Trees, Decompositions, and Knot Theory

Thèse de doctorat en informatique de l'Université Gustave Eiffel

Mathématiques et Sciences et Technologies de l'Information et de la Communication (Mstic)
École Doctorale 532
Laboratoire d'Informatique Gaspard-Monge (LIGM)

Thèse présentée et soutenue à l'Université Gustave Eiffel, le 23/09/2024, par :

Corentin Lunel

Composition du Jury

Dominique ATTALI CNRS, Gipsa-lab	Examinatrice
Eric COLIN de VERDIÈRE CNRS, Université Gustave Eiffel	Membre invité
David EPPSTEIN University of California	Rapporteur
Livio LIECHTI Université de Fribourg	Examinateur
Delphine MOUSSARD Aix-Marseille Université	Examinatrice
Lionel POURNIN Université Sorbonne Paris Nord	Examinateur
Stephan TILLMANN University of Sydney	Rapporteur

Encadrement de la thèse

Arnaud de MESMAY CNRS, Université Gustave Eiffel	Directeur de thèse
Pierre DEHORNOY Aix-Marseille Université	Co-Directeur de thèse

Acknowledgements & Remerciements

Avant tout, je tiens à exprimer ma profonde gratitude envers mes maîtres de thèse, Arnaud de Mesmay et Pierre Dehornoy. Arnaud et Pierre, merci infiniment pour votre patience, votre expertise et pour les innombrables conseils que vous m'avez donnés tout au long de la thèse. Je vous remercie sincèrement pour l'ensemble des discussions que nous avons eues. En particulier, j'ai adoré faire de la recherche avec vous : ces moments d'échanges ont été de formidables expériences scientifiques. Merci encore pour toute votre aide et votre disponibilité tout au long de la thèse. Je suis honoré et heureux d'avoir été votre étudiant.

I am profoundly thankful to David Eppstein and Stephan Tillmann for taking the time to review this thesis and showing a sincere interest in my work. Thanks again for your comments and feedback. I am deeply grateful to the members of my thesis committee: Dominique Attali, Livio Liechti, Delphine Moussard, and Lionel Pournin; thank you for attending my defence and for your wonderful questions and comments. Additionally, I want to especially thank Éric Colin de Verdière for accepting to be the formal supervisor at the beginning of my thesis and for the many discussions we had throughout my thesis.

I want to express my gratitude to all of the seasoned researchers with whom I have had formal or informal discussions over the years, whether at conferences or at Marne-la-Vallée. I am sure that I will forget some, whom I'd like to thank now. For the others, thank you to Alfredo Hubard, Kristóf Huszár, Vincent Jugé, Clément Maria, Jonathan Spreer, Mathijs Wintraecken, and the members of ANR SOS.

Now that my thoughts turn to my academic family, I feel compelled to write in english just for you Niloufar, even though we both know you should be fluent in french by now. So thank you very much for your friendship during these years. I'm really grateful for your support, for your outstanding example as a PhD student, for the countless gossip breaks, and for the mesmerising time that

we shared. Merci à Jean Chartier pour les nombreuses discussions et soirées au long de ma thèse, ce fut un plaisir que de suivre et partager toutes tes aventures. Enfin, merci beaucoup à Loïc Dubois, j'ai été très content de partager mon bureau et d'échanger à de nombreuses occasions, avec toi.

Ma passion pour les mathématiques a été éveillée par deux de mes professeurs. Merci profondément à Patrick Cabau et Vincent Bayle pour votre enseignement exceptionnel, et pour votre soutien au début de mes études supérieures en ce domaine.

Je tiens aussi à remercier Aurore Blelly, dont la compagnie et le soutien moral ont été salvateurs pendant ces années. Merci énormément d'avoir pris le temps de relire mon manuscrit, surtout dans un laps de temps aussi court... Je suis sûr que ceci aura contribué à renforcer la vision unifiée des maths que tu m'as si souvent partagée.

Merci à tous ceux qui ont pris le temps de discuter de mon travail et de relire tout ou partie de mon introduction. Merci ainsi à Clément Brochet, Damien Théret, Alexis, Stéphanie et Yves-Laurent Lunel. Je tiens aussi à remercier particulièrement Achille Fraisse pour son soutien moral et son enthousiasme lors de nos nombreuses discussions.

Je remercie chaleureusement les membres de ma famille que je n'ai pas encore mentionnés pour leur soutien et leur présence.

Enfin, merci à toi Audrey, pour l'aide inestimable que tu m'as apportée, en particulier pendant la rédaction de cette thèse, et ton soutien indéfectible.

Abstract

Graph theory and knot theory are two well established mathematical fields which present profound interactions. This thesis focuses on investigating some knot theory problems from a computational point of view and drawing inspiration from results and methods stemming from graph theory.

The first problem we address concerns the decidability of a knot invariant. The genus of a knot is a classical knot invariant: it is the minimal genus of an embedded orientable surface in the 3-dimensional space admitting the knot as its boundary. It is now fairly well understood from a computational perspective. On the contrary, no algorithm is known for its four-dimensional variants, both in the smooth and in the topologically locally flat category. We investigate a class of knots and links called Hopf arborescent links, which are obtained as the boundaries of surfaces constructed by iterated plumbings of Hopf bands. We show that for such links, computing the genus defects, which measure how much the four-dimensional genera differ from the classical genus, is decidable. Our proof is non-constructive and is obtained by proving that a containment relation on surfaces associated with Hopf arborescent links forms a well-quasi-order.

The second problem we tackle is motivated by the existence of efficient algorithms to compute many knot invariants and properties on diagrams of low treewidth. It was recently proved that there exist knots which do not admit any diagram of low treewidth, and the proof relied on intricate low-dimensional topology techniques. We initiate here a thorough investigation of tree decompositions of knot diagrams (or more generally, diagrams of spatial graphs) using ideas from structural graph theory. We define an obstruction on spatial embeddings that forbids low treewidth diagrams, and we prove that it is optimal with respect to a related width invariant. We then show the existence of this obstruction whenever an embedding into a surface with high compression-representativity exists, which is the case for torus knots. Thus, we provide a new and self-contained proof that those do not admit diagrams of low treewidth.

Finally, we shift our focus toward the complexity of knot diagrams throughout the action of Reidemeister moves. Recognising the trivial knot is a fundamental question of knot theory. A natural way to attack this problem is by applying Reidemeister moves on a diagram in a brute force or random manner until the diagram corresponds to a circle. It turns out that the number of crossings of some unknot diagrams must increase during the execution of this algorithm. No super constant lower bound is known on how many such crossings need to be added during the execution of similar algorithms. The problem of deciding whether two links are split i.e., can be separated by a sphere, is approachable in the same way and presents the same issue. We prove that there exist link diagrams that require an arbitrarily large number of added crossings to be split via this method.

Keywords: Knot theory, Computational topology, Structural graph theory, Knot diagram, Geometric topology

Résumé

La théorie des graphes et la théorie des nœuds sont deux célèbres domaines mathématiques qui présentent de profondes interactions. Cette thèse se concentre sur l'étude de certains problèmes de théorie des nœuds d'un point de vue informatique et en s'inspirant de résultats et méthodes issues de la théorie des graphes.

Le premier problème que nous abordons concerne la décidabilité d'un invariant de nœud. Le genre d'un nœud est un invariant classique : c'est le genre minimal d'une surface orientable plongée dans l'espace de dimension 3 qui est bordée par le nœud. De nos jours, il est assez bien compris d'un point de vue informatique. En revanche, aucun algorithme n'est connu pour ses variantes en dimension 4, à la fois dans la catégorie des variétés lisses et dans celle des variétés localement plates. Nous étudions une classe de nœuds et d'entrelacs appelés entrelacs arborescents de Hopf, qui sont obtenus comme bords de surfaces construites par des plombages itérés de bandes de Hopf. Nous montrons que, sur ces entrelacs, le calcul des défauts de genre, qui mesurent à quel point les genres quadridimensionnels diffèrent du genre classique, est décidable. Notre preuve est non constructive et est obtenue en prouvant qu'une relation d'inclusion sur les surfaces associées aux entrelacs arborescents de Hopf forme un bel ordre.

Le deuxième problème que nous attaquons est motivé par l'existence d'algorithmes efficaces pour calculer de nombreux invariants et propriétés de nœuds sur des diagrammes de faible largeur arborescente (*treewidth*). Il a été récemment prouvé, par de complexes résultats de topologie en basse dimension, qu'il existe des nœuds qui n'admettent pas de diagrammes de faible largeur arborescente. Nous entamons ici une étude approfondie des décompositions arborescentes des diagrammes de nœuds (ou plus généralement, des diagrammes de graphes spatiaux) en utilisant des idées de la théorie structurelle des graphes. Nous définissons une obstruction sur les plongements dans l'espace qui s'oppose à des diagrammes de faible largeur arborescente. Nous prouvons de plus que cette obstruction est optimale vis-à-vis d'un invariant de largeur que nous définissons. Nous montrons ensuite l'existence d'une telle obstruction dès qu'il existe un plongement dans une surface avec haute représentativité. Cette dernière partie est toujours vérifiée sur les nœuds toriques : nous fournissons ainsi une nouvelle preuve que ces nœuds n'admettent aucun diagramme de faible largeur arborescente.

Pour finir, nous nous concentrons sur la complexité des diagrammes de nœuds sous l'action de mouvements de Reidemeister. Reconnaître le nœud trivial est un problème fondamental de théorie des nœuds. Une façon naturelle de l'attaquer est d'appliquer des mouvements de Reidemeister sur un diagramme du nœud de manière exhaustive ou aléatoire jusqu'à ce que le diagramme corresponde à un cercle. Il s'avère que le nombre de croisements de certains diagrammes du nœud trivial doit nécessairement augmenter pendant l'exécution de cet algorithme. De plus, on ne connaît pas de borne inférieure meilleure que constante sur le nombre de croisements qui doivent être ajoutés pendant l'exécution d'algorithmes similaires. Cet algorithme peut aussi être appliqué au problème de décider si un entrelacs est séparé, c'est-à-dire, s'il existe une sphère qui sépare deux de ses composantes. Il présente le même problème qui est de devoir ajouter des croisements à un diagramme pendant son

exécution. Nous prouvons qu'il existe des diagrammes d'entrelacs qui nécessitent un nombre arbitrairement grand de croisements à ajouter pour être séparés en utilisant cet algorithme.

Mots clefs : Théorie des nœuds, Topologie algorithmique, Théorie structurelle des graphes, Diagrammes de nœuds, Topologie géométrique

Contents

Chapter 1

Introduction

This thesis falls within the domain of computational topology: a mathematical field that aims at formalising and dealing with topological questions from a computational point of view. In particular, we are interested in the rich connections appearing between graph theory and knot theory. In the following, we will broadly present these domains while highlighting the links that they offer. Then, we will present more precisely the contributions and the organisation of this thesis.

This first chapter is intended to be readable with little mathematical knowledge, although the concepts and notions presented will be more difficult when moving to the contributions.

1.1 General presentation

Topology. Topology can be defined as the branch of mathematics studying *shapes*. More precisely, it focuses on properties that are preserved by continuous operations. Such properties are called topological properties. For instance, stretching, twisting, or bending are continuous operations, while cutting or glueing are not. An example of property preserved by such continuous operations is the ability to reach any point of a given object from any other one; an object satisfying this property is called **connected**. If an object is made of exactly two disjoint parts, no amount of bending, twisting, or expanding will make us able to reach a point on the second part when starting from a point on the first part. Only when the two parts will merge i.e., will be glued (which is not a continuous operation), one will be able to join the two aforementioned points, and the object will be connected. Conversely, bending, twisting, or expanding will never change the fact that an object is connected.

