ whose endpoints define a self-intersection of γ. A **bigon** of γ is defined by two self-intersections joined by two homotopic subpaths (with fixed endpoints) $\gamma|_\sigma$ and $\gamma|_{\sigma'}$ with $\sigma, \sigma' \subset \mathbb{R}/\mathbb{Z}$. The bigon is a **singular bigon** if the defining segments σ, σ' are disjoint and a **weak bigon** otherwise. Hass and Scott prove that a curve with excess self-intersection on an orientable surface must have a singular monogon or a singular bigon (see Figure 55 for the classification of bigons).

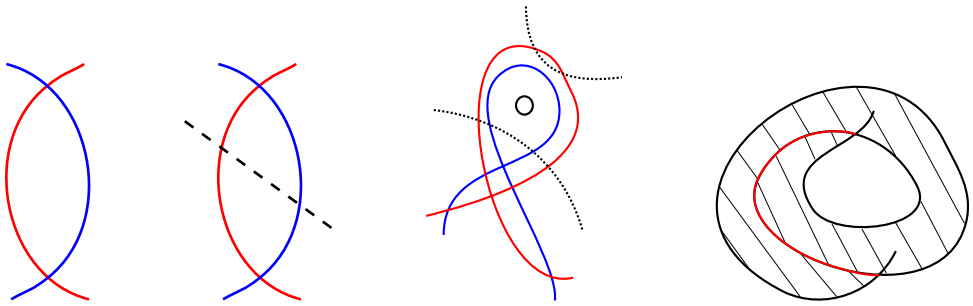

Figure 55: The bigons: an innermost, an embedded, a singular and a non-singular.

Their result directly suggests an algorithm to compute the geometric intersection number of a curve: iteratively remove monogons or untie bigons until there is no more. The final configuration must have the minimal number of self-intersections.

Designing an efficient procedure to find monogons and bigons remains the crux of this approach. In any case, the method cannot be extended to compute the geometric intersection number of two curves since Hass and Scott give two counter-examples to the fact that two curves with excess intersections should have a singular monogon or bigon. One of their counter-examples contains a curve with excess self-intersections and the other one contains a non-primitive curve. Our counter-example, Figure 56, shows that even assuming each curve to be primitive and in minimal configuration, we may have excess intersections without singular bigons.

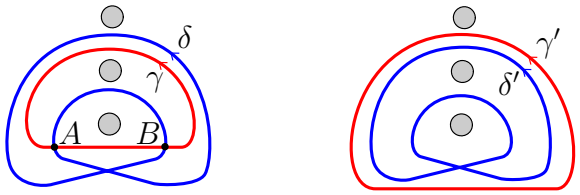

Figure 56: The plain circles represent non-contractible curves. The two curves γ and δ on the left have homotopic disjoint curves γ' and δ'. They thus have excess intersection although there is no singular bigon between the two. If $A = \gamma(0) = \delta(0)$ and $B = \gamma(u) = \delta(v)$ we nonetheless have $\delta|_{[0,v]} \sim \gamma|_{[0,1+u]}$ where $\gamma|_{[0,1+u]}$ is the concatenation of γ with $\gamma|_{[0,u]}$. In particular, $\gamma|_{[0,1+u]}$ wraps more that once around γ.

Nonetheless, it was proved by [Hass and Scott, 1994; de Graaf and Schrijver, 1997; Paterson, 2002] that starting from any configuration of curves one may reach a configuration with a minimal number of intersections by applying a finite sequence of elementary moves involving innermost monogons, bigons and trigons similar to the Reidemeister moves in knot theory. A surprising consequence was obtained by Neumann-Coto [Neumann-Coto, 2001]. Define a **cut and paste** on a family of curves by cutting the curves at some of their intersection points and glueing the resulting arcs in a different order. Neumann-Coto proves that any set of primitive curves can be brought to a homotopic set with minimal (self-)intersections by a set of cut and paste operations. Note that each intersection of two curve pieces can be re-arranged in three different manners, including the original one, by a cut and paste. Hence, if the curves have I (self-)intersections we may find a minimal configuration out of the 3^I possible re-configurations!

We also mention the algebraic approach of Gonçalves et al. [Gonçalves et al., 2005] based on previous works by Turaev [Turaev, 1979] who introduced intersection forms over the integral group ring of the fundamental group of the surface. See [Cohen and Lustig, 1987; Gonçalves et al., 2005] for more historical notes.

The last point we have to discuss is the complexity of the algorithm that we can design using those works. No explicit complexity analysis has been done but some are quite easy to evaluate. The next table summaries it, where b is the number of boundaries of the underlying surface. In our model, the entry is a map of size n and a closed walk in that map of size l. Since the considered algorithms mainly take as input a word in a representation of the π_1 it requires $O(n)$ time to obtain the word from the walk and the result can have a size of $2gl$. For the case with boundaries Birman and Series, Cohen and Lustig and [Arettines, 2015] give the best complexities. In the case with empty boundaries, the algebraic method can lead to, in the most optimistic scenario, a $O(l\ell)$ time complexity. It is interesting to note that no algorithm from the bibliography gives a polynomial time for the computation of a representative in that case.

Before	Simple	Number	Representative
$b > 0$	$O(n + (g\ell)^2)$	$O(n + (g\ell)^2)$	$O(n + (g\ell)^4)$
$b = 0$	$O(n + (g\ell)^5)$	$O(n + (g\ell)^5)$	exponential

3.3 Notations

For simplicity, we shall only consider orientable surfaces in this chapter.

Combinatorial curves Let c be an edge-path. If c is closed, we write $c(i)$, $i \in \mathbb{Z}/|c|\mathbb{Z}$, for the vertex of index i of c and $c[\bar{i}, \overline{i+1}]$ for the arc joining $c(i)$ to $c(i+1)$. For convenience we set $c[i+1, i] = c[i, i+1]^{-1}$ to allow the traversal of c in reverse direction. In order to differentiate the arcs with their occurrences we denote by $[i, i \pm 1]_c$ the corresponding occurrence of the arc $c[i, i \pm 1]$ in $c^{\pm 1}$, where c^{-1} is obtained by traversing $c^1 := c$ in the opposite direction. More generally, for any non-negative integer ℓ and any sign $\varepsilon \in \{-1, 1\}$, The sequence of indices

$$(i, i+\varepsilon, i+2\varepsilon, \ldots, i+\varepsilon\ell)$$

is called an **index path** of c of length ℓ. The index path can be **forward** ($\varepsilon = 1$) or **backward** ($\varepsilon = -1$) and can be longer than c so that an index may appear more than once in the sequence. We denote this path by $[\bar{i} \overset{\varepsilon\ell}{\rightarrow}]_c$. Its **image path** is given by the arc sequence

$$c[\bar{i} \overset{\varepsilon\ell}{\rightarrow}] := (c[i, i+\varepsilon], c[i+\varepsilon, i+2\varepsilon], \ldots, c[\varepsilon(\ell-1), \varepsilon\ell]).$$

The image path of a length zero index path is just a vertex.

Combinatorial crossings Given an immersion $\mathcal{I}$ of two combinatorial closed curves c and d we define a **double point** of (c,d) as a pair of indices $(\bar{i},\bar{j}) \in \mathbb{Z}/|c|\mathbb{Z} \times \mathbb{Z}/|d|\mathbb{Z}$ such that $c(i) = d(j)$. Likewise, a double point of c is a pair $(\bar{i},\bar{j}) \in \mathbb{Z}/|c|\mathbb{Z} \times \mathbb{Z}/|c|\mathbb{Z}$ with $i \neq j$ and $c(i) = d(j)$ The double point $(\bar{i},\bar{j})$ is a **crossing** in $\mathcal{I}$ if the pairs of arc occurrences $([i-1,i]_c,[i,i+1]_c)$ and $([j-1,j]_d,[j,j+1]_d)$ are linked in the $\preceq_{c(i)}$-order, i.e. if they appear in the cyclic order

$$\cdots [i,i-1]_c \cdots [j,j-1]_d \cdots [i,i+1]_c \cdots [j,j+1]_d \cdots ,$$

with respect to $\preceq_{c(i)}$ or the opposite order. An analogous definition holds for a self-crossing of a single curve, taking $c = d$ in the above definition. Note that the notion of (self-)crossing is independent of the traversal directions of c and d. The number of crossings of c and d and of self-crossings of c in $\mathcal{I}$ is denoted respectively by $i_{\mathcal{I}}(c,d)$ and $i_{\mathcal{I}}(c)$.

We define the **combinatorial self-crossing number** of c, denoted by $i(c)$, as the minimum of $i_{\mathcal{I}}(c')$ over all the combinatorial immersions $\mathcal{I}$ of any combinatorial curve c' freely homotopic to c. The **combinatorial crossing number** of two combinatorial curves c and d is defined the same way taking into account crossings between c and d only.

Lemma 3.3.1. *The combinatorial (self-)crossing number coincides with the geometric (self-)intersection number, i.e. for every combinatorial closed curves c and d:*

$$i(c) = i(\rho(c)) \qquad \text{and} \qquad i(c,d) = i(\rho(c),\rho(d)).$$

Proof. Every combinatorial immersion $\mathcal{I}$ of c can be realized by a continuous curve $\gamma \sim \rho(c)$ with the same number of self-intersections as the number of self-crossings of $\mathcal{I}$. To construct γ we consider small disjoint disks centered at the images of the vertices of G in S. We connect those vertex disks by disjoint strips corresponding to the edges of G. For every arc a of G we draw inside the corresponding edge strip parallel curve pieces labelled by the occurrences of a or a^{-1} in c in the left-to-right order $\preceq_a$. The endpoints of those curve pieces appear on the boundary of the vertex disks in the circular vertex orders induced by $\mathcal{I}$. It remains to connect those endpoints by straight line segments inside each disk (via a parametrization over a Euclidean disk) according to the labels of their incident curve pieces. The resulting curve γ can be homotoped to $\rho(c)$ and its self-intersections may only appear inside the vertex disks. Clearly, an intersection of two segments of γ in a disk corresponds to linked pairs of arc occurrences, i.e. to a combinatorial crossing, and vice-versa. It follows that $i(c) \geq i(\rho(c))$. To prove the reverse inequality we show that for any continuous curve γ in generic position there exists a combinatorial immersion

of a curve c' such that $\rho(c') \sim \gamma$ and the number of self-crossings c' is at most the number of self-intersections of γ. By an isotopy we can enforce the self-intersections of γ to lie inside the vertex disks. After removing the vertex disks from S, we are left with a set of disjoint and simple pieces of γ that can be isotoped inside the edge strips and vertex disks without introducing any new intersections. If a piece of curve in a strip has its endpoints on the same end of the strip we can further isotope the piece inside the incident vertex disk. This way all the curve pieces join the two ends of their strip and can be ordered from left-to-right inside each strip (assuming a preferred direction of the strip). Replacing each curve piece by an arc occurrence we thus define a combinatorial immersion of a curve c' with the required properties. The same constructions apply to an immersion of two curves showing that $i(c, d) = i(\rho(c), \rho(d))$. □

3.4 Geodesics

The canonical form of a curve c is the rightmost combinatorial geodesic homotopic to it (see Section 2.3.3). The definitions of a geodesic and of a canonical form extend trivially to paths. In particular, the canonical form of a path is its unique rightmost homotopic geodesic and is characterized as for closed curves.

Theorem 3.4.1. *The canonical form of a combinatorial path is the unique homotopic path that contains no spurs or brackets and whose turning sequence contains no -1's.*

Although we cannot claim in general the uniqueness of geodesics in a homotopy class, homotopic geodesics are almost equal and have the same length. Specifically, define a **(quad) staircase** as a planar sequence of quads obtained by stitching an alternating sequence of rows and columns of quads to form a staircase. Assuming that the staircase goes up from left to right, we define the **initial tip** of a quad staircase as the lower left vertex of the first quad in the sequence. The **final tip** is defined as the upper right vertex of the last quad. A **closed staircase** is obtained by identifying the two vertical arcs incident to the initial and final tips of a staircase.

Theorem 3.4.2. *Let c, d be two non-trivial homotopic combinatorial geodesics. If c, d are closed curves, then they label the two boundary cycles of an annular diagram composed of a unique closed staircase or of an alternating sequence of paths (possibly reduced to a vertex) and quad staircases connected through their tips. Likewise, if c, d are paths, then the closed curve $c \cdot d^{-1}$ labels the boundary of a disk diagram composed of an alternating sequence of paths (possibly reduced to a vertex) and quad staircases connected through their tips.*

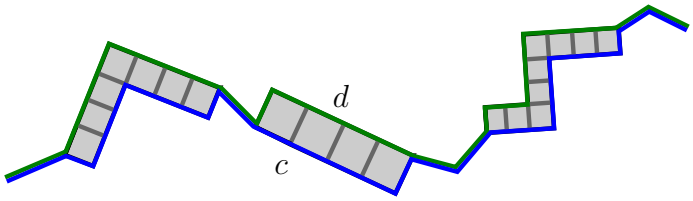

Figure 57: A disk diagram for two homotopic paths c and d composed of paths and staircases.

Proof. We only detail the proof when c, d are paths. See Figure 57. The similar case of closed curves is covered in [Erickson and Whittelsey, 2013]. By Proposition 2.3.2, $c \cdot d^{-1}$ is the label of the facial walk of the perforated face of a disk diagram Δ. This diagram has a cactus-like structure composed of 2-cells subdivided into quads and connected by trees. A vertex v such that $\Delta \backslash v$ is not connected is called a **cut vertex**. We also consider as cut vertices the endpoints of c in Δ. A 2-cell of Δ must have more than one cut vertex on its boundary. Otherwise, this boundary is entirely labelled by a subpath of either c or d and Theorem 2.3.6 implies the existence of four brackets, one of which (in fact two) must avoid the cut vertex, hence be contained in the interior of this subpath. This would contradict that c and d are geodesic. Moreover, because c and d have no spur, Δ cannot have more than two degree one vertices. It follows that Δ is an alternating sequence of paths and 2-cells. In particular, each 2-cell has exactly two cut vertices. No 2-cell in this sequence has an interior vertex. For otherwise, by the second part of Theorem 2.3.7, the boundary of this 2-cell would contain five brackets one of which would be contained in the interior of either c or d and this would again contradict that c and d are geodesic. It follows that the dual of a 2-cell, viewed as an assembling of quads, is a tree. We finally remark that this tree must be a path with a staircase shape. Indeed, any other shape would imply the existence of a bracket in either c or d. □

Corollary 3.4.3. *With the hypothesis of Theorem 3.4.2, c and d have equal length which is minimal among homotopic curves.*

Corollary 3.4.4. *A combinatorial geodesic has no non-trivial index path whose image path is contractible.*

The next two remarks follow directly from the characterization of geodesics and canonical forms in terms of spurs, brackets and turns.

Remark 3.4.5. *The image path of any index path of a combinatorial geodesic is geodesic. If the combinatorial geodesic is in canonical form, so is the image path.*

Remark 3.4.6. *Likewise, any power c^k of a combinatorial closed geodesic c is also a combinatorial geodesic. Moreover, if c is in canonical form, so is c^k.*

Let c_R and c_L^{-1} be canonical paths such that $c_R \sim c_L$. In other words, c_L is the leftmost geodesic homotopic to c_R. By Theorem 3.4.2, there is a disk diagram Δ composed of quad staircases and paths and whose left and right boundaries Δ_L and Δ_R are labelled by c_L and c_R respectively. The next technical Lemma will be used in Proposition 3.7.1 to analyze the intersection of canonical curves. A **spoke** is a non-boundary edge of Δ.

Lemma 3.4.7. *Let v, w be two vertices, one on each side of Δ. Then Δ contains a path p from v to w labelled by a canonical path. Moreover, p can be uniquely decomposed as either $\lambda.\rho$, $\rho.\lambda$, $\lambda.e.\rho$ or $\rho.e.\lambda$, where*

1. *λ is a subpath (possibly reduced to a vertex) of Δ_L or Δ_L^{-1},*
2. *ρ is a subpath (possibly reduced to a vertex) of Δ_R or Δ_R^{-1},*
3. *e is a spoke,*
4. *if λ is a subpath of Δ_L^{-1} of positive length then $p \cap \Delta_L^{-1} = \lambda$,*
5. *if ρ is a subpath of Δ_R of positive length then $p \cap \Delta_R = \rho$.*
6. *if ρ is a subpath of Δ_R^{-1} and λ is a subpath of Δ_L then either ρ is reduced to a vertex and $p \cap \Delta_R = \rho$ or λ is reduced to a vertex and $p \cap \Delta_L^{-1} = \lambda$.*

Proof. We assume that v is on the left side and w on the right side of Δ, the other case being symmetric. Let i, j be such that $v = \Delta_L(i)$ and $w = \Delta_R(j)$. We first consider the case where $j \geq i$. If Δ_L and Δ_R coincide at v, then we trivially obtain the desired decomposition as $p = \Delta_L[\bar{i} \overset{0}{\rightarrow}].\Delta_R[\bar{i} \overset{j-i}{\rightarrow}]$. Otherwise, v may be incident to $0, 1$ or 2 spokes.

- If v is not incident to a spoke then $\Delta_L(i-1)$ is either the initial tip of a staircase or is incident to a spoke e. We set $q = \Delta_L[\bar{i} \overset{-1}{\rightarrow}].\Delta_R[\overline{i-1} \overset{j-i+1}{\rightarrow}]$ in the first case and $q = \Delta_L[\bar{i} \overset{-1}{\rightarrow}].e.\Delta_R[\overline{i-1} \overset{j-i+1}{\rightarrow}]$ otherwise. The path q has no spurs or -1 turns but may start with a bracket. If not, by the characterization of Theorem 3.4.1, q is already canonical and we can set $p = q$. Otherwise, we short cut the bracket in q to obtain a canonical path p satisfying the above points 1 to 5. See Figure 58.

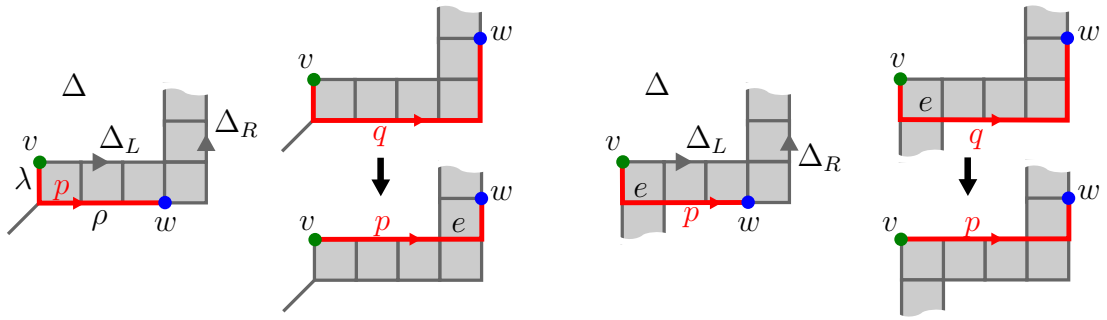

Figure 58: The canonical path p from v to w when v is not incident to a spoke.

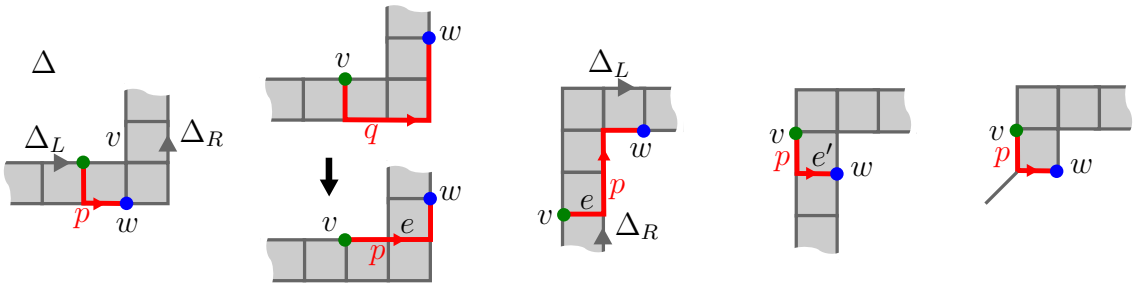

Figure 59: The canonical path p from v to w when v is incident to exactly one spoke.

- If v is incident to exactly one spoke e, then e connects v to either $\Delta_R(i-1)$ or $\Delta_R(i+1)$. In the former case we set $q = e.\Delta_R[\overline{i-1}\ ^{j-i+1}\!\!\rightarrow]$. As above, q may be canonical or starts with a bracket and we easily obtain the path p with the desired properties. See Figure 59. If e is incident to $\Delta_R(i+1)$ and $j > i$ the path $p = \Delta_L[\bar{i}\ ^{0}\!\!\rightarrow].e.\Delta_R[\overline{i+1}\ ^{j-i-1}\!\!\rightarrow]$ has the required properties. When e is incident to $\Delta_R(i+1)$ and $j = i$, then either $\Delta_L(i-1)$ is incident to a spoke e' and we must have $p = \Delta_L[\bar{i}\ ^{-1}\!\!\rightarrow].e'.\Delta_R[\bar{i}\ ^{0}\!\!\rightarrow]$, or $\Delta_L(i-1)$ must be the initial tip of a staircase and we must have $p = \Delta_L[\bar{i}\ ^{-1}\!\!\rightarrow].\Delta_R[\overline{i-1}\ ^{1}\!\!\rightarrow]$.

- If v is incident to two spokes, then one of them, e_-, connects v to $\Delta_R(i-1)$ and the other e_+ connects v to $\Delta_R(i+1)$. We can directly set $p = e_-.\Delta_R[\overline{i-1}\ ^{1}\!\!\rightarrow]$ if $j = i$ and $p = e_+.\Delta_R[\overline{i+1}\ ^{j-i-1}\!\!\rightarrow]$ if $j > i$.

For Point 6 in the Lemma, we note that we cannot have $p = \lambda.\rho$ or $p = \lambda.e.\rho$ with both λ and ρ being subpaths of positive length of Δ_L and Δ_R^{-1} respectively. Indeed, p would have a spur or a $\bar{1}$ turn in the first case and a bracket in the other case.

Moreover, if ρ is reduced to a vertex and λ is a subpath of Δ_L, we may assume that the intersection $p \cap \Delta_R$ is reduced to ρ. Otherwise, the canonical path between any other intersection point and ρ would have to follow Δ_R and we could express p so that ρ is a subpath of positive length of Δ_R. An analogous argument holds to show that we can assume $p \cap \Delta_L^{-1} = \lambda$ if λ is reduced to a vertex.

We next consider the case $i > j$. When Δ_L and Δ_R coincide at w, we obtain the desired decomposition as $p = \Delta_L[\bar{i} \overset{j-i}{\to}].\Delta_R[\bar{i} \overset{0}{\to}]$. Otherwise, w may be incident to $0, 1$ or 2 spokes. Similar arguments as in the case $j \geq i$ allow to conclude the proof. □

3.5 Our Strategy for Counting Intersections

Following Poincaré's original approach we represent the surface S as the hyperbolic quotient surface $\mathbb{D}/\Gamma$ where Γ is a discrete group of hyperbolic motions of the Poincaré disk $\mathbb{D}$. We denote by $p : \mathbb{D} \to \mathbb{D}/\Gamma = S$ the universal covering map. Any closed curve $\alpha : \mathbb{R}/\mathbb{Z} \to S$ gives rise to its infinite power $\alpha^\infty : \mathbb{R} \to \mathbb{R}/\mathbb{Z} \to S$ that wraps around α infinitely many times. A **lift** of α is any curve $\tilde{\alpha} : \mathbb{R} \to \mathbb{D}$ such that $p \circ \tilde{\alpha} = \alpha^\infty$ where the parameter of $\tilde{\alpha}$ is defined up to an integer translation (we thus identify the curves $t \mapsto \tilde{\alpha}(t + k)$, $k \in \mathbb{Z}$). Note that $p^{-1}(\alpha)$ is the union of all the images $\Gamma \cdot \tilde{\alpha}$ of $\tilde{\alpha}$ by the motions in Γ. The curve $\tilde{\alpha}$ has two limit points on the boundary of $\mathbb{D}$ which can be joined by a unique hyperbolic line L. The projection $p(L)$ wraps infinitely many times around the unique geodesic homotopic to α. In particular, the limit points of $\tilde{\alpha}$ are independent of the chosen representative in the homotopy class of α.

No two motions of Γ have a limit point in common unless they are powers of the same motion. This can be used to show that when α is primitive, its lifts are uniquely identified by their limit points [Farb and Margalit, 2011]. Let α and β be two primitive curves. We fix a lift $\tilde{\alpha}$ of α and denote by $\tau \in \Gamma$ the hyperbolic motion sending $\tilde{\alpha}(0)$ to $\tilde{\alpha}(1)$. Let $\Gamma \cdot \tilde{\beta}$ be the set of lifts of β. We consider the subset of lifts

$$B = \{\tilde{\beta}' \in \Gamma \cdot \tilde{\beta} \mid \text{ the limit points of } \tilde{\beta}' \text{ and } \tilde{\alpha} \text{ alternate along } \partial\mathbb{D}\}.$$

Lemma 3.5.1 ([Reinhart, 1962])**.**

$$i(\alpha, \beta) = |B/\tau|,$$

where B/τ is the set of equivalence classes of lifts generated by the relations $\tilde{\beta}' \sim \tau(\tilde{\beta}')$.

Proof. Put $I(\alpha,\beta) = \{(t,t') \mid t,t' \in \mathbb{R}/\mathbb{Z} \text{ and } \alpha(t) = \beta(t')\}$. define a map $\varphi : I(\alpha,\beta) \to B/\tau$ as follows. Given $(u \bmod 1, v \bmod 1) \in I(\alpha,\beta)$ there is, by the unique lifting property of coverings, a unique lift $\tilde{\beta}'$ of β that satisfies $\tilde{\beta}'(v) = \tilde{\alpha}(u)$. We set $\varphi(u \bmod 1, v \bmod 1)$ to the class of this lift. Note that changing u to $u+k$, $k \in \mathbb{Z}$, leads to the lift $\tau^k(\tilde{\beta}')$, so that φ is well-defined. The map φ is onto. Indeed, if $\tilde{\beta}' \in B$ then $\tilde{\beta}'$ and $\tilde{\alpha}$ must intersect at some point $\tilde{\beta}'(v) = \tilde{\alpha}(u)$. It follows that $\varphi(u \bmod 1, v \bmod 1)$ is the class of $\tilde{\beta}'$. As an immediate consequence,

$$|I(\alpha,\beta)| \geq |B/\tau|$$

When α and β are geodesics all their lifts are hyperbolic lines and φ is a bijection. We conclude that $|I(\alpha,\beta)|$ is minimized among all homotopic curves so that $i(\alpha,\beta) = |B/\tau|$. □

When α and β are hyperbolic geodesics, their lifts being hyperbolic lines have alternating limit points exactly when they have a non-empty intersection and they have a unique intersection point in that case. This point projects to a crossing of α and β that actually identifies the corresponding element of B/τ. When α and β are not geodesic the situation is more ambiguous and their lifts may have multiple intersection points. Those intersection points project to crossings of α and β, so that the elements of B/τ are now identified with subsets of crossing points (with odd cardinality) rather than single crossing points. The induced partition is generated by the following relation: two crossings are equivalent if they are connected by a pair of homotopic subpaths of α and β, namely one of the two subpaths of α and one of the two subpaths of β cut by the two crossings. Indeed, if the two crossings are projections of intersections of two lifts, then the paths between those intersections in each lift project to homotopic paths. Conversely, homotopic paths lift to paths with common endpoints that can be seen as subpaths of two intersecting lifts. In order to compute the above partition, we thus essentially need an efficient procedure for testing if two paths are homotopic. This homotopy test can indeed be performed in linear time according to Theorem 2.3.8. Since a combinatorial curve of length ℓ may have $O(\ell^2)$ crossings, we directly obtain an algorithm with time complexity $O(\ell^5)$ to compute the above partition.

When dealing with combinatorial (canonical) geodesics, the situation is more constrained and somehow intermediate between the ideal hyperbolic case and the most general situation. Thanks to Theorem 3.4.2, we know that homotopic paths must stay parallel and at distance at most one. See Figure 60. This allows us to identify equivalent crossings more efficiently. These ideas are formalized in the next sections.

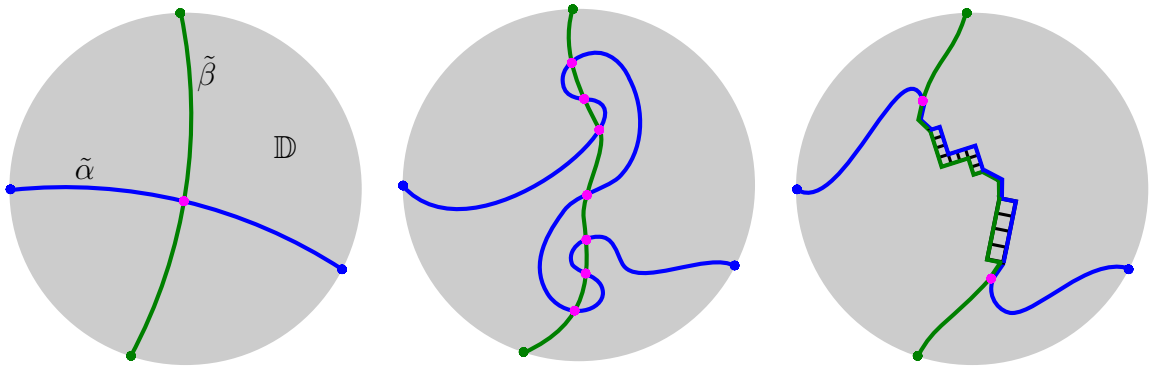

Figure 60: Left, two intersecting hyperbolic lines. Middle, two lifts of non-geodesic curves may intersect several times. Right, lifts of combinatorial geodesics.

3.6 Crossing Double-Paths

Let c, d be two combinatorial closed curves on a combinatorial surface. A **double-path** of (c, d) of length ℓ is a pair of forward index paths $([\overline{i} \overset{\ell}{\to}]_c, [\overline{j} \overset{\ell}{\to}]_d)$ with the same image path $c[\overline{i} \overset{\ell}{\to}] = d[\overline{j} \overset{\ell}{\to}]$. If $\ell = 0$ then the double path is just a double point. A double path of c is defined similarly, taking $c = d$ and assuming $i \neq j$. The next Lemma follows from Remark 3.4.5.

Lemma 3.6.1. *Let $[\overline{i} \overset{\ell}{\to}]_c$ and $[\overline{j} \overset{k}{\to}]_d$ be forward index paths of two canonical curves c and d such that the image paths $c[\overline{i} \overset{\ell}{\to}]$ and $[\overline{j} \overset{k}{\to}]$ are homotopic. Then $k = \ell$ and $([\overline{i} \overset{\ell}{\to}]_c, [\overline{j} \overset{\ell}{\to}]_d)$ is a double path.*

A double path $([\overline{i} \overset{\ell}{\to}]_c, [\overline{j} \overset{\ell}{\to}]_d)$ gives rise to a sequence of $\ell + 1$ double points $(\overline{\imath + k}, \overline{\jmath + k})$ for $k \in [0, \ell]$. A priori a double point could occur several times in this sequence. The next two lemmas claim that this is not possible when the curves are primitive. Recall that a curve is **primitive** if its homotopy class cannot be expressed as a proper power of another class.

Lemma 3.6.2. *A double path of a primitive combinatorial curve c cannot contain a double point more than once in its sequence. In particular, a double path of c must be strictly shorter than c.*

Proof. Suppose that a double path $\mathbb{P}$ of c contains two occurrences of a double point (i, j). Because the couples (i, j) and (j, i) represent the same double point there are two cases to consider.

- If $\mathbb{P}$ contains the couple (i, j) twice then it must contain a subsequence of length $|c|$ starting with (i, j). We thus have $c[\bar{i} \overset{|c|}{\to}] = c[\bar{j} \overset{|c|}{\to}]$. This implies that c is equal to some non-trivial circular permutation of itself. It is a simple exercise to check that c must then be a proper power of some other curve, contradicting that c is primitive.

- Otherwise $\mathbb{P}$ contains (i, j) and (j, i). Let ℓ be the distance between these two occurrences in $\mathbb{P}$. We thus have $c[\bar{i} \overset{\ell}{\to}] = c[\bar{j} \overset{|c|-\ell}{\to}]$ from which we deduce that c is a square (and $\ell = |c|/2$), contradicting that c is primitive.

□

Lemma 3.6.3. *Let c and d be two non-homotopic primitive combinatorial curves. A double path of (c, d) cannot contain a double point more than once in its sequence. Moreover, the length of a double path of (c, d) must be less than $|c| + |d| - 1$.*

Proof. Suppose that a double path of (c, d) contains two occurrences of a double point. After shortening the double path if necessary, we may assume that these two occurrences are the first and the last double points of the double path. Its length must accordingly be a nonzero integer multiple p of $|c|$ as well as a nonzero integer multiple q of $|d|$. It follows that for some circular permutations c' of c and d' of d we have $c'^p = d'^q$. By a classical result of combinatorics on words [Lothaire, 1997, Prop. 1.3.1] this implies that c' and d' are powers of a same curve, in contradiction with the hypotheses in the lemma. In fact, by a refinement due to Fine and Wilf [Lothaire, 1997, Prop. 1.3.5] it suffices that c'^p and d'^q have a common prefix of length $|c|+|d|-1$ to conclude that c' and d' are powers of a same curve. This proves the second part of the lemma. □

A double path whose index paths cannot be extended is a **maximal double path**. As an immediate consequence of Lemmas 3.6.2 and 3.6.3 we have:

Corollary 3.6.4. *The maximal double paths of a primitive curve or of two primitive curves in canonical form induce a partition of the double points of the curves.*

Let (i, j) and $(\overline{i + \ell}, \overline{j + \ell})$ be the first and the last double points of a maximal double path of (c, d), possibly with $c = d$. When $\ell \geq 1$ the arcs $c[i, i - 1]$, $d[j, j - 1]$, $c[i, i + 1]$ must be pairwise distinct because canonical curves have no spurs, and similarly for the three arcs $c[i + \ell, i + \ell + 1]$, $d[j + \ell, j + \ell + 1]$, $c[i + \ell, i + \ell - 1]$. We declare the maximal double path to be a **crossing double path** if the circular ordering of the first three arcs at $c(i)$ and the circular ordering of the last three arcs

at $c(i+\ell)$ are either both clockwise or both counterclockwise with respect to the rotation system of the system of quads. When $\ell = 0$, that is when the maximal double path is reduced to the double point[1] (i, j), we require that the arcs $c[i, i-1], d[j, j-1], c[i, i+1], d[j, j+1]$ are **pairwise distinct** and appear in this circular order, or its opposite, around the vertex $c(i) = d(j)$.

3.7 Counting Intersections Combinatorially

Let c, d be primitive combinatorial curves such that d is canonical and let c_R and c_L^{-1} be the canonical curves homotopic to c and c^{-1} respectively. We denote by Δ the annular diagram corresponding to c_R and c_L. When the two boundaries Δ_R and Δ_L of Δ have a common vertex we implicitly assume that c_R and c_L are indexed so that this vertex corresponds to the same index along Δ_R and Δ_L. We consider the following set of double paths:

- $\mathcal{D}_+$ is the set of crossing double paths of positive length of c_R and d,
- $\mathcal{D}_0$ is the set of crossing double paths (i, j) of zero length of c_R and d such that either
  - the two boundaries of Δ coincide at $\Delta_L(i) = \Delta_R(i)$ and $d[j-1, j] = c_L[i-1, i]$ or $d[j, j+1] = c_L[i, i+1]$, or
  - one of $d[j, j-1]$ or $d[j, j+1]$ is the label of a spoke $(\Delta_R(i), \Delta_L(i'))$ of Δ and $d[j-2, j-1] = c_L[i'-1, i']$ in the first case or $d[j+1, j+2] = c_L[i', i'+1]$ in the other case.
- $\mathcal{D}_-$ is the set of crossing double paths $([\overline{i} \xrightarrow{\ell}]_{c_L^{-1}}, [\overline{j} \xrightarrow{\ell}]_d)$ $(\ell \geq 0)$ of c_L^{-1} and d such that **none** of the following situations occurs:
  - the two boundaries of Δ coincide at $\Delta_L^{-1}(i) = \Delta_R(i')$ and $d[j-1, j] = c_R[i'-1, i']$,
  - the two boundaries of Δ coincide at $\Delta_L^{-1}(i+\ell) = \Delta_R(i')$ and $d[\overline{j+\ell}, \overline{j+\ell+1}] = c_R[\overline{i'}, \overline{i'+1}]$,
  - $d[j-1, j]$ is the label of a spoke $(\Delta_L^{-1}(i), \Delta_R(i'))$ of Δ and $d[j-2, j-1] = c_R[i'-1, i']$,

[1] Those crossing double points should not be confused with the combinatorial crossings of an immersion as defined in Section 3.3. Which notion of crossings is used should always be clear from the context.

- $d[j+\ell, j+\ell+1]$ is the label of a spoke $(\Delta_L^{-1}(i+\ell), \Delta_R(i'))$ of Δ and $d[j+\ell+1, j+\ell+2] = c_R[i', i'+1]$.

Those definitions allow the case $c \sim d$, recalling that the index paths of a double path of c must be distinct by definition. Referring to Section 4.3, we view the underlying surface of the system of quads Σ as a quotient $\mathbb{D}/\Gamma$ of the Poincaré disk. The system of quads lifts to a quadrangulation of $\mathbb{D}$ and the lifts of a combinatorial curve in Σ are combinatorial bi-infinite paths in this quadrangulation. By Remark 3.4.6, if the combinatorial curve is geodesic (resp. canonical) so are its lifts. In this case, each lift is simple by Corollary 3.4.4. We fix a lift $\tilde{c_R}$ of c_R and consider the set B/τ of Lemma 3.5.1 corresponding to the classes of lifts of d whose limit points alternate with the limit points of $\tilde{c_R}$ along $\partial\mathbb{D}$.

Proposition 3.7.1. *B/τ is in 1-1 correspondence with the disjoint union $\mathcal{D}_+ \cup \mathcal{D}_0 \cup \mathcal{D}_-$.*

Proof. Let $\tilde{c_L}$ be the lift of c_L with the same limit points as $\tilde{c_R}$. These two lifts project onto the boundaries of the annular diagram Δ and thus form an infinite strip $\tilde{\Delta}$ of width at most 1 in $\mathbb{D}$ composed of paths and quad staircases (possibly a single infinite staircase). We shall define a correspondence between B/τ and $\mathcal{D}_+ \cup \mathcal{D}_0 \cup \mathcal{D}_-$. To this end we consider a lift $\tilde{d}$ of d whose limit points alternate with those of $\tilde{c_R}$. In other words, $\tilde{d} \in B$. The lift $\tilde{d}$ must cross $\tilde{\Delta}$. Let i and j be respectively the smallest and largest index k such that $\tilde{d}(k)$ is in $\tilde{\Delta}$. By Remark 3.4.5, the corresponding subpath $\tilde{d}[i,j]$ of $\tilde{d}$ is canonical. Since $\mathbb{D}$ is simply connected, $\tilde{d}[i,j]$ is homotopic to any path joining the same extremities and we can apply Lemma 3.4.7 to show that $\tilde{d}[i,j]$ is actually contained in $\tilde{\Delta}$ and that it can be decomposed as either $\lambda.\rho$, $\rho.\lambda$, $\lambda.e.\rho$ or $\rho.e.\lambda$ where e is a spoke of $\tilde{\Delta}$, $\rho = \tilde{c_R}[\overline{a} \xrightarrow{r}]$ and $\lambda = \tilde{c_L}[\overline{b} \xrightarrow{\ell}]$ for some $a, b, r, \ell \in \mathbb{Z}$.

- If $r > 0$ then by Point 5 of Lemma 3.4.7 we have $\tilde{d} \cap \tilde{c_R} = \rho$ so that this intersection defines a maximal double path of $\tilde{d}$ and $\tilde{c_R}$. It must be crossing since $\tilde{d} \in B$. Its projection on Σ is a crossing double path of length $r > 0$ of d and c_R, hence in $\mathcal{D}_+$, to which we map $\tilde{d}$.

- If $r = 0$ and $\ell > 0$, then Point 6 of Lemma 3.4.7 implies $\tilde{d} \cap \tilde{c_R} = \rho$. As above, ρ must define a crossing double path of length zero of $\tilde{d}$ and $\tilde{c_R}$. We map $\tilde{d}$ to the projection of this crossing double point on Σ and remark that this projection is in $\mathcal{D}_0$.

- Otherwise, we must have $r \leq 0$ and $\ell \leq 0$ by Point 6 of Lemma 3.4.7. If $\ell < 0$ then Point 4 of Lemma 3.4.7 implies $\tilde{d} \cap \tilde{c_L}^{-1} = \lambda$. This intersection defines

a crossing double path of $\tilde{d}$ and $\tilde{c_L}^{-1}$ and we map $\tilde{d}$ to its projection on Σ. If $\ell = 0$ and $r < 0$ then Point 6 of Lemma 3.4.7 implies $\tilde{d} \cap \tilde{c_L^{-1}} = \lambda$, which also holds true if $r = \ell = 0$. In both cases λ corresponds to a crossing double path of length zero of $\tilde{d}$ and $\tilde{c_L^{-1}}$ and we map $\tilde{d}$ to its projection on Σ. We finally remark that this last projection or the above one belong to $\mathcal{D}_-$.

Because $\tilde{\Delta}$ is left globally invariant by τ (the hyperbolic motion that sends $\tilde{c_R}(0)$ to $\tilde{c_R}(|c_R|)$), we have $\tau(\tilde{d}) \cap \tilde{\Delta} = \tau(\tilde{d}) \cap \tau(\tilde{\Delta}) = \tau(\tilde{d} \cap \tilde{\Delta})$. It follows that $\tilde{d}$ and $\tau(\tilde{d})$ are mapped to the same crossing double path by the above rules. We thus have a well defined map $B/\tau \to \mathcal{D}_+ \cup \mathcal{D}_0 \cup \mathcal{D}_-$. The uniqueness of the decomposition in Lemma 3.4.7 implies that this map is 1-1. In order to check that the map is onto we consider a maximal crossing double path $\mathbb{P} = ([\overline{i \mod |c_R|} \xrightarrow{\ell}]_{c_R}, [\overline{j \mod |d|} \xrightarrow{\ell}]_d)$ of c_R and d in $\mathcal{D}_+$. By the unique lift property of coverings there is a unique lift $\tilde{d}$ of d such that $\tilde{d}(j) = \tilde{c_R}(i)$ and $\mathbb{P}$ lifts to a crossing double path of $\tilde{d}$ and $\tilde{c_R}$. By Lemma 3.6.1 this double path is the only intersection of $\tilde{d}$ and $\tilde{c_R}$ so these lifts must have alternating limit points. In other words $\tilde{d}$ is in B and is mapped to $\mathbb{P}$. A similar argument applies to the crossing double paths of $\mathcal{D}_0$ and $\mathcal{D}_-$. □

This leads to a simple algorithm for computing combinatorial crossing numbers.

Corollary 3.7.2. *Let c, d be primitive curves of length at most ℓ on a combinatorial surface with complexity n. The crossing numbers $i(c, d)$ and $i(c)$ can be computed in $O(n + \ell^2)$ time.*

Proof. We may assume that the surface is a system of quads. By Theorem 2.3.8 we may compute the canonical forms of c, c^{-1} and d in $O(\ell)$ time. According to Proposition 3.7.1, we have

$$i(c, d) = |\mathcal{D}_+| + |\mathcal{D}_0| + |\mathcal{D}_-|$$

The set $\mathcal{D}_+$ can be constructed in $O(\ell^2)$ time. Indeed, since the maximal double paths of c and d form disjoint sets of double points by Corollary 3.6.4, we just need to traverse the grid $\mathbb{Z}/|c|\mathbb{Z} \times \mathbb{Z}/|d|\mathbb{Z}$ and group the double points into maximal double paths. Those correspond to diagonal segments in the grid that can be computed in time proportional to the size of the grid. We can also determine which double paths are crossing in the same amount of time. Likewise, we can construct the sets $\mathcal{D}_0$ and $\mathcal{D}_-$ in $O(\ell^2)$ time. We can also compute $i(c)$ in quadratic time using that $i(c, c) = 2i(c)$. □

3.8 Non-Primitive Curves and Proof of Theorem 3.1.1

In order to finish the proof of Theorem 3.1.1, we need to tackle the case of non-primitive curves. Thanks to canonical forms, computing the primitive root of a curve becomes extremely simple.

Lemma 3.8.1. *Let c be a combinatorial curve of length $\ell > 0$ in canonical form. A primitive curve d such that c is homotopic to d^k for some integer k can be computed in $O(\ell)$ time.*

Proof. By Theorem 2.3.8, we may assume that c and d are in canonical form. By Remark 3.4.6, the curve d^k is also in canonical form. The uniqueness of the canonical form implies that $c = d^k$, possibly after some circular shift of d. It follows that d is the smallest prefix of c such that c is a power of this prefix. It can be found in $O(\ell)$ time using a variation of the Knuth-Morris-Pratt algorithm to find the smallest period of a word [Knuth et al., 1977]. □

The geometric intersection number of non-primitive curves is related to the geometric intersection number of their primitive roots. The next result is rather intuitive, see Figure 61, and is part of the folklore although we could only find references in some relatively recent papers.

Proposition 3.8.2 ([de Graaf and Schrijver, 1997; Gonçalves et al., 2005])**.** *Let c and d be primitive curves and let p, q be integers. Then,*

$$i(c^p) = p^2 \times i(c) + p - 1 \quad \text{and} \quad i(c^p, d^q) = \begin{cases} 2pq \times i(c) & \text{if } c \sim d \text{ or } c \sim d^{-1}, \\ pq \times i(c, d) & \text{otherwise.} \end{cases}$$

Proof of Theorem 3.1.1. Let c, d and Σ be the two combinatorial curves and the combinatorial surface as in the Theorem. We first assume that the Σ has negative Euler characteristic. We can compute the canonical forms of c and d in $O(\ell)$ time after $O(n)$ time preprocessing by Lemma **??**. Thanks to Lemma 3.8.1 we can further determine primitive curves c' and d' and integers p, q such that $c \sim c'^p$ and $d \sim d'^p$ in $O(\ell)$ time. We can compute $i(c', d')$ and $i(c')$ in $O(\ell^2)$ time according to Corollary 3.7.2. We finally use the formulas in the previous proposition to deduce $i(c, d)$ and $i(c)$ from $i(c', d')$ and $i(c')$.

If Σ is a sphere or a disk, then every curve is contractible and $i(c, d) = i(c) = 0$. If Σ is a cylinder, then every two curves can be made non crossing so that $i(c, d) = 0$ while $i(c) = p-1$. Finally, if Σ is a torus, the radial graph of the system of quads can be decomposed into two loops α, β such that $c \sim \alpha^x \cdot \beta^y$ and $d \sim \alpha^{x'} \cdot \beta^{y'}$. We may then use the classical formulas: $i(c) = \gcd(x, y) - 1$ and $i(c, d) = |\det\big((x, y), (x', y')\big)|$. □

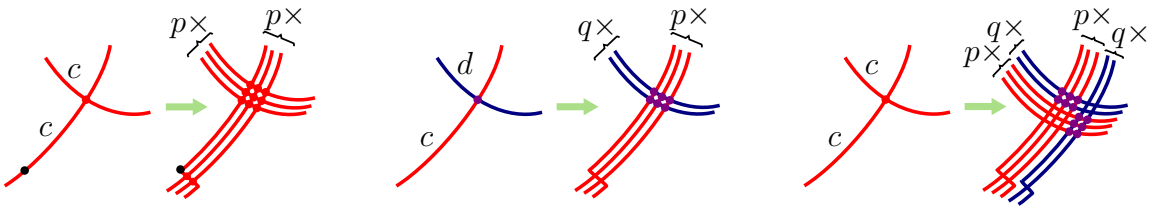

Figure 61: Left, each self-intersection of c gives rise to p^2 self-intersections of its pth power obtained by wrapping p times around c in a small tubular neighborhood. The starting basepoint also adds $p-1$ self-intersections. Middle, each intersection of c and d gives rise to pq intersections of (homotopic perturbations of) c^p and d^q. Right, when c and d are homotopic, one should count $2pq$ intersections per self-intersection of c.

3.9 Computing a Minimal Immersion

In the subsequent sections we only deal with the self-intersection number of a single curve. We thus drop the subscript c to denote an index path $[\bar{i} \xrightarrow{\ell}]$ or an arc occurrence $[\bar{i}, \overline{i+1}]$. We also implicitly consider intersections as self-intersections. By Theorem 3.1.1 we can compute the geometric intersection number of a curve efficiently. Here, we describe a way to compute a minimal immersion, that is an actual immersion with the minimal number of intersections. We refer to the combinatorial framework of Section 3.3 to describe such an immersion combinatorially. Thanks to Lemma 3.3.1, we know that this framework faithfully encodes the topological configurations.

Bigons and monogons. A **bigon** of an immersion $\mathcal{I}$ of c is a pair of index paths $([\bar{i} \xrightarrow{\ell}], [\bar{j} \xrightarrow{k}])$ whose **sides** $c[\bar{i} \xrightarrow{\ell}]$ and $c[\bar{j} \xrightarrow{k}]$ have strictly positive lengths, are homotopic, and whose **tips** (i,j) and $(i+\ell, j+k)$ are combinatorial crossings for $\mathcal{I}$. A **monogon** of $\mathcal{I}$ is an index path $[\bar{i} \xrightarrow{\ell}]$ of strictly positive length such that $(i, i+\ell)$ is a combinatorial crossing and the image path $c[\bar{i} \xrightarrow{\ell}]$ is contractible.

Proposition 3.9.1. *A combinatorial immersion of a primitive curve has excess self-crossing if and only if it contains a bigon or a monogon.*

Proof. We just need to prove the existence of a bigon or monogon for a continuous realization $\gamma : \mathbb{R}/\mathbb{Z} \to S$ of the combinatorial immersion. This bigon or monogon corresponds to a combinatorial bigon or monogon as claimed in the Proposition. To prove the topological counterpart, we first note as in Section 3.5 that the minimal

number of intersections is counted by the lifts of γ whose limit points alternate along $\partial\mathbb{D}$. It follows that γ is minimally crossing when its lifts are pairwise disjoint unless the corresponding limit points alternate in $\partial\mathbb{D}$ and they should intersect exactly once in that case. When γ is not minimally crossing there must exist two lifts intersecting more than necessary, thus forming a bigon in $\mathbb{D}$. The projection of such a bigon on S gives the desired result. Note that a statement similar to the one in the Proposition holds for an immersion of a pair of primitive curves c and d with excess crossing. □

We provide a purely combinatorial proof of the direct implication of the Proposition in the case of two curves.

Combinatorial Proof. We consider the case of an immersion $\mathcal{I}$ of two curves c and d. Suppose that $\mathcal{I}$ has p excess crossings. We shall prove the stronger claim that $\mathcal{I}$ has at least $\lceil p/2 \rceil$ bigons whose tips are pairwise distinct. Let $\mathcal{J}$ be an immersion of two curves c' and d' respectively homotopic to c and d such that $\mathcal{J}$ has no excess crossings. Consider a sequence of elementary homotopies from (c', d') to (c, d). Following the above remark, we can insert adjacent transpositions between each elementary homotopy in order to obtain a sequence of elementary moves from $\mathcal{J}$ to $\mathcal{I}$. See the above remark. The claim is trivially verified by $\mathcal{J}$. We now check that the claim remains true after each elementary move. If the move is an adjacent transposition we may assume that we exchange an arc occurrence of c with an arc occurrence of d since we only consider intersections between c and d. Without loss of generality we also assume exchanging the forward occurrences $[i, i+1]_c$ and $[j, j+1]_d$. There are three cases to consider.

1. If none of $(\bar{i}, \bar{j})$ and $(\overline{i+1}, \overline{j+1})$ is a crossing then the transposition adds two crossings that we may pair as the tips of a new bigon. We now have $p+2$ excess crossings with at least $\lceil (p+2)/2 \rceil$ bigons as required.

2. If $(\bar{i}, \bar{j})$ is a crossing but $(\overline{i+1}, \overline{j+1})$ is not and if $(\bar{i}, \bar{j})$ is paired to form the tips of a bigon, say $([\bar{i} \overset{\ell}{\to}]_c, [\bar{j} \overset{k}{\to}]_d)$, we may just replace this bigon by $([\overline{i+1} \overset{\ell-1}{\to}]_c, [\overline{j+1} \overset{k-1}{\to}]_d)$ sliding its tip $(\bar{i}, \bar{j})$ to $(\overline{i+1}, \overline{j+1})$. A similar procedure applies when $(\overline{i+1}, \overline{j+1})$ is a crossing but $(\bar{i}, \bar{j})$ is not. In each case the number of excess crossings and bigons is left unchanged.

3. It remains the case where both $(\bar{i}, \bar{j})$ and $(\overline{i+1}, \overline{j+1})$ are crossings. If none of the two is paired to form the tips of a bigon, then the transposition removes two crossings and no pairing, so that the claim remains trivially true. If exactly one of the two is paired or if the two are paired together, we loose one bigon and two crossings after the transposition so that the claim remains true. Otherwise,

there are two bigons of the form $([\overline{i} \xrightarrow{\ell}]_c, [\overline{j} \xrightarrow{k}]_d)$ and $([\overline{i+1} \xrightarrow{\ell'}]_c, [\overline{j+1} \xrightarrow{k'}]_d)$ that we can recombine to form the bigon $([\overline{i+\ell} \xrightarrow{1+\ell'-\ell}]_c, [\overline{j+k} \xrightarrow{1+k'-k}]_d)$. We again have one less bigon and two less crossings.

We now consider the application of an elementary homotopy as described before the proof. There are again three possibilities.

1. If the homotopy replaces a nonempty subpath u of a facial walk uv^{-1} by the nonempty complementary part v then no crossing may appear or disappear as we assume the immersion in good position. In particular, no crossing may use an internal vertex of u and u is either entirely included in or excluded from any side of any bigon. We can replace u by v in any bigon side where u occurs to obtain valid bigons in the new immersion after the elementary homotopy is applied. The number of excess crossings and bigons is left unchanged.

2. When u is empty in the above replacement, or when inserting a spur, we may only add crossings by pairs forming bigons with one zero-length side and the complement of u or the spur as the other side. A similar analysis as in case (2) of a transposition applies to take care of each pair.

3. When v is empty in the above replacement, or when removing a spur, we may only remove crossings by pairs and a similar analysis as in case (3) of a transposition applies to take care of each pair.

This ends the proof for the case of two curves. A similar proof holds for the existence of a bigon or a monogon in an immersion with excess self-intersection. This time the excess crossings are either paired to form bigons or left alone as the tips of monogons. □

Suppose that $\mathcal{I}$ is a combinatorial immersion of a primitive geodesic c. It cannot have a monogon by Corollary 3.4.4. Hence, according to Proposition 3.9.1 the immersion $\mathcal{I}$ is minimal unless it contains a bigon. The two sides of such a bigon may share some common part that would prevent us from swapping these two sides. In agreement with the terminology of Hass and Scott [Hass and Scott, 1985], a bigon $([\overline{i} \xrightarrow{\ell}], [\overline{j} \xrightarrow{k}])$ is said **singular** if

1. its two index paths have disjoint interiors, i.e. they do not share any arc occurrence;

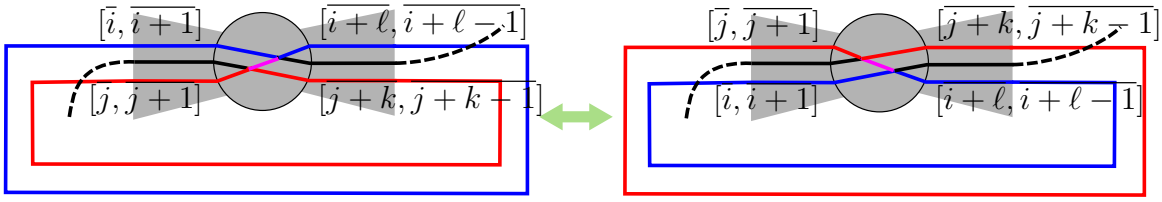

Figure 62: Right, the realization of the bigon $([\bar{i} \overset{\ell}{\to}], [\bar{j} \overset{k}{\to}])$ where $j = i + \ell$ and $[\bar{i}, \overline{i-1}]$ is in-between $[\overline{i+\ell}, \overline{i+\ell-1}]$ and $[\overline{j+k}, \overline{j+k-1}]$, and $[\overline{j+\ell}, \overline{j+\ell+1}]$ is in-between $[\bar{i}, \overline{i+1}]$ and $[\bar{j}, \overline{j+1}]$. The small purple part is at the same time the beginning of the red side of the bigon and at the end of the blue side. Swapping this bigon does not reduce the number of crossings.

2. when $j = i + \ell$ the following arc occurrences

$$[\bar{i}, \overline{i-1}], [j, j-1], [\bar{i}, \overline{i+1}], [j+k, j+k+1], [j, j+1], [j+k, j+k-1]$$

do not appear in this order or its opposite in the circular ordering induced by $\mathcal{I}$ at $c(j)$;

3. when $i = j + k$ the following arc occurrences

$$[\bar{i}, \overline{i-1}], [\bar{j}, \overline{j-1}], [\overline{i+\ell}, \overline{i+\ell-1}], [\bar{i}, \overline{i+1}], [\overline{i+\ell}, \overline{i+\ell+1}], [\bar{j}, \overline{j+1}]$$

do not appear in this order or its opposite in the circular ordering induced by $\mathcal{I}$ at $c(i)$.

Note that when c is geodesic and $k = -\ell$ there cannot be any identification between $\{i, i+\ell\}$ and $\{j, j-\ell\}$. For instance, $i = j$ implies $c[\bar{i} \overset{\ell}{\to}] \sim c[\bar{j} \overset{-\ell}{\to}] = c[\bar{i} \overset{-\ell}{\to}]$ so that $c[\overline{i-\ell} \overset{2\ell}{\to}] \sim 1$, while $j = i + \ell$ implies $c[\bar{i} \overset{\ell}{\to}] \sim c[\bar{j} \overset{-\ell}{\to}] = c^{-1}[\bar{i} \overset{\ell}{\to}]$ so that $c^2[\bar{i} \overset{\ell}{\to}] \sim 1$. In both cases c would have a monogon which would contradict Lemma 3.4.4.

When c is primitive and geodesic, we cannot have $j = i + \ell$ and $i = j + \ell$ at the same time. Otherwise, we must have $k = \ell$ by the preceding paragraph, and c would be a circular shift of $c[\bar{i} \overset{\ell}{\to}] \cdot c[\bar{j} \overset{\ell}{\to}]$ and thus homotopic to a square. Also remark that when one of the last two conditions in the definition is not satisfied, the bigon maps to a non-singular bigon in the continuous realization of $\mathcal{I}$ as in Lemma 3.3.1. See Figure 62.

When the bigon is singular we can swap its two sides by exchanging the two arc occurrences $[i + p, i + p + 1]$ and $[j + \varepsilon p, i + \varepsilon(p + 1)]$, for $0 \leq p < \ell$ and $k = \varepsilon\ell$.

Lemma 3.9.2. *Swapping the two sides of a singular bigon of an immersion of a geodesic primitive curve decreases its number of crossings by at least two.*

This is relatively obvious if one considers a continuous realization of the immersion, performs the swapping and comes back to a combinatorial immersion as in the proof of Lemma 3.3.1. We nonetheless provide a purely combinatorial proof.

Proof. Consider a singular bigon $([\overline{i} \overset{\ell}{\to}], [\overline{j} \overset{\ell}{\to}])$ of an immersion $\mathcal{I}$ of a closed primitive canonical curve c. Let $\mathcal{J}$ be the immersion after the bigon has been swapped. We shall partition the set of potential double points and show that the net change of the number of crossings with respect to $\mathcal{I}$ and $\mathcal{J}$ is non positive in each part. As the tips of the bigon are not crossings in $\mathcal{J}$, this will prove the lemma. We set for $x \in \mathbb{Z}/|c|\mathbb{Z} \setminus ([\overline{i} \overset{\ell}{\to}] \cup [\overline{j} \overset{\ell}{\to}])$ and $0 < p, q < \ell$:

$$\mathbb{D}_{p,q} = \{(\overline{i+p}, \overline{j+q}), \{(\overline{i+q}, \overline{j+p})\}, \qquad \mathbb{D}_{x,p} = \{(\overline{x}, \overline{i+p}), (\overline{x}, \overline{j+p})\},$$
$$\mathbb{D}_{p,0} = \{(\overline{i}, \overline{i+p}), (\overline{i}, \overline{j+p}), (\overline{j}, \overline{i+p}), (\overline{j}, \overline{j+p})\}, \qquad \mathbb{D}_{x,0} = \{(\overline{x}, \overline{i}), (\overline{x}, \overline{j})\},$$
$$\mathbb{D}_{p,\ell} = \{(\overline{i+\ell}, \overline{i+p}), (\overline{i+\ell}, \overline{j+p}), (\overline{j+\ell}, \overline{i+p}), (\overline{j+\ell}, \overline{j+p})\}, \quad \mathbb{D}_{x,\ell} = \{(\overline{x}, \overline{i+\ell}), (\overline{x}, \overline{j+\ell})\},$$
$$\mathbb{D}'_{p,q} = \{(\overline{i+p}, \overline{i+q}), \{(\overline{j+q}, \overline{j+p})\}, \qquad \mathbb{D}_1 = \{(\overline{y}, \overline{z}) \mid y, z \notin ([\overline{i} \overset{\ell}{\to}] \cup [\overline{j} \overset{\ell}{\to}])\},$$
$$\mathbb{D}_2 = \{(\overline{i}, \overline{i+\ell}), (\overline{i}, \overline{j+\ell}), (\overline{j}, \overline{j+\ell}), (\overline{j}, \overline{i+\ell})\}, \qquad \mathbb{D}_3 = \{(\overline{i}, \overline{j}), (\overline{i+\ell}, \overline{j+\ell})\}$$

When $i = j + \ell$ or $j = i + \ell$, the set $\mathbb{D}_2$ consists of three double points only (or zero if $c(i) \neq c(i+\ell)$). Note that we cannot have both equalities as this would imply $j = j + 2\ell$, whence $\ell = |c|/2$ and c would be a square, in contradiction with the hypothesis that c is primitive. For each double point in $\mathbb{D}_1$ the four incident arc occurrences are left in place in $\mathcal{I}$ and $\mathcal{J}$. It ensues that $\mathbb{D}_1$ has the same crossings in $\mathcal{I}$ and $\mathcal{J}$. The double points in $\mathbb{D}_3$ are the tips of the bigon. They are crossings in $\mathcal{I}$ and not in $\mathcal{J}$, whence a net change of -2 crossings. For each of $\mathbb{D}_{x,p}, \mathbb{D}_{p,q}$ or $\mathbb{D}'_{p,q}$, the first double point in their above definition is a crossing in $\mathcal{I}$ (resp. $\mathcal{J}$) if and only if the second one is a crossing in $\mathcal{J}$ (resp. $\mathcal{I}$). Their net change of crossings is thus null. For $\mathbb{D}_{x,0}$, $\mathbb{D}_{x,\ell}$, $\mathbb{D}_{p,0}$ and $\mathbb{D}_{p,\ell}$ there are a few case analysis depending on the relative ordering of the two arc occurrences incident to x (resp. $i+p$, $j+p$) with respect to the crossing $(\overline{i}, \overline{j})$ (resp. $(\overline{i+\ell}, \overline{j+\ell})$. In each case (see Figure 63) the number of crossings cannot increase from $\mathcal{I}$ to $\mathcal{J}$. For $\mathbb{D}_2$ there are three cases according to whether $i + \ell = j$, $j + \ell = i$, or none of the two identifications occurs. Recall from the paragraph before the lemma that we cannot have both identifications. The first two cases can be treated the same way. See Figure 64. The number of possible configurations is much larger in the last case, i.e. when $i + \ell \neq j$ and $j + \ell \neq i$. We essentially have to shuffle two circular orderings of length four corresponding to the two crossing tips. Without loss of generality we can assume that the path from i to

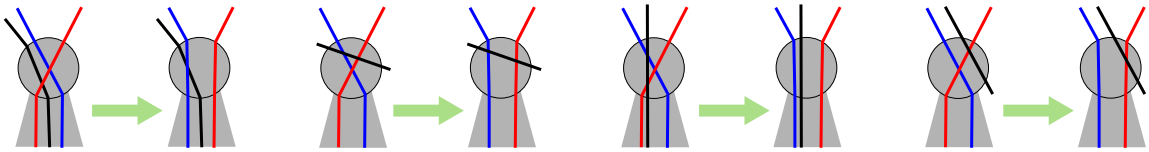

Figure 63: Up to some obvious symmetries the case $\mathbb{D}_{x,0}$, $\mathbb{D}_{x,\ell}$, $\mathbb{D}_{p,0}$ or $\mathbb{D}_{p,\ell}$ has four distinct possible configurations. The blue and red strands represent the crossing $(\overline{i}, \overline{j})$ (resp. $(\overline{i+\ell}, \overline{j+\ell})$) and each arrow links the left configuration before the bigon swap (in $\mathcal{I}$) with the right configuration after the swap is applied (in $\mathcal{J}$).

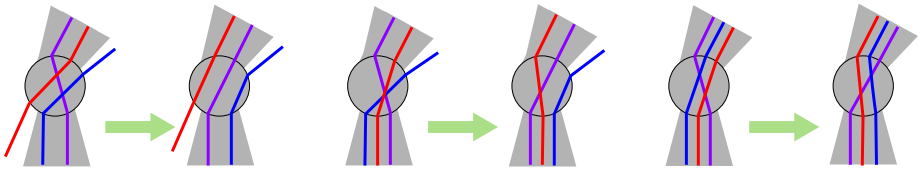

Figure 64: When $i + \ell = j$, the end of $[\overline{i} \xrightarrow{\ell}]$ overlaps with the beginning of $[\overline{j} \xrightarrow{\ell}]$ (in the purple strand). The third configuration on the right is actually forbidden by the definition of a singular bigon! It corresponds to the non-singular bigon on Figure 62.

$i + \ell$ is the right side of the bigon. We consider the following arc occurrences:

$$\begin{array}{llll} \alpha = [\overline{i}, \overline{i+1}], & \alpha' = [\overline{i}, \overline{i-1}], & \beta = [\overline{j}, \overline{j+1}], & \beta' = [\overline{j}, \overline{j-1}] \\ \gamma = [\overline{i+\ell}, \overline{i+\ell+1}], & \gamma' = [\overline{i+\ell}, \overline{i+\ell-1}], & \delta = [\overline{j+\ell}, \overline{j+\ell+1}], & \delta' = [\overline{j+\ell}, \overline{j+\ell-1}] \end{array}$$

Since $(\overline{i}, \overline{j})$ and $(\overline{i+\ell}, \overline{j+\ell})$ are crossings we must see $(\alpha, \beta, \alpha', \beta')$ in this counterclockwise order around the vertex $c(i) = c(i+\ell)$ and similarly for $(\gamma, \delta, \gamma', \delta')$. We denote by S_1 and S_2, respectively, these two circular sequences. We need to consider the effect of the bigon swapping on all the possible shuffles of S_1 and S_2. The restriction of these shuffles to α, β, γ and δ gives the 6 possible shuffles of (α, β) and (γ, δ). Among them the order $(\alpha, \delta, \gamma, \beta)$ cannot occur. Indeed, since the bigon is a thick double path the arcs $c(\alpha)$ and $c(\beta)$ either coincide or form a corner of a quad. This would force $c(\delta)$ and $c(\gamma)$ to lie in a similar configuration. In turn, the constrained order S_2 would also enforce $c(\delta')$ and $c(\delta)$ to coincide or form a corner of a quad in contradiction with the hypothesis that c has no spurs or brackets. Similar arguments show that the orders $(\alpha, \gamma, \beta, \delta)$, $(\alpha, \delta, \beta, \gamma)$, $(\alpha, \gamma, \delta, \beta)$ and $(\alpha, \beta, \delta, \gamma)$ can only occur as factors in the possible shuffles of S_1 and S_2. These 4 orders thus leads to 24 distinct shuffles of S_1 and S_2 by factoring with the 6 shuffles of (α', β') with (γ', δ'). By exchanging the roles of (α, β) and (γ, δ) and by turning clockwise instead of counterclockwise we see that $(\alpha, \delta, \beta, \gamma)$ leads to the same orders as $(\alpha, \gamma, \beta, \delta)$ and

a similar correspondence holds for $(\alpha, \gamma, \delta, \beta)$ and $(\alpha, \beta, \delta, \gamma)$. We thus only need to check the 12 configurations depicted on Figure 65. It remains to consider the order

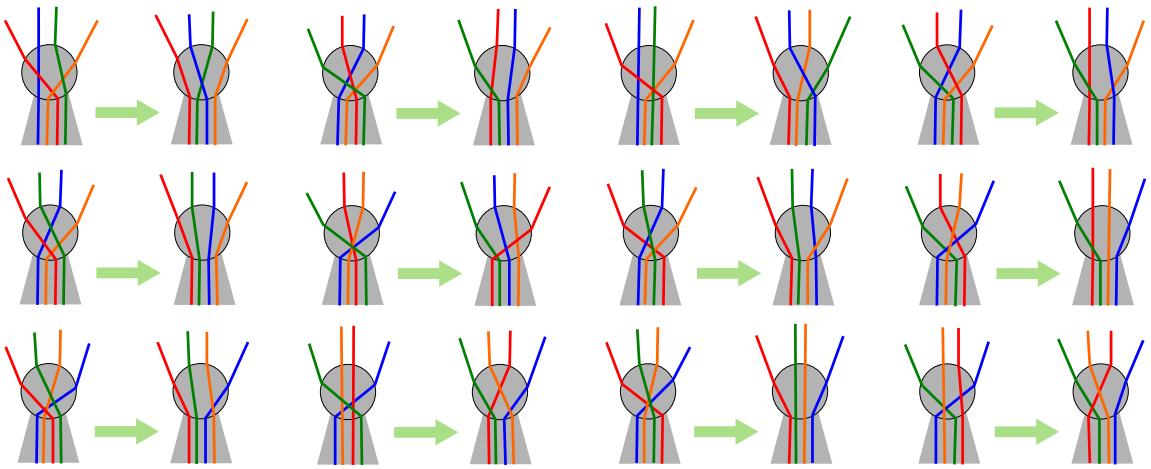

Figure 65: Effect of swapping the singular bigon $([\overline{i} \overset{\ell}{\to}], [\overline{j} \overset{\ell}{\to}])$ on $\mathbb{D}_2$ for all the orders including the factor $(\alpha, \gamma, \beta, \delta)$ or $(\alpha, \gamma, \delta, \beta)$. The blue strand represents (α, α'), the red one (β, β'), the orange one (γ, γ') and the green one (δ, δ').

$(\alpha, \beta, \gamma, \delta)$. Because of the constrained order S_2, γ' and δ' cannot lie between γ and δ. This leaves out $\binom{4}{2} = 6$ possible shuffles of (γ', δ') and $(\alpha, \beta, \gamma, \delta)$:

$$\begin{array}{lll} S'_1 : (\alpha, \beta, \gamma, \delta, \gamma', \delta') & S'_2 : (\alpha, \beta, \delta', \gamma, \delta, \gamma') & S'_3 : (\alpha, \beta, \gamma', \delta', \gamma, \delta) \\ S'_4 : (\alpha, \delta', \beta, \gamma, \delta, \gamma') & S'_5 : (\alpha, \gamma', \beta, \delta', \gamma, \delta) & S'_6 : (\alpha, \gamma', \delta', \beta, \gamma, \delta) \end{array}$$

We finally shuffle each of these orders with (α', β'). Since α' and β' cannot lie between α and β, we obtain $\binom{6}{2} = 15$ possible shuffles when considering either S'_1, S'_2 or S'_3, $\binom{5}{2} = 10$ possible shuffles with S'_4 or S'_5, and $\binom{4}{2} = 6$ possible shuffles with S'_6. By swapping left and right and turning clockwise instead of counterclockwise, we remark that S'_1 and S'_3 lead to the same orders and similarly for S'_4 and S'_5. We thus only need to consider the shuffles of (α', β') and S'_1, S'_2, S'_4 or S'_6 to complete the inspection of all the cases. Those shuffles are represented on Figures 66, 67, 68, 69 respectively.

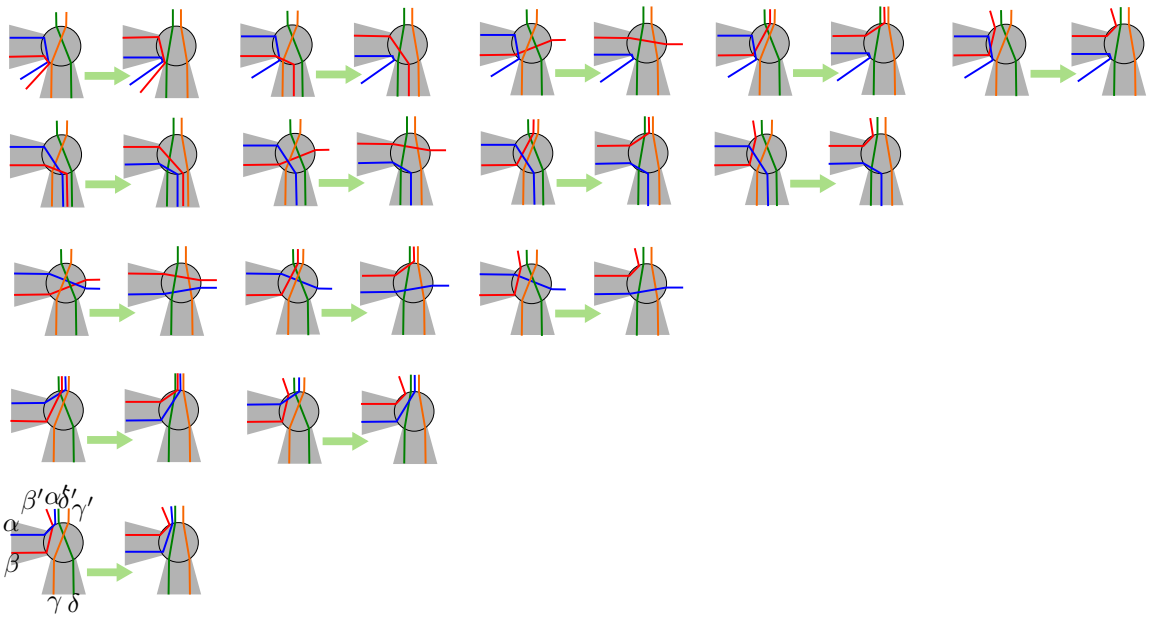

Figure 66: Effect of the bigon swapping on the 15 shuffles of S_1' and (α', β').

In each of the configurations, we trivially check that the number of crossings is not increasing. This allows to conclude the lemma when the two index paths of a singular bigon are directed the same way. A similar analysis can be made when their directions are opposite, that is when the considered bigon has the form $([\bar{i} \xrightarrow{\ell}], [\bar{j} \xrightarrow{-\ell}])$. □

Hence, by swapping singular bigons we may decrease the number of crossings until there is no more singular bigons. It follows from the next lemma that the resulting immersion has no excess crossing.

Theorem 3.9.3 (Hass and Scott [Hass and Scott, 1985, Th. 4.2])**.** *An immersion of a primitive geodesic curve has excess crossing if and only if it contains a singular bigon.*

Proof. We realize $\mathcal{I}$ by a continuous curve γ with the same construction as in Lemma 3.3.1. By Theorem 4.2 in [Hass and Scott, 1985] the curve γ has a singular bigon. In turn this bigon corresponds to a singular bigon in $\mathcal{I}$. □

In the next section we describe how to detect bigons in practice.

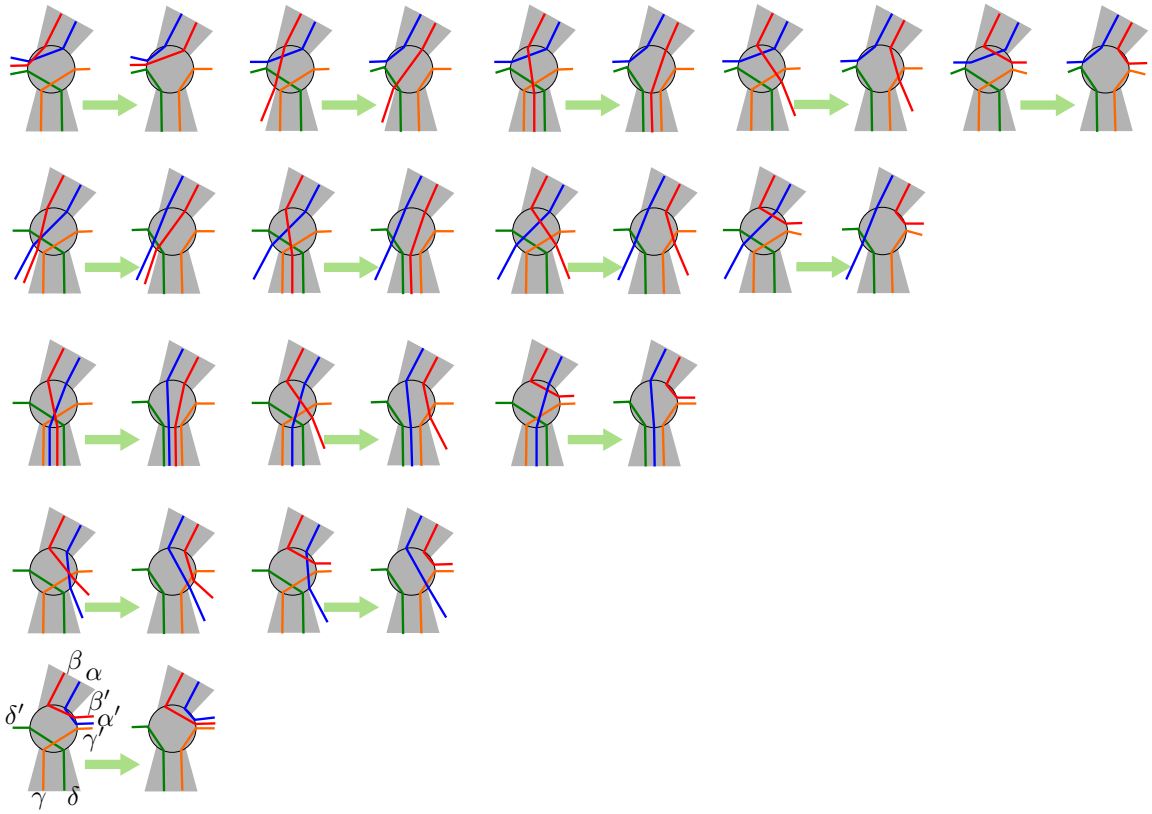

Figure 67: Effect of the bigon swapping on the 15 shuffles of S_2' and (α', β').