Historically, the birth of topology can be attributed [17, Chapter 10] to Riemann, whose work in the 1850s led to the definition of topological spaces, which are the building blocks of topological theories. However, fundamental results like Euler's formula go back to Descartes in the 17th century (this formula will be discussed later in this section in the proof of Proposition 1.1), and the key notion of continuity was already conceived by Greek philosophers like Aristotle. It is now a very active field of research whose results are prevalent in many

branches of mathematics. Famously, among the Millennium Prize Problems, which are seven renowned and highly difficult mathematical problems selected by the Clay Mathematics Institute in 2000, only one, which is of topological nature and called the Poincaré conjecture, has been solved [106, 107, 108]. Recently, with the rise of theoretical computer science, a need for the computation of topological properties of computer simulations appeared. This pushed the development of computational topology as presented above and fuelled flourishing exchanges of tools, methods, and problems to study between classical topology and theoretical computer science.

As said above, topology focuses on properties preserved by continuous operations. As such, it is natural that topology will consider two objects equivalent when they only differ by a continuous operation. Depending on the operation in question, the equivalence relation will be different. One of the most basic stems from homeomorphisms: a **homeomorphism** is a bijective and continuous map between two spaces that has a continuous inverse function. All inherent topological properties of two homeomorphic objects are the same. For instance, $\mathbb{R}$ and $\mathbb{R}_+^*$ are homeomorphic via the function $\exp : \mathbb{R} \to \mathbb{R}_+^*$. However, there is no homeomorphism between $[0, 1]$ and $\mathbb{R}$ because they are topologically different. One topological difference between $[0, 1]$ and $\mathbb{R}$ is the fact that everywhere in $\mathbb{R}$ one can move forward and backward, but that is not the case at the endpoints of $[0, 1]$ where only one of these movements is possible (see Figure 1.1).

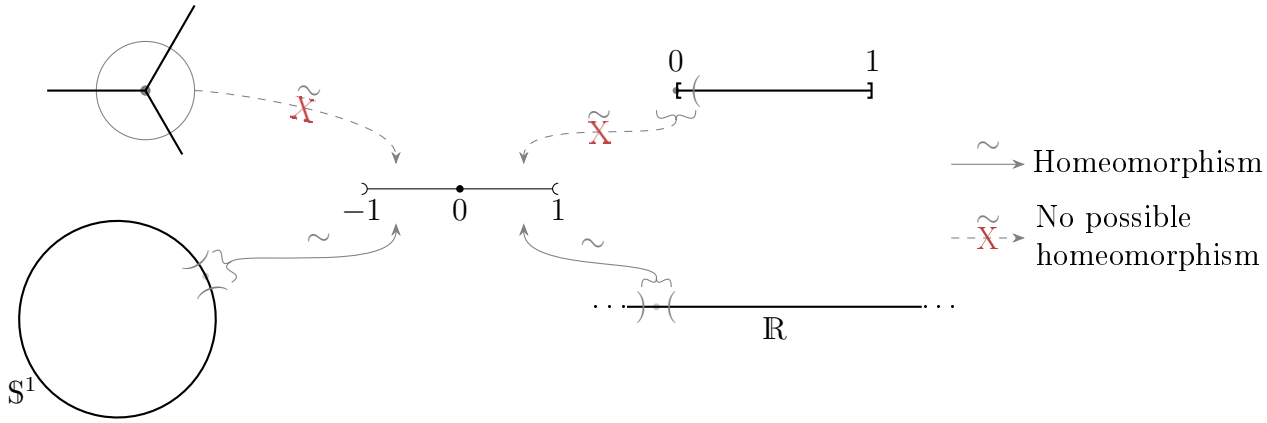

Figure 1.1: Two objects non homeomorphic to $(-1, 1)$ around the illustrated points (3 or only 1 direction for movement is possible) while $\mathbb{S}^1$ (a circle) and $\mathbb{R}$ are locally homeomorphic to $(-1, 1)$.

Intuitively, this notion of the number of directions of movement, or number of degrees of freedom around a point, is what we call the **dimension** of an object. For now this definition is flawed since it can depend on the point considered. Let us get more formal and define more precisely a common object studied in topology called **n-manifold**. An object M is an n-manifold if locally it looks like the n-dimensional space $\mathbb{R}^n$, i.e., for each point of M there exists a neighbourhood homeomorphic to the open unit n-ball of $\mathbb{R}^n$: $\{x \in \mathbb{R}^n \mid \|x\| < 1\} = \mathbb{B}^n$. This last property is naturally preserved by homeomorphisms. Hence, any object homeomorphic to an n-manifold is an n-manifold too. Thus $\mathbb{R}$ is a 1-manifold while $[0, 1]$ is not.

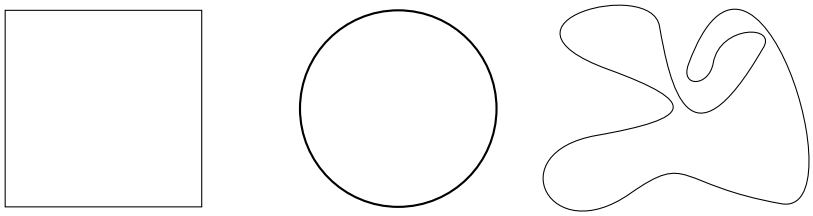

Figure 1.2: Three $\mathbb{S}^1$, that is to say three objects homeomorphic to a circle.

We can already classify connected 1-manifolds. If, when starting from a point and keeping a direction, one is able to reach again this starting point, then the 1-manifold is in fact homeomorphic to a circle denoted $\mathbb{S}^1$, commonly described by $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$, and we will say that the manifold is a $\mathbb{S}^1$ (see Figure 1.2). Otherwise, the previous walk will never end, and the manifold is a line i.e., it is homeomorphic to $\mathbb{R}$. More generally, as it is common for topologists, we will shorten the expression "is homeomorphic to" by "is". Hence, a n-sphere denoted $\mathbb{S}^n$ is any n-manifold homeomorphic to the unit Euclidean sphere of $\mathbb{R}^{n+1}$: $\{x \in \mathbb{R}^{n+1} \mid \|x\| = 1\}$. In particular, we choose the unit sphere as a common representative, but any radius would be equivalent since increasing and reducing the radius are continuous operations.

Before moving to surfaces, which are the 2-manifolds, let us introduce a second topological property, which will be a key concept in the rest of this section: embeddability. A fundamental topological operation is that of **embedding**: an embedding $j : X \to Y$ of an object X into an object Y is a homeomorphism f from X onto its image $f(X)$. If such a map exists, X is said to be **embeddable** into Y. Since the composition of two homeomorphisms is a homeomorphism, it follows that the notion of embeddability is a property preserved by homeomorphisms.

Since $[0, 1]$ can be embedded in $\mathbb{R}$ via the natural inclusion and $\mathbb{R}$ into $[0, 1]$ via $\frac{1}{2} + \frac{1}{\pi}\arctan$, we cannot distinguish $[0, 1]$ and $\mathbb{R}$ via embeddability. Indeed, the objects they can be embedded into, or the objects which can be embedded in them, are the same. However, such a notion allows us to distinguish between the two 1-manifolds $\mathbb{R}$ and $\mathbb{S}^1$: $\mathbb{R}$ embeds into $\mathbb{S}^1 = \left\{e^{i\theta} \mid \theta \in [0, 2\pi)\right\}$[1] $\subset \mathbb{C}$ via $x \mapsto e^{i \arctan(x)}$ but there is no embedding from $\mathbb{S}^1$ into $\mathbb{R}$ since such an embedding would require two disjoint paths in $\mathbb{R}$ between the images of two disjoint points of $\mathbb{S}^1$.

Surfaces. Moving up one dimension, we now focus on 2-manifolds, which are called **surfaces**. First, let us discuss some examples. The plane $\mathbb{R}^2$ and the open unit 2-ball are both surfaces: for each point of both these objects, there is a neighbourhood homeomorphic to an open 2-ball (take any small enough disc). As can be seen in Figure 1.3, and up to bending a little, the same is true for infinite tubes $\mathbb{R} \times \mathbb{S}^1$ and the torus $\mathbb{T} = \mathbb{S}^1 \times \mathbb{S}^1$, which can be seen as the surface of a doughnut.

[1]Notice that this description of $\mathbb{S}^1$ by $\phi : [0, 2\pi) \to \mathbb{S}^1$ such that $\phi(\theta) = e^{i\theta}$ is an example of a continuous bijection that is not an embedding: the inverse map ϕ^{-1} is not continuous at $\phi(0)$. Indeed the limit of ϕ^{-1} at this point, by continuity, is 0 from one side and 2π from the other one.

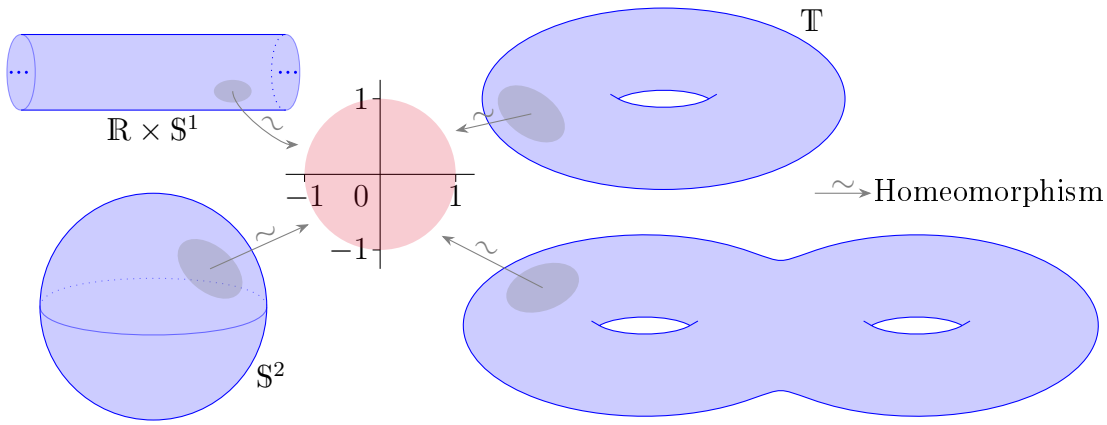

Figure 1.3: Some examples of surfaces: the surroundings of each of their points are homeomorphic to the open unit disc depicted in the middle.

Let us emphasise again that we usually consider objects up to homeomorphism. Hence, we do not distinguish between a cube and a 2-sphere or between an open disc and a triangle as shown in Figure 1.4. A common description of this phenomenon is to picture every object as if they were made of rubber.

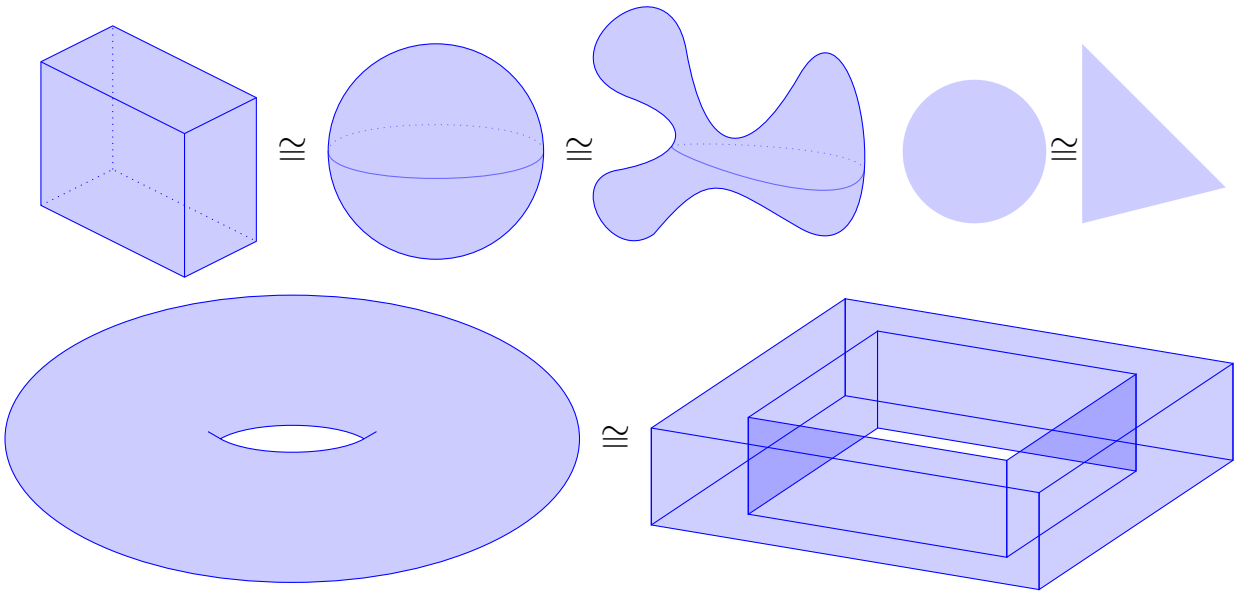

Figure 1.4: Some surfaces homeomorphic to: the 2-sphere (top left), the open unit 2-ball (top right), and the torus (bottom). The homeomorphism relation is denoted by $\cong$.