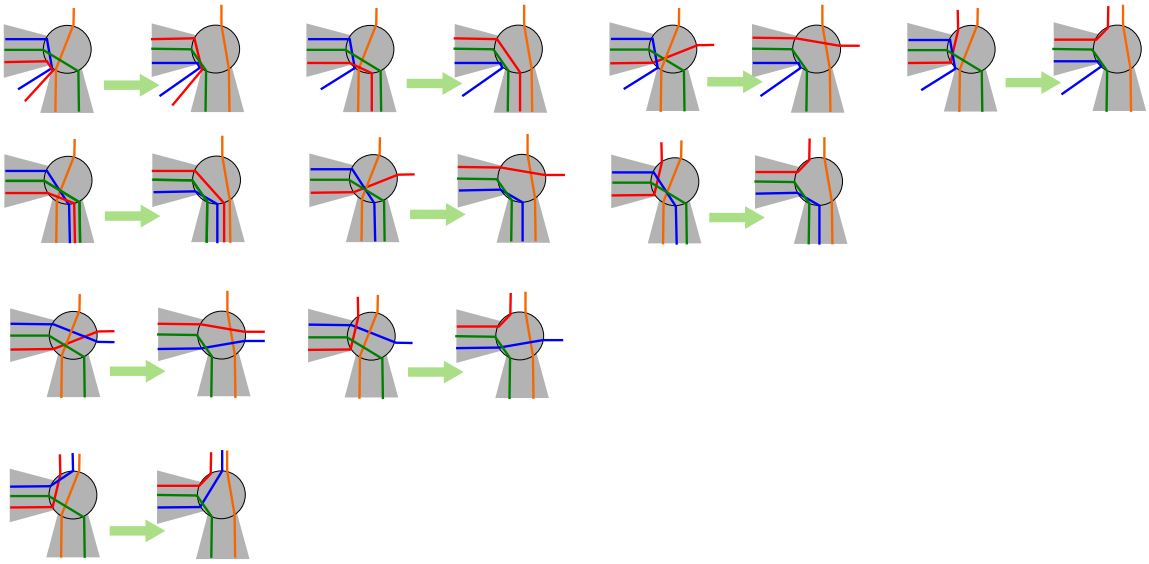

Figure 68: Effect of the bigon swapping on the 10 shuffles of S'_4 and (α', β').

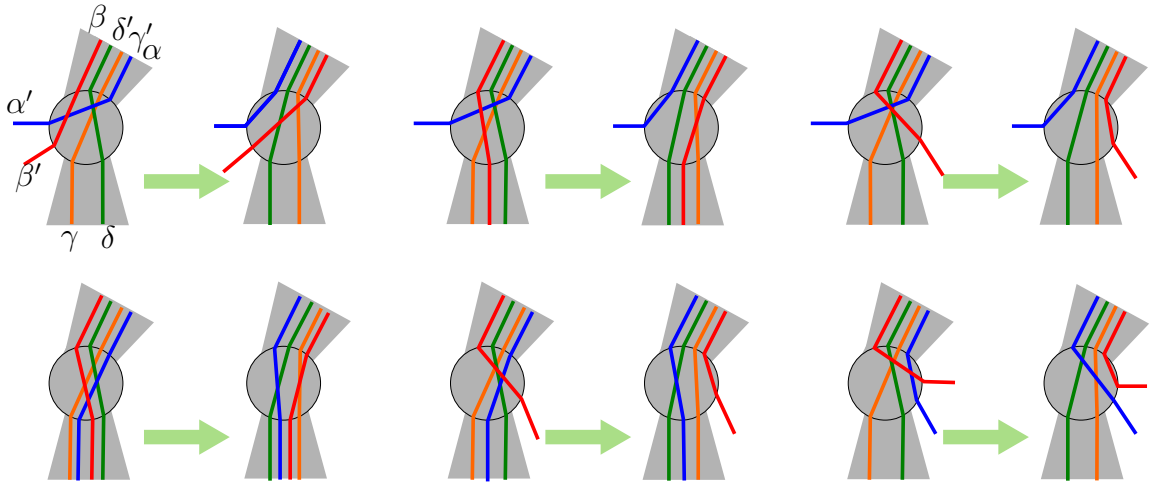

Figure 69: Effect of the bigon swapping on the 6 shuffles of S'_6 and (α', β').

3.10 Finding Bigons

If an immersion $\mathcal{I}$ of a primitive geodesic c has a bigon then its two sides have equal length according to Corollary 3.4.3 and form a sequence of paths and staircases following the description of Theorem 3.4.2. In particular, the vertices of the two sides can be put in 1-1 correspondence and corresponding vertices are at distance at most one. Hence, when looking for bigons in $\mathcal{I}$, we can start from any crossing (i, j) and walk in parallel along the two sides from i and j, checking that the two sides stay at distance one. However, we do not know a priori if the two sides will cross again to form a bigon and the search may be unsuccessful. In order to analyze the complexity of the bigon search we thus consider pair of paths that could potentially be part of a bigon, but that are not necessarily so. Formally, a pair $([\overline{i} \xrightarrow{\ell}], [\overline{j} \xrightarrow{\varepsilon\ell}])$ of distinct index paths such that $c[\overline{i} \xrightarrow{\ell}]$ and $c[\overline{j} \xrightarrow{\varepsilon\ell}]$ are corresponding subpaths of homotopic geodesic paths is called a **thick double path** of length ℓ. In other words, a thick double path is such that its image paths can be extended to form homotopic geodesic paths. As for a double path, a thick double path gives rise to a sequence of $\ell + 1$ index pairs $(\overline{i + k}, \overline{j + \epsilon k})$ for $k \in [0, \ell]$. Let Δ be a disk diagram whose boundary is labelled by homotopic geodesic paths extending the thick double path as illustrated on Figure 70.

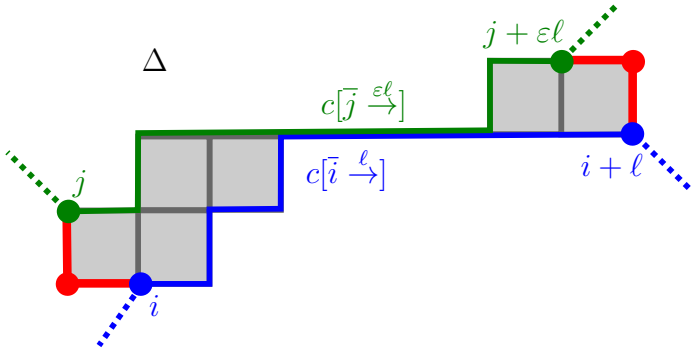

Figure 70: A thick double path extends to a pair of homotopic geodesics and bound a partial diagram.

The restriction of Δ to the part delimited by the thick double path is called a **partial diagram**. In a partial diagram each index pair $(\overline{i + k}, \overline{j + \epsilon k})$ may be in at most one of five **configurations**: either the corresponding vertices coincide in the partial diagram, or they appear diagonally opposite in a quad; in turn this quad may be labelled by one of the (at most) four quads incident to either $c[i, i+1]$ or $c[j, j+\varepsilon]$. A partial diagram is a **maximal partial diagram** if it cannot be extended to form

a partial diagram of a longer thick double path.

Lemma 3.10.1. *If c is a primitive geodesic curve each configuration of an index pair may occur at most once in a partial diagram.*

Proof. Suppose for a contradiction that the index pair (i, j) occurs twice in the same configuration in a partial diagram of a thick double path of c. Exchanging the roles of i and j if necessary, we may assume that the indices i, j are represented by integers such that $0 < j - i \leq |c|/2$. The two occurrences of (i, j) are separated by a path whose length is an integer multiple of $|c|$, say $k|c|$. We first assume that the two index paths of the thick double path are both forward. In the partial diagram, $c(i)$ and $c(j)$ label vertices that are either identical or opposite in a quad. They are thus connected by a path p of length zero or two. We infer that

$$c[\bar{i} \overset{k|c|}{\rightarrow}] \cdot p \sim p \cdot c[\bar{j} \overset{k|c|}{\rightarrow}] \sim p \cdot c[i, j]^{-1} \cdot c[\bar{i} \overset{k|c|}{\rightarrow}] \cdot c[i, j]$$

This can be written as $c'^k \cdot q \sim q \cdot c'^k$ where $c' = c[\bar{i} \overset{|c|}{\rightarrow}]$ is a cyclic permutation of c and $q = p \cdot c[i, j]^{-1}$ is a (closed) path of length at most $j - i + 2$. Since the homotopy classes of the loops c'^k and q are commuting they must admit a common primitive root [Reinhart, 1962]. We thus have *(i)* $c'^k \sim u^s$ and *(ii)* $q \sim u^t$ for some primitive curve u and some $s, t \in \mathbb{Z}$. Relation *(i)* implies that c' and u^s are also commuting so that c' is homotopic to a power of the primitive curve u. But c' being also primitive we actually have $c' \sim u^{\pm 1}$. By Remark 3.4.6, c'^t is geodesic, hence has minimal length in its homotopy class. From *(ii)* we obtain $j - i + 2 \geq |t|.|c|$. This is only possible if $t = 0$ or if $|t| = 1$ and $j - i = |c| - 2$ (indices in a pair have the same parity).

- If $t = 0$, then $q \sim 1$, so that $p \sim c[i, j]$. Since c is geodesic, we must have $j - i \leq |p| \leq 2$, hence $j = i + 2$ and the corresponding vertices must be diagonally opposite in a quad. Two homotopic paths each of length two bound a disk diagram composed of either a length two path or a single quad. In other words, $c[i, j]$ and p are either equal or bound the first quad in the partial diagram. In either case we infer that $c[i, j + 1] = c[i, j].c[j, j + 1]$ contains a spur or a bracket, leading to a contradiction.

- If $|t| = 1$ and $j - i = |c| - 2$, then $|c| = 4$ (recall that $0 < j - i \leq |c|/2$). The proof is delayed in Lemma 3.10.2.

We now assume that the thick double path is composed of a forward and a backward index paths. Similarly to the previous case, we infer that $c' \sim q \cdot (c')^{-1} \cdot q^{-1}$, where $c' = c[\bar{i} \overset{|c|}{\rightarrow}]$ and $q = p \cdot c[i, j]^{-1}$. Equivalently, c and its inverse are freely homotopic. This is however impossible unless c is contractible. □

Lemma 3.10.2. *Let c be a primitive geodesic of length four and let k be a positive integer. Then $c[\overline{0} \overset{4k}{\to}]$ and $c[\overline{2} \overset{4k}{\to}]$ cannot bound a partial diagram where the index pairs $(0,2)$ and $(4k, 4k+2)$ are in the same configuration.*

Proof. Consider a partial diagram Δ for the thick double path $\mathbb{P} = ([\overline{0} \overset{4k}{\to}], [\overline{2} \overset{4k}{\to}])$ and suppose that the index pairs $(0,2)$ and $(4k, 4k+2)$ are in the same configuration. In order to reach a contradiction we will use the fact that $c[\overline{0} \overset{2}{\to}] = c[\overline{4k} \overset{2}{\to}]$ label the first arcs of the $[\overline{0} \overset{4k}{\to}]$ side of $\mathbb{P}$ as well as the last arcs of the $[\overline{2} \overset{4k}{\to}]$ side, and similarly for $c[\overline{2} \overset{2}{\to}] = c[\overline{4k-2} \overset{2}{\to}]$.

If the first and last index pairs of $\mathbb{P}$ correspond to the same vertex in Δ, then there are five possible configurations for the first two arcs on each side of $\mathbb{P}$ (depicted on Figure 71, b1-5) and five possible configurations for the last two arcs on each side of $\mathbb{P}$ (depicted on Figure 71, e1-5). Clearly, Δ cannot start with a path part of length two (case b1), as this would imply $c[\overline{0} \overset{2}{\to}] = c[\overline{2} \overset{2}{\to}]$ and in turn $c = c[\overline{0} \overset{2}{\to}].c[\overline{2} \overset{2}{\to}]$ would be a square; this would contradict the fact that c is primitive. Case b2 leads to the same contradiction noting that $c[\overline{0} \overset{2}{\to}] \sim c[\overline{2} \overset{2}{\to}]$. Cases e1 and e2 can be dealt with analogously. As in Lemma 3.10.6, we denote by (x, y) the conjunction of configurations x and y.

- If the orientations of the first and last staircases in Δ coincide, we easily obtain that $c(1)$ has degree 2 for the conjunctions $(b3, e3)$, $c(3)$ has degree 2 for $(b4, e4)$, $c(1)$ has degree 3 for $(b3, e4)$ and $(b4, e3)$ and $c(1)$ has degree 4 for $(b3, e5)$, $(b4, e5)$, $(b5, e3)$ and $(b5, e4)$. In case $(b5, e5)$ a quad has two sides labelled by the same arc, which is forbidden by Lemma 3.10.4.

- If the orientations of the first and last staircases in Δ are opposite then each conjunction of $b3, 4, 5$ with $e3, 4, 5$ implies that a quad has two sides labelled by the same arc, which is again forbidden.

When the first and last index pairs of $\mathbb{P}$ correspond to diagonally opposite vertices in a quad, we get eight possible configurations for the first two arcs on each side of $\mathbb{P}$ (depicted on Figure 71, B1-8) and eight possible configurations for the last two arcs on each side of $\mathbb{P}$ (depicted on Figure 71, E1-8). All the pink quads in the figure denote the same quad with the same orientation, recalling that the configuration for the first and last index pairs of $\mathbb{P}$ is the same. We easily get in each case that a vertex has degree at most 5, or that a quad has two sides labelled by the same arc, or that c is a square (case $(B2, E2)$). In each case, we have thus reached a contradiction. □

Corollary 3.10.3. *Each configuration of an index pair may occur at most once in the set of maximal partial diagrams.*

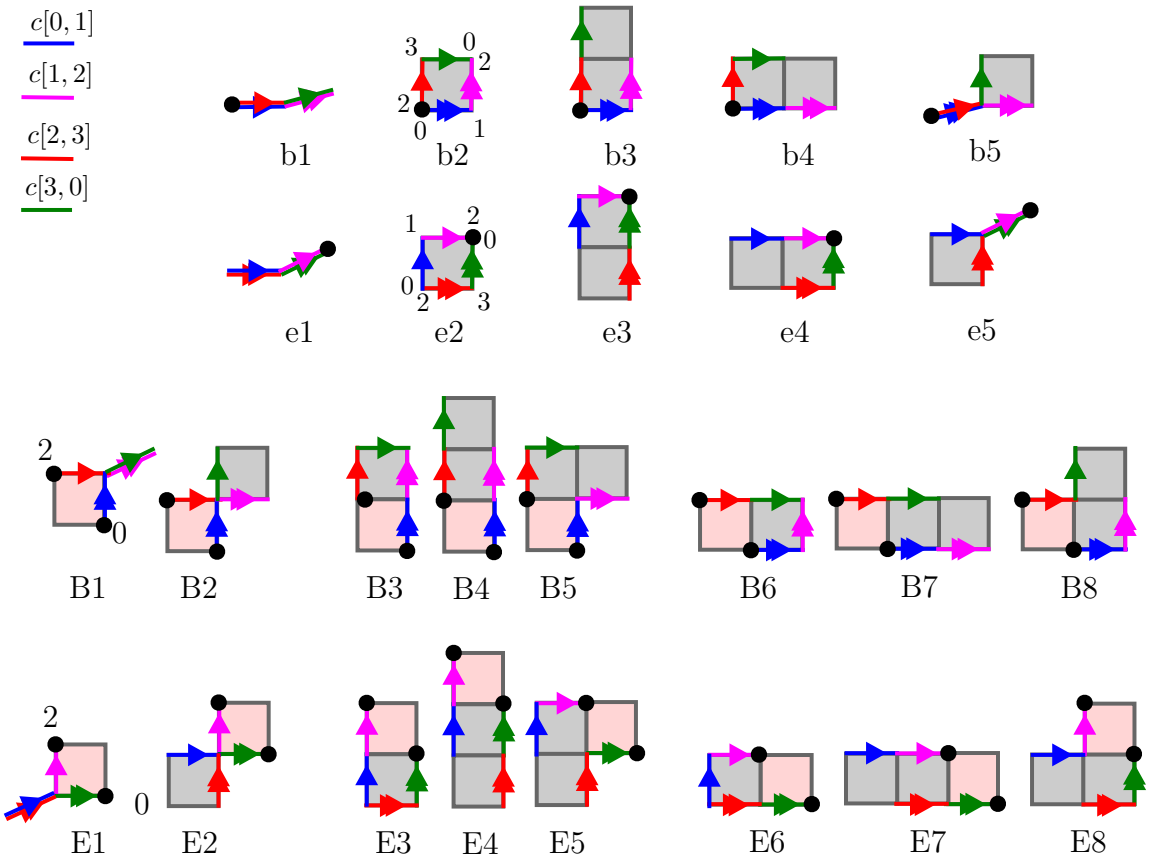

Figure 71: The possible configurations for the first two arcs and the last two arcs on each side of $\mathbb{P}$ in the corresponding partial diagram.

Proof. By the preceding Lemma, we only need to check that a configuration of an index pair cannot occur twice in distinct partial diagrams of maximal thick double paths. This is essentially the unique lift property of coverings. □

Proof of Theorem 3.1.2. By Theorem 2.3.8 we can assume that c is geodesic. We first consider the case where c is primitive. Let $\mathcal{I}$ be a randomly chosen immersion of c. We store the induced vertex orderings of the arc occurrences in arrays so that we can test if a double point is a crossing in constant time using pointer arithmetic. By Theorem 3.9.3, if $\mathcal{I}$ has excess crossing it has a singular bigon. This singular bigon defines a thick double path that must appear in some maximal partial diagram of $\mathcal{I}$ as two boundary paths with common extremities. Hence, in order to find a singular bigon we just need to scan the index pairs in every maximal partial diagram and test whether the two sides between consecutive index pairs define a singular bigon. By the very definition of a singular bigon, each test can be easily performed in constant time. The number of tests is itself bounded by the number of index pairs in all the maximal partial diagrams. This is $O(\ell^2)$ by Corollary 3.10.3. Once a singular bigon is found we swap its sides in $O(\ell)$ time. This preserves the geodesic character and, by Lemma 3.9.2, the number of crossings is reduced by at least two. Since $\mathcal{I}$ may have $O(\ell^2)$ excess crossings, we need to repeat the above procedure $O(\ell^2)$ times and we may conclude the theorem in the case of primitive curves. When $c = d^p$, with d primitive, we first find a minimally crossing immersion for d by the above procedure. We further traverse the immersion p times duplicating d as many times. As we start each traversal we connect the last arc occurrence of the previously traversed copy with the first occurrence of the next copy. We continue this copy by duplicating each traversed arc occurrence to its right. It is easily seen that the number of crossings of the final immersion of c, after connecting the last traversed arc with the first one, satisfies the formula in Proposition 3.8.2. □

We end this section with some refinements of Lemma 3.10.1 and its Corollary. We first highlight a simple property of the system of quads.

Lemma 3.10.4. *The facial walk of a quad cannot contain an arc twice, either with the same or the opposite orientations.*

Proof. If the facial walk of a quad contains two occurrences of an arc with the same orientation, then their identification creates a Möbius strip in the system of quads, in contradiction with its orientability. If the facial walk contains two consecutive occurrences of an arc with opposite orientations, then their common endpoint must have degree one in the system of quads. This contradicts the minimal degree 8 in our system of quads. Finally, if the facial walk contains non consecutive occurrences of

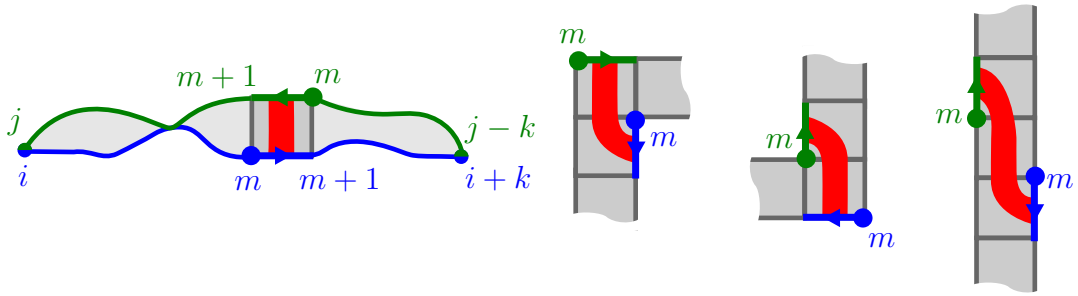

Figure 72: When the index pair (m, m) corresponds to distinct vertices (the thick m dots) in the diagram, the arc $c[m, m+1]$ labels two arcs that may be in one of four possible configurations. The whole diagram is only represented in the left case. The red strip projects to a Möbius strip in the system of quads.

an arc with opposite orientations, then their identification creates a cylinder bounded by loop edges. However, the radial graph being bipartite, this cannot occur. □

Lemma 3.10.5. *Let c be a primitive geodesic curve. If a forward index path $[\bar{i} \overset{k}{\rightarrow}]$ and a backward index path $[\bar{j} \overset{-k}{\rightarrow}]$ have homotopic image paths $c[\bar{i} \overset{k}{\rightarrow}] \sim c[\bar{j} \overset{-k}{\rightarrow}]$, then they cannot share any index.*

Proof. Remark that we cannot have $i = j$ for otherwise $c[\overline{j-k} \overset{2k}{\rightarrow}] = (c[\bar{j} \overset{-k}{\rightarrow}])^{-1}.c[\bar{i} \overset{k}{\rightarrow}] = 1$ in contradiction with Corollary 3.4.4. Likewise, we cannot have $i+k = j-k$. By way of contradiction, suppose that the index paths $[\bar{i} \overset{k}{\rightarrow}]$ and $[\bar{j} \overset{-k}{\rightarrow}]$ have a common index $i + r = j - t$ for some integers $0 \leq r, t \leq k$. By the previous remark we have $0 < \frac{r+t}{2} < k$. It follows that (m, m) is an interior index pair of $([\bar{i} \overset{k}{\rightarrow}], [\bar{j} \overset{-k}{\rightarrow}])$, where $m = i + \frac{r+t}{2} = j - \frac{r+t}{2}$ (recall that i and j have the same parity). Consider a disk diagram for $c[\bar{i} \overset{k}{\rightarrow}] \sim c[\bar{j} \overset{-k}{\rightarrow}]$. The index pair (m, m) cannot correspond to the same vertex in the diagram since the path $c[i, j] = c[i, m].c[m, j]$ would then label a closed path in the diagram, hence be contractible. This would again contradict Corollary 3.4.4. It ensues that (m, m) corresponds to opposite vertices in a quad of the diagram. This quad is part of a staircase where the arc $c[m, m+1]$ labels two arcs on either sides of the partial diagram. Figure 72 depicts the four possible configurations. In each case we infer that the system of quads would contain a Möbius strip, in contradiction with its orientability. □

Lemma 3.10.6. *Two distinct index paths of a non-trivial primitive geodesic curve c having homotopic image paths have length distinct from $|c|$ and at most $|c|+1$.*

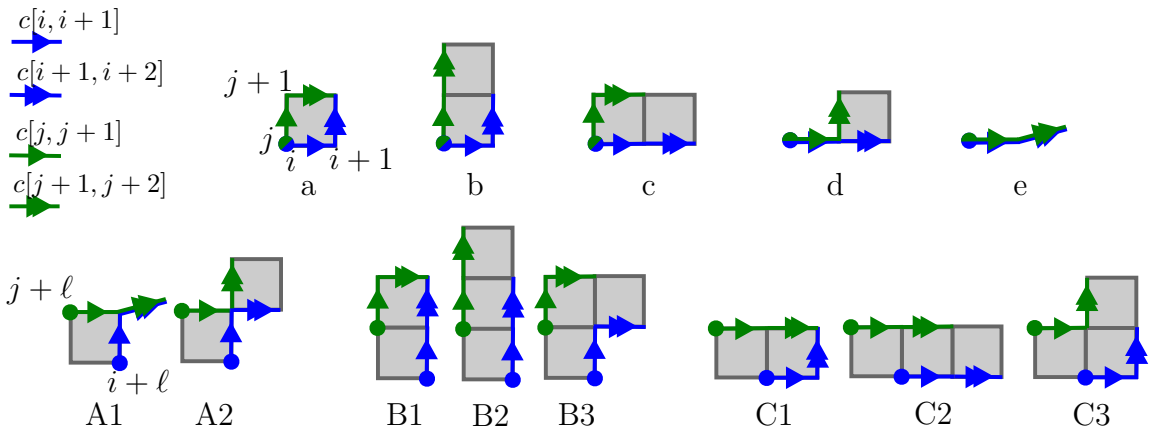

Figure 73: Upper row, the five possible configurations at the beginning of the disk diagram for $c[\bar{i} \xrightarrow{k}] \sim c[\bar{j} \xrightarrow{k}]$. Lower row, the eight possible configurations in the neighborhood of the index pair (i, j) in this diagram.

Proof. When the index paths have opposite orientations, the previous lemma directly implies the result. Let $\ell = |c|$ and let $[\bar{i} \xrightarrow{k}], [\bar{j} \xrightarrow{k}]$ be two index paths such that $c[\bar{i} \xrightarrow{k}] \sim c[\bar{j} \xrightarrow{k}]$. Clearly, $k \neq \ell$ since otherwise c would be homotopic to a conjugate of itself by the shorter curve $c[i, j]$, in contradiction with the primitivity of c. For the sake of a contradiction, suppose that $k > \ell + 1$. Consider a diagram Δ with sides Δ_R and Δ_L for the homotopic paths $c[\bar{i} \xrightarrow{k}] \sim c[\bar{j} \xrightarrow{k}]$. By Lemma 3.10.1 the vertices $\Delta_R(i + \ell)$ and $\Delta_L(j + \ell)$ must be diagonally opposite in a quad of Δ. Using that $\Delta_R[\bar{i} \xrightarrow{2}]$ and $\Delta_R[\overline{i + \ell} \xrightarrow{2}]$ are labelled by the same edges and similarly for $\Delta_L[\bar{j} \xrightarrow{2}]$ and $\Delta_L[\overline{j + \ell} \xrightarrow{2}]$ we easily deduce that the system of quads has a quad with two occurrences of a same arc in its facial walk, or has a vertex of degree at most 5, or is non-orientable. This would contradict Lemma 3.10.4 or the hypotheses on the system of quads. In details, Figure 73 depicts the five possible configurations (a,b,c,d) for $\Delta_R[\bar{i} \xrightarrow{2}]$ and $\Delta_L[\bar{j} \xrightarrow{2}]$ in the diagram Δ and the eight possible configurations (A1-2, B1-3, C1-3) for $\Delta_R[\overline{i + \ell} \xrightarrow{2}]$ and $\Delta_L[\overline{j + \ell} \xrightarrow{2}]$. Remark that cases C1-3 are symmetric to cases B1-3, so that we only need to consider the five configurations A1-2, B1-3. We denote by (x,Y) $\in$ {a,b,c,d,e} $\times$ {A1-2, B1-3} the conjunction of configurations x and Y in Δ. The quads in configurations a,b,c,d are part of an initial staircase in Δ that may have a positive or negative orientation according to whether or not its orientation is consistent with some default orientation of the quad system and similarly for the final staircase containing $\Delta_R(i + \ell)$ and $\Delta_L(j + \ell)$.

- If the orientations of the two staircases have the same sign, then the face

to the right of $\Delta_L[j, j+1]$ has the same facial walk as the face to the right of $\Delta_L(j+\ell, j+\ell+1)$. The induced identifications imply that for (x,Y) $\in$ {a,b,c,d}$\times${A1-2, B1-3}$\cup${e}$\times${A1-2,B1} a quad is bounded twice by the same arc (with or without the same orientation) in contradiction with Lemma 3.10.4. If (x,Y) $\in$ {e} $\times$ {B2-3}, we easily deduce that $c(i+1) = c(j+1)$ has degree four in the system of quads, again a contradiction.

- When the initial and final staircases have opposite orientations the faces to the right of $\Delta_L[j, j+1]$ and to the right of $\Delta_L(j+\ell, j+\ell+1)$ (resp. to the left of $\Delta_R[i, i+1]$ and to the left $\Delta_R[i+\ell, i++\ell+1]$) correspond to the two faces incident to $c[j, j+1]$ (resp. $c[i, i+1]$). We further split case $A2$ into two variants A2′ and A2″ according to whether the two quads (see Figure 73) have consistent orientations or not. By the induced identifications we derive the following forbidden situations: two occurrences of a same arc in a quad for (x,Y) $\in$ {a,d,e} $\times$ {A1, A2″} $\cup$ {e} $\times$ {B1}, a degree two vertex (namely $c(j+1)$) for (x,Y) $\in$ {a,c}$\times${B1,B3}, a degree three vertex for (x,Y) $\in$ {b}$\times${B1-3}$\cup${a,c}$\times${B2}$\cup${(e,B1)}, a degree four vertex for (x,Y) $\in$ {b,c}$\times$ {A1}$\cup${e}$\times${B2-3}$\cup${(a,A2′),(d,B1)}, a degree five vertex for (x,Y) $\in$ {b,c}$\times$ {A2′} $\cup$ {d} $\times$ {B2-3}, and finally a Möbius band for (x,Y) $\in$ {b,c} $\times$ {A2″}.

□

When the initial and final staircases in the above proof have consistent orientations it is sufficient to use that $\Delta_R[\bar{i} \overset{1}{\rightarrow}]$ and $\Delta_R[\overline{i+\ell} \overset{1}{\rightarrow}]$ are labelled by the same edge, and similarly for $\Delta_L[\bar{j} \overset{1}{\rightarrow}]$ and $\Delta_L[\overline{j+\ell} \overset{1}{\rightarrow}]$, in order to reach a contradiction for cases $A1$ and $A2$. In other words, the two homotopic paths $c[\bar{i} \overset{k}{\rightarrow}] \sim c[\bar{j} \overset{k}{\rightarrow}]$ cannot have length $\ell+1$. This leads to the following refinement.

Lemma 3.10.7. *A bigon* $([\bar{i} \overset{k}{\rightarrow}], [\bar{j} \overset{k}{\rightarrow}])$ *of an immersion of a primitive geodesic curve c without intermediate crossings (i.e., no index pair* $(i+r, j+r)$, $1 < r < k$, *is a combinatorial crossing) has length* $k < |c|$.

Lemma 3.10.8. *If two forward index paths of a canonical curve have homotopic image paths then they actually have the same image paths, i.e. they form a double path. In particular, bigons composed of two forward paths must be flat.*[2]

Proof. By Theorem 3.4.2 the image paths bound a disk diagram composed of paths and staircase. Remark that a staircase with both sides directed forward must have a

[2]In a preliminary version of this work posted on arXiv the authors erroneously claimed an analogous property for the case of a bigon composed of a forward and a backward index paths.

$\bar{1}$ turn on its left side. Since the curve is canonical we infer that the diagram cannot have any staircase, implying that the two sides coincide. □

3.11 The Unzip Algorithm

We now turn to the original problem of Poincaré [Poincaré, 1904, §4], deciding whether a given curve c is homotopic to a simple curve. In the affirmative we know by Lemma 3.9.2 and Theorem 3.9.3 that some geodesic homotopic to c must have a (combinatorial) **embedding**, i.e. an immersion without crossings. Rather than swapping the sides of a singular bigon as in Lemma 3.9.2 we can choose to switch one side along the other side. This will also decrease the number of crossings if the bigon contains no other interior bigons. By considering interior-most bigons only, we can thus enforce a given edge of c to stay fix as we remove crossings. This suggests an incremental computation of an embedding in which the image of the first arc occurrence is left unchanged: we assume that c is canonical and consider the trivial embedding of its first arc occurrence $[0, 1]$. We next insert the successive arc occurrences incrementally to obtain an embedding of the path formed by the already inserted arcs. When inserting the occurrence $[i, i+1]$ we need to compare its left-to-right order with each already inserted arc occurrence β of its supporting arc. If $\beta \neq [0, 1]$ we can use the comparison of the occurrence $[i-1, i]$ with the occurrence γ preceding β (or succeeding β if it is a backward occurrence). If $[i-1, i]$ and γ have the same supporting arc, we just propagate their relative order to $[i, i+1]$ and β. Otherwise, we use the circular ordering of the supporting arcs of $[i-1, i]$, γ and $[i, i+1]$ in order to conclude. When $\beta = [0, 1]$, we cannot use the occurrence preceding $[0, 1]$ as it is not yet inserted. We rather compare $[i, i+1]$ and $[0, 1]$ as follows. In the Poincaré disk, we consider two lifts $\tilde{d}_i$ and $\tilde{d}_0$ of c such that $\tilde{d}_i[i, i+1] = \tilde{d}_0[0, 1]$. We decide to insert $[i, i+1]$ to the left (right) of $[0, 1]$ if one of the limit points of $\tilde{d}_i$ lies to the left (right) of $\tilde{d}_0$. Note that when c is homotopic to a simple curve the two limit points of $\tilde{d}_i$ should lie on the same side of $\tilde{d}_0$.

After comparing $[i, i+1]$ with all the occurrences of its supporting arc, we can insert it in the correct place. If no crossings were introduced this way, we proceed with the next occurrence $[i+1, i+2]$. It may happen, however, that no matter how we insert $[i, i+1]$ in the left-to-right order of its supporting arc, the resulting immersion of $[\bar{0} \overset{i+1}{\to}]$ will have a combinatorial crossing. In order to handle this case, we first check if $[i, i+1]$ is **switchable**, i.e. if for some $k \geq 0$ and some turns t, u the subpath $p := c[\bar{i} \overset{k+2}{\to}]$ has turn sequence $t2^k1u$ and the index path $[\bar{i} \overset{k+2}{\to}]$ does not contain the arc occurrence $[0, 1]$. When $[i, i+1]$ is switchable we can switch p to a new subpath

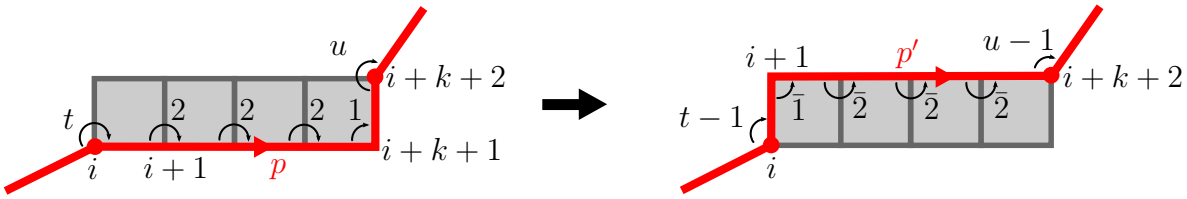

Figure 74: The arc $[i, i+1]$ is switchable.

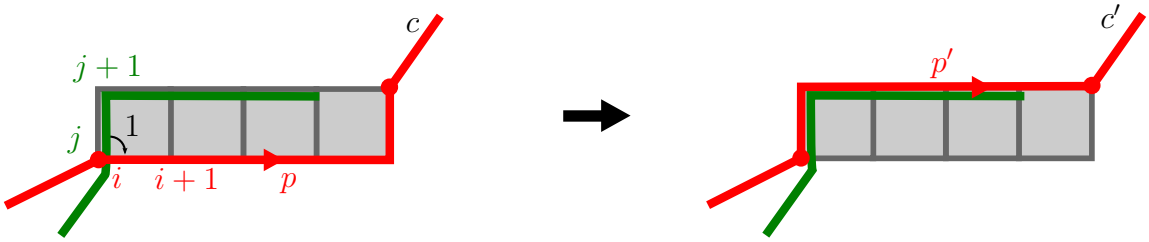

Figure 75: A switch may avoid a crossing.

p' with turn sequence $(t-1)\bar{1}\bar{2}^k(u-1)$ such that p and p' bound a diagram composed of a single horizontal staircase. See Figure 74. We indeed perform the switch if some intersection is actually avoided this way. More formally, this happens when there is a crossing (i, j) such that the turn of $(c[j, j+1], c[i, i+1])$ or of $(c[j, j-1], c[i, i+1])$ is one as illustrated on Figure 75. Note that replacing p by p' leads to a curve c' that is still geodesic and homotopic to c. Moreover, **the part of c' that remains to be embedded is in canonical form** as it contains no $\bar{1}$ turns. We then insert the arc occurrence $[i, i+1]$ and proceed with the algorithm using c' in place of c. The successive switches in the course of the computation untangle c incrementally and we call our embedding procedure the **unzip algorithm**. For further reference we make some observations that easily follows from the above switch procedure and the fact that c is initially canonical.

Remark 3.11.1. *If $[i, i+1]$ is switchable but not switched at the time of being processed it remains true until the end of the algorithm, meaning that the vertex of index i is directly followed by a subpath p with a turn sequence of the form 2^k1 (k may get smaller if some other arc of p is switched). Similarly, if $[i, i+1]$ is not switchable because of an inappropriate turn sequence, it remains true until the end of the algorithm.*

Remark 3.11.2. *Let c' be a curve homotopic to c obtained after a certain number of switches during the execution of the unzip algorithm. Then c' cannot contain a*

subpath with turn sequence $\bar{1}\bar{2}^*$ *that ends at index* 0*, nor can it contain a subpath with turn sequence* $\bar{2}^*\bar{1}$ *that starts at index* 0*.*

Lemma 3.11.3. *The unzip algorithm applied to a canonical primitive curve* c *of length* ℓ *can be implemented to run in* $O(\ell \log^2 \ell)$ *time.*

Proof. We first traverse c in reverse direction to mark all the switchable arcs. In the course of the algorithm, each time a switchable arc triggers the switch of a subpath p we unmark all the arcs in p except the last one that may become switchable (depending on the last turn of p and on the status of the arc following p). It easily follows that an arc can be switched at most twice and that the amortized cost for the switches is linear. Alternatively, we could use the run-length encoded turn sequence of c (as defined in [Erickson and Whittelsey, 2013]) to detect each switchable arc and update the turn sequence in constant time per switches. In a preprocessing phase we also compute the relative order of $[0, 1]$ with all the other occurrences of the same arc in c as follows. If $c[i, i+1] = c[0, 1]$, the corresponding arc occurrences form a double path of length one and we compute their relative order by extending maximally this double path in the backward direction. Looking at the tip of this double path, say $(\bar{j}, \bar{k})$, we can decide which side is to the left of the other. Indeed, the three arcs $c[\bar{j}, \overline{j-1}]$, $c[\bar{k}, \overline{k-1}]$ and $c[\bar{j}, \overline{j+1}] = c[\bar{k}, \overline{k+1}]$ must be pairwise distinct and their circular order about their common origin vertex in the system of quads provides the necessary information as follows from Lemma 3.10.8.

The computation of the maximal extensions in the backward direction amounts to evaluate the longest common prefix of c^{-1} with all its circular shifts. Overall, this can be done in $O(\ell)$ time thanks to a simple variation of the Knuth-Morris-Pratt algorithm. We next initialize an empty (balanced) binary search tree for every edge of the system of quads. Assuming a default orientation of each edge, its search tree should eventually contain the set of occurrences of the corresponding arc using the left-to-right order induced by the embedding.

We now traverse c in the forward direction starting with $[0, 1]$ and insert each traversed arc occurrence $\alpha = [i, i + 1]$ in its tree. If the search tree is empty, we just insert α at the root. Otherwise, either α is switchable and needs to be switched to a new arc a, or it should be inserted into a non-empty search tree. The former case can be easily detected in $O(\log \ell)$ time using a dichotomy over the occurrences of a: α **needs to be switched when one of these occurrences defines a crossing with** α **at their common endpoint**. In the latter case, we perform the usual tree insertion: each time α must be compared with some already inserted arc occurrence $\beta \neq [0, 1]$ we can use the comparison of the occurrence preceding α with the occurrence preceding β (or succeeding β if it is a backward occurrence). This

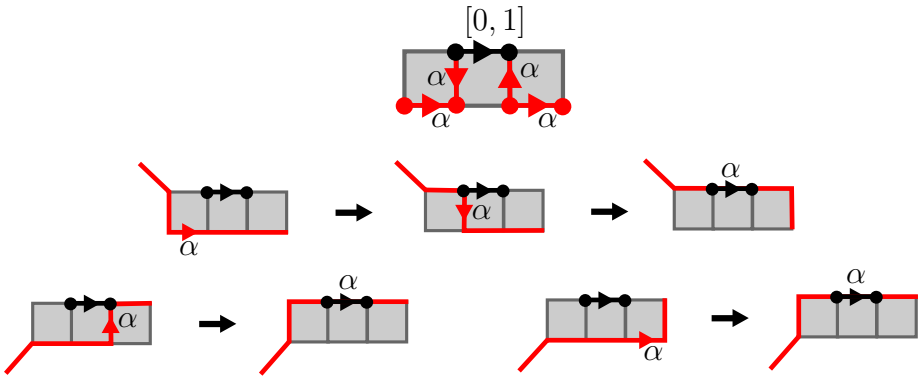

Figure 76: Top: the arc $c(\alpha)$ may be in one of four possible configurations with respect to the first arc $c[01]$. Middle and bottom row: from its initial canonical position, α is switched to the same arc as $[0, 1]$.

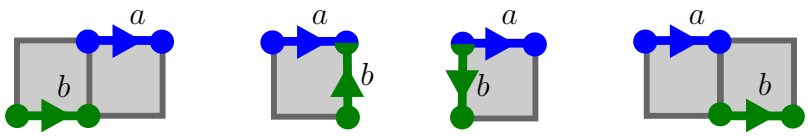

Figure 77: The relative positions of a and b.

takes $O(\log \ell)$ time. When $\beta = [0, 1]$ we can use our precomputed comparisons unless α was previously switched, thus not compared with $[0, 1]$ during the preprocessing phase. Since c is canonical, arcs may only be switched to their left in one of four possible configurations as on Figure 76 and we infer that α must lie to the right of $[0, 1]$ as justified by Lemma 3.11.4 below. In conclusion, each comparison costs $O(\log \ell)$ time so that α can be inserted in its search tree in $O(\log^2 \ell)$ time.

Note that after α is inserted we do not try to determine if the current immersion has a crossing or not. This will be checked in a second step after the unzip algorithm is completed. □

Lemma 3.11.4. *Let $\tilde{c}$ and $\tilde{d}$ be two lifts of a primitive canonical geodesic in the Poincaré disk $\mathbb{D}$. The limit points of $\tilde{c}$ cut the boundary circle $\partial\mathbb{D}$ into two pieces. By the **right piece**, we mean the piece of $\partial\mathbb{D}$ that bounds the part of $\mathbb{D} \setminus \tilde{c}$ to the right of c. Suppose that $\tilde{c}$ contains an arc a and that $\tilde{d}$ contains an arc b such that a and b are in one of the four relative positions depicted on Figure 77. Then one of the limit points of $\tilde{d}$ is in the right piece of $\partial\mathbb{D}$.*

Proof. We first note that b itself lies to the right of $\tilde{c}$. Indeed, $\tilde{c}$ would have to use a $\bar{1}$ turn to see b on its left or to pass along b, in contradiction with its canonicity. If

both limit points of $\tilde{d}$ were to the left of $\tilde{c}$, then $\tilde{c}$ and $\tilde{d}$ would form a bigon with one staircase part containing b. The $\tilde{c}$ side of this staircase would thus see b on its right, which is impossible since $\tilde{c}$ is canonical (here, we have not used that $\tilde{d}$ is canonical but only geodesic). □

Proposition 3.11.5. *If $i(c) = 0$ the unzip algorithm returns an embedding of a geodesic homotopic to c.*

Proof. Let $\mathcal{I}$ be the immersion computed by the unzip algorithm. We denote by c_k the geodesic homotopic to c resulting from the first switches in the algorithm up to the insertion of the arc occurrence $[k, k+1]$. Note that $k > r$ implies $c_k[\overline{0} \overset{r+1}{\rightarrow}] = c_r[\overline{0} \overset{r+1}{\rightarrow}]$. Suppose that $\mathcal{I}$ has a crossing. For the rest of this proof, we denote by i the smallest index such that the insertion of $[i, i+1]$ creates a crossing, i.e. the restriction of $\mathcal{I}$ to $[\overline{0} \overset{i+1}{\rightarrow}]$ has a crossing double point while its restriction to $[\overline{0} \overset{i}{\rightarrow}]$ is an embedding. By convention we set $i = \ell$ if the crossing appeared after the last arc insertion. We shall show that $c_{\ell-1}$ has two lifts whose limit points are alternating on the boundary of the Poincaré disk. It will follow from Section 4.3 that $i(c) = i(c_{\ell-1}) > 0$ thus proving the Proposition. We first establish some preparatory claims.

Claim 1. *If (i, j) is a crossing of $\mathcal{I}$ with $0 < j < i$, then the backward arc $c_i[j, j-1]$ and the forward arc $c_i[j, j+1]$ are distinct from the supporting arc $c_i[i, i+1]$ of $[i, i+1]$. Moreover, if $[i, i+1]$ was switched just before its insertion, then $c_i[j, j-1]$ and $c_i[j, j+1]$ are also distinct from the supporting arc $c_{i-1}[i, i+1]$ of $[i, i+1]$* ***just before its switch****.*

Proof. The first part of the claim follows directly from our insertion procedure and the fact that the restriction of $\mathcal{I}$ to $[\overline{0} \overset{i}{\rightarrow}]$ is an embedding. Moreover, assume that $[i, i+1]$ was switched just before its insertion. By the insertion procedure, this means that the insertion of $[i, i+1]$ before the switch would have induced a crossing (i, k) where $0 < k < i$ and $c_{i-1}[k, k+1], c_{i-1}[k, k-1] \neq c_{i-1}[i, i+1]$ by the first part of the claim. By the same first part we have $c_i[j, j-1], c_i[j, j+1] \neq c_i[i, i+1]$. Since the restriction of $\mathcal{I}$ to $c_i[\overline{0} \overset{i}{\rightarrow}]$ has no crossing, the length 2 path $c_{i-1}[\overline{k-1} \overset{2}{\rightarrow}] = c_i[\overline{k-1} \overset{2}{\rightarrow}]$ separates $c_i[\overline{j-1} \overset{2}{\rightarrow}]$ from $c_{i-1}[i, i+1]$, implying $c_i[j, j-1], c_i[j, j+1] \neq c_{i-1}[i, i+1]$ as desired. □

Claim 2. *If c_i has a $\bar{1}$ turn at index $k+1$, with $0 < k < i$, then there is an index r with $0 < r < k$ such that*

- $c_k[k, k+1] = c_k[r, r-1]$,

- *the arc* $c_i[r, r-1]$ *lies to the right of* $c_i[k, k+1]$,
- $c_k[r, r-1] \neq c_{k-1}[k, k+1]$.

Proof. A $\bar{1}$ turn can only occur at the destination of an arc that has been switched and such that the next arc was not switched (otherwise, we would get a $\bar{2}$ turn). If c_i has a $\bar{1}$ turn at $k+1$ then so has c_k since the algorithm never backtracks. It follows that the arc $c_k[k, k+1]$ must have been switched just before its insertion, thus witnessing the existence of an arc a of the form $c_k[u, u-1]$ or $c_k[u, u+1]$, with $0 < u < k$, that lies parallel to and to the right of $c_k[k, k+1]$ (with respect to $\mathcal{I}$). Moreover, according to Claim 1 we can also assume $c_k[u, u-1]$ and $c_k[u, u+1]$ to be distinct from the supporting arc of $[k, k+1]$ just before it was switched, namely $c_{k-1}[k, k+1]$. In the case $a = c_k[u, u-1]$, we may set $r = u$ and conclude. Otherwise, $a = c_k[u, u+1]$ and, recalling that the restriction of $\mathcal{I}$ to $[\overline{0} \overset{k+1}{\rightarrow}]$ has no crossing, it must be that $c_k[\overline{u} \overset{2}{\rightarrow}]$ also lies parallel to and to the right of $c_k[\overline{k} \overset{2}{\rightarrow}]$. So c_i has a $\bar{1}$ turn at index $u+1$, and we are back to the hypothesis of the claim, decreasing the value of k to u. We can repeat the same arguments inductively, each time decreasing the value of k in the claim. Since $k > 1$, the process must stop, implying that we have reached the former case. □

Recall that i is the smallest index such that the restriction of $\mathcal{I}$ to $[\overline{0} \overset{i+1}{\rightarrow}]$ has a crossing. We denote by $c' = c_{\ell-1}$ the geodesic homotopic to c resulting from all the switches in the course of the algorithm execution. Let (i, j) be a crossing of $\mathcal{I}$ with $0 < j < i$. We consider two lifts $\tilde{d}_i$ and $\tilde{d}_j$ of c' in the Poincaré disk such that $\tilde{d}_i(i) = \tilde{d}_j(j)$. We first suppose $i < \ell$.

Claim 3. *We can choose* j *so that* $\tilde{d}_i[\overline{i} \overset{+\infty}{\rightarrow}]$ *has no crossing with* $\tilde{d}_j$. *(Crossings are defined with respect to the lift of* $\mathcal{I}$ *in the Poincaré disk.)*

Proof. Fix any j such that (i, j) is a crossing of $\mathcal{I}$ $(0 < j < i)$ and suppose that $\tilde{d}_i[\overline{i} \overset{+\infty}{\rightarrow}]$ and $\tilde{d}_j$ cross. Let s be the smallest positive integer such that $\tilde{d}_i$ and $\tilde{d}_j$ crosses at $\tilde{d}_i(i+s)$. We thus have a bigon of the form $([\overline{i} \overset{s}{\rightarrow}], [\overline{j} \overset{\varepsilon s}{\rightarrow}])$ for some $\varepsilon \in \{-1, 1\}$. Moreover, this bigon has no intermediate crossings by the choice of s. Let Δ be a disk diagram for this bigon, oriented consistently with the system of quads. Δ must start with a staircase part by Claim 1. In particular, the turn t between $c'[j, j+\varepsilon]$ and $c'[i, i+1]$ should be ± 1. The unzip algorithm may have run across four possible situations at step i.

1. Either $[i, i+1]$ was switchable just before its insertion but was not switched,

2. or it was switchable and switched,

3. or it was not switchable because of an inappropriate turn sequence,

4. or it was not switchable because the part to be switched contains $[0, 1]$.

In the first situation, we know by Remark 3.11.1 that $c'(i)$ is followed by a turn sequence of the form 2^*1. Hence, t is exactly 1; but this contradicts the fact that $[i, i+1]$ was not switched though switchable. The third situation together with Remark 3.11.1 also lead to a contradiction as the inappropriate turn sequence prevents $c[\bar{i} \overset{s}{\rightarrow}]$ from being part of any staircase. Thanks to Claim 1, the second situation equally prevents $c[\bar{i} \overset{s}{\rightarrow}]$ from being part of any staircase. It remains to consider the fourth situation. We first suppose $\varepsilon = -1$, i.e. that the diagram Δ corresponds to the bigon $([\bar{i} \overset{s}{\rightarrow}], [\bar{j} \overset{-s}{\rightarrow}])$. By Lemma 3.10.7, we have $s < \ell$ and the fourth situation implies that the $[\bar{i} \overset{s}{\rightarrow}]$ side of the bigon contains $[0, 1]$. These two properties imply that $0 < i + s - \ell < i$. It ensues that the $[\bar{j} \overset{-s}{\rightarrow}]$ side contains index i, for otherwise $(i + s, j - s)$ would be a crossing occurring before step i. However, the occurrence of i on both sides of the bigon contradicts Lemma 3.10.5. Hence, it must be that $\varepsilon = 1$, i.e. that Δ is bounded by two forward paths. Since the $[\bar{i} \overset{s}{\rightarrow}]$ side of Δ contains $[0, 1]$ and since $i \neq \ell$, it ensues from Remark 3.11.2 that $\mathcal{I}$ has a $\bar{1}$ turn at index $j + 1$. Claim 2 ensures the existence of a smaller index r such that $c_j[j, j+1] = c_j[r, r-1]$ and $c_j[r, r-1] \neq c_{j-1}[j, j+1]$. Since $c_j[j, j+1]$ makes a one turn with both $c_{j-1}[j, j+1]$ and $c'[i, i+1]$, these last arcs are equal and, since (j, r) is not a crossing, (i, r) must be a crossing of $\mathcal{I}$. We replace j by r to obtain another lift $\tilde{d}_r$ that crosses $\tilde{d}_i$ at the same point $\tilde{d}_i(i) = \tilde{d}_r(r)$ with the additional property that $\tilde{d}_r[r, r-1]$ and $\tilde{d}_i[i, i+1]$ make a one turn. We finally observe that $\tilde{d}_r$ satisfies the claimed property. Indeed, the first staircase of Δ must lie to the left of its $[\bar{i} \overset{s}{\rightarrow}]$ side by Remark 3.11.2, hence the other side of Δ must start with $[r, r-1]$. This is however impossible by the preceding argument for $\varepsilon = -1$. It follows that $\tilde{d}_i[\bar{i} \overset{+\infty}{\rightarrow}]$ and $\tilde{d}_r$ have no crossing. □

We next consider the smallest positive integer r such that $\tilde{d}_i$ and $\tilde{d}_j$ crosses at $\tilde{d}_i(i - r)$. If no such r exists, then by the above Claim 3, $\tilde{d}_j$ and $\tilde{d}_i$ have a unique intersection point and we may conclude that their limit points alternate. We can thus assume the existence of a bigon $([\overline{i-r} \overset{r}{\rightarrow}], [\overline{j-\varepsilon r} \overset{\varepsilon r}{\rightarrow}])$ with $r > 0$ and $\varepsilon \in \{-1, 1\}$, and without intermediate crossings. We first examine the case $\varepsilon = -1$ where the bigon $([\overline{i-r} \overset{r}{\rightarrow}], [\overline{j+r} \overset{-r}{\rightarrow}])$ has oppositely oriented index paths. We must have $r \geq i - j$ for otherwise the tip $(i - r, j + r)$ of the bigon would define a crossing occurring before step i, contradicting the choice of i. Hence, $[\overline{i-r} \overset{r}{\rightarrow}]$ contains

j. This is however impossible by Lemma 3.10.5. We now look at the case of two forwards index paths ($\varepsilon = 1$). We must have $r \geq j$ for otherwise the tip $(i-r, j-r)$ of the bigon would define a crossing occurring before step i, again contradicting the choice of i. It follows that $[\overline{j-r}\ \overset{r}{\to}]$ contains $[0,1]$ and that the forward branches $\tilde{d}_j[\overline{i-j}\ \overset{+\infty}{\to}]$ and $\tilde{d}_i[\overline{0}\ \overset{+\infty}{\to}]$ have a unique intersection point. The bigon labels a disk diagram Δ composed of paths and staircases as described in Theorem 3.4.2.

- If $[0,1]$ and $\alpha := [i-j, i-j+1]$ label the same arc of a path part in Δ there are two possibilities: either it holds initially that $c[0,1] = c(\alpha)$ or α was switched in the course of the algorithm. In the former case, we know by the preprocessing phase and Lemma 3.10.8 that the left-to-right order of $[0,1]$ and α is coherent with its extension in the backward direction. This implies that ultimately in the backward direction $\tilde{d}_i$ lies on the same side of $\tilde{d}_j$ as does α. In the latter case, as described in the end of the proof of Lemma 3.11.3, we know by the insertion procedure that at least one of the limit points of $\tilde{d}_i$ lies on the same side of $\tilde{d}_j$ as does α. Since (i,j) is the only crossing in the forward direction, we conclude in both cases that $\tilde{d}_j$ and $\tilde{d}_i$ have alternating limit points.

- Otherwise, $[0,1]$ and α label two distinct arcs, say c_0 and a_{i-j}, of a staircase part σ of Δ. Let $[i-v, i-u]$ and $[j-v, j-u]$, $0 \leq u < v \leq r$ be the index paths corresponding to the sides σ_L and σ_R of σ. By Remark 3.11.2, $[j-u, j-v]$ must label the right side σ_R while $[i-v, i-u]$ cannot contain $[0,1]$. It follows that $v < i$, whence $[i-v, i-u] \subset [1, i]$. If a_{i-j} belongs to a horizontal part of σ_L, the first vertex in this part has a $\bar{1}$ turn and we let m be the index of this vertex. Otherwise, a_{i-j} belongs to a vertical part whose last vertex has a $\bar{1}$ turn. It follows that the vertex of index $i-j+1$ had a $\bar{1}$ turn at step $i-j$ of the algorithm and we set $m = i-j+1$. By Claim 2, there is an arc occurrence $[x, x-1]$, with $x < m-2$, that lies to the right of a_{i-j} and such that the turn at $x+1$ is not one. We view Δ as a subset of the Poincaré disk so that Δ_L and Δ_R can be seen as portions of $\tilde{d}_i$ and $\tilde{d}_j$ respectively. Let $\tilde{q}$ be the lift of c' that extends the above occurrence $[x, x-1]$ in $\mathbb{D}$. See Figure 78. We denote by $q_+ := \tilde{q}[\overline{x}\ \overset{+\infty}{\to}]$ and $q_- := \tilde{q}[\overline{x}\ \overset{-\infty}{\to}]$ the portion of $\tilde{q}$ respectively after and before its vertex with index x.

  We claim that q_+ cannot cross $\tilde{d}_i$. Otherwise, we would get a pair of homotopic paths, one piece of q_+ starting at index x and one piece of $\tilde{d}_i$ ending at index m, with opposite orientations. By Lemma 3.10.5, the q_+ piece would not contain index m and would thus have length at most $m - x < m$. In turn, this would imply that the $\tilde{d}_i$ piece would only contain vertices with indices in $[1, m]$. The crossing of q_+ and $\tilde{d}_i$ would thus occur before step i, a contradiction.

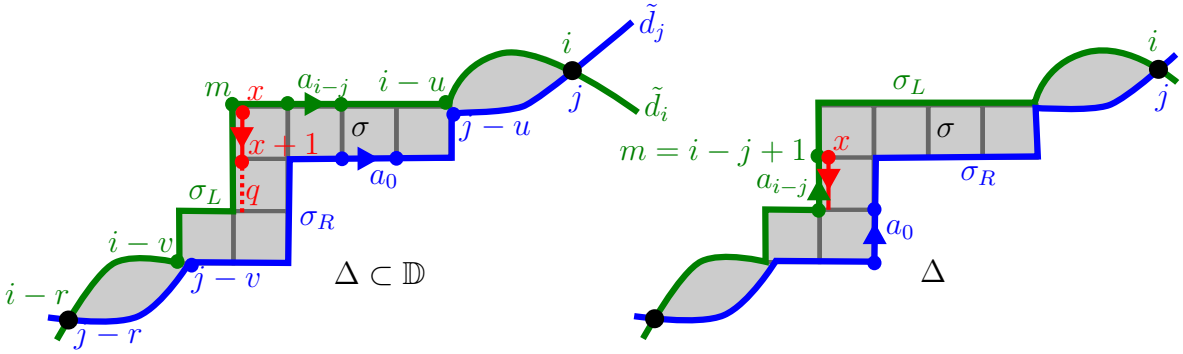

Figure 78: a_{i-j} may belong to a horizontal (left figure) or vertical (right figure) part of σ_L.

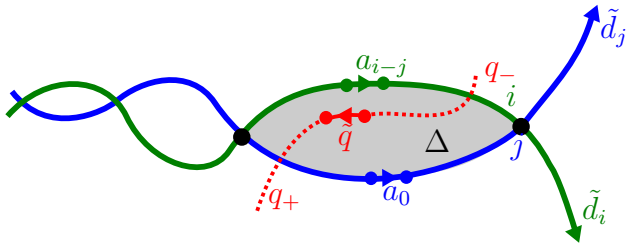

Figure 79: q_+ cannot cross $\tilde{d}_i$, q_- cannot cross $\tilde{d}_j$, and $\tilde{d}_j$ cannot cross $\tilde{d}_i[\bar{i} \overset{+\infty}{\rightarrow}]$.

We next claim that q_- cannot cross $\tilde{d}_j$. To see this, first note that the index path corresponding to the part of q_- inside Δ must contain 0. Indeed, q_- must cross Δ_L or the part of Δ_R after c_0 and, in any case, the index along q_- of the corresponding crossing should be at least i, since no crossing occurs before step i.

It follows from the above claims that $\tilde{q}$ and $\tilde{d}_i$ have alternating limit points. See Figure 79.

It remains to consider the case $i = \ell$ where the first crossing appears after the last arc insertion. Since all the arcs have been inserted without introducing crossings it means that the crossings of the computed immersion $\mathcal{I}$ must have the form $(0, j)$. We first claim that each bigon of $\mathcal{I}$ must have its two sides oriented the same way. Otherwise, the tips of the bigon must have the form $(0, j)$ and $(j, 0)$ for some index $j \neq 0$, implying that $c_\ell[\overline{0} \overset{\pm j}{\rightarrow}]$ is contractible. This would contradict Corollary 3.4.4. Let $\tilde{d}_\ell$ and $\tilde{d}_j$ be two intersecting lifts of c' in the Poincaré disk. If $\tilde{d}_\ell$ and $\tilde{d}_j$ intersect only once, then we are done as they must have alternating limit points. We now suppose that $\tilde{d}_\ell$ and $\tilde{d}_j$ intersect at least twice and we consider the bigon Δ between

their last two intersections in the forward direction. By Lemma 3.10.7, the length of Δ is smaller than ℓ. Let $(0, j)$ and $(\ell - j, 0)$ be the tips of Δ, so that $\tilde{d}_\ell[\overline{0} \overset{j}{\rightarrow}] \sim \tilde{d}_j[\overline{\ell - j} \overset{j}{\rightarrow}]$.

- If $[0, 1]$ and $\alpha = [\ell - j, \ell - j + 1]$ label the same arc $\tilde{d}_\ell[0, 1] = \tilde{d}_j(\alpha)$ then, by the insertion procedure, one of the limit points of $\tilde{d}_\ell$ is on the same side of $\tilde{d}_j$ as does α. The same argument as in the general case $i < \ell$ allows to conclude that $\tilde{d}_\ell$ and $\tilde{d}_j$ have alternating limit points.

- We finally suppose that $[0, 1]$ and α label distinct arcs, say c_0 and $a_{\ell-j}$, of a staircase part σ of Δ. As in the general case $i < \ell$, c_0 must see σ on its left, while $a_{\ell-j}$ must see σ on its right. Hence, the $\tilde{d}_j$ side of σ is canonical while the side along $\tilde{d}_\ell$ is not and must have been switched. By Claim 2, there is an arc occurrence β to the right of $\tilde{d}_\ell[\ell - j, \ell - j + 1]$ with the opposite orientation. Let $\tilde{q}$ be the lift of c' that extends β. As in the general case $i < \ell$, we can show that the part of $\tilde{q}$ after β cannot cross $\tilde{d}_\ell$, while the part before β cannot cross $\tilde{d}_j$. Using that $(0, j)$ is the last crossing along $\tilde{d}_\ell$ and $\tilde{d}_j$, we equally conclude that $\tilde{d}_\ell$ and $\tilde{q}$ have alternating limit points. See Figure 79 with $i = \ell$.

□

Proof of Theorem 3.1.3. Let c be a combinatorial curve of length ℓ on a combinatorial surface of size n. We compute its canonical form in $O(n + \ell)$ time and check in linear time that c is primitive. In the negative, we conclude that either c is contractible, hence reduced to a vertex, or that c has no embedding by Proposition 3.8.2. In the affirmative, we apply the unzip algorithm to compute an immersion $\mathcal{I}$ of some geodesic c' homotopic to c. According to Proposition 3.11.5, we have $i(c) = 0$ if and only if $\mathcal{I}$ has no crossings. This is easily verified in $O(n + \ell)$ time by checking for each vertex v of the system of quads that the set of paired arc occurrences with v as middle vertex form a well-parenthesized sequence with respect to the local ordering $\prec_v$ induced by $\mathcal{I}$. We conclude the proof thanks to Lemma 3.11.3. □

A related problem was tackled by Chang et al. [Chang et al., 2015, Th. 8.2] who try to find an embedding of a given closed path on a combinatorial surface. In their formulation, though, the path is not authorized to be modified. In our terminology, they only look for the existence of a combinatorial immersion without crossings. They suggest a linear time complexity for this problem and it seems likely that we could also eliminate the $\log^2 \ell$ factor in our complexity.

3.12 Concluding Remarks

The existence of a singular bigon claimed in Theorem 3.9.3 relies on Theorem 4.2 of Hass and Scott [Hass and Scott, 1985]. As noted by the authors themselves this result is "surprisingly difficult to prove". Except for this result and the recourse to some hyperbolic geometry in the general strategy of Section 4.3 our combinatorial framework allows to provide simple algorithms and to give simple proofs of results whose known demonstrations are rather involved. Concerning Proposition 3.9.1, the existence of an immersion without bigon could be achieved in our combinatorial viewpoint by showing that if an immersion has bigons, then one of them can be swapped to reduce the number of crossings. (The example in Figure 56 shows that such a bigon need not be singular.) What is more, if those swappable bigons could be found easily this would provide an algorithm to compute a minimally crossing immersion of two curves by iteratively swapping bigons as in Section 3.9 for the case of a single curve. However, we were unable to show the existence or even an appropriate definition of a swappable bigon. Note that in the analogous approaches using Reidemeister-like moves by de Graff and Schrijver [de Graaf and Schrijver, 1997] or by Paterson [Paterson, 2002], the number of moves required to reach a minimal configuration is unknown. Comparatively, the number of bigon swapping would be just half the excess crossing of a given immersion.

Although the geometric intersection number of a combinatorial curve of length ℓ may be $\Omega(\ell^2)$, it is not clear that the complexity in Theorem 3.1.1 is optimal. In particular, it would be interesting to see if the unzip algorithm of Section 3.11 yields minimally crossing curves even with curves that are not homotopic to simple curves, thus improving Theorem 3.1.2. This would lead to an algorithm that is faster to find an immersion with the minimal number of crossings than to actually count them! It is also tempting to check whether the unzip algorithm applies to compute the geometric intersection number of two curves rather than a single curve. Another intriguing question related to the computation of minimally crossing immersions comes from the fact that they are not unique. Given a combinatorial immersion we can construct a continuous realization as described in the proof of Lemma 3.3.1. Say that two realizations are in the **same configuration** if there is an ambient isotopy of the surface where they live that brings one realization to the other. It was shown by Neumann-Coto [Neumann-Coto, 2001] that every minimally crossing immersion is in the configuration of shortest geodesics for some Riemaniann metric μ, but Hass and Scott [Hass and Scott, 1999] gave counterexamples to the fact that we could always choose μ to be hyperbolic.

Is there an algorithm to construct or detect combinatorial immersions that have a

realization in the configuration of geodesics for some hyperbolic metric?

Chapter 4

Barnette's Conjecture

4.1 Introduction

A **splitting cycle** on a topological surface Σ is a simple closed curve that cuts the surface into two non-trivial pieces, none of which is homeomorphic to a disk. Refer to Fig. 80. One can easily see that the genus of Σ is split into two numbers summing to the genus of Σ. A torus does not have any splitting cycle. By the classification of surfaces any closed surface (orientable or not) of genus at least two admits a splitting cycle. Given a combinatorial surface homeomorphic to Σ it is natural to ask whether G contains a cycle whose image in Σ is a splitting cycle. **Here, a cycle in a graph is a closed walk without any repeated vertex.** This is the key point. To ask for simplicity in a combinatorial map is strongly different than its continuous analogue since it is obviously relative to the map and not only the underlying surface. We thus remark that the fact of being simple has a very different sense speaking of surfaces or combinatorial maps. This must be taken in consideration for any tentative of translation of a continuous problem in a discrete setting. It is known to be NP-hard

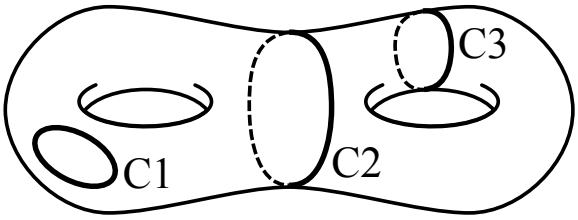

Figure 80: A cycle may be null-homotopic (C1) or separating but non null-homotopic (C2) or neither null-homotopic nor separating (C3). C2 is also called a splitting cycle.

to decide whether a combinatorial surface contains a splitting cycle or not [Cabello et al., 2011; Chambers et al., 2006]. However, it was conjectured by Barnette that

Conjecture 4.1.1 (Barnette '1982 [Mohar and Thomassen, 2001, p. 166])**.** *Every triangulation of a surface of genus at least 2 has a splitting cycle.*

Two splitting cycles on a topological surface have the same **type** if there exists a self-homeomorphism of the surface that maps one cycle to the other one. On an orientable surface of genus g there are $\lfloor g/2 \rfloor$ possible types corresponding to splittings into components of respective genus k and $g-k$, $1 \leq k \leq g/2$. This type is denoted by $(k, g-k)$ or simply by $min(k, g-k)$ when the genus is clear from the context. A stronger version of the above conjecture was later proposed.

Conjecture 4.1.2 ([Mohar and Thomassen, 2001, p. 167])**.** *Every triangulation of an orientable surface of genus at least 2 has a splitting cycle of every possible type.*

We prove that Conjecture 4.1.1 holds for the embeddings of the complete graphs K_n described by Ringel and Youngs [Ringel, 1974] or by Gross and Tucker [Gross and Tucker, 1987] when $n \equiv 7$ modulo 12. We next present counter-examples to Conjecture 4.1.2 that also disprove a stronger conjecture of Zha and Zhao [Zha and Zhao, 1993]. Let M_{19} be one of the embeddings of K_{19}, the complete graph on 19 vertices, given by Lawrencenko et al. [Lawrencenko et al., 1994]. From the Euler characteristic it is easily seen that M_{19} is a triangulation of genus 20. Indeed $\chi(M_{19}) = 19 - 19 \times 18/2 + 2/3 \times 19 \times 18/2$ because its number of edges is clearly $19 * 18/2$ and any triangulation verifies $2e = 3f$, so $\chi(M_{19}) = 19 - 1/3 * 19 * 9 = -38 = 2 - 2 * 20$ and so $g = 20$. A brute force approach to test if any of the cycles of K_{19} is splitting in M_{19} would lead to years of computations. Thanks to a simple branch and cut heuristic I was able to check on a computer that Conjecture 4.1.2 fails for M_{19}. Only 4 of the 10 possible types occur and, in particular, it is not possible to split M_{19} into two pieces of equal genus. We first describe the previous work on the subject before discussing our new results.

4.2 State of the Art

4.2.1 Splitting Cycles on Triangulations

If T' is obtained from a triangulation T by an edge contraction, then every splitting cycle in T' is the contraction of at least one splitting cycle of the same type in T. See Figure 81. It follows that Conjectures 4.1.1 and 4.1.2 can be equivalently restricted to

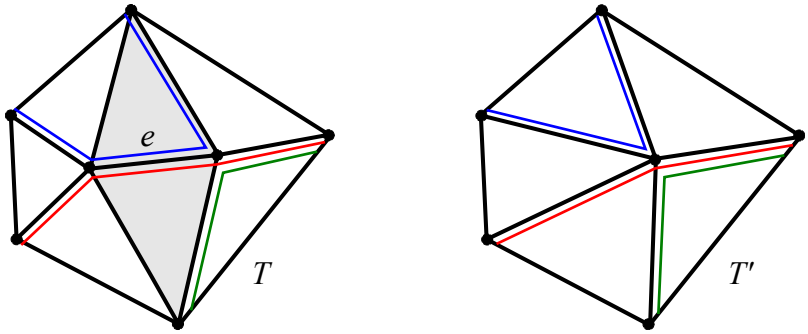

Figure 81: The contraction of edge e in T and the resulting triangulation T'. Every simple path on T' is the contraction of (at least) a simple path on T.

irreducible triangulations. On the other hand, Barnette and Edelson [Barnette and Edelson, 1988; Barnette and Edelson, 1989] proved that the number of irreducible triangulations of a given surface is finite. Nakamoto and Ota [Nakamoto and Ota, 1995] further showed that the number of vertices in an irreducible triangulation is at most linear in the genus of the surface. (This result has been extended to surfaces with boundaries by Boulch et al. [Boulch et al., 2013]. The best upper bound known to date is due to Joret and Wood [Joret and Wood, 2010] who proved that this number is at most $\max\{13g-4, 4\}$.) In theory, one can thus list all irreducible triangulations with fixed genus. This makes conjectures 4.1.1 and 4.1.2 decidable for fixed genus. Indeed, we can consider every irreducible triangulation in turn to test whether one of its cycles is splitting and compute its type as the case may be. Sulanke [Sulanke, 2006a] describes an algorithm for generating all the irreducible triangulations with given genus and was able to list the irreducible triangulations of the orientable surface of genus 2 and of the non-orientable surfaces up to genus 4. According to Sulanke there are already 396784 irreducible triangulations of the orientable surface of genus 2 and 6297982 irreducible triangulations of the non-orientable surface of genus 4. In practice, the number of irreducible triangulations is growing too fast and the technique cannot be used for higher genera. Thanks to its enumeration Sulanke could conclude by brute force computation that Conjecture 4.1.1 is true for the orientable surface of genus 2 [Sulanke, 2006b]. A formal and highly technical proof seems to have appeared in Jennings' thesis [Jennings, 2003]. Sulanke [Sulanke, 2006b] also gives a simple counter-example to an extension of Conjecture 4.1.2 to the non-orientable case of genus 3. It is constructed from a triangulation of a torus and a triangulation of a projective plane. Remove a triangle from each of those surfaces and glue them along their boundary. Let C be the joining cycle in the resulting non-orientable triangulation of genus 3. This surface cannot be cut by any cycle into a

perforated Klein bottle and a projective plane since the cycle would have to cross the length three cycle C at least four times. A similar argument holds when reversing the roles of the torus and the Klein bottle, that is gluing a Klein bottle with a projective plane. To our knowledge, no progress has been made on conjectures 4.1.1 and 4.1.2 since then.

4.2.2 Splitting Cycles on Maps with Large Face-Width

The splitting cycle Conjectures for triangulations have their counterpart for maps with face-width 3:

Conjecture 4.2.1 ([Zha and Zhao, 1993])**.** *Every map of genus at least 2 and face-width at least 3 has a splitting cycle.*

Conjecture 4.2.2 ([Mohar and Thomassen, 2001])**.** *Every map of genus at least 2 and face-width at least 3 has splitting cycles of every possible type.*

Since a triangulation has no loop or multiple edge, its face-width is at least 3. The above conjectures are thus stronger than conjectures 4.1.1 and 4.1.2 respectively. It was proved by Zha and Zhao [Zha and Zhao, 1993] that a map of genus ≥ 2 with face-width at least 6 in the orientable case and at least 5 in the non-orientable case has a splitting cycle. Each of their constructions leads to splitting cycles of type 1, which seems to be the most occurring type. For maps of genus 2, those conditions were lowered to face-width 4 in the orientable case [Ellingham and Zha, 2003] and face-width 3 otherwise [Robertson and Thomas, 1991]. This last case is somehow buried in a paper related to the computation of the genus of a graph. The proof involves Menger's theorem and is quite different from the other approaches on the subject. So we describe it in our vocabulary and in a ad-hoc way.

Theorem 4.2.3 ([Robertson and Thomas, 1991])**.** *Every non-orientable map of genus 2 and face-width at least 3 has a splitting cycle.*

In fact, the condition on the face-width can be lowered to the condition that any closed curve homotopic to some fixed non-separating two-sided simple loop ℓ intersects the graph of the map at least three times.