We will now explain why, among the previous examples of surfaces, the torus is topologically more complex. This will be illustrated by the continuous transformations which can be applied to a **curve**: a $\mathbb{S}^1$ continuously mapped to the considered surface (it need not be an embedding). If such a curve can be contracted to a point, it is said to be **contractible**; otherwise, it is non-contractible. On a 2-sphere S, every curve is contractible (see Figure 1.5).

Indeed, if we picture them as rubber bands moving continuously on S, we can move these curves so that they will stay on a cap of the sphere that will eventually taper to a point (picture the equator sliding toward the north pole of the sphere of Figure 1.5).

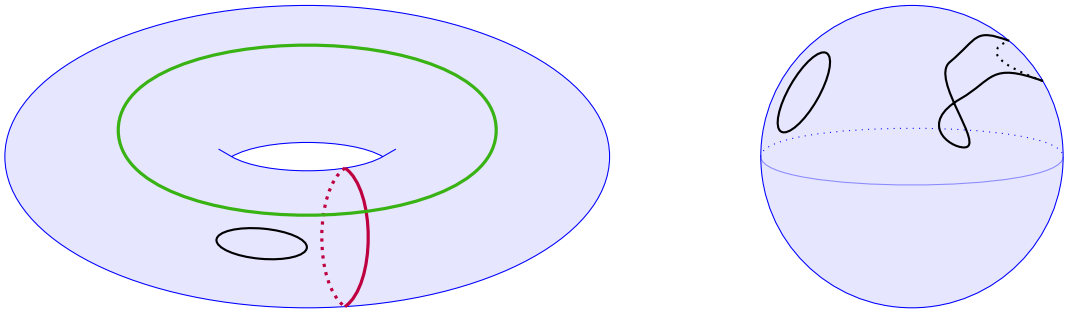

Figure 1.5: Some non-contractible curves of the torus, a meridian in purple and a latitude in green. The other curves, in black, can be continuously reduced to a point.

On the torus, there exist non-contractible curves, two of which can be seen in Figure 1.5, in purple and green. From our point of view, which considers everything up to continuous deformation on the surface, these non-contractible curves are non trivial: they are not equivalent to points, while contractible curves are. Such a non-contractible curve attests to the existence of a hole in the torus that does not exist in 2-spheres. Among the surfaces of Figure 1.3, only the surface on the bottom right, and $\mathbb{T}$, the torus, admit non-contractible curves. It is worth noticing that the surface on the bottom right has intuitively two holes (two distinct and disjoint non-contractible curves) while the torus has only one.

The previous examples of surfaces present another topological difference: some of them are *compact*, like the 2-sphere and the torus, while $\mathbb{R}^2$, the open 2-ball $\mathbb{B}^2$, or the infinite tube $\mathbb{S}^1 \times \mathbb{R}$ are not. These last surfaces present some kind of infinite behaviour that the first ones do not. The precise definition of a compact manifold would require a level of detail exceeding the ambitions of this section. Hence, we will settle for the following: a surface is not compact if we can find a sequence of points of the manifold that converges "outside" of the manifold, like the sequence $(1-\frac{1}{n}, 0)_{n\in\mathbb{N}}$ in the unit open 2-disc, which converges toward $(1, 0)$, "outside" of the disc. From our computational point of view, the compactness property is especially useful. Indeed, the compact surfaces can be described by closed triangles glued on their boundaries: they can be *triangulated* by a finite number of triangles [130]. For instance, tori are often represented by a rectangle whose opposite sides are identified (see bottom of Figure 1.6): gluing two of these sides results in a tube. The last stage is gluing its ends to get the torus.

Compact surfaces can be input into computers, or into their abstract equivalent called *Turing machine*, through the description of a finite number of triangles and how they are glued on their boundaries. Such a discrete representation could be surprising since topology focuses on continuous structures and operations, which seem to be opposed to a discrete description. However, that is not the case: these discrete representations are equivalent to the continuous one. Thus, they encapsulate the entirety of the topological properties of surfaces and are far more convenient to manipulate algorithmically. Hence, computational topologists tend to prefer compact manifolds to non-compact ones.

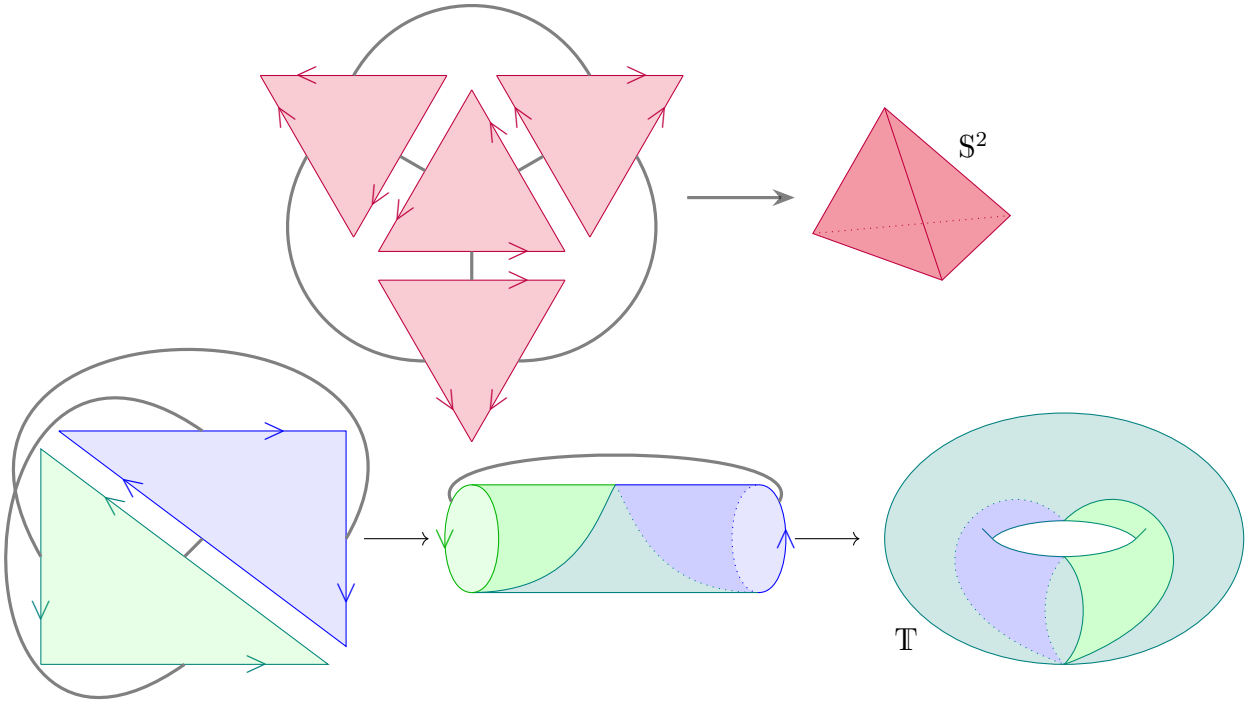

Figure 1.6: Representation of a 2-sphere and a torus by gluing pairs of segments on the boundaries of triangles according to the grey links and orientations.

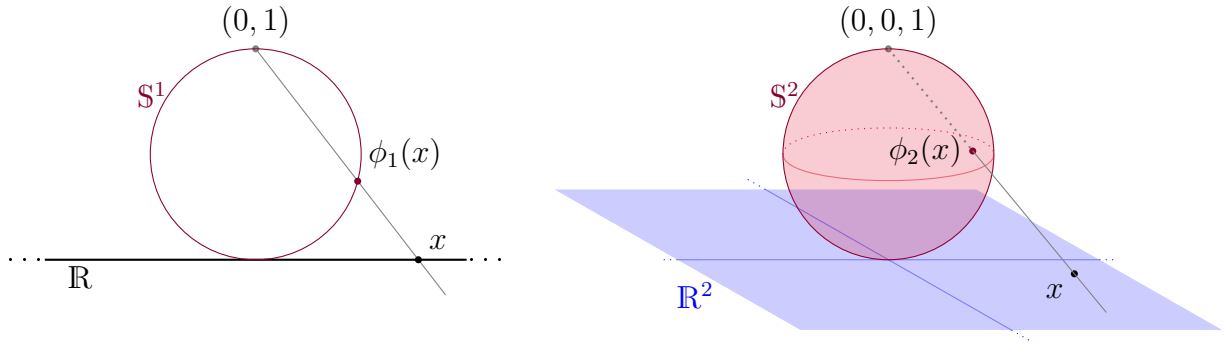

Figure 1.7: Left: the real line $\mathbb{R}$ compactified into $\mathbb{S}^1$, here the unit sphere of $\mathbb{R}^2$, by ϕ_1. Right: the plane $\mathbb{R}^2$ compactified into $\mathbb{S}^2$, here the unit sphere of $\mathbb{R}^3$, by ϕ_2.

We have seen earlier that $\mathbb{S}^2$ and $\mathbb{R}^2$ both have no non-contractible curves. In fact, they have a very similar topology, so much so that they are the same up to one point: there is a homeomorphism that we will denote by ϕ_n between each $\mathbb{R}^n$ and $\mathbb{S}^n \setminus \{N\}$ where N is the "north pole" of $\mathbb{S}^n$. Let us detail the construction of this homeomorphism, which is called the **stereographic projection** (see Figure 1.7). We first put $\mathbb{S}^n$ on top of $\mathbb{R}^n$, seen as a hyperplane of $\mathbb{R}^{n+1}$. Then, for each point x of $\mathbb{R}^n$, we define the half-line ℓ_x starting at $N = (0, \dots, 0, 1)$ and passing across x. There is exactly one intersection between ℓ_x and $\mathbb{S}^n$, this is $\phi_n(x)$. Notice that, intuitively, the "infinity" of $\mathbb{R}^n$ is sent to N by continuity. Since $\mathbb{S}^n$ is compact while $\mathbb{R}^n$ is not, we will say $\mathbb{S}^n$ is the *Alexandrov compactification* of $\mathbb{R}^n$. The

stereographic projection allows us to reduce problems in $\mathbb{R}^n$ to problems in $\mathbb{S}^n$, which are manifolds we can handle more easily from a computational point of view.

Finally, the fact that every 2-manifold can be triangulated is also verified by 3-manifolds [96, 130]. In the latter case, the building blocks are not triangles but tetrahedra (the 2-sphere, pictured on the top right of Figure 1.6, if filled on its inside, is a tetrahedron) that are glued on their triangular faces. It means that 3-manifolds can also be input to Turing machines and studied by theoretical computer science. This essential property is not satisfied by n-manifolds when $n \geq 4$, and this fact shows that low-dimensional spaces are particularly amenable to computational questions. On the contrary, in high dimensions, many questions cannot even be put in an algorithmic framework.

Triangulations of surfaces can actually be seen on old 3D computer graphics where triangles can be outlined on the surfaces used to depict objects. Nowadays these triangles are harder to delineate. They are small enough to create the illusion of a continuous smooth surface. However, the underlying triangulations are what make the computations and simulations possible. In addition, in a similar way that non-contractible curves are essential to define holes in surfaces, hence being constitutive of the topological behaviour of surfaces, surfaces within 3-manifolds are constitutive of their topological properties. Thus, the computations involving 3-manifolds exploit surfaces. It goes without saying that computations involving 3-manifolds are major since the space we live in is, on our scale, a 3-manifold: it follows that real-world simulations need surfaces and a computational understanding thereof.