Proof. Consider a map as in the theorem and let Σ be the Klein bottle on which its graph is embedded. With a little abuse of notations we write G for the graph as well as its embedding. Choose ℓ as above on Σ that further minimizes the number $k \geq 3$ of intersections with G. After cutting along ℓ we get an annulus bounded by two copies ℓ' and ℓ'' of ℓ with opposite orientations. See Figure 82. The graph G is also cut by ℓ,

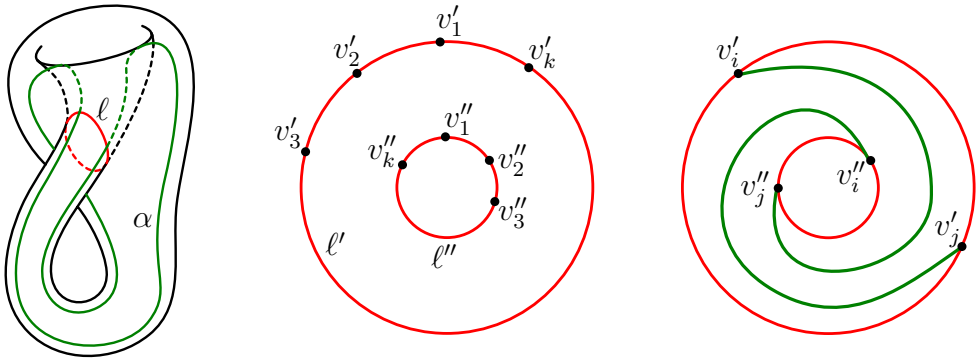

Figure 82: Left, a Klein bottle with a non-separating two-sided curve ℓ and a splitting cycle α. Middle, the annulus after cutting along ℓ. Right, The two paths in the annulus merge to a splitting cycle α on the Klein bottle.

yielding a graph G' with k vertices $v'_1, \ldots, v'_k$ in cyclic order along ℓ' that correspond to k other vertices $v''_1, \ldots, v''_k$ on ℓ''. By the minimal property of ℓ, every closed curve in the annulus separating ℓ' from ℓ'' cuts G' at least k times. By Menger's theorem it easily follows that G' contains k vertex disjoint paths connecting v'_i to $v''_{\sigma(k+1-i)}$ $(i = 1 \ldots k)$ for some circular shift σ of $[1, k]$. The permutation $i \mapsto \sigma(k + 1 - i)$ is an involution that cannot be the identity as $k \geq 3$. Any of its 2-cycle (i, j) provides a splitting cycle by merging the path from v'_i to v''_j and the path from v'_j to v''_i. See Figure 82. □

4.3 Experimental Approach

In order to get more insight on the conjecture we start with some simple examples. This is not easily done by freehand drawing. Since every cycle in a torus is either contractible or non-separating, the simplest surface to consider is the genus 2 surface. Freehand drawing of combinatorial maps of genus 2 is rather tricky especially for maps of small face-width as required by Zha and Zhao's result [Zha and Zhao, 1993]. Looking at such drawings for splitting cycles is even more intricate so that we can hardly get any intuition on the problem. An implementation on a computer can thus be helpful in order to test the conjecture even against small examples. In practice I chose to implement the flag data-structure described in Section 2.2.3 in C++. The running time tables described below were obtained on a quad core laptop with 8Gb RAM. Implementation is available at `http://vdespre.free.fr/VincentDESPRE.html`.

4.3.1 First Exploration

As noted in the state of the art it is sufficient to test the conjecture against every irreducible triangulation of a fixed genus for proving the conjecture for all triangulations of the corresponding genus. Sulanke has computed the full list of irreducible triangulations of genus 2. There are approximately 400 000 of them. For each of them he found at least one splitting cycle and computed the length of the shortest one. Those results are shown in the next table. The irreducible triangulations of genus 2 have between 10 and 17 vertices. The entry in row i and column j stores the number of maps with i vertices whose shortest splitting cycle has length j. For instance we see on the first row that 681 irreducible triangulations of genus 2 with 10 vertices have their shortest splitting cycle of length 6. The entries of the last column give the average length of a minimum splitting cycle over all the maps considered in their row.

•	3	4	5	6	7	8	Average
10		2	51	681	130	1	6.09
11	2	58	2249	16138	7818	11	6.21
12	25	1516	20507	72001	22877	121	6.00
13	710	13004	50814	78059	16609	9	5.61
14	8130	30555	12308	3328	205	1	4.21
15	36794	1395	3	1	2		3.04
16	661	3					3.01
17	5						3.00

My implementation leads to the exact same results. Looking at the previous table we realize that triangulations with few vertices have longer shortest splitting cycles than those with more vertices. This made us doubting of the correctness of the conjecture and we decided to try to look at bigger triangulations of higher genus. In order to see if my intuition was correct I first tried all the 59 non isomorphic triangulations of K_{12} on S_6. In all maps I found a splitting cycle. I thus decided to look for the type of the splitting. This study highlighted the fact that it is more difficult (meaning needs longer cycles) to split into pieces of genera 3 and 3 than 1 and 5. The algorithm stores for each map the length of the shortest splitting cycle of each type. The average on that value and the worst cases on the set of the 59 non-isomorphic embeddings of K_{12} are shown on the next table.

Type	$(1,5)$	$(2,4)$	$(3,3)$
Average	7.58	9.41	10.32
Worst case	8	10	12

For instance there is an embedding of K_{12} whose only splitting cycle of balanced type is a Hamiltonian cycle. It suggests that complete graph may not all have that kind of splitting cycle. We then decided to study complete graphs of higher genus to see if that phenomenon is confirmed.

4.3.2 Complete Graphs

Details on embeddings of complete graphs can be found in the next section. By trying to find explicit constructions of complete triangulations we found that Ellingham and Stephens [Ellingham and Stephens, 2005] had proposed as an application of their work the test of the conjectures about splitting cycles. In fact, it is not easy in practice to search for splitting cycles. There is not efficient algorithm for general maps (which is excepted since the problem is NP-complete). Thus the only algorithm we know consists in testing every cycles of the entry map. The number of cycles of a graph with n vertices is $\sum_{k=3}^{n} \frac{1}{2}\binom{n}{k}(k-1)!$ which is morally growing like $n!$. Using that basic exhaustive search algorithm we obtain the following time of computation for complete graphs with Ringel and Youngs embeddings. The last row shows the final computation time obtained after designing our final algorithm.

n	12	15	16	19	31	43
basic	2 s.	1 h.	12 h.			
final	5 s.	6 s.	6 s.	10 sec.	15 m.	4 h.

The growth of computation times of the naive algorithm is similar to $n!$. We got a theoretical computation time of approximately 10 years for $n = 19$. Let us see the result we obtain for K_{15} and K_{16}. Recall that the the embeddings considered are those of Ringel and Youngs. K_{15} triangulates S_{11} and K_{16}, S_{13}. For instance, the embedding of K_{15} we consider has a shortest splitting cycle of type $(3, 8)$ of length 12.

Type	$(1,\cdot)$	$(2,\cdot)$	$(3,\cdot)$	$(4,\cdot)$	$(5,\cdot)$	$(6,\cdot)$
K_{15}	8	11	12	13	14	
K_{16}	8	10	12	14	14	15

Again, it appears that the more balanced is the type, the more vertices are required to build a splitting cycle of the corresponding type. Considering larger n may just lead to something interesting. Note that this is correlated with the following lemma:

Lemma 4.3.1. *Let M be an embedding of a complete graph K_n on the surface of genus $g = \frac{(n-3)(n-4)}{12}$. Then a splitting cycle of type $(\lfloor \frac{g}{2} \rfloor, \lceil \frac{g}{2} \rceil)$ has length at least $\left\lceil \frac{5+\sqrt{2n^2-14n+25}}{2} \right\rceil$.*

Proof. Since K_n is a complete graph, it is not possible that the two sides of a splitting cycles have an interior vertex. Hence, after cutting along a splitting cycle, there is a map with one boundary and no interior vertex of genus at least the smaller value of the type. Let S be a splitting of type $(\lfloor \frac{g}{2} \rfloor, \lceil \frac{g}{2} \rceil)$. Let k be the length of S and M' be the map without interior vertices. M' has genus at least $\lfloor \frac{g}{2} \rfloor$ and so $\chi(M') \leq 2 - 2\lfloor \frac{g}{2} \rfloor - 1 \leq 2 - (g-1) - 1 = 2 - g$. M' has k vertices, $e \leq \frac{k(k-1)}{2}$ edges and f faces. The double counting of the number of edges gives $3f = 2e - k$ because all the edges are on exactly 2 faces except the k on the boundary. So $\chi(M') = k - e + 2\frac{e}{3} - \frac{k}{3} = \frac{2k-e}{3} \geq \frac{4k-k(k-1)}{6} = \frac{5k-k^2}{6}$. By putting together the two inequalities we obtain: $2-g \geq \frac{5k-k^2}{6}$ leading to $k^2-5k+6-6g \geq 0$. $\Delta = 25-4(6-6g) = 1+24g$ and so $k \geq \frac{5+\sqrt{1+24g}}{2} = \frac{5+\sqrt{1+2(n-3)(n-4)}}{2} = \frac{5+\sqrt{2n^2-14n+25}}{2}$. □

For K_{15} and K_{19} this gives that a balanced splitting cycle has length at least 11 and 14 respectively. When n is big enough then a splitting cycle of balanced type has length greater than $\frac{\sqrt{2}}{2}n$ and thus uses at least 70% of the vertices. It suggests that balanced splittings are difficult to build and may not exits for n big enough. In order to confirm that intuition it is needed to design a more efficient algorithm in our particular setting. The description of the algorithm we used is given in section 4.6. This algorithm is of critical importance in this study. Indeed thanks to that efficiency we discovered a counter-example to conjecture 4.1.2 but since it allows to go up to $n = 43$ easily it gives an intersecting view on the behavior of splitting cycles in bigger triangulations (see section 4.7.2). Note that it is also very efficient to find one splitting rather than an exhaustive search. We were able to find a splitting cycle in K_{1207} for instance in a few minutes. This last calculus made us realize that we can prove the existence of splitting cycles in Ringel and Youngs embeddings by finding a feature in the output cycles we obtained for big n (see section 4.5). Our implementation contains routine to create any embedding of K_n of Ringel and Youngs or Gross and Tucker for $n = 12s + 7$ by choosing the s. However, the other cases are made one by one.

Without the efficient algorithm it may also have been possible to find counter-examples by testing other embeddings of K_{15} which was the limit size in terms of time of computation. Uniform sampling is not known for embeddings of K_{15} (see chapter 5) so the best option would have been to use an approach similar to Ellingham and Stephens one.

4.4 Embeddings of Complete Graphs

As noted in Section 4.2.1, one only needs to test conjectures 4.1.1 and 4.1.2 against irreducible triangulations. Any triangulation whose graph is complete is irreducible. On the other hand, Sulanke's experimentation [Sulanke, 2006a] on irreducible genus two triangulations suggests that denser graphs, i.e. triangulations with fewer vertices, have longer — hence potentially fewer — splitting cycles. It is thus legitimate to confront the conjectures with complete graphs.

4.4.1 Existence

The existence of embeddings of complete graphs as triangulations has appeared quite early. It was a key point in the proof of Heawood's conjecture:

Conjecture 4.4.1 ([Heawood, 1890])**.** *Let Σ be a surface without boundary of characteristic χ and different from the sphere. Lets define:*

$$\gamma(\chi) = \left\lfloor \frac{7 + \sqrt{49 - 24\chi}}{2} \right\rfloor$$

Then any simple graph drawn on Σ can be colored using $\gamma(\chi)$ colors and this is tight.

The formula is not absolutely natural but it comes very easily from Euler formula. Since it involves very common ideas it is interesting to develop it with modern formulation. The main result of Heawood is in fact the following:

Proposition 4.4.2 ([Heawood, 1890])**.** *Let Σ be a surface without boundary of characteristic χ. Let G be a simple graph embedded on Σ. Then G has a vertex of degree less than $max(5, \gamma(\chi) - 1)$.*

Proof. Let k be the minimum degree of G. G has n vertices, e edges and f faces. by the usual vertex-edge incidence double counting, $e \geq \frac{kn}{2}$. If G embeds as a triangulation then $2e = 3f$ and in the general case $2e \geq 3f$ since edges can be added to obtain a triangulation. So $f \leq \frac{2e}{3}$. We obtain $\chi = n - e + f \leq n - e + \frac{2e}{3} = n - \frac{e}{3}$. So $\chi \leq n(1 - \frac{k}{6})$. If $k \geq 6$ then $n(1 - \frac{k}{6})$ is a decreasing function of n. So $n \geq k + 1$ implies that $\chi \leq k + 1 - \frac{k(k+1)}{6}$ leading to $k^2 - 5k + 6\chi - 6 \leq 0$. The discriminant of the equation is $\Delta = 25 - 4(6\chi - 6) = 49 - 24\chi$. So k is an integer between $\frac{5-\sqrt{49-24\chi}}{2}$ and $\frac{5+\sqrt{49-24\chi}}{2}$. Finally $k \leq \left\lfloor \frac{5+\sqrt{49-24\chi}}{2} \right\rfloor = \gamma(\chi) - 1$. □

Using this result we can prove the conjecture without the tightness. Indeed, it is possible to iteratively remove vertices of degree less than $\gamma(\chi)-1$ and then coloring them in reverse order. Unfortunately it is not interesting for the sphere since it is far from giving the bound of 4. This part of the result has been generalized in higher dimension as a conjecture by Kühnel [Kühnel, 1994]. Here, I do not pretend to make a careful study of higher dimension problems but I give the definitions required to understand the generalization. It uses the **Betti numbers** β_i that are defined as the dimension of the i-th homology group. For instance the β_1 of a surface is the dimension of its H_1, which is $2g$ for an orientable surface without boundary. A manifold is k connected if it is non-empty, arc-connected and such that its k first homotopy groups π_i are trivial. The π_i are defined in a similar way as the fundamental group of a surface except it involves higher dimensional objects (see [Hatcher, 2002]). We can now state the conjecture.

Conjecture 4.4.3 ([Kühnel, 1994])**.** *If the k-skeleton of a n-simplex embeds in a compact, $k-1$-connected $2k$-manifold then: $\binom{n-k-1}{k+1} \leq \binom{2k+1}{k+1}\beta_k$.*

The best result known so far on that conjecture is given by Goaoc et al. [Goaoc et al., 2015]. They show that with the same assumptions but the connectedness, the following inequality holds: $n \leq 2\beta_k\binom{2k+2}{k} + 2k + 5$.

The tightness in conjecture 4.4.1 consists in constructing a $\gamma(\chi)$-colorable graph embedded on the corresponding surface that is not $(\gamma(\chi)-1)$-colorable. This has been done by using complete graphs thus completing the proof of the conjecture for all surfaces except the Klein bottle. All the details can be found in the book of Ringel and Youngs [Ringel, 1974] (see also [Gross and Tucker, 1987, Sec. 5.1.5]). Note that the case of the Klein bottle is quite strange and a one page technical proof is given in that book. Here, we focus on the cases where those complete graphs embed as triangulations. We first state some arithmetic restriction on the size of the complete graph.

Lemma 4.4.4. *If the complete graph on n vertices K_n embeds as a triangulation then $n \equiv 0,\ 3,\ 4$ or 7 mod 12 for orientable surfaces and $n \not\equiv 2,\ 5$ mod 6 otherwise. In addition the genus of the triangulated surface is $g = \frac{(n-3)(n-4)}{12}$ in the orientable case and $g = \frac{(n-3)(n-4)}{6}$ otherwise.*

Proof. Let K_n be a complete graph that triangulates an orientable surface Σ of genus g. It has n vertices, $n(n-1)/2$ edges and $2/3 * n(n-1)/2$ faces. So we have $\chi(\Sigma) = n - n(n-1)/2 + 2/3 * n(n-1)/2 = \frac{6n-n(n-1)}{6} = \frac{7n-n^2}{6} = 2-2g$ so $g = \frac{n^2-7n+12}{12} = \frac{(n-3)(n-4)}{12}$. Since g is an integer, it implies that $(n-3)(n-4) \equiv 0$ mod 12

leading to the announced cases. The non-orientable case is done similarly but with $\chi(\Sigma) = 2 - g$. □

All the work described in the proof of the map color theorem by Ringel and Youngs consists in showing that in all the possible cases such a triangulation exists. Indeed, since for such embeddings of K_n we get $\chi = 2 - \frac{(n-3)(n-4)}{6}$ then, the Heawood bound gives $\gamma(\chi) = n$, which is obviously the minimum number of colors required to properly color K_n.

4.4.2 Non-Isomorphic Embeddings

Lawrencenko et al. [Lawrencenko et al., 1994] identify three triangular embeddings of the complete graph K_{19} which they prove to be non-isomorphic with the help of a computer (see [Korzhik and Voss, 2001] for a non-computer proof). Each occurs as a covering of one of the orientable genus two **base** maps as shown Figure 83. Those

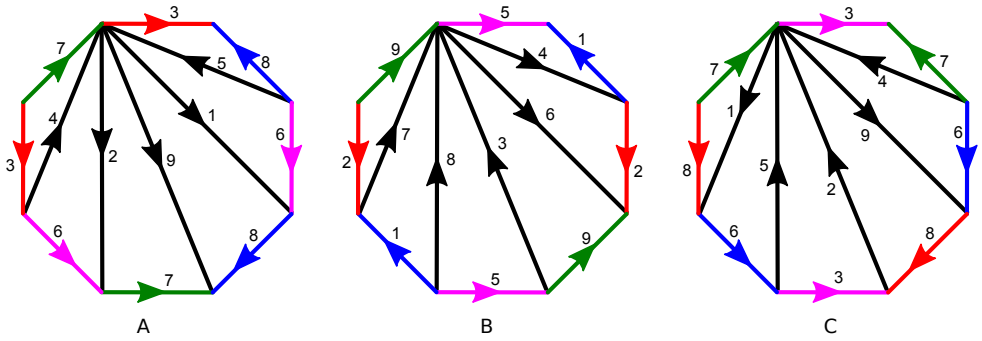

Figure 83: Three genus two maps with voltages in $\mathbb{Z}_{19}$. (The arcs opposite to those represented must receive opposite voltages.) For each octagon the sides should be pairwise identified according to their voltages. This results in each case A, B, C in a map with 9 edges and a single vertex. The corresponding coverings provide triangular embeddings of the complete graph K_{19}. Case A is from Ringel and Youngs contruction and case B from Gross and Tucker [Gross and Tucker, 1987].

coverings are generated from assignments of the directed edges (or arcs) of each base map to elements of $\mathbb{Z}_{19}$, the cyclic group of order 19. Two opposite arcs should receive opposite assignments and the sum of the assignments along an oriented facial cycle should be zero, whence the name of **voltage** given to such assignments [Gross and Tucker, 1987, Ch. 2]. The three voltages in [Lawrencenko et al., 1994] happen to be injective so that the arcs can be identified with their assignment. This leads to the following simple description of each covering. Given one of the base maps

with its voltage, we label the vertices of K_{19} with $\mathbb{Z}_{19}$ and declare (i, j, k) to be a triangle if and only if the three arcs with respective assignments $j - i, k - j$ and $i - k$ form a facial cycle in the base map. This construction can be generalized to produce embeddings of K_{12s+7} for every positive integer s. Gross and Tucker [Gross and Tucker, 1987] use the base map reproduced on Figure 84 as a $4(s+1)$-gon whose sides are pairwised identified according to their voltage. The resulting surface has genus $s+1$ and is covered by a triangular embedding M_{12s+7} of K_{12s+7} whose genus is $1 + s(12s + 7)$. As for M_{19}, we may identify the vertices of K_{12s+7} with $\mathbb{Z}_{12s+7}$ so that (i, j, k) is a triangle of M_{12s+7} if and only if $j - i, k - j$ and $i - k$ label the three sides of a triangle on the left Figure 84.

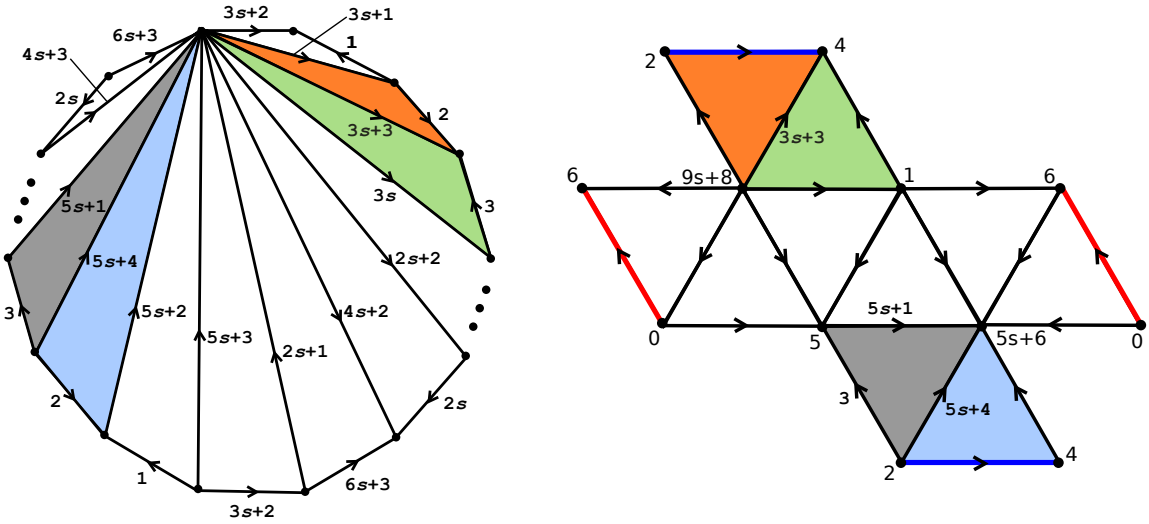

Figure 84: Left, a base map and its voltage for constructing the Gross and Tucker's embedding of K_{12s+7}. Note that this base map has a single vertex and is not a simplicial triangulation. Right, the 10 triangles form a sub-surface of genus one with one boundary component in M_{12s+7}. Each arc (u, v) is assigned the voltage $v - u$. Hence, the arc $(9s + 8, 4)$ receives the voltage $3s + 3 \mod (12s + 7)$. Some triangles are colored to show the corresponding covered triangles in the base map.

Recent studies indicate that the number of non-isomorphic triangulations of a surface by the same complete graph is actually quite high [Korzhik and Voss, 2001; Ellingham and Stephens, 2005; Grannell and Knor, 2012; Grannell and Knor, 2010]. It is far from being exhaustive but it contains all the different main ideas. We discuss those article in order to list the corresponding concepts. First, Korzhik and Voss show that many non-isomorphic embeddings can be constructed in Ringel and Youngs' way. They find a way to ensure that some transformations of the voltages

lead to non-isomorphic embeddings. This strategy only works for K_n with $n \equiv 4$ mod 12 and $n \equiv 7 \mod 12$, which corresponds to the simplest cases of Ringel and Youngs' embeddings. When $n = 12s+4$ or $n = 12s+7$, they prove that there exists at least 4^s non-isomorphic triangular embeddings of K_n. Their construction is far from giving all possible embeddings. For instance, it gives 4 different embeddings for K_{19}. Ellingham and Stephens provides a good idea of the number of non-isometric embeddings. They construct a canonical representation of a triangulation in a similar way that the one of Steiner Triple Systems. Steiner Triple System are of interest when considering triangular embeddings. Given a set of basis elements a Steiner Triple System is two sets of set of basis elements of size 2 and 3 such that each pair appears in exactly one triplet. Then a biembedding of a triple system leads to an embedding of the graph designed by the basis elements as vertices, the pairs as the edges and the triplet as the triangles. It only allowed to described two face-colorable embeddings (embeddings whose dual is bipartite). It is possible to use directly that kind of systems but Ellingham and Stephens proposed a description more appropriate that suits all embeddings and not 2-colorables. Then by brute-force they find 182, 200 non-isomorphic embedding of K_{12} on its corresponding non-orientable surface and 243, 088, 286 for K_{13}. The choice of 12 and 13 is strictly fixed by computation time. It is possible to handle smaller complete graphs directly (Altshuler, Bokowski and Schuchert show that there are exactly 59 non-isomorphic orientable triangulations of K_{12} [Altshuler et al., 1996]) and it becomes unreachable for bigger values. For purely theoretical counting Grannell and Knor improve the numbers of Korzhik and Voss by applying surgery tools and complete tripartite graph embeddings. Finally Grannell and Knor obtain n^{an^2} embeddings for an infinite but restrictive family of n and with a a constant. Remark that they also give an upper bound of $n^{n^2/3}$ for 2-face colorable embbedings. This suggests that it is the actual value for all the cases but nothing stronger seems to be known to date.

4.5 Many Splitting Cycles

Our implementation for searching splitting cycles lead us to discover that every triangular embedding M_{12s+7} given by Ringel and Youngs has a splitting cycle of type 1. This confirms Conjecture 4.1.1 for those triangulations. Note that by gluing irreducible triangulations along the boundary of a triangle we may obtain arbitrarily large irreducible triangulations that all possess a splitting cycle. However, the embeddings of complete graphs are not constructed this way and thus provide non-trivial confirmations of Conjecture 4.1.1. For $s \geq 3$, M_{12s+7} has indeed a splitting

cycle of length 8 given by the circular sequence of vertices:

$$\gamma_s = (0, 5, 2, 9s + 8, 6, 1, 4, 5s + 6)$$

This cycle bounds a perforated torus made of ten triangles as pictured on the right in Figure 85. On this figure the two copies of edge $(2, 4)$ should be identified, as well as the two copies of $(0, 6)$.

By construction, adding the same integer to the label of each vertex leaves the map unchanged. It provides a group action of $\mathbb{Z}_{12s+7}$ on the cycles of M_{12s+7}. The orbit of γ_s by this action provides us with $12s + 7$ distinct splitting cycles of type 1. Viewing the base map on the left Figure 85 as a fan of $4s + 2$ triangles, we can decompose the ten triangles of the perforated torus into two sub-fans of length five in the base map. See Figure 85, left. This construction can be generalized to exhibit many splitting cycles of type 1 by shifting the two fans, leading to the perforated tori as on Figure 85. Each of those cycles can be translated by the action of $\mathbb{Z}_{12s+7}$

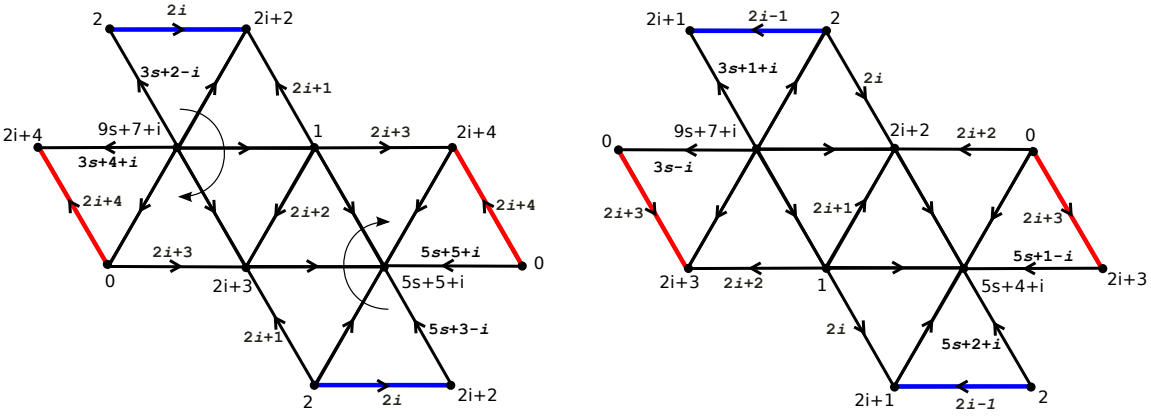

Figure 85: Left, the ten triangles of perforated torus in M_{12s+7} composed of two sub-fans glued along an edge with even voltage $2i + 2$. Right, a similar construction using a gluing edge with odd voltage.

to give the following $2(s - 1)(12s + 7)$ distinct splitting cycles of type 1:

$$\gamma_{s,i,k} = k + (0, 2i + 3, 2, 9s + 7 + i, 2i + 4, 1, 2i + 2, 5s + 5 + i)$$
$$\gamma'_{s,i,k} = k + (0, 2i + 2, 2, 5s + 4 + i, 2i + 3, 1, 2i + 1, 9s + 7 + i)$$

for $i \in [1, s - 1]$ and $k \in \mathbb{Z}_{12s+7}$. We can further generalize the construction by gluing larger fans of length $4j + 1$ as on Figure 86 to obtain splitting cycles of type j for $j = 1 \dots \lceil \frac{s-1}{2} \rceil$. Hence, the embeddings M_{12s+7} constitute an infinite family of irreducible triangulations that admit splitting cycles.

A similar construction applies to the embeddings of K_{12s+7} by Ringel and Youngs.

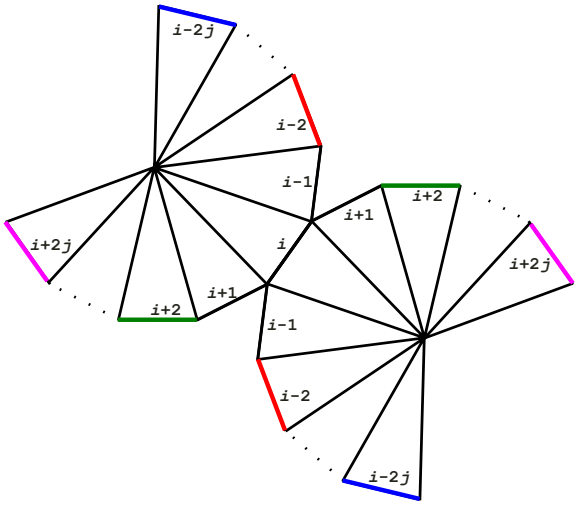

Figure 86: A punctured surface of genus j in M_{12s+7}. The thick edges should be glued according to their voltage.

4.6 Testing Algorithm for Conjecture 4.1.2

Here we provide the details of our implementation for testing Conjecture 4.1.2 on any triangular embedding M_n of the complete graph K_n (for a relevant n), not necessarily the embedding given by Ringel and Youngs or by Gross an Tucker. A straightforward approach consists in checking for every possible cycle in K_n whether it is splitting or not in M_n and computing its type in the former case. This assumes that we can list all the $\sum_{k=3}^{n} \frac{1}{2}\binom{n}{k}(k-1)!$ undirected cycles of K_n. For $n = 19$ this is already more than 9×10^{16} cycles to test, which is out of reach of current computers.

4.6.1 Pruning the Cycle Trees

Labelling the vertices of K_n with $\mathbb{Z}_n$, we identify a directed cycle with the sequence of its vertex labels starting with the smallest label. The (directed) cycles can be organized in $n-2$ rooted trees where the parent of a cycle is obtained by deleting its last vertex. Before exploring those **cycle trees**, we make two simple observations.

Remark 4.6.1. *If the automorphism group of M_n acts transitively on the set of its vertices we only need to consider the tree of cycles through vertex* 0.

Indeed, an automorphism of M_n does not change the type of a cycle. This remark applies to the three embeddings of K_{19} by Lawrencenko et al. [Lawrencenko et al.,

1994] and to all the orientable embeddings of K_{12s+7} by Ringel and Youngs and by Gross and Tucker. In each case $\mathbb{Z}_{12s+7}$ acts transitively by addition on vertex labels.

Remark 4.6.2. *We can assume that a splitting cycle does not contain two consecutive edges bounding a same triangle. We could otherwise replace the two edges by the third one in the triangle.*

This allows us to decrease by two the degree of the (non root) nodes in the cycle trees. A much more efficient pruning of the cycle trees is provided with the following simple heuristic. Suppose we are given a directed splitting cycle γ on M_n. We view an edge as a pair of oppositely directed arcs. Color in red or blue all the interior edges of the components of $M_n \setminus \gamma$ respectively to the left or right of γ. The resulting coloring satisfies that (i) every arc not in γ receives the same color as its opposite and (ii) for every vertex, the set of arcs directed inward that vertex is either monochromatic or composed of a red and a blue nonempty sequences separated by two arcs of γ. This leads to the following coloring test. As we go down a cycle tree from cycle $(v_0, \ldots, v_k)$ to $(v_0, \ldots, v_k, v_{k+1})$ we color the arcs pointing outward v_k and distinct from $v_k v_{k-1}$ and from $v_k v_{k+1}$ in red or blue according to whether they lie to the left or right of the directed subpath[1] (v_{k-1}, v_k, v_{k+1}). Each time an arc $v_k v$ is colored we check that

- $v v_k$ is either colored as $v_k v$ or has not been visited yet,
- the arc preceding $v_k v$ around v is not colored with the other color,
- the arc following $v_k v$ around v is not colored with the other color,
- if another arc uv has the same color as $v_k v$, then all the arcs entering v with the other color lie on the same side of the path (v_k, v, u).

If any of these tests fails the current partial coloring cannot be extended to fulfill the above conditions (i) and (ii). We can thus stop exploring the cycle subtree rooted at $(v_0, \ldots, v_{k+1})$. Together with remarks 4.6.1 and 4.6.2 those simple tests happens to be very effective and to reduce drastically the number of cycles to consider. See Section 4.7 for experimental results.

[1] In the non-orientable case, the left and right orientation should be propagated along the path $(v_0, \ldots, v_{k+1})$.

4.6.2 Computing the Type of a Cycle

When the cycle $\sigma := (v_0, \ldots, v_{k+1})$ passes the above tests it remains to check if σ is splitting and to compute its type. To this end we temporarily color the arcs pointing outward v_0 and the arcs pointing outward v_{k+1} in a way similar to that of the other v_i, $i = 1, \ldots, k$. We also perform the above tests and reject the cycle if anyone fails. At this point all the arcs outward the vertices of σ have been colored. We say that a vertex of $K_n - \sigma$ is **partially monochromatic** if all its colored inward arcs received the same color.

Lemma 4.6.3. *σ is separating if and only if all the vertices of $K_n - \sigma$ are partially monochromatic and all of the same color.*

Proof. If σ is separating then one of the two components of $M_n \setminus \sigma$ has no interior vertex. Otherwise, there would be an interior vertex in each component. However, those vertices would be connected by an edge of the (complete) graph of M_n leading to a contradiction. The direct implication in the Lemma easily follows. For the reverse implication, suppose that every vertex of $K_n - \sigma$ is partially monochromatic and that none of the arcs pointing to $K_n - \sigma$ has color c for some $c \in \{\text{blue, red}\}$. Then, every arc a with color c must connect two vertices of σ. The sides of the two triangles of M_n incident to a are thus either in σ or have color c. It ensues that the set of triangles of M_n each of whose sides is either in σ or colored with c forms a subsurface of M_n whose boundary is σ, proving that σ is separating. □

When M_n is orientable and σ is separating we can directly compute its type. With the notations of the preceding proof this amounts to compute the genus g' of the c-colored component of $M_n \setminus \sigma$. Denote by A_c the number of c-colored arcs and let $|\sigma|$ be the number of edges of σ. Using the Euler characteristic and double counting of the edge-triangle incidences, we easily obtain $g' = (A_c - 2|\sigma| + 6)/12$ when the c-colored component is orientable and $g' = (A_c - 2|\sigma| + 6)/6$ otherwise.

In practice, we maintain the following information for every vertex of $v \in M_n$: whether it belongs to the current cycle, the number of incident blue edges, the number of incident red edges, and a data-structure to store the colored arcs inward v. The updating of this information as well as the color tests can be performed in $O(\log n)$ time per newly colored edge using a simple data-structure that allows to find the next or previous colored arc around v and to insert or remove an arc in $O(\log n)$ time. Traversing the cycle trees in depth first order our algorithm thus spends $O(n \log n)$ time per cycle in the cycle trees. When M_n is non-orientable we can compute the Euler characteristic of the c-colored component in the same amount of time. If the c-colored component is orientable the other component must be non-orientable

and we are done. Otherwise we may need $O(|K_n|) = O(n^2)$ time to determine the orientability of the other component.

4.7 Results

In this section we detail the results for testing our algorithm on embeddings of complete graphs, leading to counter-examples to Conjectures 4.1.2 and 4.2.2. Remark 4.6.1 applies to each of the tested embeddings so that we only need to explore the cycle tree rooted at vertex 0.

4.7.1 Embeddings of K_{19}

Our smallest counter-examples to Conjecture 4.1.2 are provided by the three embeddings of K_{19} described in Section 4.4. We refer to them as A,B and C in accordance with Figure 83. A splitting cycle is said to have type k if it cuts the surface into components of respective genus k and $20 - k$. The next table shows for each embedding and each type the number NSC of splitting cycles of that type found as we traverse the cycle tree with root vertex 0 and the minimum length of any of those NSC cycles. Note that this minimum length would be the same if we would not take Remark 4.6.2 into account for pruning the cycle tree. For instance, we may note that every splitting cycle of type 4 in embedding B is Hamiltonian.

	Type	1	2	3	4	5-10
A	NSC	450	545	79	18	0
	Min Length	11	14	16	18	$\perp$
B	NSC	468	494	130	19	0
	Min Length	10	14	18	19	$\perp$
C	NSC	355	257	17	36	0
	Min Length	11	15	17	18	$\perp$

None of the three embeddings admits a splitting cycle of type 5 or more, thus disproving Conjecture 4.1.2. In particular, these embeddings cannot be split in a balanced way into two punctured surfaces of genus 10. As a side remark, the fact that the numbers in the table are distinct for the three embeddings is another confirmation that they are not isomorphic [Lawrencenko et al., 1994]. The next table indicates the proportion of contractible and splitting directed cycles among the visited nodes (that passes the color tests of Section 4.6.1) in the cycle tree. Since every cycle appears with both directions in the cycle tree, the number of splitting cycles is twice the sum in the corresponding NSC row in the previous table.

	# visited nodes	# contractible	# splitting
A	250221	36	2164
B	244229	36	2222
C	210808	36	1330

Hence, thanks to our pruning heuristic, less than 3×10^5 of the 1.8×10^{17} directed cycles of K_{19} are visited. We remark that by a Dehn type argument in the universal cover of the triangulation (see [Stillwell, 1993, Sec. 6.1.3]) a contractible cycle that does not contain two consecutive edges of any triangle must be the link (i.e. the boundary of the star) of a vertex. Since every cycle in the cycle tree must contain vertex 0 this leaves 18 link cycles as contractible cycles and explains the 36 found in each row of the previous table.

4.7.2 More Counter-Examples

Table 1 summarizes our results for the Ringel and Youngs embeddings of K_n with $n \in \{15, 16, 19, 27, 28, 31, 39, 40, 43\}$. Each entry in the table gives the length of the smallest splitting cycle of a given type for some K_n, if any. The last row indicates the largest possible type of a splitting cycle in the Ringel and Youngs embedding of K_n. It took less than 10 seconds to explore the pruned cycle tree for K_{19}, less than one hour for K_{31} and about half a day for K_{43}. We have also tested some of the Gross and Tucker embeddings of K_{12s+7}. Some results are listed in Table 2.

4.8 Conclusion

Our counter-examples to Conjecture 4.1.2 were checked with the help of a computer. Can we give a formal proof that would not recourse to a computer (at least for M_{19})? Although we could not find such a proof, three points seem relevant to this purpose.

- Suppose there is a splitting cycle of type 10 with k edges in M_{19}. As noted in the proof of Lemma 4.6.3, on one side of the cut surface we have no interior vertex. Let f and e be the number of faces and edges on this side (the number of vertices is k). By double counting incidences we have $3f = 2e - k$, and by Euler's formula: $k - e + f = 2 - 20 - 1 = -19$. It ensues that $e = 2k + 57$. Because the graph is simple we also have $e \leq \binom{k}{2}$. This implies $k^2 - 5k - 114 > 0$ and in turn $14 \leq k \leq 19$. A similar computation shows that a splitting cycle of type half the genus of M_{12s+7} has length at least $(5 + \sqrt{(1 + (24s + 7)^2)/2})/2$.

Type \ K_n	K_{15}	K_{16}	K_{19}	K_{27}	K_{28}	K_{31}	K_{39}	K_{40}	K_{43}
1	8	10	11	12	12	8	12	10	8
2	11	12	14	16	17	13	15	15	11
3	12	14	16	19	18	15	20	18	12
4	13	16	18	20	$\perp$	17	24	19	15
5	14	16	$\perp$	27	$\perp$	20	26	24	18
6		16	$\perp$	$\perp$	$\perp$	21	30	26	20
7			$\perp$	$\perp$	$\perp$	23	32	28	21
8			$\perp$	$\perp$	$\perp$	24	$\perp$	30	23
9			$\perp$	$\perp$	$\perp$	28	$\perp$	33	24
10			$\perp$	$\perp$	$\perp$	28	$\perp$	35	25
11				$\perp$	$\perp$	29	$\perp$	36	27
12				$\perp$	$\perp$	$\perp$	$\perp$	38	29
13				$\perp$	$\perp$	$\perp$	$\perp$	40	30
14				$\perp$	$\perp$	$\perp$	$\perp$	$\perp$	31
$\vdots$				$\perp$	$\perp$	$\perp$	$\perp$	$\perp$	$\vdots$
29						$\perp$	$\perp$	$\perp$	42
30						$\perp$	$\perp$	$\perp$	$\perp$
max type	5	6	10	23	25	31	52	55	65

Table 1: Minimal size of splitting cycles of Ringel and Youngs embeddings according to their type.

- Every splitting cycle leads to an arrangement of $12s+7$ splitting cycles thanks to the action of $\mathbb{Z}_{12s+7}$ on M_{12s+7}. Is it possible to take advantage of this arrangement to obtain an "impossible" dissection of M_{12s+7} when assuming the existence of a splitting cycle of type $\lfloor g(M_{12s+7})/2 \rfloor$?

- Finally, we saw in Section 4.5 that M_{12s+7} contains many short splitting cycles. On the other hand, the first above point tells that a splitting cycle γ of type half the genus is relatively long, hence must cut many of the short splitting cycles. Being separating, γ has to cut every other cycle an even number of times. Would this enforce γ to have length larger than $12s+7$, leading to a contradiction?

Our counter-examples show that it is not always possible to split a genus g triangulation into two genus $g/2$ triangulations. Looking at the tables in Section 4.7.2,

Type \ K_n	K_{19}	K_{31}	K_{43}
1	10	8	8
2	14	13	11
3	18	17	12
4	19	19	15
5	$\perp$	20	18
6	$\perp$	25	23
7	$\perp$	26	26
8	$\perp$	26	26
9	$\perp$	$\perp$	28
10	$\perp$	$\perp$	34
11		$\perp$	34
12		$\perp$	35
13		$\perp$	35
14		$\perp$	$\perp$
max	10	31	65

Table 2: Minimal size of splitting cycles of Gross and Tucker embeddings according to their type.

it seems that the proportion of the types of the splitting cycles of M_{12s+7} is roughly decreasing as s grows. This leads to the following conjecture in replacement of Conjecture 4.1.2.

Conjecture 4.8.1. *For any positive real number $\alpha \leq 1/2$, there exists a triangulation (or a graph embedding of face-width at least 3) of arbitrarily large genus g that has no splitting cycle of type larger than αg.*

Chapter 5

Bijections for Toroidal Triangulations

5.1 Introduction

In this chapter, we consider graphs embedded on a surface with no contractible cycle of size 1 or 2 ,i.e. no contractible loop-edges and no multiple edges forming a contractible cycle. Note that this is a assumption weaker than the graph being simple. We denote by n the number of vertices, m the number of edges and f the number of faces of a given map all along this chapter.

Poulalhon and Schaeffer introduced in [Poulalhon and Schaeffer, 2006] a method (called here PS method for short) to linearly encode a planar triangulation with a binary word of length $\log_2 \binom{4n}{n} \sim n \log_2(\frac{256}{27}) \approx 3.2451\, n$ bits. This is asymptotically optimal since it matches the information theory lower bound. The method is the following. Given a planar triangulation G and its minimal Schnyder wood of G, a special depth-first search algorithm is applied by "following" ingoing edges and "cutting" outgoing ones. The algorithm outputs a rooted spanning tree with exactly two leaves (also called stems) on each vertex from which the original triangulation can be recovered in a straightforward way. This tree can be encoded very efficiently. Besides its interesting encoding properties, this method gives a bijection between planar triangulations and a particular type of planar trees. A very nice application of this work is that it allows to have a linear time uniform sampling for plane triangulation. : it is possible to generate a random tree corresponding to a plane triangulation with n vertices in time linear in n. The transformation from the tree to the triangulation is also linear. The distribution of the random trees is uniform. No so efficient uniform sampling is known for higher genus.

Castelli Aleardi, Fusy and Lewiner [Castelli Aleardi et al., 2010] adapt the PS method to encode planar triangulations with boundaries. A consequence of this work is that any triangulation of an oriented surface can be encoded by cutting the surface along non-contractible cycles and see the surface as a planar map with boundaries. This method is a first attempt to generalize the PS algorithm to higher genus. The algorithm obtained is asymptotically optimal (in terms of number of bits) but it is not linear, nor bijective. Generalizations to plane maps that allow α-orientations for $\alpha > 3$ can be found in [Albenque and Poulalhon, 2013; Bernardi, 2007; Duchi et al., 2013]. Interestingly, this kind of generalization can also be imagined in the torus case by using the same kind of tools we use in this chapter.

The goal of this chapter is to present a new generalization of the PS algorithm to higher genus based on some strong structural properties. Applied on the canonical Schnyder woods of a toroidal triangulation, what remains after the execution of the algorithm is a rooted unicellular map that can be encoded optimally using $3.2451\,n$ bits. Moreover, the algorithm runs in linear time and leads to a new bijection between toroidal triangulations and a particular type of unicellular maps.

The main issue while trying to extend the PS algorithm to higher genus is the **accessibility**, i.e. the existence of an oriented path from every vertex to the special face. Accessibility toward the outer face is given almost for free in the planar case because of Euler's formula that sums to a strictly positive value. For an oriented surface of genus $g \geq 1$ new difficulties occur. Already in genus 1, even if the orientation is minimal and accessible the PS algorithm can visit all the vertices but not all the corners of the map because of the existence of non-contractible cycles. We can show that the canonical orientation that we define avoids this problem. In genus $g \geq 2$ things get even more difficult.

Another problem is to recover the original map after the execution of the algorithm. If what remains after the execution of PS method is a spanning unicellular map then the map can be recovered with the same simple rules as in the plane. Unfortunately, for many minimal orientations the algorithm leads to a spanning unicellular embedded graph that is not a map (the only face is not a disk) and it is not possible to directly recover the original map. Here again, the choice of the canonical orientation ensures that this never happens.

Finally the method presented here can be implemented in linear time. Clearly the execution of the PS algorithm is linear but the difficulty is to provide the algorithm with the appropriate input orientation. Computing the minimal Schnyder wood of a planar triangulation can be done in linear time quite easily by using a so-called shelling order (or canonical order, see [Kant, 1996]). Other similar ad-hoc linear algorithms can be found for other kinds of α-orientations of planar maps (see for

example [Fusy, 2007, Chapter 3]). Such methods are not known in higher genus. We solve this problems by first computing an orientation in the canonical lattice of Section 2.4.4 and then go down in the lattice to find the minimal orientation. All this can be performed in linear time.

In Section 5.2, we introduce a reformulation of the Poulalhon and Schaeffer original algorithm that is applicable to any orientation of any map on an oriented surface. The main theorem of this chapter is proved in Section 5.3, that is, for a toroidal triangulation given with an appropriate root and orientation, the output of the algorithm is a toroidal spanning unicellular map. In Section 5.4, we show how one can recover the original triangulation from the output. This output is then used in Section 5.5 to optimally encode a toroidal triangulation. The linear time complexity of the method is discussed in Section 5.6. In Section 5.7, we exhibit a bijection between appropriately rooted toroidal triangulations and rooted toroidal unicellular maps. A similar bijection with non-rooted maps is described in Section 5.9. To obtain the non-rooted bijection, further structural results concerning the particular Schnyder woods considered in this chapter are given in Section 5.8. Finally, a possible generalization to higher genus is discussed in Section 5.10.

5.2 The Poulalhon and Schaeffer Algorithm on Oriented Surfaces

In this section we introduce a reformulation of the Poulalhon and Schaeffer original algorithm. This version is more general in order to be applicable to any orientation of any map on an oriented surface. The execution slightly differs from the original formulation, even on planar triangulations. In [Poulalhon and Schaeffer, 2006], the authors first delete some outer edges of the triangulation before executing the algorithm. We do not consider any edge as special here since we want to apply the algorithm on any surface, but the core of the algorithm is the same. We show general properties of the algorithm in this section before considering toroidal triangulations in the forthcoming sections.

Algorithm PS

Input : A map whose graph G is oriented, a root vertex v_0 and a root edge e_0 incident to v_0.

Output : An embedded graph U with stems.

1. Let $v := v_0$, $e := e_0$, $U := \emptyset$.

2. Let v' be the extremity of e different from v.

Case 1 : e **is unmarked and entering** v**.** Add e to U and set $v := v'$.

Case 2 : e **is unmarked and leaving** v**.** Add a stem to U incident to v and corresponding to e.

Case 3 : e **is already marked and entering** v**.** Do nothing.

Case 4 : e **is already marked and leaving** v**.** Set $v := v'$.

3. Mark e.

4. Let e be the next edge around v in counterclockwise order after the current e.

5. While $(v, e) \neq (v_0, e_0)$ go back to 2.

6. Return U.

We insist on the fact that the output of Algorithm PS is a graph embedded on the same surface as the input map but that this embedded graph is not necessarily a map (i.e some faces may not be homeomorphic to open disks). In the following section we show that in our specific case the output U is a unicellular map.

Consider any oriented map G on an oriented surface given with a root vertex v_0 and a root edge e_0 incident to v_0. When Algorithm PS is considering a couple (v, e) we see this like it is considering the angle at v that is just before e in clockwise order. The particular choice of v_0 and e_0 is thus in fact a particular choice of a root corner c_0 that automatically defines a root vertex v_0, a root edge e_0, as well as a root face f_0. From now on we consider that the input of Algorithm PS is an oriented map plus a root angle (without specifying the root vertex, face and edge).

The **corner graph** of G, is the graph defined on the corners of G and where two corners are adjacent if and only if they are consecutive around a vertex or around a face. An execution of Algorithm PS can be seen as a walk in the corner graph. Figure 87 illustrates the behavior of the algorithm corresponding to Case 1 to 4. In each case, the algorithm is considering the corner in top left position and depending on the marking of the edge and its orientation the next corner that is considered is the one that is the end of the magenta arc of the corner graph. The cyan edge of Case 1 represents the edge that is added to U by the algorithm. The stems of U added in Case 2 are not represented in cyan, in fact we will represent them later by an edge in the dual. Indeed seeing the execution of Algorithm PS as a walk in the corner graph enables us to show that Algorithm PS behaves exactly the same in the primal or in the dual map (as explained later).

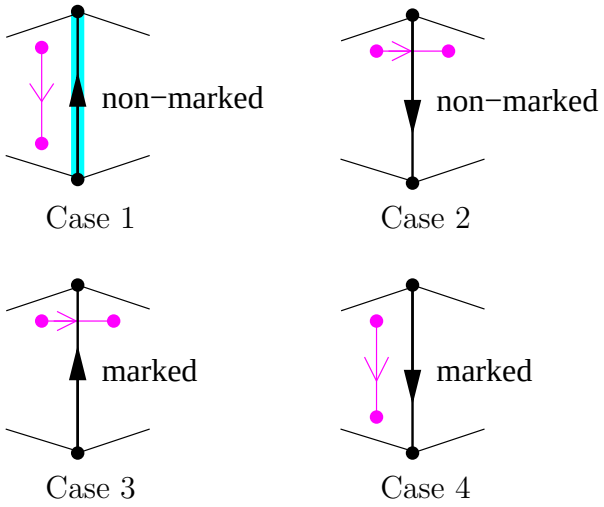

Figure 87: The four cases of ALGORITHM PS .

On Figure 88, we give an example of an execution of ALGORITHM PS on the orientation corresponding to the minimal HTC Schnyder wood of K_7 of Figure 52.

Let c be a particular corner of the map G. It is adjacent to four other corners in the **corner graph** (see Figure 89). Let v, f be such that c is a corner of vertex v and face f. The **next-vertex** (resp. **previous-vertex**) corner of c is the angle appearing just after (resp. before) c in counterclockwise order around v. Similarly, the **next-face** (resp. **previous-face**) angle of c is the corner appearing just after (resp. before) c in clockwise order around f. These definitions enable one to orient consistently the edges of the corner graph like in Figure 89 so that for every oriented edge (c, c'), c' is a next-vertex or next-face corner of c.

The different cases depicted in Figure 87 show that an execution of ALGORITHM PS is just an oriented walk in the corner graph (i.e. a walk that is following the orientation of the edges described in Figure 89). The condition in the while loop ensures that when the algorithm terminates, this walk is back to the root corner. The following proposition shows that the algorithm actually terminates:

Proposition 5.2.1. *Consider an oriented map G on an oriented surface and a root corner c_0. The execution of* ALGORITHM PS *on (G, c_0) terminates and corresponds to a cycle in the corner graph.*

Proof. We consider the oriented walk W in the corner graph corresponding to the execution of ALGORITHM PS . Note that W may be infinite. The walk W starts with c_0, and if it is finite it ends with c_0 and contains no other occurrence of c_0 (otherwise

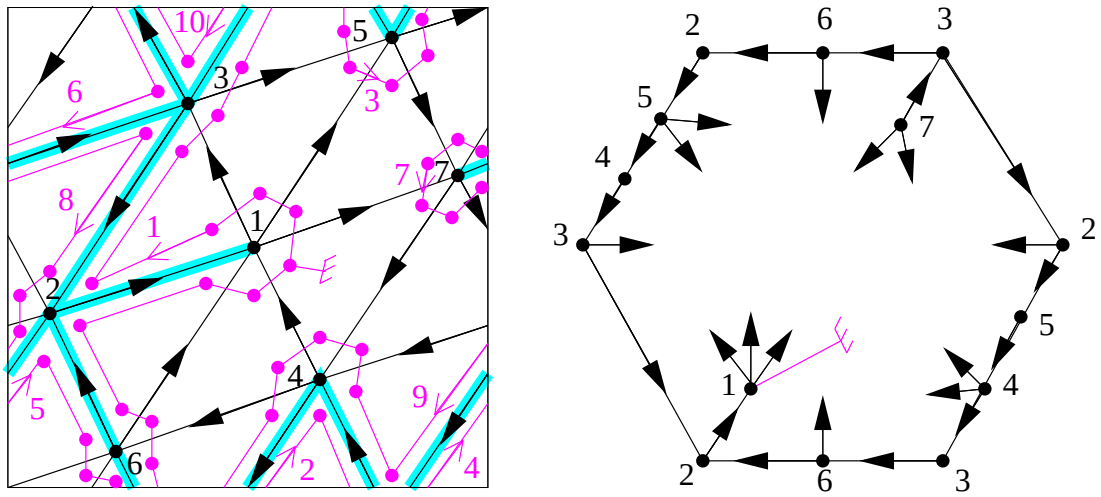

Figure 88: An execution of ALGORITHM PS on K_7 given with the orientation corresponding to the minimal HTC Schnyder wood of Figure 52. Vertices are numbered in black. The root corner is identified by a root symbol and chosen in the face for which the orientation is minimal (i.e. the shaded face of Figure 52). The magenta arrows and numbers are here to help the reader to follow the cycle in the corner graph. The output U is a toroidal unicellular map, represented here as an hexagon where the opposite sides are identified.

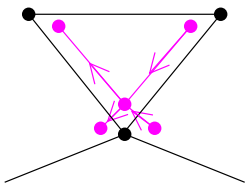

Figure 89: Orientation of the edges of the corner graph.

the algorithm should have stopped earlier). Toward a contradiction, suppose that W is not simple (i.e. some corners different from the root corner c_0 are repeated). Let $a \neq c_0$ be the first corner along W that is met for the second time. Let c_1, c_2 be the corners appearing before the first and second occurrence of c in W, respectively. Note that $c_1 \neq c_2$ by the choice of c.

If c_1 is the previous-vertex corner of c, then c_2 is the previous-face corner of c. When the algorithm considers c_1, none of c and c_2 are already visited, thus edge e is not marked. Since the execution then goes to c after c_1, we are in Case 2 and

the edge e between c and c_1 is oriented from v, where v is the vertex incident to c. Afterward, when the algorithm reaches c_2, Case 3 applies and the algorithm cannot go to c, a contradiction. The case where c_1 is the previous-face corner of c is similar.

So W is simple. Since the corner graph is finite, W is finite. So the algorithm terminates, thus W ends on the root corner and W is a cycle. □

In the next section we see that in some particular cases the cycle in the corner graph corresponding to the execution of the PS algorithm (Proposition 5.2.1) can be shown to be Hamiltonian like on Figure 88.

By Proposition 5.2.1, a corner is considered at most once by ALGORITHM PS. This implies that the corners around an edge can be visited in different ways depicted on Figure 90. Consider an execution of ALGORITHM PS on G. Let C be the cycle formed in the corner graph by Proposition 5.2.1. Let P be the set of edges of the output U (without the stems) and Q be the set of dual edges of edges of G corresponding to stems of U. These edges are represented on Figure 90 in cyan for P and in yellow for Q. They are considered with their orientation (recall that the dual edge e^* of an edge e goes from the face on the left of e to the face on the right of e). Note that C does not cross an edge of P or Q, and moreover P and Q do not intersect (i.e. an edge can be in P or its dual in Q but both cases cannot happen).

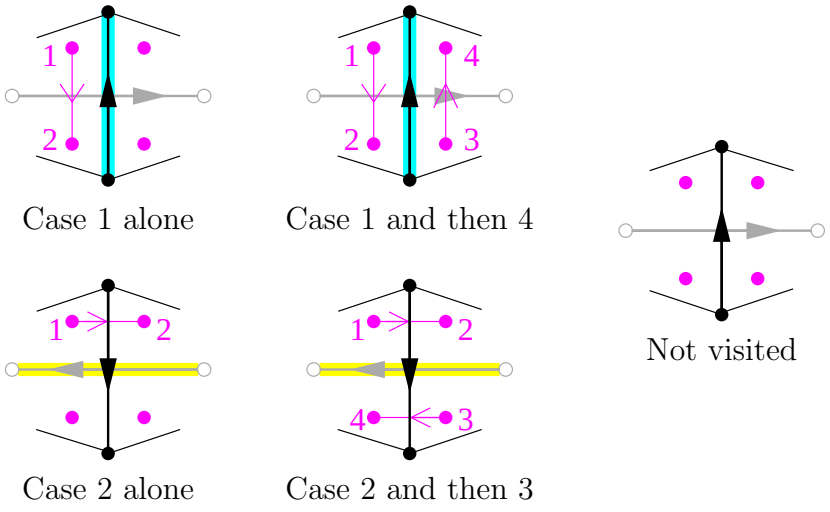

Figure 90: The different cases of ALGORITHM PS seen in a dual way. The number of the corners gives the order in which the algorithm visits them (unvisited corners are not numbered). The edges of P and Q are respectively cyan and yellow.

One can remark that the cases of Figure 90 are dual of each other. One can

see that ALGORITHM PS behaves exactly the same if applied on the primal map or on the dual map. The only modifications to make is to start the algorithm with the face f_0 as the root vertex, the dual of edge e_0 as the root edge and to replace counterclockwise by clockwise at Line 4. Then the cycle C formed in the corner graph is exactly the same and the output is Q with stems corresponding to P (instead of P with stems corresponding to Q). Note that this duality is also illustrated by the fact that the minimality of the orientation of G w.r.t. the root face is nothing else than the accessibility of the dual orientation toward the root face. Indeed, a clockwise null-homologous oriented subgraph of G w.r.t f_0 corresponds to a directed cut of the dual where all the edges are oriented from the part containing f_0. The following lemma shows the connectivity of P and Q:

Lemma 5.2.2. *At each step of the algorithm, for every vertex v appearing in an edge of P (resp. Q), there is an oriented path from v to v_0 (resp. f_0) consisting only of edges of P (resp. Q). In particular P and Q are connected.*

Proof. If at a step a new vertex is reached then it correspond to Case 1 and the corresponding edge is added in P and oriented from the new vertex, so the property is satisfied by induction. As observed earlier the algorithm behaves similarly in the dual map. □

Let $\overline{C}$ be the set of corners of G that are not in C. Any edge of G is bounded by exactly 4 corners. Since C is a cycle, the 4 corners around an edge are either all in C, all in $\overline{C}$ or 2 in each set (see Figure 90). Moreover, if they are 2 in each set, these sets are separated by an edge of P or an edge of Q. Hence the frontier between C and $\overline{C}$ is a set of edges of P and Q. Moreover this frontier is an union of oriented closed walks of P and of oriented closed walks of Q. In the next section we study this frontier in more details to show that $\overline{C}$ is empty in the case considered there.

5.3 From Triangulations to Unicellular Maps

Let G be a toroidal triangulation. In order to choose appropriately the root corner c_0, we have to consider separating triangles. A **triangle** is a closed walk of size 3 (it is not necessarily a cycle since non-null-homologous loops are allowed and it is not necessarily null-homologous). A **separating triangle** is a null-homologous triangle that is different from a face of G. We say that a corner is **in the strict interior of a separating triangle** if it is in its null-homologous region and not incident to a vertex of the triangle. We choose as root corner c_0 any corner that is not in the strict interior of a separating triangle. One can easily see that such a corner c_0 always

exists. Indeed the interiors of two null-homologous triangles are either disjoint or one is included in the other. So, the corners that are incident to a null-homologous triangle whose interior is maximal by inclusion satisfy the property.

The choice of a root corner that is not in the interior of a separating triangle is mandatory to apply Poulalhon and Schaeffer method. Indeed, in a 3-orientation of a toroidal triangulation, by Euler's formula, all the edges that are incident to a separating triangle and in its interior are oriented towards the triangle. Thus if one apply ALGORITHM PS from a corner in the strict interior of a triangle, it will be stucked in its interior and will not visit all the vertices.

A subgraph of a graph is **spanning** if it is covering all the vertices. The main result of this section is the following theorem (see Figure 88 for an example):

Theorem 5.3.1. *Consider a toroidal triangulation G, a root corner c_0 that is not in the strict interior of a separating triangle and the orientation of the edges of G corresponding to the minimal HTC Schnyder wood with respect to the root face f_0 containing c_0. Then the output U of* ALGORITHM PS *applied on (G, c_0) is a toroidal spanning unicellular map.*

Consider a toroidal triangulation G, a root corner c_0 that is not in the strict interior of a separating triangle and the orientation of the edges of G corresponding to the minimal HTC Schnyder wood with respect to the root face f_0 containing c_0. Let U be the output of ALGORITHM PS applied on (G, c_0). We use the same notation as in previous section: the cycle in the corner graph is C, the set of corners that are not in C is $\overline{C}$, the set of edges of U is P, the dual edges of stems of U is Q. The frontier is defined as follows. Add to C and $\overline{C}$ all the edges between adjacent corners. Then the frontier is the set of edges of G and G^* that are not crossing by one of this edges.

Lemma 5.3.2. *The frontier between C and $\overline{C}$ contains no oriented closed walk of Q.*

Proof. Suppose by contradiction that there exists such a walk W. Then along this walk, all the dual edges of W are edges of G oriented from the region containing C toward $\overline{C}$ as one can see in Figure 90. If W is non-null-homologous, then W contains an oriented non-null-homologous cycle, a contradiction to Lemma 2.4.23. So W is null-homologous. So it contains an oriented null-homologous cycle W', and then either C is in the null-homologous region delimited by W', or not. The two case are considered below:

Suppose first that C lies in the non-null-homologous region of W'. Then consider the plane map G' obtained from G by keeping only the vertices and edges that lie

(strictly) in the null-homologous region delimited by W'. Let n' be the number of vertices of G'. All the edges incident to G' that are not in G' are entering G'. So in G' all the vertices have outdegree 3 as we are considering 3-orientations of G. Thus the number of edges of G' is exactly $3n'$, contradicting the fact that the maximal number of edges of planar map on n vertices is $3n-6$ by Euler's formula.

Suppose now that C lies in the null-homologous region of W'. All the dual edges of W' are edges of G oriented from its null-homologous region toward its exterior. Consider the graph G_{out} obtained from G by removing all the edges that are cut by W' and all the vertices and edges that lie in the null-homologous region of W'. As G is a map, the face of G_{out} containing W' is homeomorphic to an open disk. Let F be its facial walk (in G_{out}) and let k be the length of F. We consider the map obtained from the facial walk F by putting back the vertices and edges that lied inside. We transform this map into a plane map G' by duplicating the vertices and edges appearing several times in F, in order to obtain a triangulation of a cycle of length k. Let n', m', f' be the number of vertices, edges and faces of G'. Every inner vertex of G' has outdegree 3, there are no other inner edges, so the total number of edges of G' is $m' = 3(n'-k)+k$. All the inner faces have size 3 and the outer face has size k, so $2m' = 3(f'-1)+k$. By Euler's formula $n'-m'+f' = 2$. Combining the three equalities gives $k = 3$ and F is hence a separating triangle of G. This contradicts the choice of the root corner, as it should not lie in the strict interior of a separating triangle. □

Lemma 5.3.3. *The cycle C is a Hamiltonian cycle of the corner graph, all the edges of G are marked exactly twice, the subgraph Q of G^* is spanning, and, if $n \geq 2$, the subgraph P of G is spanning.*

Proof. Suppose for a contradiction that $\overline{C}$ is non empty. By Lemma 5.3.2 and Section 5.2, the frontier T between C and $\overline{C}$ is an union of oriented closed walks of P. Hence a face of G has either all its corners in C or all its corners in $\overline{C}$. Moreover T is a non-empty union of oriented closed walk of P that are oriented clockwise according to the set of faces containing $\overline{C}$ (see the first case of Figure 90). This set does not contain f_0 since c_0 is in f_0 and C. As in Section 2.4.4, let $\mathcal{F}$ be the set of counterclockwise facial walks of G and F_0 be the counterclockwise facial walk of f_0. Let $\mathcal{F}' = \mathcal{F} \setminus F_0$, and $\mathcal{F}_{\overline{C}} \subseteq \mathcal{F}'$ be the set of counterclockwise facial walks of the faces containing $\overline{C}$. We have $T = -\sum_{F \in \mathcal{F}_{\overline{C}}} \partial(F)$. So T is a clockwise non-empty null-homologous oriented subgraph with respect to f_0. This contradicts Lemma 2.4.22 and the minimality of the orientation with respect to f_0. So $\overline{C}$ is empty, thus C is Hamiltonian and all the edges of G are marked twice.

Suppose for a contradiction that $n \geq 2$ and P is not spanning. Since the algorithm

starts at v_0, P is not covering a vertex v of G different from v_0. Then the corners around v cannot be visited since by Figure 90 the only way to move from a corner of one vertex to a corner of another vertex is through an edge of P incident to them. So P is spanning. The proof is similar for Q (note that in this case we have $f \geq 2$). □

Lemma 5.3.4. *The first cycle created in P (resp. in Q) by the algorithm is oriented.*

Proof. Let e be the first edge creating a cycle in P while executing ALGORITHM PS and consider the steps of ALGORITHM PS before e is added to P. So P is a tree during all these steps. For every vertex of P we define $P(v)$ the unique path from v to v_0 in P (while P is empty at the beginning of the execution, we define $P(v_0) = \{v_0\}$). By Lemma 5.2.2, this path $P(v)$ is an oriented path. We prove the following

Claim 4. *Consider a step of the algorithm before e is added to P and where the algorithm is considering a vertex v. Then all the corners around the vertices of P different from the vertices of $P(v)$ are already visited.*

Proof. Suppose by contradiction that there is such a step of the algorithm where some corners around the vertices of P different from the vertices of $P(v)$ have not been visited. Consider the first such step. Then clearly we are not at the beginning of the algorithm since $P = P(v) = \{v_0\}$. So at the step just before, the conclusion holds and now it does not hold anymore. Clearly at the step before we were considering a vertex v' distinct from v, otherwise $P(v)$ and P have not changed and we have the conclusion. So from v' to v we are either in Case 1 or Case 4 of ALGORITHM PS. If v has been considered by Case 1, then $P(v)$ contains $P(v')$ and the conclusion holds. If v has been considered by Case 4, then since P is a tree, all the corners around v' have been considered and v' is the only element of $P \setminus P(v)$ that is not in $P \setminus P(v')$. Thus the conclusion also holds. □

Consider the iteration of ALGORITHM PS where e is added to P. The edge e is added to P by Case 1, so e is oriented from a vertex u to a vertex v such that v is already in P or v is the root vertex v_0. Consider the step of the algorithm just before u is added to P. By Claim 4, vertex u is not in $P \setminus P(v)$ (otherwise e would have been considered before and it would be a stem). So $u \in P(v)$ and $P(v) \cup \{e\}$ induces an oriented cycle of G. The proof is similar for Q. □

Lemma 5.3.5. *P is a spanning unicellular map of G and Q is a spanning tree of G^*. Moreover one is the dual of the complement of the other.*

Proof. Suppose that Q contains a cycle, then by Lemma 5.3.4 it contains an oriented cycle of G^*. This cycle is null-homologous by Lemma 2.4.23. Recall that by Lemma 5.3.3, C is a Hamiltonian cycle, moreover it does not cross Q, a contradiction. So Q contains no cycle and is a tree.

By Lemma 5.3.3, all the edges of G are marked at the end. So every edge of G is either in P or its dual in Q (and not both). Thus P and Q are the dual of the complement of each other. So P is the dual of the complement of a spanning tree of G^*. Thus P is a spanning unicellular map of G. □

Theorem 5.3.1 is then a direct reformulation of Lemma 5.3.5 by the definition of P and Q:

A toroidal unicellular map on n vertices has exactly $n+1$ edges: $n-1$ edges of a tree plus 2 edges corresponding to the size of a basis of the homology (i.e. plus $2g$ in general for an oriented surface of genus g). Thus a consequence of Theorem 5.3.1 is that the obtained unicellular map U has exactly n vertices, $n+1$ edges and $2n-1$ stems since the total number of edges is $3n$. The orientation of G induces an orientation of U such that the stems are all outgoing, and such that while walking clockwise around the unique face of U from c_0, the first time an edge is met, it is oriented counterclockwise according to this face, see Figure 91 where all the tree-like parts and stems are not represented. There are two types of toroidal unicellular maps depicted on Figure 91. Two cycles of U may intersect either on a single vertex (square case) or on a path (hexagonal case). The square can be seen as a particular case of the hexagon where one side has length zero and thus the two corners of the hexagon are identified.

In Figure 92, we give several examples of executions of Algorithm PS on minimal 3-orientations. These examples show how important is the choice of the minimal HTC Schnyder wood in order to obtain Theorem 5.3.1. In particular, the third example shows that Algorithm PS can visit all the corners of the triangulation (i.e. the cycle in the corner graph is Hamiltonian) without outputting an unicellular map.

Note that the orientations of Figure 92 are not Schnyder woods. One may wonder if the fact of being a Schnyder wood is of any help for our method. This is not the case since there are examples of minimal Schnyder woods that are not HTC and where Algorithm PS does not visit all the vertices. One can obtain such an example by replicating 3 times horizontally and then 3 times vertically the second example of Figure 92 to form a 3×3 tiling and starts Algorithm PS from the same root corner. Conversely, there are minimal Schnyder woods that are not HTC where Algorithm PS does output a toroidal spanning unicellular map (the Schnyder wood of Figure 53 can serve as an example while starting from a corner of the only face oriented clockwise).

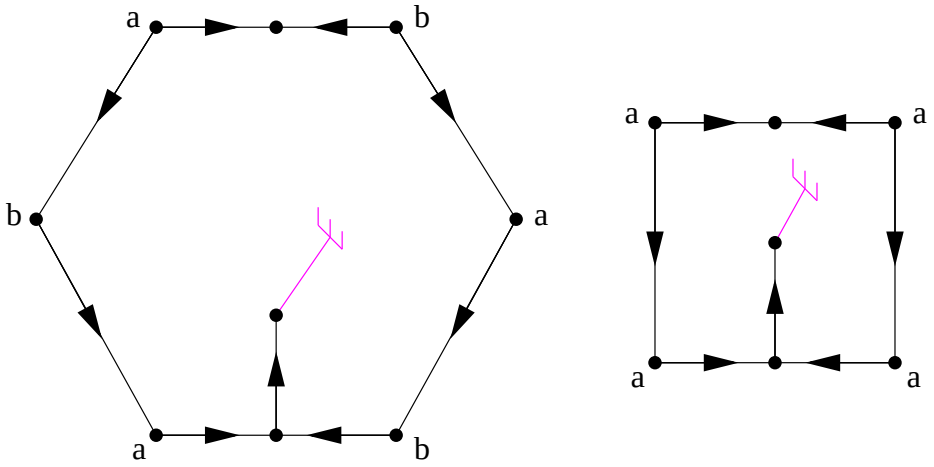

Hexagon Square

Figure 91: The two types of rooted toroidal unicellular maps.

5.4 Recovering the Original Triangulation

This section is dedicated to show how to recover the original triangulation from the output of ALGORITHM PS. The method is very similar to [Poulalhon and Schaeffer, 2006] since like in the plane the output has only one face that is homeomorphic to an open disk (i.e. a tree in the plane and an unicellular map in general).

Theorem 5.4.1. *Consider a toroidal triangulation G, a root corner c_0 that is not in the strict interior of a separating triangle and the orientation of the edges of G corresponding to the minimal HTC Schnyder wood with respect to the root face f_0 containing c_0. From the output U of* ALGORITHM PS *applied on (G, c_0) one can reattach all the stems to obtain G by starting from the root corner c_0 and walking along the face of U in counterclockwise order (according to this face): each time a stem is met, it is reattached in order to create a triangular face on its left side.*

thm 5.4.1 is illustrated on Figure 93 where one can check that the obtained toroidal triangulation is K_7 (like on the input of Figure 88).

In fact in this section we define a method, more general than the one described in thm 5.4.1, that will be useful in next sections.

Let $\mathcal{U}_r(n)$ denote the set of toroidal unicellular maps U rooted on a particular corner, with exactly n vertices, $n+1$ edges and $2n-1$ stems satisfying the following property. A vertex that is not the root, has exactly 2 stems if it is not a corner, 1 stem if it is the corner of a hexagon and 0 stem if it is the corner of a square. The

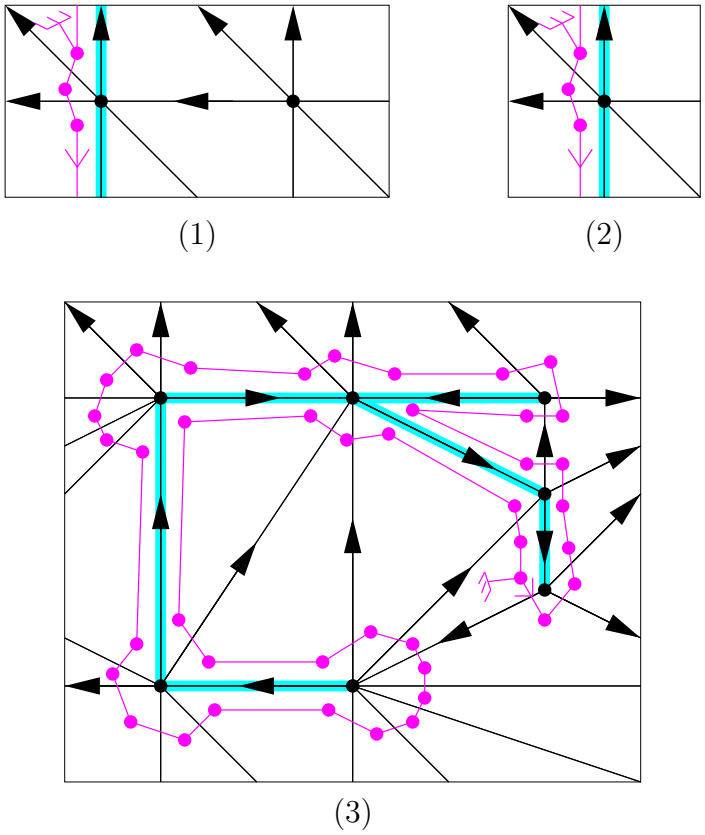

Figure 92: Examples of minimal 3-orientations that are not HTC Schnyder woods and where ALGORITHM PS respectively: (1) does not visit all the vertices, (2) visits all the vertices but not all the corners, and (3) visits all the corners but does not output an unicellular map.

root vertex has 1 additional stem, i.e. it has 3 stems if it is not a corner, 2 stems if it is the corner of a hexagon and 1 stem if it is the corner of a square. Note that the output U of ALGORITHM PS given by Theorem 5.3.1 is an element of $\mathcal{U}_r(n)$.

Similarly to the planar case [Poulalhon and Schaeffer, 2006], we define a general way to reattached step by step all the stems of an element U of $\mathcal{U}_r(n)$. Let $U_0 = U$, and, for $1 \leq k \leq 2n-1$, let U_k be the map obtained from U_{k-1} by reattaching one of its stem (we explicit below which stem is reattached and how). The **special face** of U_0 is its only face. For $1 \leq k \leq 2n-1$, the **special face** of U_k is the face on the right of the stem of U_{k-1} that is reattached to obtain U_k. For $0 \leq k \leq 2n-1$, the border of the special face of U_k consists of a sequence of edges and stems. We

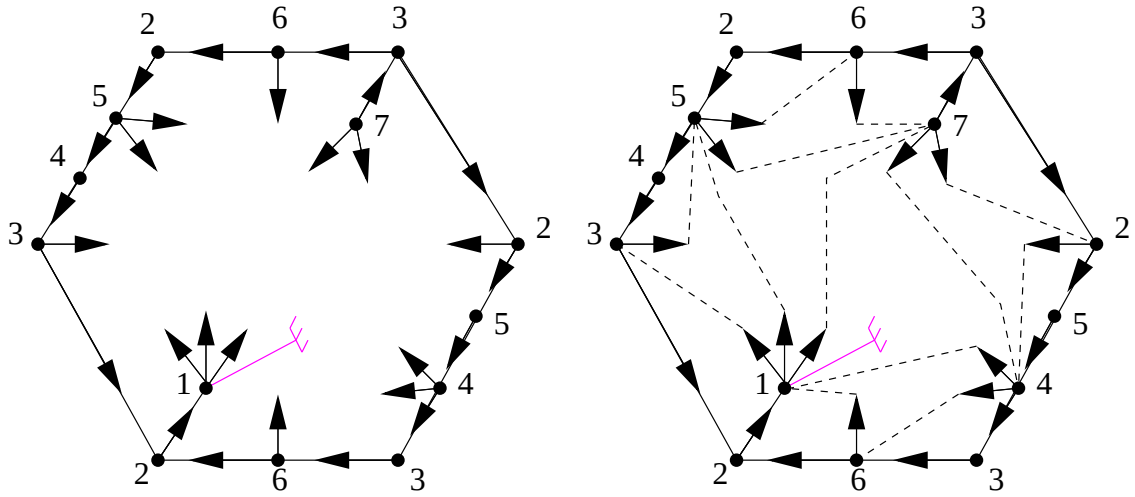

Figure 93: Example of how to recover the original toroidal triangulation K_7 from the output of ALGORITHM PS .

define an **admissible triple** as a sequence (e_1, e_2, s), appearing in counterclockwise order along the border of the special face of U_k, such that $e_1 = (u, v)$ and $e_2 = (v, w)$ are edges of U_k and s is a stem attached to w. The **closure** of the admissible triple consists in attaching s to u, so that it creates an edge (w, u) oriented from w to u and so that it creates a triangular face (u, v, w) on its left side. The **complete closure** of U consists in closing a sequence of admissible triple, i.e. for $1 \leq k \leq 2n-1$, the map U_k is obtained from U_{k-1} by closing any admissible triple.