Graphs. As we explained, discrete structures are highly convenient for computational purposes. It is then natural to shift our focus to graphs, which are a major part of both mathematics and computer science. Their study is the purpose of graph theory. A graph consists of two pieces of information: on the one hand, a set of points called **vertices**, and on the other hand, the links between them, which are called **edges**. Edges and vertices can both be supplemented with further information, like weights, colours, orientation, or labels. Hence, as mathematical structures, they are highly flexible to encompass information about objects and relations between them and have a wide range of applications as models. Any company organisation chart, subway map, network, or dependency graph is a graph of some sort. Graphs are often described visually by depicting vertices by points and edges by segments (see Figure 1.8 for example). In fact, we have already pictured some graphs, since Figure 1.6 can be seen as a graph whose vertices are the segments composing the boundary of each triangle and edges are the grey links representing their gluing.

The seminal paper of graph theory is a paper published in 1736, written by Leonhard Euler, about the Seven Bridges of Königsberg [39]. Its goal was to settle the unsolvability of an old mathematical challenge of the inhabitants of Königsberg. That challenge, stated as whether a tour of the city using each bridge exactly once exists, could not be proved until formalised within the framework of graph theory. As we want to emphasise links between graph theory and topology, it is meaningful that this paper is also often quoted as fundamental for the birth of topology since its concerns are of topological nature.

Since graphs are fairly understandable objects, some famous graph results are also approachable. The 4-color theorem is one of them: it states that any geographic map can be coloured using four distinct colours in such a way that any two regions sharing a border have

different colours. This theorem, which can naturally be expressed in the framework of graph theory, withstood attempts at proofs for more than 100 years before a computer-assisted proof succeeded in 1969 [59]. Graph theory is currently a very active field of research which presents significant interactions with other fields of mathematics and computer science, like algebra, probability, data science, and, more importantly for us, topology.

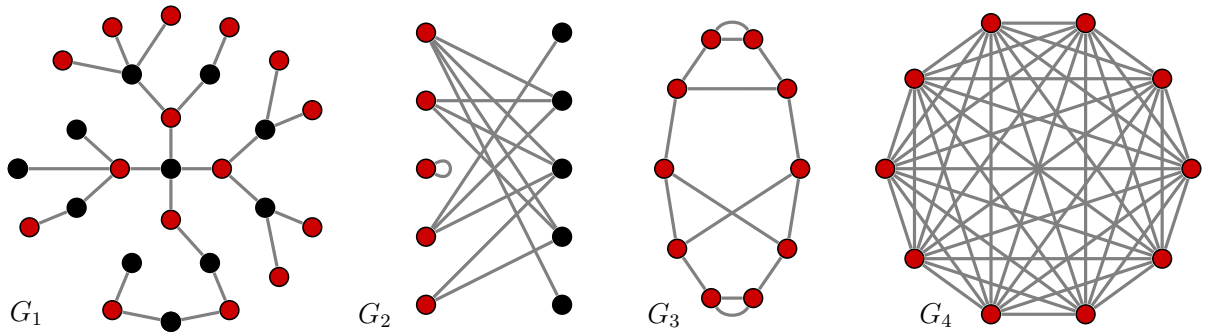

Figure 1.8: Some examples of graphs. Among them, G_1 is a tree and G_4 is the complete graph K_{10}.

Let us now dive into some graph definitions and properties using the 4 graphs of Figure 1.8. A **path** on a graph is a sequence of distinct vertices such that each one of them is linked to the next one by an edge. As in topology, we say that a graph where any pair of points can be joined with a path is **connected**. All graphs of Figure 1.8 are connected, except G_2 where one vertex is isolated. A **cycle** is a path except on its endpoints, which must be the same. When seen as a continuous space, that is to say, if edges are seen as segments that are glued on their endpoints when they share a vertex, a cycle is a non-contractible curve embedded on the graph. Some problems can be reduced to the search of a specific cycle in a given graph: if one looks for an exciting, exhaustive tour of a country (visiting the same city twice is boring), it amounts to looking for a cycle visiting each vertex exactly once.

Let us digress about computational complexity. We can endow the previous edges with the corresponding length of the road they represent. In this graph, finding a cycle of minimal length passing through each vertex is known as a particularly hard to solve computational problem called the Travelling Salesman Problem. Here, hard does not mean finding a way to compute the solution, since there exists a rather naive algorithm solving the problem: this algorithm proceeds by enumerating all the possible orders of visit of the vertices and remembers which one is the shorter. The *running time* of an algorithm is the number of operations that is performed by the algorithm, seen as a function of the size of the entry. The greater the number, the more time will be required to run the algorithm. Finding an algorithm with the best possible running time and classifying problems depending on this best possible running time is one of the major goals of computer science. For our naive algorithm, since there are $n!$ visiting orders where n is the number of vertices, it will take a huge amount of time to run as each order is considered. The problem is considered hard because we do not know, yet, of any algorithm with a time complexity better than exponential which solves this problem in the general case.

Coming back to cycles in graphs, the graph G_4, called the **complete graph** on 10 vertices, contains an edge between each pair of vertices, and thus a lot of cycles. We will denote by

K_n the complete graph on n vertices. On the contrary, graphs may also have no cycles. See G_1 for instance. A graph that is both connected and without a cycle is called a **tree**. In that regard, trees have a very simple topology. Hence, they are highly interesting and useful: similarly to spheres, they have no non-contractible curves. Another of their properties is that there is a unique path between each pair of their vertices. For example, trees are highly used in computer science as an efficient data structure. One can be interested in summarising tasks to accomplish for a process in a dependency graph: vertices are tasks to accomplish, and edges are dependencies between them. Here, edges are oriented: if a first task is linked by an oriented arrow to another one, it means the first task must be completed before the next one starts. It is then crucial that no cycle appears during the creation of such a dependency graph.

Graphs can have several edges between two vertices, like G_3, or even an edge coming from a vertex to itself, like the isolated vertex of G_2: such an edge is called a **self loop**. Graphs can be **bipartite**: their vertices can be split into two sets, such that there is no edge between two vertices of the same set. Both G_1 and G_2 are bipartite, the colours of vertices describing the aforementioned splitting. Every tree is bipartite: if we first fix a vertex v in the tree and colour any other one according to the parity of the number of edges on the path between it and v, it yields the desired partition. Notice that cycles of bipartite graphs must be of even length. Since there are as many edges as there are vertices in cycles, and because these vertices alternate between the two sets of the partition, there is an even number of vertices and an even number of edges too. Finally, we can define the complete bipartite graph $K_{n,m}$, it is the bipartite graph where A has n vertices, B has m vertices, and there is an edge between each vertex of A and each vertex of B (see Figure 1.9).

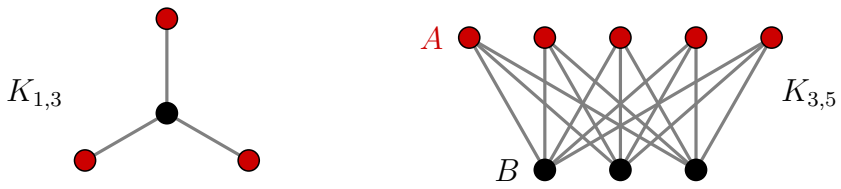

Figure 1.9: Two examples of bipartite complete graphs.

From a computational point of view, the structural property of being bipartite has some algorithmic benefits. For example, it is easier to compute a maximum matching i.e., a maximum sized set of edges that do not share vertices, in a bipartite graph. Similarly, many algorithmic problems are far easier on trees, since they contain few edges and no cycles. A fruitful idea, from a computational perspective, has been to identify graphs that look like trees and generalise algorithms on trees to them. Hence, many measures of how close to a tree a graph is have been developed. The *treewidth* is one such measure. Intuitively, elements of a graph of treewidth k can be arranged in *bags* of size $k+1$ and presented in the shape of a tree to form a *tree decomposition*. Figure 1.10 shows a tree decomposition of size 3.

Although tree decompositions must satisfy several properties, we do not want to delve too much into technical definitions here. For the purpose of this introduction, it is enough to say that the treewidth is the minimum width of a tree decomposition and to focus on the following examples. Some graphs have high treewidth, like complete graphs or grids (see the

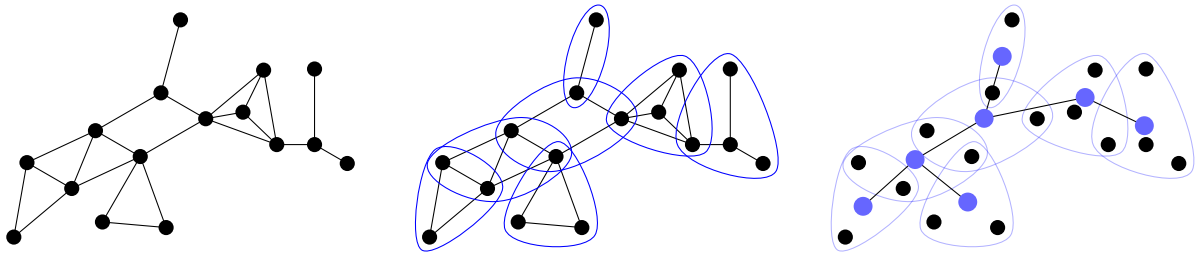

Figure 1.10: Left: a graph G. Middle: a decomposition of vertices of G in bags. Right: A tree decomposition of G.

right part of Figure 1.11): they are far from trees. Some other graphs, like trees or graphs like the one on the left part of Figure 1.11 have low treewidth.

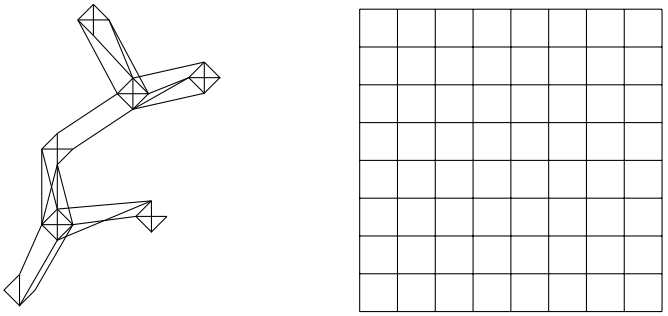

Figure 1.11: Left: A graph with low treewidth. Right: a graph with high treewidth.

When faced with an algorithmic problem on a graph of treewidth k, one can first find a tree decomposition of width k. Then, one can solve the problem on each bag and try to recompose a solution for the whole graph using the tree. This approach has proven to be very fruitful for many computational problems.

Planar graphs. A very natural issue appears to anyone who tries to draw a graph such that the drawing is easy to read. An obstacle to readability is crossings between edges so that one usually tries to draw the graph with no crossings. For example, the two drawings of Figure 1.12 represent the same graph, but the one on the right has no crossings, and therefore is easier to read.

Hence the following definition: a graph is said to be **planar** if there exists a representation of this graph in the plane such that no two edges cross. Here, a class of graphs is defined by a topological property, namely being embeddable in the plane or, since it is equivalent, in the 2-sphere.

This property is particularly useful for some applications when edges represent connections that must be constructed. Such examples are printed circuits, where vertices are components and edges are circuits etched on a sheet, or more simply road design, where vertices are cities and edges are the roads between them. In both cases, one wants to avoid crossings since they are costly to handle; in our examples, they are handled by, respectively, an electrical component or a bridge.