Note that, for $0 \leq k \leq 2n-1$, the special face of U_k contains all the stems of U_k. The closure of a stem reduces the number of edges on the border of the special face and the number of stems by 1. At the beginning, the unicellular map U_0 has $n+1$ edges and $2n-1$ stems. So along the border of its special face, there are $2n+2$ edges and $2n-1$ stems. Thus there is exactly three more edges than stems on the border of the special face of U_0 and this is preserved while closing stems. So at each step there is necessarily at least one admissible triple and the sequence U_k is well defined. Since the difference of three is preserved, the special face of U_{2n-2} is a quadrangle with exactly one stem. So the reattachment of the last stem creates two faces that are triangles and at the end U_{2n-1} is a toroidal triangulation. Note that at a given step there might be several admissible triples but their closure are independent and the order in which they are performed does not modify the obtained triangulation U_{2n-1}.

We now apply the closure method to our particular case. Consider a toroidal triangulation G, a root corner c_0 that is not in the strict interior of a separating triangle and the orientation of the edges of G corresponding to the minimal HTC Schnyder wood with respect to the root face f_0. Let U be the output of ALGORITHM PS applied on (G, c_0).

Lemma 5.4.2. *When a stem of U is reattached to form the corresponding edge of G, it splits the (only) face of U into two faces. The root corner of U is in the face that is on the right side of the stem.*

Proof. By Lemma 5.3.3, the execution of ALGORITHM PS corresponds to an Hamiltonian cycle $C = (c_0, \dots, a_{2m}, c_0)$ in the corner graph of G. Thus C defines a total order $<$ on the corners of G where $c_i < c_j$ if and only if $i < j$. Let us consider now the corners on the face of U. Note that such a corner corresponds to several corners of G, that are consecutive in C and that are separated by a set of incoming edges of G (those incoming edges corresponding to stems of U). Thus the order on the corners of G defines automatically an order on the corners of U. The corners of U considered in clockwise order along the border of its face, starting from the root corner, correspond to a sequence of strictly increasing corners for $<$.

Consider a stem s of U that is reattached to form an edge e of G. Let c_s be the corner of U that is situated just before s (in clockwise order along the border of the face of U) and c'_s be the corner of U where s should be reattached. If $c'_s < c_s$, then when ALGORITHM PS consider the corner c_s, the edge corresponding to s is already marked and we are not in Case 2 of ALGORITHM PS . So $c_s < c'_s$ and c_0 is on the right side of s. □

Recall that U is an element of $\mathcal{U}_r(n)$ so we can apply on U the complete closure procedure described above. We use the same notation as before, i.e. let $U_0 = U$ and for $1 \leq k \leq 2n-1$, the map U_k is obtained from U_{k-1} by closing any admissible triple. The following lemma shows that the triangulation obtained by this method is G:

Lemma 5.4.3. *The complete closure of U is G, i.e. $U_{2n-1} = G$.*

Proof. We prove by induction on k that every face of U_k is a face of G, except for the special face. This is true for $k = 0$ since $U_0 = U$ has only one face, the special face. Let $0 \leq k \leq 2n-2$, and suppose by induction that every non-special face of U_k is a face of G. Let (e_1, e_2, s) be the admissible triple of U_k such that its closure leads to U_{k+1}, with $e_1 = (u, v)$ and $e_2 = (v, w)$. The closure of this triple leads to a triangular face (u, v, w) of U_{k+1}. This face is the only "new" non-special face while going from U_k to U_{k+1}.

Suppose, by contradiction, that this face (u, v, w) is not a face of G. Let c_v (resp. c_w) be the corner of U_k at the special face, between e_1 and e_2 (resp. e_2 and s). Since G is a triangulation, and (u, v, w) is not a face of G, there exists at least one stem of U_k that should be attached to c_v or c_w to form a proper edge of G. Let s' be such a stem that is the nearest from s. In G the edges corresponding so s and s' should be incident to the same triangular face. Let x be the origin of the stem s'. Let $z \in \{v, w\}$ such that s' should be reattached to z. If $z = v$, then s should be reattached to x to form a triangular face of G. If $z = w$, then s should be reattached to a common neighbor of w and x located on the border of the special face of U_k in counterclockwise order between w and x. So in both cases s should be reattached to a vertex y located on the border of the special face of U_k in counterclockwise order between w and x (with possibly $y = x$). To summarize s goes from w to y and s' from x to z, and z, x, y, w appear in clockwise order along the special face of U_k. By Lemma 5.4.2, the root corner is on the right side of both s and s', this is not possible since their right sides are disjoint, a contradiction.

So for $0 \leq k \leq 2n-2$, all the non-special faces of U_k are faces of G. In particular every face of U_{2n-1} except one is a face of G. Then clearly the (triangular) special face of U_{2n-1} is also a face of G, hence $U_{2n-1} = G$. □

Lemma 5.4.3 shows that one can recover the original triangulation from U with any sequence of admissible triples that are closed successively. This does not explain how to find the admissible triples efficiently. In fact the root corner can be used to find a particular admissible triple of U_k:

Lemma 5.4.4. *For $0 \leq k \leq 2n-2$, let s be the first stem met while walking counterclockwise from c_0 in the special face of U_k. Then before s, at least two edges are met and the last two of these edges form an admissible triple with s.*

Proof. Since s is the first stem met, there are only edges that are met before s. Suppose by contradiction that there is only zero or one edge met before s. Then the reattachment of s to form the corresponding edge of G is necessarily such that the root corner is on the left side of s, a contradiction to Lemma 5.4.2. So at least two edges are met before s and the last two of these edges form an admissible triple with s. □

Lemma 5.4.4 shows that one can reattach all the stems by walking once along the face of U in counterclockwise order. Thus we obtain thm 5.4.1.

Note that U is such that the complete closure procedure described here never **wraps over the root corner**, i.e. when a stem is reattached, the root corner is always on its right side (see Lemma 5.4.2). The property of never wrapping over

the root corner is called a **balanced corner** in [Albenque and Poulalhon, 2013]. Let $\mathcal{U}_{r,b}(n)$ denote the set of elements of $\mathcal{U}_r(n)$ that are balanced. So the output U of Algorithm PS given by thm 5.3.1 is an element of $\mathcal{U}_{r,b}(n)$. We exhibit in Section 5.7 a bijection between appropriately rooted toroidal triangulations and a particular subset of $\mathcal{U}_{r,b}(n)$.

The possibility to close admissible triples in any order to recover the original triangulation is interesting comparing to the simpler method of Theorem 5.4.1 since it enables to recover the triangulation even if the root corner is not given. This property is used in Section 5.9 to obtain a bijection between toroidal triangulations and some unrooted unicellular maps.

Moreover if the root corner is not given, then one can simply start from any corner of U, walk twice around the face of U in counterclockwise order and reattached all the admissible triples that are encountered along this walk. Walking twice ensure that at least one complete round is done from the root corner. Since only admissible triples are considered, we are sure that no unwanted reattachment is done during the process and that the final map is G. This enables to reconstruct G in linear time even if the root corner is not known. This property is used in Section 5.5.

5.5 Optimal Encoding

The results presented in the previous sections allow us to generalize the encoding of planar triangulations, defined by Poulalhon and Schaeffer [Poulalhon and Schaeffer, 2006], to triangulations of the torus. The construction is direct and it is hence really different from the one of [Castelli Aleardi et al., 2010] where triangulations of surfaces are cut in order to deal with planar triangulations with boundaries. Here we encode the unicellular map outputted by Algorithm PS by a plane rooted tree with n vertices and with exactly two stems attached to each vertex, plus $O(\log(n))$ bits. As in [Castelli Aleardi et al., 2010], this encoding is asymptotically optimal and uses approximately $3.2451n$ bits. The advantage of our method is that it can be implemented in linear time. Moreover we believe that our encoding gives a better understanding of the structure of triangulations of the torus. It is illustrated with new bijections that are obtained in Sections 5.7 and 5.9.

Consider a toroidal triangulation G, a root corner c_0 that is not in the strict interior of a separating triangle and the orientation of the edges of G corresponding to the minimal HTC Schnyder wood with respect to the root face f_0. Let U be the output of Algorithm PS applied on (G, c_0). As already mentioned at the end of Section 5.4, to retrieve the triangulation G one just needs to know U without the information of its root corner (by walking twice around the face of U in counter-

clockwise order and reattached all the admissible triples that are encountered along this walk, one can recover G). Hence to encode G, one just has to encode U without the position of the root corner around the root vertex (see Figure 94.(a)).

By Lemma 5.2.2, the unicellular map U contains a spanning tree T which is oriented from the leaves to the root vertex. The tree T contains exactly $n-1$ edges, so there is exactly 2 edges of U that are not in T. We call these edges the **special edges** of U. We cut these two special edges to transform them into stems of T (see Figures 94.(a) and (b)). We keep the information of where are the special stems in T and on which corner of T they should be reattached. This information can be stored with $O(\log(n))$ bits. One can recover U from T by reattaching the special stems in order to form non-null-homologous cycles with T (see Figure 94.(c)).

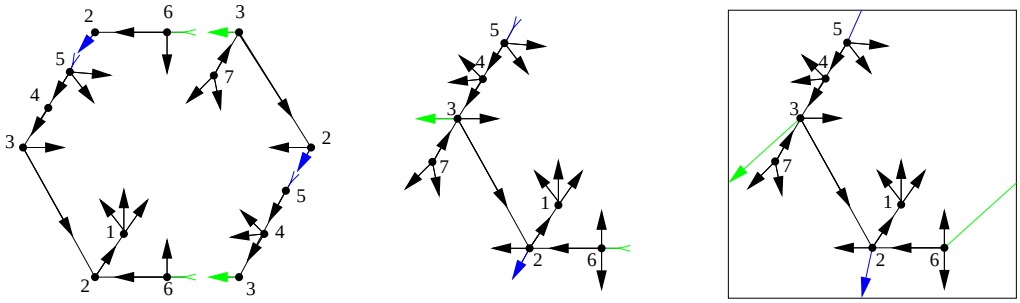

(a) (b) (c)

Figure 94: From unicellular maps to trees with special stems and back.

So T is a plane tree on n vertices, each vertex having 2 stems except the root vertex v_0 having three stems. Choose any stem s_0 of the root vertex, remove it and consider that T is rooted at the angle where s_0 should be attached. The information of the root enables to put back s_0 at its place. So now we are left with a rooted plane tree T on n vertices where each vertex has exactly 2 stems (see Figure 95.(a)).

This tree T can easily be encoded by a binary word on $6n-2$ bits: that is, walking in counterclockwise order around T from the root corner, writing a "1" when going down along T, and a "0" when going up along T (see Figure 95.(a)). As in [Poulalhon and Schaeffer, 2006], one can encode T more compactly by using the fact that each vertex has exactly two stems. Thus T is encoded by a binary word on $4n-2$ bits: that is, walking in counterclockwise order around T from the root corner, writing a "1" when going down along an edge of T, and a "0" when going up along an edge or along a stem of T (see Figure 95.(b) where the "red 1's" of Figure 95.(a) have been removed). Indeed there is no need to encode when going down along stems, this information can be retrieved afterward. While reading the binary word to recover

T, when a "0" is met, we should go up in the tree, except if the vertex that we are considering does not have already its two stems, then in that case we should create a stem (i.e. add a "red 1" before the "0"). So we are left with a binary word on $4n-2$ bits with exactly $n-1$ bits "1" and $3n-1$ bits "0".

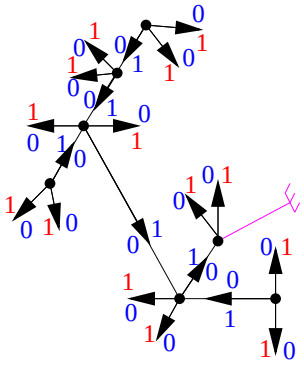

10101110111010010100101101000101011010000

(a)

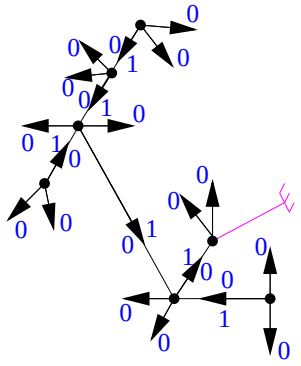

00110110000000100000010000

(b)

Figure 95: Encoding a rooted tree with two stems at each vertex.

Similarly to [Poulalhon and Schaeffer, 2006], using [Bonichon et al., 2003, Lemma 7], this word can then be encoded with a binary word of length $\log_2 \binom{4n-2}{n-1} + o(n) \sim n \log_2(\frac{256}{27}) \approx 3.2451\, n$ bits. Thus we have the following thm whose linearity is discussed in Section 5.6:

Theorem 5.5.1. *Any toroidal triangulation on n vertices, can be encoded with a binary word of length $3.2451n + o(n)$ bits, the encoding and decoding being linear in n.*

5.6 Linear Complexity

In this section we show that the encoding method described in this chapter, that is encoding a toroidal triangulation via an unicellular map and recovering the original triangulation, can be performed in linear time. The only difficulty lies in providing Algorithm PS with the appropriate input it needs in order to apply thm 5.3.1. Then clearly the execution of Algorithm PS, the encoding phase and the recovering of the triangulation are linear. Thus we have to show how one can find in linear time a root corner c_0 that is not in the strict interior of a separating triangle, as well as the minimal HTC Schnyder wood with respect to the root face f_0.

Consider a toroidal triangulation G. Let us see how one can build a Schnyder wood of G in linear time. The **contraction** of a non-loop-edge e of G is the operation consisting of continuously contracting e until merging its two ends, as shown on Figure 96. Note that only one edge of each pair of edges forming a null-homologous 2-cycle is preserved (edges e_{wx} and e_{wy} on the figure).

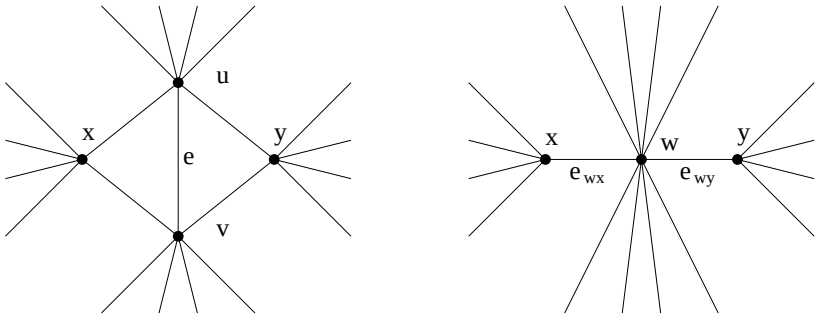

Figure 96: The contraction operation.

An edge uv is said to be **contractible** if it is not a loop and if it is not on a separating triangle (i.e. if after contracting uv one obtains a triangulation that is still without contractible 1- or 2-cycles). In [Gonçalves and Lévêque, 2013] the existence of crossing Schnyder wood is proved by contraction. Unfortunately this proof cannot easily be transformed into a linear algorithm because of the crossing property that has to be maintained during the contraction process. Nevertheless we use contractions to obtain non-necessarily crossing Schnyder woods. If the triangulation obtained after contracting a contractible edge admits a Schnyder wood it is then easy to obtain a Schnyder wood of G. The rules for decontracting an edge in the case of toroidal triangulations are depicted on [Gonçalves and Lévêque, 2013, Figure 21] where for each case one can choose any of the proposed colorings. For any toroidal triangulation, one can find contractible edges until the toroidal map has only one vertex (see [Mohar, 1996]). A Schnyder wood of the toroidal map on one vertex is depicted on the right of Figure 49. Thus one can obtain a Schnyder wood of any toroidal triangulation by this process. Nevertheless, to maintain linearity we have to be more precise since it is not trivial to find contractible edges.

Consider an edge uv of G with incident faces uvx and vuy such that these vertices appear in clockwise order around the corresponding face (so we are in the situation of Figure 96). If u and v have more common neighbors, then consider their second common neighbor going clockwise around u from uv (the first one being x, and the last being y) and call it x'. Call y' their second common neighbor going counter-clockwise around u from uv. Then either uvx' or uvy' is a separating triangle or edge

uv is contractible. We consider these two cases below:

- If uv is contractible, then we contract it and apply the procedure recursively to obtain a Schnyder wood of the contracted graph. Then we update the Schnyder wood as described above. Note that this update is done in constant time.
- If uvx' (resp. uvy') is a separating triangle, one can remove its interior, recursively obtain a toroidal Schnyder wood of the remaining toroidal triangulation, build a planar Schnyder wood of the planar triangulation inside uvx' (resp. uvy'), and then superimpose the two (by eventually permuting the colors) to obtain a Schnyder wood of the whole graph. Note that computing a planar Schnyder wood can be done in linear time using a canonical ordering (see [Kant, 1996]).

The difficulty here is to test if uvx' or uvy' are contractible triangles. For that purpose, one first need to compute a basis (B_1, B_2) for the homology. Consider a spanning tree of the dual map G^*. The map obtained from G by removing those edges is unicellular, and removing its treelike parts one obtains two cycles (B_1, B_2) (intersecting on a path with at least one vertex) that form a basis for the homology. This can be computed in linear time for G and then updated in constant time when some edge is contracted or when the interior of some separating triangle is removed. Then a closed walk W, given with an arbitrary orientation, is contractible if and only if W crosses B_i from right to left as many times as W crosses B_i from left to right, for $i \in \{1, 2\}$. This test is linear in $|W|$ hence constant time for the triangles uvx' and uvy'. Vertex u is fixed during the whole process so the total running time to compute a Schnyder wood of G is linear.

From this Schnyder wood, one can compute in linear time a root corner c_0 not in the strict interior of a separating triangle. First note that in a 3-orientation of a toroidal triangulation, the edges that are inside a separating triangle and that are incident to the three vertices on the border are all oriented toward these three vertices by Euler's formula. Thus an oriented non-contractible cycle cannot enter in the interior of a separating triangle. Now follow any oriented monochromatic path of the Schnyder wood and stop the first time this path is back to a previously met vertex v_0. The end of this path forms an oriented monochromatic cycle C containing v_0. If C is contractible then Euler's formula is violated in the contractible region. Thus C is an oriented non-contractible cycle and cannot contain some vertices that are in the interior of a separating triangle. So v_0 is not in the interior of a separating triangle and we can choose as root corner c_0 any corner incident to v_0.

In [Gonçalves et al., 2015] it is proved how one can transform any 3-orientation (hence a Schnyder wood) of a toroidal triangulation into a half-crossing (hence HTC)

Schnyder wood. The method consists in computing a so called "middle-path" (a directed path where the next edge chosen is the one leaving in the "middle") and reversing some non-contractible "middle-cycles". Clearly the method is linear even if not explicitly mentioned in [Gonçalves et al., 2015]. Let D_0 be the corresponding obtained orientation of G.

It remains to compute the minimal HTC Schnyder wood with respect to the root face f_0. There is a generic known method [Meunier, 2015] (see also [Ueckerdt, 2011, p.23]) to compute in linear time a minimal α-orientation of a planar map as soon as an α-orientation is given. This method also works on oriented surfaces and can be applied to obtain the minimal HTC Schnyder wood in linear time. We explain the method briefly below.

It is much simpler to compute the minimal orientation $D_{\min}$ homologous to D_0 in a dual setting. The first observation to make is that two orientations D_1, D_2 of G are homologous if and only if there dual orientations D_1^*, D_2^* of G^* are equivalent up to reversing some directed cuts. Furthermore $D_1 \leq_{f_0} D_2$ if and only if D_1^* can be obtained from D_2^* by reversing directed cuts oriented from the part containing f_0. Let us compute $D_{\min}^*$ which is the only orientation of G^*, obtained from D_0^* by reversing directed cuts, and without any directed cut oriented from the part containing f_0. For this, consider the orientation D_0^* of $G^* = (F, E^*)$ and compute the set $X \subseteq F$ of vertices of G^* that have an oriented path toward f_0. Then $(X, F \setminus X)$ is a directed cut oriented from the part containing f_0 that one can reverse. Then update the set of vertices that can reach f_0 and go on until $X = F$. It is not difficult to see that this can be done in linear time. Thus we obtain the minimal HTC Schnyder wood with respect to f_0 in linear time.

5.7 Bijection with Rooted Unicellular Maps

Given a toroidal triangulation G with a root corner c_0, we have defined a unique associated orientation: the minimal HTC Schnyder wood with respect to the root face f_0. Suppose that G is oriented according to the minimal HTC Schnyder wood. If c_0 is not in the strict interior of a separating triangle then thms 5.3.1 and 5.4.1 show that the execution of Algorithm PS on (G, c_0) gives a toroidal unicellular map with stems from which one can recover the original triangulation. Thus there is a bijection between toroidal triangulations rooted from an appropriate angle and their image by Algorithm PS. The goal of this section is to describe this image.

Recall from Section 5.4 that the output of Algorithm PS on (G, c_0) is an element of $\mathcal{U}_{r,b}(n)$. One may hope that there is a bijection between toroidal triangulations rooted from an appropriate corner and $\mathcal{U}_{r,b}(n)$ since this is how it works in the

planar case. Indeed, given a planar triangulation G, there is a unique orientation of G (the minimal Schnyder wood) on which ALGORITHM PS, performed from an outer corner, outputs a spanning tree. In the toroidal case, things are more complicated since the behavior of ALGORITHM PS on minimal HTC Schnyder woods does not characterize such orientations.

Figure 97 gives an example of two (non-homologous) orientations of the same triangulation that are both minimal with respect to the same root face. For these two orientations, the execution of ALGORITHM PS from the same root corner gives two different elements of $\mathcal{U}_{r,b}(2)$ (from which the original triangulation can be recovered by the method of thm 5.4.1). Thus we have to exhibit a particular property of HTC Schnyder woods that can be used to characterize which particular subset of $\mathcal{U}_{r,b}(n)$ is in bijection with appropriately rooted toroidal triangulations.

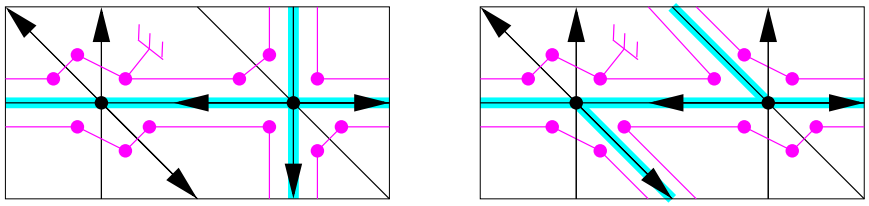

HTC orientation non-HTC orientation

Figure 97: A graph that can be represented by two different unicellular maps.

Let us use the γ_0 property on $\mathcal{U}_r(n)$. Consider an element U of $\mathcal{U}_r(n)$ whose edges and stems are oriented with respect to the root corner as follows: the stems are all outgoing, and while walking clockwise around the unique face of U from c_0, the first time an edge is met, it is oriented counterclockwise with respect to the face of U. Then one can compute γ on the cycles of U (edges and stems count). We say that an unicellular map of $\mathcal{U}_r(n)$ satisfies the γ_0 property if γ equals zero on its (non-contractible) cycles. Let us call $\mathcal{U}_{r,b,\gamma_0}(n)$ the set of elements of $\mathcal{U}_{r,b}(n)$ satisfying the γ_0 property. So the output of ALGORITHM PS given by thm 5.3.1 is an element of $\mathcal{U}_{r,b,\gamma_0}(n)$.

Let $\mathcal{T}_r(n)$ be the set of toroidal triangulations on n vertices rooted at a corner that is not in the clockwise interior of a separating triangle. Then we have the following bijection:

Theorem 5.7.1. *There is a bijection between $\mathcal{T}_r(n)$ and $\mathcal{U}_{r,b,\gamma_0}(n)$.*

Proof. Consider the mapping g that associates to an element of $\mathcal{T}_r(n)$, the output of ALGORITHM PS executed on the minimal HTC Schnyder wood with respect to the

root face. By the above discussion the image of g is in $\mathcal{U}_{r,b,\gamma_0}(n)$ and g is injective since one can recover the original triangulation from its image by thm 5.4.1.

Conversely, given an element U of $\mathcal{U}_{r,b,\gamma_0}(n)$ with root corner c_0, one can build a toroidal map G by the complete closure procedure described in Section 5.4. The number of stems and edges of U implies that G is a triangulation. Recall that c_0 defines an orientation on the edges and stems of U. Consider the orientation D of G induced by this orientation. Since U is balanced, the execution of Algorithm PS on (G, c_0) corresponds to the cycle in the corner graph of U obtained by starting from the root corner and walking clockwise in the face of U. Thus the output of Algorithm PS executed on (G, c_0) is U. It remains to show that G is appropriately rooted and that D corresponds to the minimal HTC Schnyder wood with respect to this root, then g will be surjective.

First note that by definition of $\mathcal{U}_r(n)$, the orientation D is a 3-orientation.

Suppose by contradiction that c_0 is in the strict interior of a separating triangle. Then, since we are considering a 3-orientation, by Euler's formula, the edges in the interior of this triangle and incident to its border are all entering the border. So Algorithm PS started from the strict interior cannot visit the vertices on the border of the triangle and outside. Thus the output of Algorithm PS is not a toroidal unicellular map, a contradiction. So c_0 is not in the strict interior of a separating triangle.

The γ_0 property of U implies that γ equals zero on two cycles of U. Hence these two cycles considered in G also satisfy γ equals 0 and form a basis for the homology. So D is a hTC Schnyder wood.

Suppose by contradiction that D is not minimal. Then, by Lemma 2.4.22, it contains a clockwise (non-empty) null-homologous oriented subgraph with respect to f_0. With the notations of Section 2.4.4, let T be such a subgraph with $T = -\sum_{F\in\mathcal{F}'} \lambda_F \partial(F)$, with $\lambda \in \mathbb{N}^{|\mathcal{F}'|}$. Let $\lambda_{F_0} = 0$, and $\lambda_{\max} = \max_{F\in\mathcal{F}} \lambda_F$. For $0 \leq i \leq \lambda_{\max}$, let $X_i = \{F \in \mathcal{F} \mid \lambda_F \geq i\}$. For $1 \leq i \leq \lambda_{\max}$, let T_i be the oriented subgraph such that $T_i = -\sum_{F\in X_i} \partial(F)$. Then we have $T = \sum_{1\leq i\leq\lambda_{\max}} T_i$. Since T is an oriented subgraph, we have $T \in \{-1, 0, 1\}^{|E(G)|}$. Thus for any edge of G, incident to faces F_1 and F_2, we have $(\lambda_{F_1} - \lambda_{F_2}) \in \{-1, 0, 1\}$. So, for $1 \leq i \leq \lambda_{\max}$, the oriented graph T_i is the frontier between the faces with λ value equal to i and $i-1$. So all the T_i are edge disjoint and are oriented subgraphs of D. Since T is non-empty, we have $\lambda_{\max} \geq 1$, and T_1 is non-empty. All the edges of T_1 have a face of X_1 on their right and a face of X_0 on their left. Since U is an unicellular map, and T_1 is a (non-empty) null-homologous oriented subgraph, at least one edge of T_1 corresponds to a stem of U. Let s be the last stem of U corresponding to an edge of T_1 that is reattached by the complete closure procedure. Consider the step where s is reattached. As the

root corner (and thus f_0) is in the special face (see the terminology of Section 5.4), the special face is in the region defined by X_0. Thus it is on the left of s when it is reattached. This contradicts the fact that U is balanced. Thus D is the minimal HTC Schnyder wood with respect to f_0. □

5.8 The Lattice of HTC Schnyder Woods

In this section, we push further the study of HTC Schnyder woods in order to remove the root and the balanced property of the unicellular maps considered in thm 5.7.1 and obtain a simplified bijection in thm 5.9.1 of Section 5.9. The aim is to prove that returning triangles is always sufficient to travel into the lattice. More detailed proofs can be found in [Lévêque, 2016].

Consider a toroidal triangulation G given with a crossing Schnyder wood. Let D_0 be the corresponding 3-orientation of G. Let f_0 be any face of G. Recall from Section 2.4.4 that $O(G)$ denotes the set of all the orientations of G that are homologous to D_0. The elements of $O(G)$ are the HTC Schnyder woods of G and $(O(G), \leq_{f_0})$ is a distributive lattice.

We need to reduce the graph G. We call an edge of G **rigid** with respect to $O(G)$ if it has the same orientation in all the elements of $O(G)$. Rigid edges do not play a role for the structure of $O(G)$. We delete them from G and call the obtained embedded graph $\widetilde{G}$. Note that this graph is embedded but it is not necessarily a map, as some faces may not be homeomorphic to open disks. Note also that $\widetilde{G}$ might be empty if all the edges are rigid, i.e. $|O(G)| = 1$ and $\widetilde{G}$ has no edge but a unique face that is all the surface.

Lemma 5.8.1. *Given an edge e of G, the following are equivalent:*

1. *e is non-rigid*
2. *e is contained in a null-homologous oriented subgraph of D_0*
3. *e is contained in a null-homologous oriented subgraph of any element of $O(G)$*

Proof. $(1 \implies 3)$ Let $D \in O(G)$. If e is non-rigid, then it has a different orientation in two elements D', D'' of $O(G)$. Then we can assume by symmetry that e has a different orientation in D and D' (otherwise in D and D'' by symmetry). Since D, D' are homologous to D_0, they are also homologous to each other. So $T = D \setminus D'$ is a null-homologous oriented subgraph of D that contains e.

$(3 \implies 2)$ Trivial since $D_0 \in O(G)$

($2 \implies 1$) If an edge e is contained in a null-homologous oriented subgraph T of D_0. Then let D be the element of $O(G)$ such that $T = D_0 \setminus D$. Clearly e is oriented differently in D and D_0, thus it is non-rigid. □

By Lemma 5.8.1, one can build $\widetilde{G}$ by keeping only the edges that are contained in a null-homologous oriented subgraph of D_0. Note that this implies that all the edges of $\widetilde{G}$ are incident to two distinct faces of $\widetilde{G}$. Denote by $\widetilde{\mathcal{F}}$ the set of oriented subgraphs of $\widetilde{G}$ corresponding to the boundaries of faces of $\widetilde{G}$ considered counterclockwise. Let $\widetilde{f}_0$ be the face of $\widetilde{G}$ containing f_0 and $\widetilde{F}_0$ be the element of $\widetilde{\mathcal{F}}$ corresponding to the boundary of $\widetilde{f}_0$. Let $\widetilde{\mathcal{F}}' = \widetilde{\mathcal{F}} \setminus \widetilde{F}_0$. The elements of $\widetilde{\mathcal{F}}'$ are sufficient to generate the entire lattice $(O(G), \leq_{f_0})$, i.e. two elements D, D' of $O(G)$ are linked in the Hasse diagram of the lattice, with $D \leq_{f_0} D'$, if and only if $D \setminus D' \in \widetilde{\mathcal{F}}'$.

Lemma 5.8.2. *For every element $\widetilde{F} \in \widetilde{\mathcal{F}}$ there exists D in $O(G)$ such that $\widetilde{F}$ is an oriented subgraph of D.*

Proof. Let $\widetilde{F} \in \widetilde{\mathcal{F}}$. Let D be an element of $O(G)$ that maximize the number of edges of $\widetilde{F}$ that have the same orientation in $\widetilde{F}$ and D (i.e. that maximize the number of edges of D oriented counterclockwise on the border of the face of $\widetilde{G}$ corresponding to $\widetilde{F}$). Suppose by contradiction that there is an edge e of $\widetilde{F}$ that does not have the same orientation in $\widetilde{F}$ and D. Edge e is in $\widetilde{G}$ so it is non-rigid. Let $D' \in O(G)$ such that e is oriented differently in D and D'. Let $T = D \setminus D'$. There exists edge-disjoint oriented subgraphs $T_1, \ldots, T_k$ of D such that $T = \sum_{1 \leq i \leq k} T_i$, and, for $1 \leq i \leq k$, there exists $\widetilde{X}_i \subseteq \widetilde{\mathcal{F}}'$ and $\epsilon_i \in \{-1, 1\}$ such that $T_i = \epsilon_i \sum_{\widetilde{F}' \in \widetilde{X}_i} \partial(\widetilde{F}')$. Without loss of generality, we can assume that e is an edge of T_1. Let D'' be the element of $O(G)$ such that $T_1 = D \setminus D''$. The oriented subgraph T_1 intersects $\widetilde{F}$ only on edges of D oriented clockwise on the border of $\widetilde{F}$. So D'' contains strictly more edges oriented counterclockwise on the border of the face $\widetilde{F}$ than D, a contradiction. So all the edges of $\widetilde{F}$ have the same orientation in D. So $\widetilde{F}$ is a null-homologous oriented subgraph of D. □

By Lemma 5.8.2, for every element $\widetilde{F} \in \widetilde{\mathcal{F}}'$ there exists D in $O(G)$ such that $\widetilde{F}$ is an oriented subgraph of D. Thus there exists D' such that $\widetilde{F} = D \setminus D'$ and D, D' are linked in the Hasse diagram of the lattice. Thus the elements of $\widetilde{\mathcal{F}}'$ form a minimal set that generates the lattice.

Let $D_{\max}$ (resp. $D_{\min}$) be the maximal (resp. minimal) element of $(O(G), \leq_{f_0})$.

Lemma 5.8.3. *$\widetilde{F}_0$ (resp. $-\widetilde{F}_0$) is an oriented subgraph of $D_{\max}$ (resp. $D_{\min}$).*

Proof. By Lemma 5.8.2, there exists D in $O(G)$ such that $\widetilde{F}$ is an oriented subgraph of D. Let $T = D \setminus D_{\max}$. Since $D \leq_{f_0} D_{\max}$, $T = \sum_{\widetilde{F} \in \widetilde{\mathcal{F}}'} \lambda_F \partial(\widetilde{F})$ with $\lambda \in \mathbb{N}^{|\mathcal{F}'|}$. So T is disjoint from $\widetilde{F}_0$. Thus $\widetilde{F}_0$ is an oriented subgraph of $D_{\max}$. The proof is similar for $D_{\min}$. □

Note that the above three lemmas hold in a more general context than just $O(G)$. Actually they hold for any lattice of homologous orientations on an oriented surface (see [Lévêque, 2016]). From now on we use some specific properties of the object considered in this chapter, i.e. HTC Schnyder woods.

Lemma 5.8.4. *Consider an orientation D in $O(G)$ and a closed walk W of $\widetilde{G}$. If on the left (resp. right) side of W, there is no incident edges of $\widetilde{G}$, and no outgoing incident edges of D, then W is a contractible triangle with its contractible region on its left (resp. right) side.*

Proof. Consider a closed walk W of $\widetilde{G}$ such that on its left side there is no incident edge of $\widetilde{G}$, and no outgoing incident edges of D. Let k be the length of W. Let W_{left} be the edges of D that are incident to the left side of W. By assumption they are all entering W. Note that W cannot cross itself otherwise it has at least one incident edge of $\widetilde{G}$ on its left side. However it may have repeated vertices but in that case it intersects itself tangencially on the right side.

Suppose first that W is non-contractible. Then consider the closed walk W^* of the dual orientation D^* that is obtained by considering all the dual edges of W_{left} with their corresponding orientation. Since all the edges of W_{left} are entering W we have that W^* is an oriented closed walk. Moreover it is non-contractible and thus contains an oriented non-contractible cycle, a contradiction to Lemma 2.4.23. So W is contractible. Since W can intersect itself only tangencially on the right side, the region delimited by W and located on its left side is connected.

Suppose that W has its contractible region on its left side. Consider the graph G' obtained from G by keeping only the vertices and edges that lie in the contractible region delimited by W, including W. The vertices of W appearing several times are duplicated so that G' is a plane triangulation of a k-cycle. Let n', m', f' be the number of vertices, edges and faces of G'. By Euler's formula, $n' - m' + f' = 2$. All the inner faces have size 3 and the outer face has size k, so $2m' = 3(f' - 1) + k$. All the inner vertices have outdegree 3 as we are considering a 3-orientation of G. All the edges of W_{left} are oriented toward W, and there are k outer edges, so $m' = 3(n' - k) + k$. Combining these three equalities gives $k = 3$, i.e. W is a triangle and the lemma holds.

Suppose now that W has its contractible region on its right side. Then similarly as above, consider the graph G' obtained from G by keeping all the vertices and edges that lie in the contractible region delimited by W, including W. This time the vertices of W appearing several times are not duplicated. Since W can intersect itself only tangencially on the right side, we have that G' is a plane map whose outer face boundary is W and whose interior is triangulated. As above, let n', m', f' be the number of vertices, edges and faces of G'. By Euler's formula, $n' - m' + f' = 2$. All the inner faces have size 3 and the outer face has size k, so $2m' = 3(f' - 1) + k$. Since there is no outgoing incident edges of D on the left side of W, all the vertices of G' have outdegree 3 and $m' = 3n'$. Combining these three equalities gives $k = -3$, a contradiction. □

The boundary of a face of $\widetilde{G}$ may be composed of several closed walks. Let us call **quasi-contractible** the faces of $\widetilde{G}$ that are homeomorphic to a disk or to a disk with punctures. Note that such a face may have several boundaries (if there is some punctures) and then the face is not contractible, but exactly one of these boundaries contains all the other in its contractible region. Let us call **outer facial walk** this special boundary. Then we have the following:

Lemma 5.8.5. *All the faces of $\widetilde{G}$ are quasi-contractible and their outer facial walk is a (contractible) triangle.*

Proof. Suppose by contradiction that there is a face $\widetilde{f}$ of $\widetilde{G}$ that is not quasi-contractible or whose outer facial walk is not a contractible triangle. Let $\widetilde{F}$ be the element of $\widetilde{\mathcal{F}}$ corresponding to the boundary of $\widetilde{f}$. By Lemma 5.8.2, there exists an orientation D in $O(G)$ such that $\widetilde{F}$ is an oriented subgraph of D.

All the faces of G, are contractible triangles. Thus $\widetilde{f}$ is not a face of G and contains in its interior at least one edge of G. Start from any such edge e and consider the **left-walk** $W = (e_i)_{i \geq 0}$ of D obtained by the following: if the edge e_i is entering a vertex v, then e_{i+1} is choosen among the three edges leaving v as the edge that is on the left coming from e_i (i.e. the first one while going clockwise around v). Suppose that for $i \geq 0$, edge e_i is entering a vertex v that is on the border of $\widetilde{f}$. Recall that by definition $\widetilde{F}$ is oriented counterclockwise according to its interior, so either e_{i+1} is in the interior of $\widetilde{f}$ or e_{i+1} is on the border of $\widetilde{f}$. Thus W cannot leave $\widetilde{f}$.