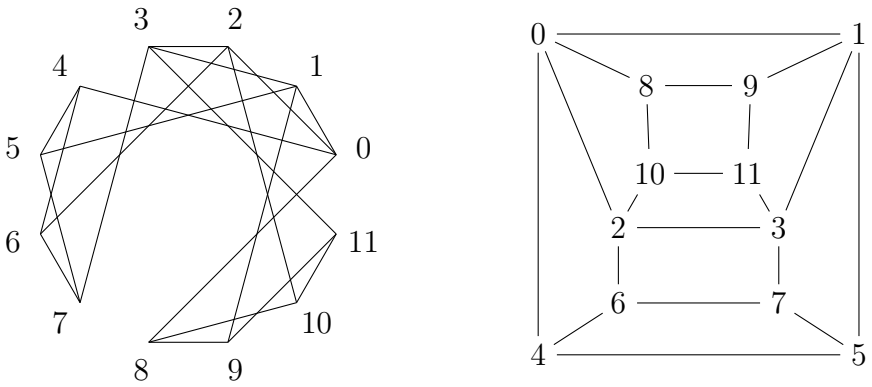

Figure 1.12: Two representations of the same graph, one being planar. Two vertices are adjacent if they differ by only one bit when expressed in binary.

Here is a famous mathematical puzzle called "the three utilities problem", first recorded in 1913 by H. Dudeney [35] who was already calling it an old problem. The aim is to connect, without crossings, three houses to three utilities, for example, gas, electricity, and water.

We encourage the reader who has never tried to solve this problem to try and doodle on a sheet of paper to solve this problem, at least until some doubt about the feasibility of the problem appears. In fact, any attempt will end up looking similar to the left part of Figure 1.13, where all edges except one are drawn. At this point, the last edge cannot be added between the leftmost house and the gas facility. Indeed, the first house is enclosed into a circle made of 4 edges (the ones between the middle and rightmost houses and the water and electricity facilities). Indeed, the famous Jordan [67] theorem states that any path between a point outside of the circle (the gas facility) and a point inside (the first house) will necessarily cross an edge.

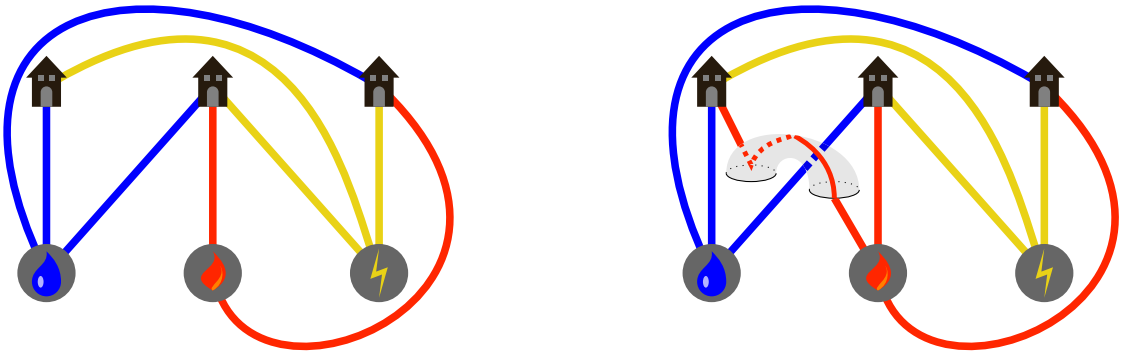

Figure 1.13: On the left, an attempt to solve the three utilities problem. On the right, the only way to solve this problem: by attaching a tunnel to the plane.

However, there is a way to solve this problem by working not on a sheet of paper but on a mug. Once again we encourage the reader who has never tried to solve this problem on a mug and who is not afraid to wash it afterwards to draw on it with a marker. The key is to have the stems of the handle lying on distinct faces. That way, in a manner similar to the

right part of Figure 1.13, where a tunnel is attached on its boundaries to two holes made on the plane, one will be able to draw the last edge without it crossing the remaining edges. Here, by modifying the surface on which we try to draw the graph, the problem becomes solvable.

In topology, the formulation of this problem is: "is $K_{3,3}$ embeddable on a 2-sphere?". And topology provides tools to answer this problem. The arguments presented above are not really a proof, since they rely on "any attempt will end up like this". Let us try a more formal approach and prove Proposition 1.1.

Proposition 1.1. *$K_{3,3}$ is non planar: it cannot be embedded on a 2-sphere.*

Proof. Assume that an embedding of $K_{3,3}$ exists on the plane. We call faces the connected components of the plane when the embedding is removed. Let us denote by V, E, and F the number of vertices, edges, and faces of the embedding, respectively. Euler's formula [40] (a list of proofs is available in [38]), states that, for any embedding of a connected graph in the plane, one has $V - E + F = 2$. In our case, since $V = 6$ and $E = 9$, it follows from this formula that $F = 5$. If, for each face f, we count the number δ_f of edges that bound it, each edge is counted twice. Since any cycle in the graph will have at least 4 edges because the graph is bipartite, we deduce then $18 = 2E = \sum_f \delta_f \geq 4F = 20$, which is absurd. Hence, $K_{3,3}$ is non planar, and the three utilities problem has no solution in the plane. □

We have here established that there exist planar graphs and non planar ones. A natural question that follows is whether there exists a characterisation of planarity for graphs. Kuratowski proved that such a characterisation exists in 1930 [77], and Wagner proved in 1937 a similar one [139]: a graph is planar if and only if one cannot "find" K_5 or $K_{3,3}$ in it (see Figure 1.14). The "find" here is voluntarily unclear; in fact, there is a subtlety between its definition in the two theorems. The graphs whose absence characterises the planar property of graphs in Wagner's theorem are called **forbidden minors**. These theorems are highly interesting for us since a structural property, that is to say, information about substructures present in graphs, characterises one of their topological properties.

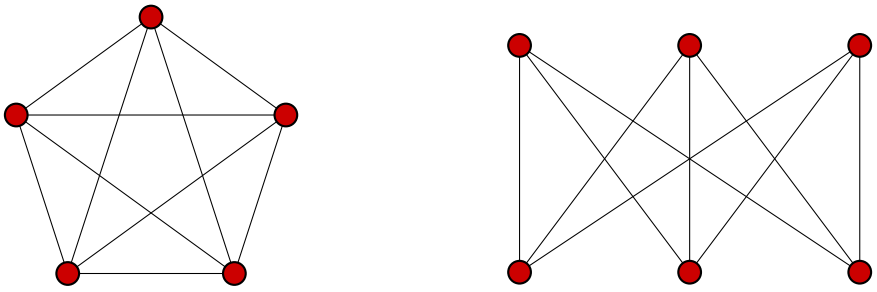

Figure 1.14: The two forbidden minors that characterise planar graphs: K_5 and $K_{3,3}$.

When a planar graph is embedded in $\mathbb{R}^2$, one can define the **dual graph** of this embedding in the following way: its vertices are the faces, and for each edge shared between two faces, there is a dual edge between the faces, crossing it (see Figure 1.15). Such a definition relies deeply on the embedding. For instance, the graph depicted in Figure 1.15 has a vertex incident to only one edge. This vertex involves a self-loop in the dual depending on which

face this vertex is embedded on (and to verify that the obtained graph is different, one can check the maximal number of edges incident to a face-vertex). The dual graph of a graph, as can be seen, is planar too.

Furthermore, the property of being planar has strong algorithmic implications. The idea is that such graphs can be cut into smaller pieces, of roughly equivalent size, in such a way that some solutions to problems can be deduced from the solutions on the smaller pieces. They have in fact treewidth roughly $\mathcal{O}(\sqrt{n})$, where n is the number of vertices, so that the recursive method presented earlier often provides efficient algorithms. It follows that some problems that can even be hard to solve in the general case admit an efficient solution when the input is planar. The existence of a small decomposition is in itself a strong result.

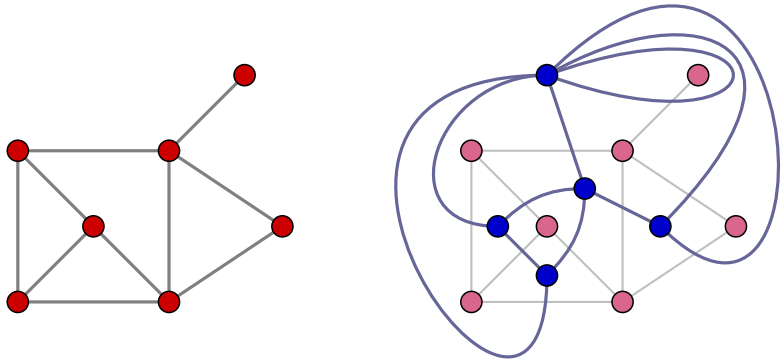

Figure 1.15: A planar embedding of a graph in $\mathbb{R}^2$ and a dual embedding in blue.

Let us remark that cycles of the dual graph are curves of $\mathbb{R}^2$ which intersect the initial graph on its edges. If one looks for a set of edges to remove to get the graph disconnected, it is then enough to find a cycle in the dual. This behaviour is a simple example of back-and-forth that can be done between a graph and its dual. In fact, the existence of a dual for planar graphs and the fact that this dual is embedded on $\mathbb{R}^2$ allow for quick computations of the aforementioned balanced decomposition in smaller pieces. Thus, planarity and duality have a strong impact on the computational characteristics of these graphs.

Linkless graphs Pushing further on topological properties that define classes of graphs, we will present here the class of linkless graphs. Two objects in the space are said to be **unlinked** if there exists a sphere such that each of these objects lies on a different side of the sphere; they are said to be **linked** otherwise. For example, two circles can be unlinked or linked (see left and right side of Figure 1.16 respectively).

Similarly to planar graphs, which are graphs which admit an embedding on a sphere, **linkless graphs** are graphs that can be embedded in space i.e., $\mathbb{R}^3$ or $\mathbb{S}^3$, in such a way that no two disjoint cycles are linked. For instance, all planar graphs are linkless graphs since any embedding on a sphere cannot have two linked disjoint cycles (a small perturbation of the sphere can separate the cycles). This is another class of graph that is defined by a topological property, and as for planar graphs, we will see that these graphs are characterised by structural properties.

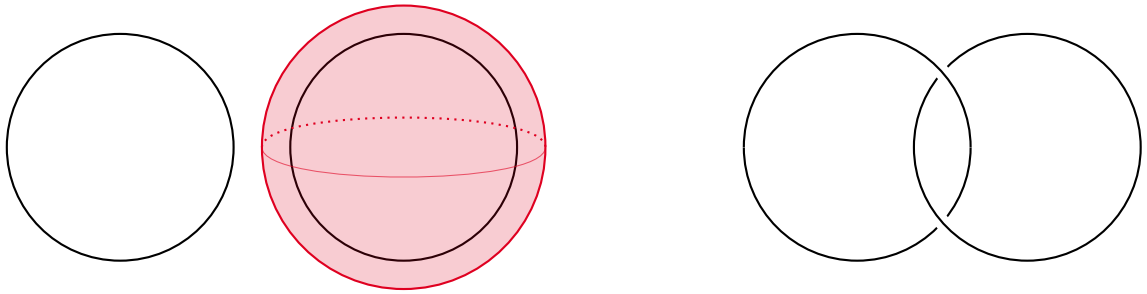

Figure 1.16: Left: two unlinked circles in the space, they are separated by a red sphere. Right: two linked circles.

In the following, we will loosely show that the complete graph on 6 vertices, K_6, as well as some graphs related to it, are all not linkless. In other words, every embedding of K_6 in space will have at least a pair of disjoint cycles that are linked. In Figure 1.17, we present an embedding of K_6, denoted $\mathcal{K}$, for which there is exactly one pair of disjoint cycles that are linked. The presentation used here is what we will loosely call a *projection*: a drawing from a fixed point of view in space of the embedding of the graph in space where ambiguities about which part is above the other are removed by additional information. We present in this figure the only pairs of disjoint cycles that could be linked since such a pair must have a crossing in the projection.

To proceed with our proof, we will take inspiration from [25, 125], and define, for each pair of disjoint cycles $\{C, C'\}$ in a fixed projection, $\delta(C, C')$: the number of crossings between C and C' where C is above.

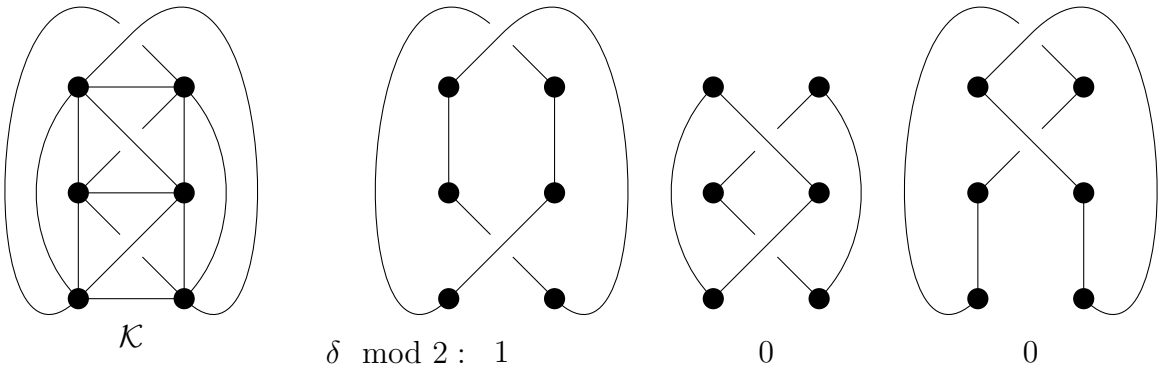

Figure 1.17: A drawing of an embedding, $\mathcal{K}$, of K_6 with only one pair of cycles linked.

Proposition 1.2. *The graph K_6 is not linkless.*

Sketch of proof: In the following we will use $\mathcal{K}$ to denote both the embedding described in Figure 1.17 and its projection. Similarly, we will now consider an embedding of K_6, and denote both this embedding and its projection by $\mathcal{K}'$. There are $\frac{1}{2}\binom{6}{3} = 10$ pairs of distinct cycles in K_6. For a projection K, we will use a quantity that we will denote $D(K) \equiv \sum_{C_1,C_2} \delta(C_1, C_2) \mod 2$ where $\{C_1, C_2\}$ are pairs of disjoint cycles of K. Only the pairs

of cycles $\{C_1, C_2\}$ of Figure 1.17 might satisfy $\delta(C_1, C_2) \mod 2$. By summing on them, we conclude that $D(\mathcal{K}) \equiv 1$.

Now, let us consider a continuous transformation, a *homotopy*, that transforms $\mathcal{K}'$ into $\mathcal{K}$. This transformation is allowed to make a pair of edges cross and edges cross themselves. Up to small perturbations, we also assume that these crossings do not happen simultaneously, that the projection is *regular* throughout the transformation (crossings on the projection can happen between 3 edges at most, and such crossings involving 3 edges are immediately resolved), and that edges do not cross vertices. If this last case happens, we push the crossing edge so that it crosses edges incident to the vertex instead.

Let us study how $D(\mathcal{K}')$ is modified whenever two non-incident edges e, e' cross during this transformation; let us call this event an edge-cross. These two edges take part in two pairs of cycles $\{C_1, C_2\}, \{C_1', C_2'\}$, depending on which one of the two remaining vertices is associated with which edge.

The edge-cross will modify $\delta(C_1, C_2)$ by ± 1 by definition. Every cycle not containing e or e' will be unaffected by the edge-cross. Hence, $D(\mathcal{K}')$ will not be modified by the modification (-2 or 2 is added to the sum mod 2). It is also not modified when the edge-cross happens between two incident edges. Since $D(\mathcal{K}') \equiv 1 \equiv D(\mathcal{K})$ at the end of the end transformation, and $D(\mathcal{K}')$ was not modified by it, we conclude that at least a pair of disjoint cycles of $\mathcal{K}'$ was linked at the beginning of this transformation. Hence, any embedding of K_6 in space has at least a pair of disjoint cycles that is linked: K_6 is not linkless. □

In [125], H. Sachs mentions transformations stemming from electrical engineering, which he calls "star-triangle-transformation", that preserve the property of a graph of being linkless. This operation, pictured in Figure 1.18, and more commonly called ΔY-transformation in graph theory, consists in replacing a 3-clique (a triangle) in the graph by a claw $K_{1,3}$ (a "Y") and conversely.

Figure 1.18: The ΔY-transformation.

The ΔY-transformation preserves the number of edges in the initial graph. Hence there is a limited number of graphs that can be obtained by applying ΔY-transformations to a graph. In the case of K_6, the graphs obtained that way are called the Petersen family (see Figure 1.19) since the famous Petersen graph is part of it (bottom graph of Figure 1.19). All these graphs are linkless, one can check that the relevant properties of $\mathcal{K}$ for the proof of Proposition 1.2 are also satisfied by all the drawings of Figure 1.19 so that the proof can be adapted. Indeed, there is a bijection between pairs of disjoint cycles before and after a ΔY-transformation. The main argument being that either the pair is disjoint from the subgraph affected by the transformation, or at most one cycle of the pair uses an edge of the considered subgraph; furthermore, if the starting subgraph is $K_{1,3}$, it necessarily uses 2 edges.

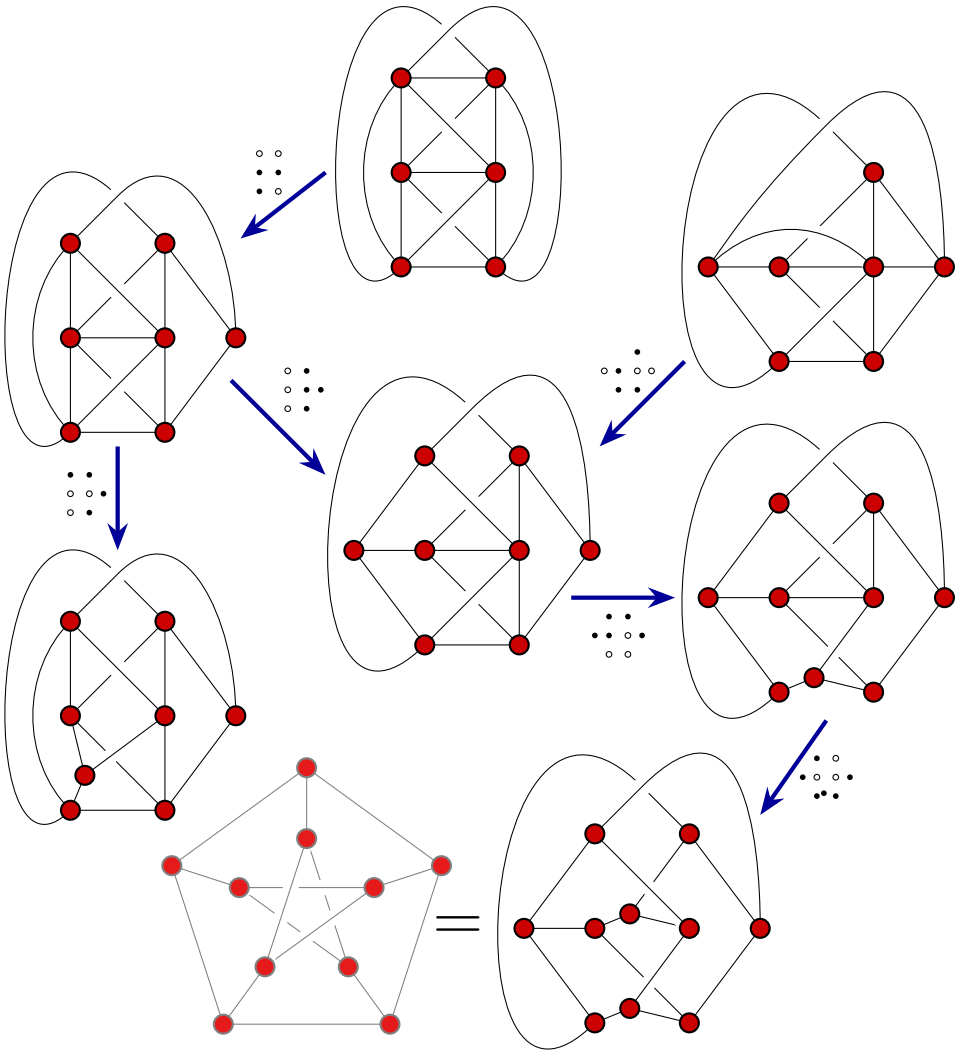

Figure 1.19: The Petersen family: graphs that can be obtained from K_6 by ΔY-transformations. The 3 white vertices near each blue arrow indicate which 3-cycle is transformed into a $K_{1,3}$ graph. The Petersen graph at the bottom is also presented in its "usual" form in lighter colours.

Remark 1.3. *As said above, Figure 1.19 presents embeddings of graphs of the Petersen family fitted for both adapting the proof of Proposition 1.2 and understanding which triangle is transformed in a claw. However, it is not clear from this picture that choosing a different triangle for each ΔY-transformation would yield the same graph. This is true, except for the top left graph, where the 3-cycle induced by the 3 vertices of degree 5 yields a graph non-homeomorphic to the one induced by the other triangles. In fact, graphs of this family present a high number of intrinsic symmetries, which explain the symmetric role played by almost all the triangles. A presentation of this family highlighting these symmetries can be seen here [2].*

Robertson, Seymour, and Thomas announced [113] and proved [114] that in fact, this

family is exactly the set of forbidden minors that characterises linkless graphs. This is another example of a class of graphs where the characterising set of forbidden minors is known.

From a computational point of view, classes of graphs characterised by forbidden minors are interesting. We have seen that planar graphs admit useful decompositions that allow some problems to be solved efficiently on them. In fact, this property is verified on every class of graphs characterised by forbidden minors [110]. Nonetheless, compared to planar graphs, fewer problems are known to admit a significant speed-up for their running time on graphs characterised by forbidden minors. On linkless graphs, for example, the topological properties associated with the embeddings are harder to exploit (a notable exception being [132]).

Knot theory The last class of graphs we considered is defined by the topological properties of embeddings of graphs in $\mathbb{S}^3$ or $\mathbb{R}^3$. We now dive deeper into this point of view by studying **knots**, which are embeddings of $\mathbb{S}^1$ into $\mathbb{R}^3$ (or equivalently $\mathbb{S}^3$) considered up to some continuous deformations we call *ambient isotopy* (see Figure 1.20 for an example of such embeddings). An ambient isotopy is, loosely speaking, a non-degenerate continuous transformation which never makes the embedding cross itself.

Intuitively, if we take a rope, tie a physical knot, and consider the continuous transformations that can be applied to it, we could untie it by reversing what has been done. Hence, we glue the extremities of the rope so that distinct knots appear: it is no more possible to untie every such physical knot. The idea is that a knot remains the same while we play with it: we can expand, bend, and distort the rope as long as we do not cut it to glue it back later (this last bit is a discontinuous operation that modifies the knot).

Knots are all homeomorphic to circles $\mathbb{S}^1$, and now is a proper time to emphasise that all intrinsic topological properties of two homeomorphic objects are the same. However, knots are embeddings on which the continuous deformations we consider are distinct from homeomorphisms: these deformations preserve topological properties of the embeddings, which are not intrinsic properties of a circle.

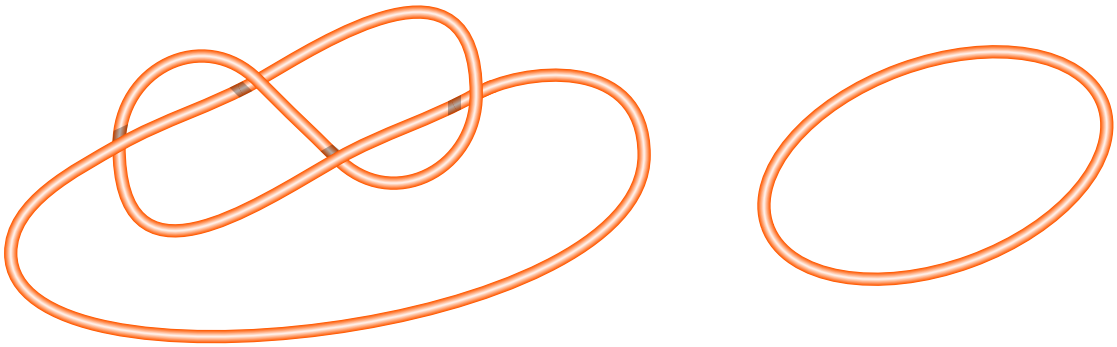

Figure 1.20: Two knots, i.e., two embeddings of $\mathbb{S}^1$ in $\mathbb{R}^3$, the knot on the left is called the figure-eight knot.