Since G has a finite number of edges, some edges are used several times in W. Consider a minimal subsequence $W' = e_k, \ldots, e_\ell$ such that no edge appears twice and $e_k = e_{\ell+1}$. Thus W ends periodically on the sequence of edges $e_k, \ldots, e_\ell$. By Lemma 5.8.4, all the closed walks that are part of $\widetilde{F}$ have some outgoing incident

edges of D on their left side. Thus we have that W' contains at least one edge that is not an edge of $\widetilde{F}$, thus it contains at least one rigid edge.

By construction, all the edges on the left side of W' are entering. Suppose that W' is not contractible. Then the oriented closed walk of the dual orientation D^* that is obtained by considering all the dual edges of its incident edges on the left side gives a contradiction to Lemma 2.4.23. So W' is contractible. So it is a null-homologous oriented subgraph of D, thus all its edges are non-rigid by Lemma 5.8.1, a contradiction. □

By Lemma 5.8.5, every face of $\widetilde{G}$ is quasi-contractible and its outer facial walk is a contractible triangle. So $\widetilde{G}$ contains all the contractible triangles of G whose interiors are maximal by inclusion, i.e. it contains all the edges that are not in the interior of a separating triangle. In particular, $\widetilde{G}$ is non-empty and $|O(G)| \geq 2$. The status (rigid or not) of an edge lying inside a separating triangle is determined as in the planar case: such an edge is rigid if and only if it is in the interior of a separating triangle and incident to this triangle. Thus an edge of G is rigid if and only if it is in the interior of a separating triangle and incident to this triangle.

Since $(O(G), \leq_{f_0})$ is a distributive lattice, any element D of $O(G)$ that is distinct from $D_{\max}$ and $D_{\min}$ contains at least one neighbor above and at least one neighbor below in the Hasse diagram of the lattice. Thus it has at least one face of $\widetilde{G}$ oriented counterclockwise and at least one face of $\widetilde{G}$ oriented clockwise. Thus by Lemma 5.8.5, it contains at least one contractible triangle oriented counterclockwise and at least one contractible triangle oriented clockwise. Next lemma shows that this property is also true for $D_{\max}$ and $D_{\min}$.

Lemma 5.8.6. *In $D_{\max}$ (resp. $D_{\min}$) there is a counterclockwise (resp. clockwise) contractible triangle containing f_0, and a clockwise (resp. counterclockwise) contractible triangle not containing f_0.*

Proof. By Lemma 5.8.5, $\widetilde{f_0}$ is quasi-contractible and its outer facial walk is a contractible triangle T. By lemma 5.8.3, $\widetilde{F_0}$ is an oriented subgraph of $D_{\max}$. Thus T is oriented counterclockwise and contains f_0. The second part of the lemma is clear since $|O(G)| \geq 2$ so $D_{\max}$ has at least one neighbor below in the Hasse diagram of the lattice. Similarly for $D_{\min}$. □

Thus by above remarks and Lemma 5.8.6, all the HTC Schnyder woods have at least one triangle oriented counterclockwise and at least one triangle oriented clockwise. Note that this property does not characterize HTC Schnyder woods. Figure 53 gives an example of a Schnyder wood that is not HTC but satisfies the property. Note also that not all Schnyder woods satisfy the property. The right of

Figure 51 is an example of a Schnyder wood that is no HTC and has no oriented contractible triangle.

Lemma 5.8.6 is used in the next section to obtained a bijection with unrooted unicellular maps.

5.9 Bijection with Unrooted Unicellular Maps

To remove the root and the balanced property of the unicellular maps considered in Theorem 5.7.1, we have to root the toroidal triangulation more precisely than before. We say that a corner is not **in the clockwise interior of a separating triangle** if it is not in its contractible region, or if it is incident to a vertex v of the triangle and situated just before an edge of the triangle in counterclockwise order around v (see Figure 98).

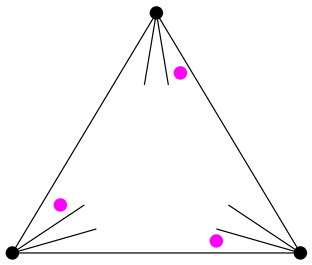

Figure 98: Angles that are in a separating triangle but not in its clockwise interior.

Consider a toroidal triangulation G. Consider a root corner c_0 that is not in the clockwise interior of a separating triangle. Note that the choice of c_0 is equivalent to the choice of a root vertex v_0 and a root edge e_0 incident to v_0 such that none is in the interior of a separating triangle. Consider the orientation of the edges of G corresponding to the minimal HTC Schnyder wood with respect to the root face f_0. By Lemma 5.8.6, there is a clockwise triangle containing f_0. Thus by the choice of c_0, the edge e_0 is leaving the root vertex v_0. This is the essential property used in this section. Consider the output U of Algorithm PS on (G, c_0). Since e_0 is leaving v_0 and c_0 is just before e_0 in counterclockwise order around v_0, the execution of Algorithm PS starts by Case 2 and e_0 corresponds in U to a stem s_0 attached to v_0. We call this stem s_0 the **root stem**.

The recovering method defined in thm 5.4.1 says that s_0 is the last stem reattached by the procedure. So there exists a sequence of admissible triples of U (see the terminology and notations of Section 5.4) such that s_0 belongs to the last admissible triple. Let $U_0 = U$ and for $1 \leq k \leq 2n-2$, the map U_k is obtained from U_{k-1} by

closing any admissible triple that does no contain s_0. As noted in Section 5.4, the special face of U_{2n-2} is a quadrangle with exactly one stem. This stem being s_0, we are in the situation of Figure 99.

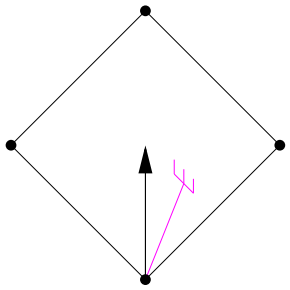

Figure 99: The situation just before the last stem (i.e. the root stem) is reattached

Consequently, if one removes the root stem s_0 from U to obtain an unicellular map U' with n vertices, $n+1$ edges and $2n-2$ stems, one can recover the graph U_{2n-2} by applying a complete closure procedure on U' (see example of Figure 100). Note that then, there are four different ways to finish the closure of U_{2n-2} to obtain an oriented toroidal triangulation. This four cases correspond to the four ways to place the (removed) root stem in a quadrangle, they are obtained by pivoting Figure 99 by 0, $\frac{\pi}{2}$, π and $\frac{3\pi}{2}$. Note that only one of this four cases leads to the original rooted triangulation G, except if there are some symmetries (like in the example of Figure 100).

Let $\mathcal{U}(n)$ denote the set of (non-rooted) toroidal unicellular maps, with exactly n vertices, $n+1$ edges and $2n-2$ stems satisfying the following: a vertex has exactly 2 stems if it is not a corner, 1 stem if it is the corner of a hexagon and 0 stem if it is the corner of a square. Note that the output of thm 5.3.1 on an appropriately rooted toroidal triangulation is an element of $\mathcal{U}(n)$ when the root stem is removed.

Note that an element U' of $\mathcal{U}(n)$ is non-rooted so we cannot orient automatically its edges with respect to the root corner like in Section 5.7. Nevertheless one can still orient all the stems as outgoing and compute γ on the cycles of U' by considering only its stems in the counting (and not the edges nor the root stem anymore). We say that an unicellular map of $\mathcal{U}(n)$ satisfies the γ_0 property if γ equals zero on its (non-contractible) cycles. Let us call $\mathcal{U}_{\gamma_0}(n)$ the set of elements of $\mathcal{U}(n)$ satisfying the γ_0 property.

A surprising property is that an element U' of $\mathcal{U}(n)$ satisfies the γ_0 property if and only if any element U of $\mathcal{U}_r(n)$ obtained from U' by adding a root stem anywhere in U' satisfies the γ_0 property (note that in U we count the edges and the root stem to compute γ). One can see this by considering the unicellular map of Figure 101.

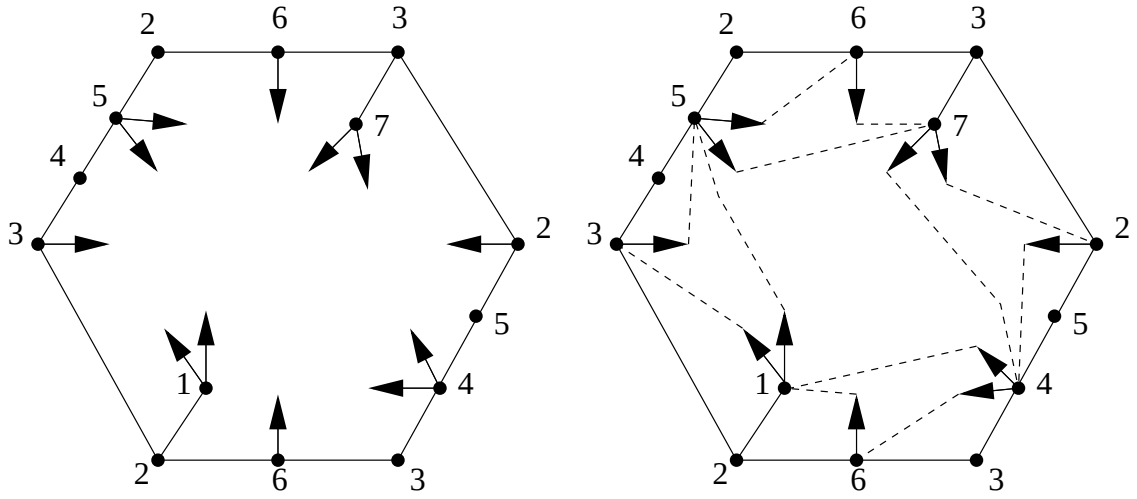

Figure 100: Example of K_7 where the root corner, the root stem and the orientation with respect to the root corner have been removed from the output of Figure 88. The complete closure procedure leads to a quadrangular face.

It represents the general case of the underlying rooted hexagon of U. The edges represent in fact paths (some of which can be of length zero). One can check that it satisfies γ equals zero on its (non-contractible) cycles. It corresponds exactly to the set of edges that are taken into consideration when computing γ on U but not when computing γ on U'. Thus it does not affect the counting (the tree-like parts are not represented since they do not affect the value γ). So the output of thm 5.3.1 on an appropriately rooted toroidal triangulation is an element of $\mathcal{U}_{\gamma_0}(n)$ when the root stem is removed.

For the particular case of K_7, the difference between the rooted output of Figure 88 and the non-rooted output of Figure 100 is represented on Figure 102 (one can superimpose the last two to obtain the first). One can check that these three unicellular maps (rooted, non-rooted and the difference) all satisfy γ equals zero on their cycles.

There is an "almost" four-to-one correspondence between toroidal triangulations on n vertices, given with a root corner that is not in the clockwise interior of a separating triangle, and elements of $\mathcal{U}_{\gamma_0}(n)$. The "almost" means that if the automorphism group of an element U of $\mathcal{U}_{\gamma_0}(n)$ is not trivial, some of the four ways to add a root stem in U are isomorphic and lead to the same rooted triangulation. In the example of Figure 100, one can root in four ways the quadrangle but this gives

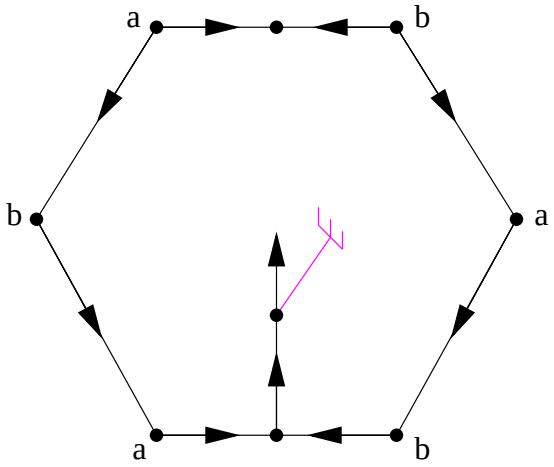

Figure 101: The parts of the unicellular map showing the correspondence while computing γ with or without the orientation with respect to the root plus the root stem.

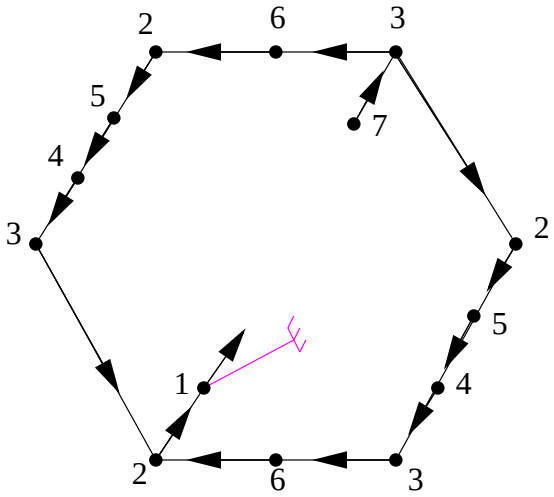

Figure 102: The difference between the rooted output of Figure 88 and the non-rooted output of Figure 100.

only two different rooted triangulations (because of the symmetries of K_7). We face this problem by defining another class for which we can formulate a bijection.

Let $\mathcal{T}(n)$ be the set of toroidal maps on n vertices, where all the faces are triangles, except one that is a quadrangle and which is not in a separating triangle. Then we have the following bijection:

Theorem 5.9.1. *There is a bijection between $\mathcal{T}(n)$ and $\mathcal{U}_{\gamma_0}(n)$.*

Proof. Let a (for "add") be an arbitrarily chosen mapping defined on the maps G' of $\mathcal{T}(n)$ that adds a diagonal e_0 in the quadrangle of G' and roots the obtained toroidal triangulation G at a vertex v_0 incident to e_0 (this defines the root corner c_0 situated just before e_0 in counterclockwise order around v_0). Note that the added edge cannot create a separating 2-cycle, since otherwise the quadrangle would be in a separating triangle. Moreover the root corner of G is not in the clockwise interior of a separating triangle. Thus the image of a is in $\mathcal{T}'_r(n)$, the subset of $\mathcal{T}_r(n)$ corresponding to toroidal triangulations rooted at an corner that is not in the clockwise interior of a separating triangle.

Let $\mathcal{U}'_{r,b,\gamma_0}(n)$ be the elements of $\mathcal{U}_{r,b,\gamma_0}(n)$ that have their root corner just before a stem in counterclockwise order around the root vertex. Consider the mapping g, defined in the proof of thm 5.7. By above remarks and thm 5.7, the image of g restricted to $\mathcal{T}'_r(n)$ is in $\mathcal{U}'_{r,b,\gamma_0}(n)$. Let r (for "remove") be the mapping that associates to an element of $\mathcal{U}'_{r,b,\gamma_0}(n)$ an element of $\mathcal{U}_{\gamma_0}(n)$ obtained by removing the root corner and its corresponding stem. Finally, let $h = r \circ g \circ a$ which associates to an element of $\mathcal{T}(n)$ an element of $\mathcal{U}_{\gamma_0}(n)$. Let us show that h is a bijection.

Consider an element G' of $\mathcal{T}(n)$ and its image U' by h. The complete closure procedure on U' gives G' thus the mapping h is injective.

Conversely, consider an element U' of $\mathcal{U}_{\gamma_0}(n)$. Apply the complete closure procedure on U'. At the end of this procedure, the special face is a quadrangle whose corners are denoted $\alpha^1, \ldots, \alpha^4$. We denote also by $\alpha^1, \ldots, \alpha^4$ the corresponding corners of U'. For $i \in \{1, \ldots, 4\}$, let U^i be the element of $\mathcal{U}_r(n)$ obtained by adding a root stem and a root corner in the corner α^i of U', with the root corner just before the stem in counterclockwise order around the root vertex. Note that by the choice of α^i, the U^i are all balanced. By above remarks they also satisfy the γ_0 property and thus they are in $\mathcal{U}'_{r,b,\gamma_0}(n)$.

By the proof of thm 5.7.1, the complete closure procedure on U^i gives a triangulation G^i of $\mathcal{T}_r(n)$ that is rooted from a corner c_0^i not in the strict interior of a separating triangle and oriented according to the minimal HTC Schnyder wood with respect to the root face. Moreover the output of Algorithm PS applied on (G^i, c_0^i) is U^i. Since in U^i, the root stem is present just after the root corner, the first edge seen by the execution of Algorithm PS on (G^i, c_0^i) is outgoing. So c_0 is not in the clockwise interior of a separating triangle (in a 3-orientation, all the edges that are in the interior of a separating triangle and incident to the triangle are entering the triangle). Thus the G^i are appropriately rooted and are elements of $\mathcal{T}'_r(n)$. Removing the root edge of any G_i, gives the same map G' of $\mathcal{T}(n)$. Exactly one of the G_i is the image of G' by the mapping a. Thus the image of G' by h is U' and the mapping h

is surjective. □

A nice aspect of Theorem 5.9.1 comparing to Theorem 5.7.1 is that the unicellular maps that are considered are much simpler. They have no root nor balanced property anymore. It would be great to use Theorem 5.9.1 to count and sample toroidal triangulations. The main issue comparing to the planar case seems to be the γ_0 property.

5.10 Conclusion

Note that the work presented here is related to a work of Bernardi and Chapuy [Bernardi and Chapuy, 2011] (their convention for the orientation of the edges is the reverse of ours). Consider a map G (not necessarily a triangulation) on an oriented surface of genus g, rooted at a particular corner c_0. An orientation of G is **right** if for each edge e, the **right-walk** starting from e (when entering a vertex, the next chosen edge is the one leaving on the right) reaches the root edge e_0 via the root vertex v_0. A consequence of [Bernardi and Chapuy, 2011] is that Algorithm PS applied on an orientation of (G, c_0) outputs a spanning unicellular submap U if and only if the considered orientation is right. Note that in this characterization, the submap U is not necessarily a map of genus g, its genus can be any value in $\{0, \dots, g\}$. In the particular case of toroidal triangulations we show that by considering minimal HTC Schnyder woods the output U is a toroidal spanning unicellular map. Hence by the above characterization, minimal HTC Schnyder woods are right. But here, the fact that U and G have the same genus is of particular interest as it yields a simple bijection.

The key property that makes U and G have same genus is the conclusion of Lemma 2.4.23 (no oriented non-null-homologous cycle in the dual orientation). Recently, [Albar et al., 2014] proved the following:

Theorem 5.10.1 ([Albar et al., 2014])**.** *A simple triangulation on a genus $g \geq 1$ orientable surface admits an orientation of its edges such that every vertex has outdegree at least* 3*, and divisible by* 3*.*

Theorem 5.10.1 is proved for simple triangulation but we believe it to be true for all triangulations. Moreover we hope for a possible generalization satisfying the conclusion of Lemma 2.4.23:

Conjecture 5.10.2. *A triangulation on a genus $g \geq 1$ orientable surface admits an orientation of its edges such that every vertex has outdegree at least* 3*, divisible by* 3*, and such that there is no oriented non-null-homologous cycle in the dual orientation.*

If Conjecture 5.10.2 is true, one can consider a minimal orientation satisfying its conclusion and apply Algorithm PS to obtain a unicellular map of the same genus as G. Note that more efforts should be made to obtain a bijection since there might be several minimal elements satisfying the conjecture and a particular one has to be identified (as the minimal HTC Schnyder wood in our case).

Conclusion

In this thesis, we looked at three different problems. We step back from our results and give some thoughts about it. We also recall some open problems relative to each subject. A more detailed description is given at the end of each chapter.

Geometric Intersection Numbers

The study of systems of quads led us to the algorithms of Chapter 3. Those algorithms are very simple and can be easily implemented using classical data-structures. Seen from above the idea is naturally described in a continuous setting. Let us assume that we have an input curve on some surface of genus at least two. The usual approach is to first provide the underlying surface with a hyperbolic metric. In a second time, the input curve is tightened with respect to this metric. What we obtain is a curve that minimizes the number of self-intersections. In the discrete setting the hyperbolic metric is replaced by a systems of quads with angles as described in the preliminary Section 2.3.5. The discrete tightening uses a canonical representative as already introduced by Lazarus and Rivaud for the test of homotopy.

At this point, the input curve is represented by a canonical combinatorial geodesic. As opposed to the continuous case, however, the most involved part is to count the number of crossings of our geodesic representative. The point is that the combinatorial geodesic is far from being in general position; a crossing may correspond to a pair of homotopic subpaths with alternate ends (see Section 3.6). The exact counting relies on a precise study of the geodesic paths in the system of quads. In particular we observe that all the homotopic geodesics are contained in a *tube* of width at most one edge. The two sides of the tube are the two canonical representative of the curve according to its orientation. It gives an intuitive notion of thick geodesic that retains the main properties of hyperbolic geodesics.

This leads to our quadratic algorithm for counting the geometric intersection number of a curve or of a pair of curves. Although the geometric intersection num-

ber could obviously be quadratic, it is not clear that this algorithm is optimal. In particular, one could expect an output sensitive algorithm. This is somehow the case when dealing with curves homotopic to simple curves. Indeed, we are able to decide in quasi-linear time whether the geometric intersection number of a curve is 0 or not. Moreover, our algorithm provides a combinatorial embedding (i.e. an immersion without crossing) in the affirmative. Computing such a minimal immersion in the general case seems even more delicate. Our quartic algorithm for computing a minimal immersion of a single curve can probably be improved. Using our approach with bigon swapping, it is not clear how to reduce the complexity or even to extend the computation to a pair of curves. In this respect, we note that no polynomial time algorithm is known. A natural and promising approach is to check whether our algorithm for simple curves would output minimal immersions in the general case. We leaves this as an open problem.

Splitting Cycles

In some sense, our counter-example to Mohar and Thomassen conjecture was not hidden. Trying any one of the triangulations by a big enough complete graph would have worked. The point is that the *big enough* part is a real issue that may require years of computation. Our coloring heuristic made the result. The idea of testing the conjecture against complete triangulations was already suggested by Ellingham and Stephens in their 2005 paper. The fact that no one tried to test the conjecture clearly shows that the critical point is the algorithm complexity. Our heuristic happened to be very efficient for searching splitting cycles in complete triangulations. However, the behavior of our algorithm when applied to general maps is not known. It is probably a good choice in many cases, in particular for dense graphs. We conjecture a polynomial time algorithm parametrized by some function of the minimum degree of the graph such as the ratio between the total number of vertices and the minimum degree.

The interest of the chapter is not only the counter-example but also the structural view of triangulations of complete graphs it provides. An informal conjecture can be given: a triangulation of a complete graph cannot have *many* splitting cycles of type more or less balanced. More precisely, we expect that the number of splitting cycles of type $t \in [\alpha\frac{g}{2}, \frac{g}{2}]$ for some constant α to be determined is very small. Considering our experimental results that there exists an embedding of K_{19} with a balanced splitting cycle. In the affirmative, it would be interesting to give an explicit construction.

It suggests an idea to deal with that problem. It seem plausible that there is only a very small proportion of embeddings of K_n on a genus $\frac{(n-3)(n-4)}{12}$ surface

than are obtained by gluing two embeddings of genus around $\frac{(n-3)(n-4)}{24}$ with one boundary. Some convergence estimations are known for the size of many families of triangulated maps. This alone cannot be enough, it is necessary to take into account that a splitting cycle of a triangulation of a complete graph must cut a part without interior vertices, in particular all the edges of this part are either boundary edges or chords. The other part should also contain many chords so that most of the gluing of such parts will create multiple edges. This would disprove the conjecture by a purely theoretical approach. However, it cannot lead to a proof that a specific triangulation avoids a given type of splitting cycle.

Bijection for Toroidal Triangulations

The key point of toroidal Schnyder woods is not their definition or the proof of their existence. It is the definition of the canonical element as the minimal element of the HTC or γ_0 lattice. In addition, it is possible to compute this minimal element in linear time. Note that all this work has been possible because Daniel Gonçalves and Benjamin Lévêque kept on trying to find new proofs of existence of Schnyder woods and of each of their results. The good definitions and proofs that can be carried to genus higher than 1 are not known by now but the various approaches they described can all be useful.

In our proof of the bijection for the torus we have tried to avoid arguments specific to genus 1. In the planar case, it has been proved that the canonical triangulation have all the properties that make Poulhalon and Schaeffer algorithm run as expected. In the torus case, we prove that the canonical element also makes Poulhalon and Schaeffer algorithm work by analogous arguments as in the plane with the additional requirement that there is no oriented cycles in the dual. In this thesis, I tried to analyze which properties are preserved or not when going from the plane to the torus. I think that it remains to understand the meaning of the vertices of outdegree greater than 3 to obtain a direct generalization of all the work on the torus to maps of higher genus.

For the torus case there remains many things to do. I claim that the bijection of Chapter 5 is the certificate that most of the implications of Schnyder woods in the planar case can be extended to the torus. The first thing to do is probably to count toroidal triangulations via our bijection with decorated unicellular maps. This would make possible to do uniform sampling as in the plane. There is probably a price to pay for partitioning and counting toroidal triangulations in order to obtain an accurate uniform sampling with the appropriate coefficients. Those coefficients can be precomputed. If the cost is too high, an approximation formula can be envisioned.

Bibliography

Albar, B., Gonçalves, D., and Knauer, K. (2014). Orienting triangulations. manuscript, arXiv:1412.4979.

Albenque, M. and Poulalhon, D. (2013). Generic method for bijections between blossoming trees and planar maps. arXiv:1305.1312.

Altshuler, A., Bokowski, J., and Schuchert, P. (1996). Neighborly 2-manifolds with 12 vertices. *Journal of Combinatorial Theory, Series A*, 75(1):148–162.

Arettines, C. (2015). A combinatorial algorithm for visualizing representatives with minimal self-intersection. *Journal of Knot Theory and Its Ramifications*, 24(11):1550058.

Barnette, D. W. and Edelson, A. L. (1988). All orientable 2-manifolds have finitely many minimal triangulations. *Israel Journal of Mathematics*, 62(1):90–98.

Barnette, D. W. and Edelson, A. L. (1989). All 2-manifolds have finitely many minimal triangulations. *Israel Journal of Mathematics*, 67(1):123–128.

Bernardi, O. (2007). Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electronic Journal of Combinatorics*.

Bernardi, O. and Chapuy, G. (2011). A bijection for covered maps, or a shortcut between harer-zagier's and jackson's formulas. *Journal of Combinatorial Theory A*, pages 1718–1748.

Birman, J. S. and Series, C. (1984). An algorithm for simple curves on surfaces. *J. London Math. Soc.*, 29(2):331–342.

Bonichon, N. (2002). Aspects algorithmiques et combinatoires des réaliseurs des graphes plans maximaux. *PhD thesis, Université de Bordeaux*.

Bonichon, N., Gavoille, C., and Hanusse, N. (2003). An information-theoretic upper bound of planar graphs using triangulation. *Proc. of the 20th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2003)*, pages 499–510.

Boulch, A., Colin de Verdière, É., and Nakamoto, A. (2013). Irreducible triangulations of surfaces with boundary. *Graphs and Combinatorics*, 29(6):1675–1688.

Cabello, S., Colin de Verdière, É., and Lazarus, F. (2011). Finding cycles with topological properties in embedded graphs. *SIAM J. Discrete Math.*, 25(4):1600–1614.

Castelli Aleardi, L., Fusy, E., and Lewiner, T. (2010). Optimal encoding of triangular and quadrangular meshes with fixed topology. *Proc. of the 22nd Canadian Conference on Computational Geometry (CCCG 2010).*

Chambers, E., Colin de Verdière, E., Erickson, J., Lazarus, F., and Whittlesey, K. (2006). Splitting (complicated) surfaces is hard. In *Proc. 22nd annual Symp. Comput. Geom.*, pages 421–429. ACM.

Chang, H.-C., Erickson, J., and Xu, C. (2015). Detecting weakly simple polygons. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1655–1670.

Chapuy, G. (2009). Combinatoire bijective des cartes de genre superieur. *PhD thesis, École Polytechnique.*

Chas, M. (2014). Self-intersection numbers of length-equivalent curves on surfaces. *Experimental Mathematics*, 23(3):271–276.

Chas, M. and Lalley, S. P. (2012). Self-intersections in combinatorial topology: statistical structure. *Inventiones mathematicae*, 188(2):429–463.

Chillingworth, D. R. (1969). Simple closed curves on surfaces. *Bulletin of the London Mathematical Society*, 1(3):310–314.

Chillingworth, D. R. (1971). An algorithm for families of disjoint simple closed curves on surfaces. *Bulletin of the London Mathematical Society*, 3(1):23–26.

Chillingworth, D. R. (1972). Winding numbers on surfaces. II. *Mathematische Annalen*, 199(3):131–153.

Cohen, M. and Lustig, M. (1987). Paths of geodesics and geometric intersection numbers: I. In *Combinatorial group theory and topology*, volume 111 of *Ann. of Math. Stud.*, pages 479–500. Princeton Univ. Press.

de Fraysseix, H. and Ossona de Mendez, P. (2001). On topological aspects of orientations. *Discrete Mathematics*.

de Graaf, M. and Schrijver, A. (1997). Making curves minimally crossing by Reidemeister moves. *Journal of Combinatorial Theory, Series B*, 70(1):134–156.

de La Harpe, P. (2010). Topologie, théorie des groupes et problèmes de décision: célébration d'un article de Max Dehn de 1910. *Gazette des mathématiciens*, 125:41–75.

Dehn, M. (1912). Transformation der kurven auf zweiseitigen flächen. *Mathematische Annalen*, 72(3):413–421.

Duchi, E., Poulalhon, D., and Schaeffer, G. (2013). Uniform random sampling of simple branched coverings of the sphere by itself. *Proc. of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, Society for Industrial and Applied Mathematics*, pages 294–304.

Ellingham, M. N. and Stephens, C. (2005). Triangular embeddings of complete graphs (neighborly maps) with 12 and 13 vertices. *Journal of Combinatorial Designs*, 13(5):336–344.

Ellingham, M. N. and Zha, X. (2003). Separating cycles in doubly toroidal embeddings. *Graphs and Combinatorics*, 19(2):161–175.

Erickson, J. and Whittelsey, K. (2013). Transforming curves on surfaces redux. In *Proc. of the 24rd Annual ACM-SIAM Symposium on Discrete Algorithms*.

Farb, B. and Margalit, D. (2011). *A Primer on Mapping Class Groups (PMS-49)*. Princeton University Press.

Felsner, S. (2004). Lattice structures from planar graphs. *Electronic Journal of Combinatorics*.

Fusy, E. (2007). Combinatoire des cartes planaires et applications algorithmiques. *PhD thesis*.

Goaoc, X., Mabillard, I., Paták, P., Patáková, Z., Tancer, M., and Wagner, U. (2015). On Generalized Heawood Inequalities for Manifolds: A Van Kampen-Flores-type Nonembeddability Result. In Arge, L. and Pach, J., editors, *31st International Symposium on Computational Geometry (SoCG 2015)*, volume 34 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 476–490, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

Gonçalves, D., Knauer, K., and Lévêque, B. (2015). Structure of Schnyder labelings on orientable surfaces. manuscript, 2015, arXiv:1501.05475.

Gonçalves, D. L., Kudryavtseva, E., and Zieschang, H. (2005). An algorithm for minimal number of (self-)intersection points of curves on surfaces. In *Proceedings of the Seminar on Vector and Tensor Analysis*, volume 26, pages 139–167.

Gonçalves, D. and Lévêque, B. (2013). Toroidal maps : Schnyder woods, orthogonal surfaces and straight-line representations. *Discrete and Computational Geometry*, pages 67–131.

Grannell, M. J. and Knor, M. (2010). A lower bound for the number of orientable triangular embeddings of some complete graphs. *Journal of Combinatorial Theory, Series B*, 100(2):216–225.

Grannell, M. J. and Knor, M. (2012). On the number of triangular embeddings of complete graphs and complete tripartite graphs. *Journal of Graph Theory*, 69(4):370–382.

Gross, J. L. and Tucker, T. W. (1987). *Topological graph theory*. Dover, reprint 2001 from Wiley edition.

Hass, J. and Scott, P. (1985). Intersections of curves on surfaces. *Israel Journal of Mathematics*, 51(1-2):90–120.

Hass, J. and Scott, P. (1994). Shortening curves on surfaces. *Topology*, 33(1):25–43.

Hass, J. and Scott, P. (1999). Configurations of curves and geodesics on surfaces. *Geometry and Topology Monographs*, 2:201–213.

Hatcher, A. (2002). *Algebraic topology*, volume 606.

Heawood, P. (1890). Map-color theorem. *Quart. J. Math. Oxford Ser. 24*.

Jennings, D. L. (2003). *Separating Cycles in Triangulations of the Double Torus*. PhD thesis, Vanderbilt University.

Joret, G. and Wood, D. R. (2010). Irreducible triangulations are small. *Journal of Combinatorial Theory, Series B*, 100(5):446–455.

Kampen, G. (1976). Orienting planar graphs. *Discrete Mathematics*, 14(4):337–341.

Kant, G. (1996). Drawing planar graphs using the canonical ordering. *Algorithmica*, pages 4–32.

Knuth, D. E., Morris, Jr, J. H., and Pratt, V. R. (1977). Fast pattern matching in strings. *SIAM journal on computing*, 6(2):323–350.

Korzhik, V. P. and Voss, H.-J. (2001). On the number of nonisomorphic orientable regular embeddings of complete graphs. *Journal of Combinatorial Theory, Series B*, 81(1):58–76.

Kühnel, W. (1994). Manifolds in the skeletons of convex polytopes, tightness, and generalized heawood inequalities. In *Polytopes: Abstract, Convex and Computational*, pages 241–247. Springer.

Lackenby, M. (2013). A polynomial upper bound on reidemeister moves. *arXiv preprint arXiv:1302.0180*.

Lawrencenko, S., Negami, S., and White, A. T. (1994). Three nonisomorphic triangulations of an orientable surface with the same complete graph. *Discrete Mathematics*, 135(1):367–369.

Lazarus, F., Pocchiola, M., Vegter, G., and Verroust, A. (2001). Computing a canonical polygonal schema of an orientable triangulated surface. In *Proceedings of the seventeenth annual symposium on Computational geometry*, pages 80–89. ACM.

Lazarus, F. and Rivaud, J. (2012). On the homotopy test on surfaces. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 440–449. IEEE.

Lévêque, B. (2016). Generalization of schnyder woods to orientable surfaces. *HDR thesis, Université de Grenoble*.

Lothaire, M. (1997). *Combinatorics on words*. Cambridge University Press.

Lustig, M. (1987). Paths of geodesics and geometric intersection numbers: II. In *Combinatorial group theory and topology*, volume 111 of *Ann. of Math. Stud.*, pages 501–543. Princeton Univ. Press.

Lyndon, R. C. (1966). On dehn's algorithm. *Mathematische Annalen*, 166(3):208–228.

Meunier, F. (2015). personal communication.

Mirzakhani, M. (2008). Growth of the number of simple closed geodesies on hyperbolic surfaces. *Annals of Mathematics*, 168(1):97–125.

Mirzakhani, M. (2016). Counting mapping class group orbits on hyperbolic surfaces. *arXiv preprint arXiv:1601.03342*.

Mohar, B. (1996). Straight-line representations of maps on the torus and other flat surfaces. *Discrete Mathematics*, pages 173–181.

Mohar, B. and Thomassen, C. (2001). *Graphs on Surfaces*. Studies in the Mathematical Sciences. Johns Hopkins University Press.

Nakamoto, A. and Ota, K. (1995). Note on irreducible triangulations of surfaces. *Journal of Graph Theory*, 20(2):227–233.

Neumann-Coto, M. (2001). A characterization of shortest geodesics on surfaces. *Algebraic and Geometric Topology*, 1:349–368.

Ossona de Mendez, P. (1994). Orientations bipolaires. PhD thesis, École des Hautes Études en Sciences Sociales, Paris.

Paterson, J. (2002). A combinatorial algorithm for immersed loops in surfaces. *Topology and its Applications*, 123(2):205–234.

Poincaré, H. (1904). Cinquième complément à l'analysis situs. *Rendiconti del Circolo Matematico di Palermo*, 18(1):45 –110.

Poulalhon, D. and Schaeffer, G. (2006). Optimal coding and sampling of triangulations. *Algorithmica*.

Propp, J. (1993). Lattice structure for orientations of graphs. *arXiv:math/0209005.*

Rademacher, H. and Steinitz, E. (1934). Vorlesungen uber die theorie der polyeder.

Reinhart, B. L. (1962). Algorithms for jordan curves on compact surfaces. *Annals of Mathematics*, pages 209–222.

Richmond, L. B. and Wormald, N. C. (1995). Almost all maps are asymmetric. *Journal of Combinatorial Theory, Series B*, 63(1):1–7.

Ringel, G. (1974). *Map color theorem*, volume 209. Springer.

Robertson, N. and Thomas, R. (1991). On the orientable genus of graphs embedded in the klein bottle. *Journal of graph theory*, 15(4):407–419.

Sapir, J. (2016). Bounds on the number of non-simple closed geodesics on a surface. *International Mathematics Research Notices*, page rnw032.

Schaeffer, G. (1998). *Conjugaison d'arbres et cartes combinatoires aléatoires*. PhD thesis, Université de Bordeaux.

Schnyder, W. (1989). Planar graphs and poset dimension. *Order*, pages 323–343.

Schnyder, W. (1990). Embedding planar graphs on the grid. In *SODA*, volume 90, pages 138–148.

Stillwell, J. (1993). *Classical Topology and Combinatorial Group Theory*. Graduate Texts in Mathematics. Springer.

Sulanke, T. (2006a). Generating irreducible triangulations of surfaces. arXiv:math/0606687.

Sulanke, T. (2006b). Irreducible triangulations of low genus surfaces. arXiv:math/0606690.

Turaev, V. G. (1979). Intersections of loops in two-dimensional manifolds. *Mathematics of the USSR-Sbornik*, 35(2):229.

Tutte, W. T. (1962). A census of planar triangulations. *Canad. J. Math*, 14(1):21–38.

Tutte, W. T. (1963). A census of planar maps. *Canad. J. Math*, 15(2):249–271.

Ueckerdt, T. (2011). Geometric representations of graphs with low polygonal complexity. *PhD thesis, Mathematiker Torsten Ueckerdt*.

Van Kampen, E. R. (1933). On some lemmas in the theory of groups. *American Journal of Mathematics*, 55(1):268–273.

Zha, X. and Zhao, Y. (1993). On non-null separating circuits in embedded graphs. *Contemporary Mathematics*, 147:349–349.

Zieschang, H. (1965). Algorithmen für einfache kurven auf flächen. *Mathematica Scandinavica*, 17:17–40.

Zieschang, H. (1969). Algorithmen für einfache kurven auf flächen II. *Mathematica Scandinavica*, 25:49–58.

Index