Similarly to graphs, which can model an incredibly high range of both physical and abstract objects, knots are the natural model to consider when it comes to understanding

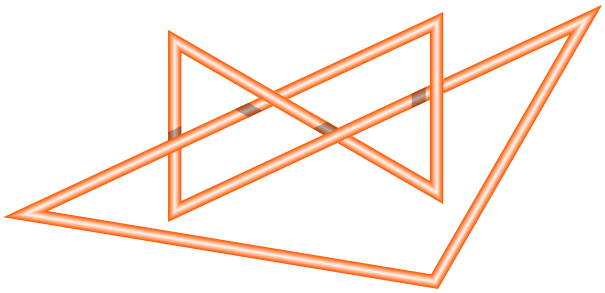

Figure 1.21: A polygonal embedding of the figure-eight knot.

and studying any object of the real world that "looks like" a 1-manifold. Thus, knot theory finds application in many fields of science. For example, DNA can be studied through that lens since it consists of a double helix of nucleotides that get knotted by the action of enzymes [111, 63] whose folding can also be studied by knot theory [64]. In physics, field lines can be seen as knots, and quantum field theory turns out to benefit from understanding knots [68]. The list could go on with 3D printing, chemistry, or material design. Going back to our more theoretical focus, understanding knots is a first step toward understanding embeddings of objects in higher dimensions, which is a basis of many other problems.

However, the study of knots is not easy. From their very definition comes the fundamental question of knot theory: are two knots the same? In other words, given two embeddings of $\mathbb{S}^1$, does an ambient isotopy from one to the other exist? This question would immediately appear highly difficult to someone tasked to decide whether two given balls of wool represent the same knot (assuming the ends are glued together). Indeed, one would need to manipulate the threads long enough to have both physical knots match. And that would only work if they represent indeed the same knot. Indeed, if they do not, how can one be sure that a longer manipulation would not eventually make both knots match? Finding an ambient isotopy between two knots is mathematically challenging. In fact, it is also hard for computers. Furthermore, that last assertion raises a concern: how does one input a knot to a Turing machine or a computer?

Similarly to surfaces that can be encoded by a finite number of triangles, knots can be equivalent to a finite representation: a finite closed broken line in space. For example, the figure-eight knot of Figure 1.20 can be transformed into the broken line of Figure 1.21. Such an embedding is called a **polygonal embedding**. It is a common restriction in knot theory to study only knots admitting a polygonal embedding: we will do so in the following.

Polygonal embeddings are a convenient representation of knots, and in fact, knots admitting such embeddings also admit **diagrams**. Intuitively, diagrams are unambiguous drawings of knots in the plane. The drawing is obtained by a *regular* projection (such a projection can be seen in Figure 1.22), which is a projection where points with multiple preimages are in a finite number and come from crossing strands. The unambiguity comes from added information of which strand is above the other one, usually pictured by a blank space of the lower strand around each crossing (as can be seen at the bottom of Figure 1.22, and was already used in Figure 1.17 to picture embeddings in space). A diagram can be seen as a planar embedded graph such that each vertex has 4 incident edges (vertices come from

crossing strands in the projection); such a graph is called 4-**regular** (for instance, the graph G_3 of Figure 1.8 is 3-regular).

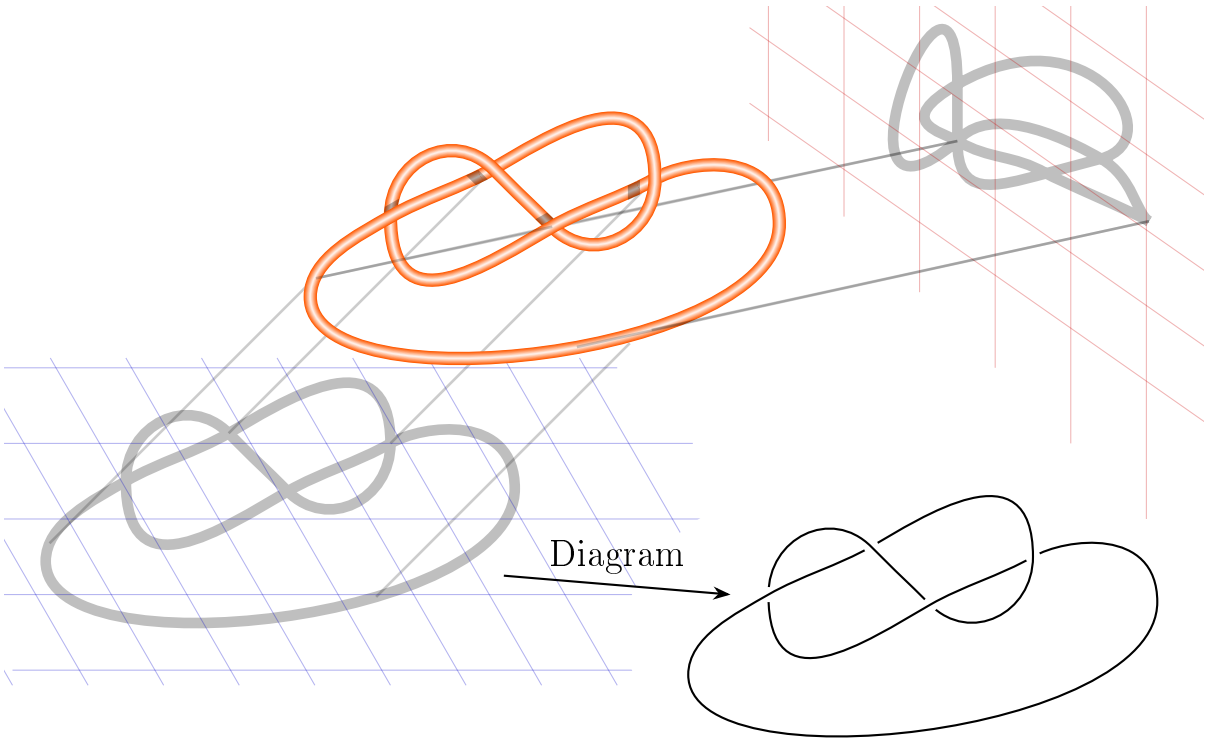

Figure 1.22: A regular projection of the figure-eight knot on the blue plane and a non-regular one on the red plane: two points of the projection have more than 2 preimages. Bottom right: the diagram arising from the regular projection.

Reidemeister proved in 1927 [138] the eponymous theorem: two knots, represented by two diagrams, are ambient isotopic if and only if a sequence of **Reidemeister moves** can turn one of the diagrams (seen as an embedded planar graph) into the other. The Reidemeister moves are local modifications of the strands that are pictured in Figure 1.23. Intuitively, RI twists or untwists a strand, RII creates or removes an overlap between two strands, and RIII represents the fact that a strand can move back and forth over another crossing.

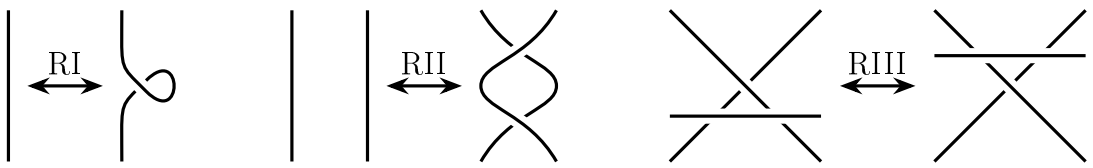

Figure 1.23: The three Reidemeister moves.

This theorem is really powerful since it connects the 3D nature of knots with their possible diagrams that are planar graphs. It also opens directions on how to solve the initial question of identifying whether two knots are the same: one can provide a sequence of Reidemeister moves to apply that certifies the equivalence between the two knots. In that description,

all that remains is finding such a sequence. It is, in fact, a troublesome question. There is no simple way to do so. Intuitively one is tied to trying out every possible move and hoping for the diagrams to eventually match. However, this search makes sense only if we limit the possible moves, otherwise, it is always possible to have more and more crossings via the repetition of RI or RII moves, eventually yielding a very complex diagram. The natural approach is then to try to only decrease the complexity of the diagram, corresponding intuitively to its number of vertices. However, it is known that some diagrams will require first increasing their complexity before they can be reduced [21]. This behaviour illustrates how hard this problem can be. Currently, the best known algorithms to detect whether two knots are the same do not use Reidemeister moves and instead exploit the topology of the "outside" of the knot. However, the running time of these algorithms is incredibly long. To echo our previous parts, let us emphasise that these algorithms use surfaces to carry this study.

Hence knot theorists tend to search for ways of distinguishing knots: ways to certify that two knots are different. And that is done by studying knot **invariants**. Invariants are quantities, or objects, that remain the same for every possible presentation of a knot. Thus if two embeddings do not share an invariant, we know they are different. The converse is not true in general; two knots sharing an invariant can be different. The invariance of some quantities can be clearly derived from their definition. The **crossing number** is one of them: it is the minimum number of crossings a knot diagram can have. Since it does not depend on a fixed embedding but on the equivalence class as a whole, it is clear that this quantity is an invariant. Some other invariants are defined by a calculation or property of a fixed embedding or diagram, therefore, proving invariance requires proving that the calculation yields the same result for every embedding or diagram.

Let us develop one simple example: tricolorability. A strand of a diagram is a continuous arc of a diagram when drawn using the convention that the lower part of the knot is replaced by a blank at each crossing as in Figures 1.23 and 1.22 (in this one, the figure-eight knot has 4 strands). A knot is said to be **tricolorable** if one of its diagrams is **3-colorable**, i.e., if it is possible to colour, using three colours, the strands of a diagram such that each crossing has either all 3 colours incident to it or only 1 (the three colours must also be used at least once). The Reidemeister theorem shows that every pair of diagrams of the same knot are connected by a sequence of Reidemeister moves. Thus, it is enough to show that if one diagram is 3-colorable, then every diagram of the sequence, including the last one, is 3-colorable too. By contradiction, it is clear that if a knot diagram is not 3-colorable, none of its diagrams will be 3-colorable.

Proposition 1.4. *Tricolorability is a knot invariant.*

Proof. Let us proceed with the proof outlined above and consider an initial 3-colouring of a knot diagram and a Reidemeister move. We will only modify the colour of the strands created by the Reidemeister move so that the properties of a 3-colouring will still be satisfied at uninvolved crossings. In particular, the colours of the strands stretching outside of each Reidemeister move pattern will be the same. In fact, the proof can be summed up by Figure 1.23 when only one colour is involved and Figure 1.24 otherwise.

•If the move is RI, a strand is separated in two, or two are merged. In both cases, they need to have the same colour (see the leftmost part of Figure 1.23 where the colour is black, by symmetry).

•If the move is RII, let us assume that more than one colour is involved in the move (otherwise there is nothing to say). Then the picture is, up to switching colours, the top left of Figure 1.24. If some crossings are involved, they satisfy the property of a 3-coloring.

•If the move is RIII, let us assume again that more than one colour is involved in the move (otherwise there is nothing to say). We will refer to the relevant part of Figure 1.24. Set first the colour of the over strand (say green in our case); it will remain the same during the move. The middle part of the knot is made of two strands that have either one or two colours (top right and bottom case, respectively). The former case is straightforward, so let us focus on the latter. Either the crossing not involving the upper strand has only one colour, in which case it is bottom left, or it has 3 and the case is bottom right. In all these cases a 3-colouring is preserved.

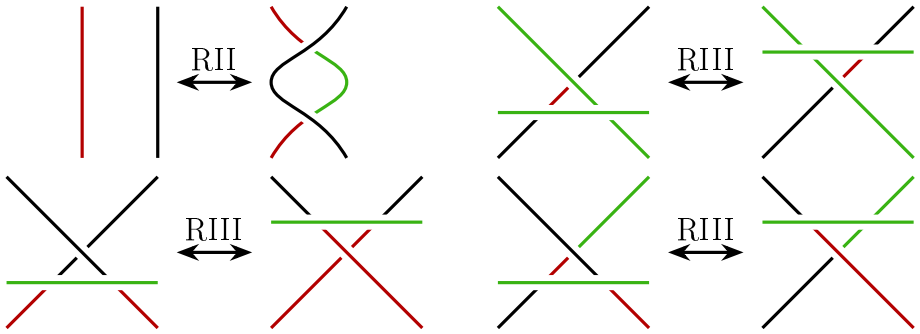

Figure 1.24: Maintaining a 3-colouring through Reidemeister moves.

□

As we have seen, the invariance of tricolorability admits a rather elementary and accessible proof. Let us use it in Figure 1.25 where the three simplest knots are pictured: from left to right, the **unknot** (the standard embedding of the circle), the trefoil knot (only knot with crossing number 3), and the figure-eight knot (only knot with crossing number 4). Among them, the trefoil knot can be distinguished from both the unknot and the figure-eight knot using this invariant, since the diagram of the unknot has only one colour: it is not 3-colorable. Similarly, pairs of strands of the figure-eight knot in the provided diagram have two crossings in common. Since there are 4 strands, at least two of them will have the same colour, and a contradiction is obtained at one of their common crossings since the remaining strand needs to be of a different colour. As said above, Proposition 1.4 and Figure 1.25 prove that the trefoil knot is different from the unknot and the figure-eight knot, but this invariant is not enough to distinguish these two last knots.

Going back to a computational point of view, knot diagrams are planar graphs. Hence, it is a natural idea to try to apply the results and methods of graph theory to understand knots via their diagrams. For instance, some knot invariants which are hard to compute in the general case can be computed efficiently when one of their diagrams looks like tree.

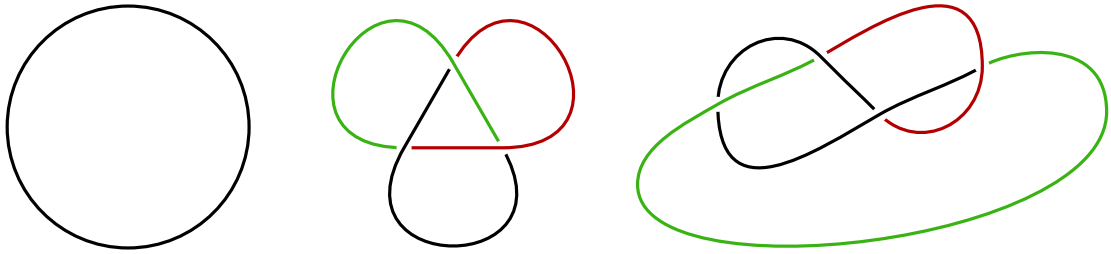

Figure 1.25: Three knot diagrams. Among them, only the trefoil knot in the middle, is tricolorable.

We presented earlier a method to handle algorithmic problems on graphs using tree decompositions. Let us sketch an algorithm to decide the tricolorability of a knot of low treewidth. Naively, if the diagram has n crossings, it will possess n strands, which can therefore be coloured in 3^n ways. Checking the validity of each colouring leads to an exponential algorithm. Then, let us assume that the diagram has treewidth k to do better. A property of planar graphs is that it is possible to find a tree decomposition where bags are bounded by circles in the plane (such a decomposition is shown in Figure 1.10). Then one can enumerate at most 3^k colourings for the strands of each bag and check the compatibility with the colouring of the neighbouring bags of the tree along the circles of the decomposition. This broadly sketched algorithm is polynomial when k is fixed, as opposed to the exponential naive one.[2]

Theme of this thesis. Earlier, we showed how topological properties of graphs define interesting classes of graphs. We also described how the topological properties of planarity and linklessness on graphs are characterised by substructures in graphs. These properties and their consequences are part of a subfield of graph theory called **structural graph theory**. We then emphasised how graphs can enlighten our understanding of knot theory. This thesis is about this last part: drawing inspiration from results and methods of structural graph theory in order to develop our understanding of knot theory.

Let us broadly present some of our main results. As we have illustrated just above with tricolorability, when a knot admits a diagram of low treewidth, this tree-like structure can be leveraged to provide efficient algorithms to compute invariants. However, every knot admits diagrams with high treewidth since it is always possible to manipulate some part of the knot to have it appearing like a grid as in Figure 1.26.

Therefore, a natural question, asked by Burton [22], Makowski, and Mariño [86], is whether there exist knots that do not admit diagrams with low treewidth. In other words, they asked whether there exist knots for which algorithms exploiting low treewidth cannot be used. This question was answered positively in 2018 [28] using complex results [57] on the shape of a decomposition of the space by surfaces. We develop a more elementary theory reproving this result that is inspired by techniques associated with treewidth, hence coming from structural graph theory. More precisely, we define a measure of how close to a

[2]In fact, deciding tricolorability can always be done in polynomial time by computing the value of the Alexander polynomial at -1 [100]. However, tricolorability can be generalised to colouring invariants whose computation is provably impossible in polynomial time (modulo standard conjectures) [76].

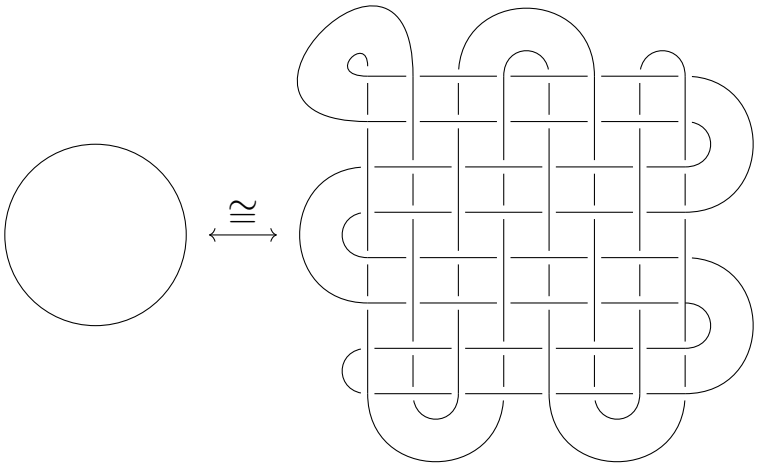

Figure 1.26: A diagram of the unknot with high treewidth.

tree a knot is (this is also a knot invariant) and an *obstruction* to it. The invariant, called *spherewidth*, is what links the knot to the treewidth of its diagram: if the spherewidth is high, the treewidth of any diagram is high. The spherewidth informally quantifies the best way to sweep the space with spheres nested in a tree-like fashion while minimising their number of intersections with the knot. The obstruction, on the other hand, provides a lower bound for the spherewidth.

In addition, we prove that our obstruction exists whenever knots can be embedded on a surface. To do so we exploit the interactions between the spheres of our sweepouts and this surface, and especially the topology arising from these intersections. This allows us to prove that all diagrams of a certain family of knots have high treewidth. Such a family is called the family of torus knots (a knot can be seen on Figure 1.27). Furthermore, our results apply as well to **links** and **spatial graphs**. Both are generalisations of knots. Links are a disjoint union of knots, each of those knots is called a **link component**. For example, the right part of Figure 1.16 pictures a link called the Hopf link, which contains two link components, each one being an unknot. Spatial graphs are embeddings of graphs in $\mathbb{S}^3$. These embeddings can be seen as a generalisation of knots and links: they locally look like knots except at vertices where they can branch off. It follows that their study is at least as complex as that of knots and links.

Our obstruction can also be used to address another problem that might in appearance seem unrelated. We expressed before how one might need to first increase the number of crossings before expecting progress on the process of untangling an unknot diagram by Reidemeister moves. In fact, that behaviour can appear on any problem dealt with by the iterated application of Reidemeister moves. However, no result is known on how many crossings need to be added while solving a problem that way. Another one of our contributions is to leverage our obstruction to provide examples of diagrams requiring an arbitrarily large number of added crossings to be simplified.

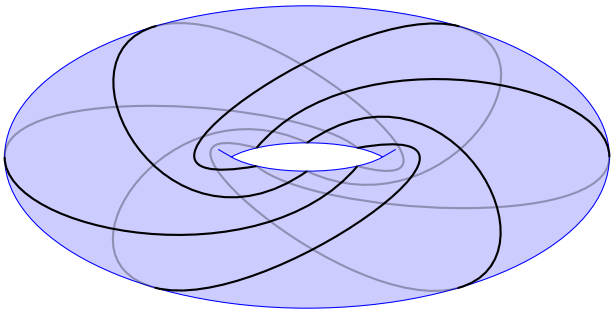

Figure 1.27: An example of a torus knot.

1.2 Contributions of this thesis

While we intended the previous section to be approachable with little knowledge in topology and graph theory, the same is not true for the following. We refer to the textbooks of Schultens on 3-manifolds [130] (it also covers basics about surfaces), Diestel on graph theory [30], Rolfsen [120] for knot theory, and Cormen, Leiserson, Rivest, and Stein for basics on algorithms [26].

We now move on to a more precise description of the contributions of this thesis without diving too deep into their specifics. Every object defined here will be redefined properly in the relevant chapter. Our work makes contributions to knot theory using various notions of tree-like decompositions relying on surfaces and how these surfaces interact with knots. Furthermore, inspiration and methods used in our results stem from structural graph theory.

Decidability of the genus defect on Hopf arborescent links. As explained in the previous section, the problem of deciding whether two knots are equivalent is hard, both from a theoretical and computational point of view. Hence, knot theory resorted to the use of invariants. Among them, a classical one is the knot *genus*. It is the minimal possible genus among its *Seifert surfaces* i.e., oriented surfaces embedded in $\mathbb{S}^3$ having the knot as their boundary. For example, the unknot is the only knot with genus 0, that is to say that the unknot is the boundary of an embedded disc (this property is a common definition for the unknot). Several algorithms to compute the genus of a knot are known, and the complexity of its computation is fairly understood: we know that the problem is both in **NP** and in **co-NP** [4, 82].

Seeing $\mathbb{S}^3$ as the boundary of $\mathbb{B}^4$, we consider a generalisation of knot genus that we will call 4-*genus*: loosely speaking, it is the minimal genus of an oriented surface embedded in $\mathbb{B}^4$ that has the knot, embedded in $\partial\mathbb{B}^4$, as its boundary. In the context of 4-dimensional topology, smooth and topologically flat embeddings are different. This gives naturally rise to two different notions of 4-genus: we refer to Chapter 3 for precise definitions. Since our methods and results apply equally well to both, we will continue to speak of 4-genus in the remaining of this introduction.

From a computational point of view, topology in dimension 4 is hard and not well understood. On the one hand, many fundamental topological problems, like deciding homeomorphisms between manifolds, are known to be undecidable in dimension 4 [89]. On the other

hand, the decidability of many fundamental problems, such as recognising the 4-dimensional sphere, is wide open, and no general framework is known for tackling these questions. The 4-genus is part of these problems whose decidability is not known in the general case. This invariant is essential to the study of *slice* knots i.e., knots of 4-genus 0, and to major knot theory conjectures like the slice-ribbon conjecture [41].

We investigate a class of knots and *links* (which are unions of knots) that we call *Hopf arborescent links*. They are knots and links that can be decomposed into a tree-like *plumbing* of Hopf bands, where a Hopf band is an embedding of an annulus that has a Hopf link as a boundary (see left of Figure 1.28, and right for such a decomposition). The *defect* is the difference between the genus and the 4-genus. We prove Theorem A which states that, for all k, there exists an algorithm to decide if a Hopf arborescent link has defect at most k. Hence, we settle the decidability of computing the defects on this class of knots and links.

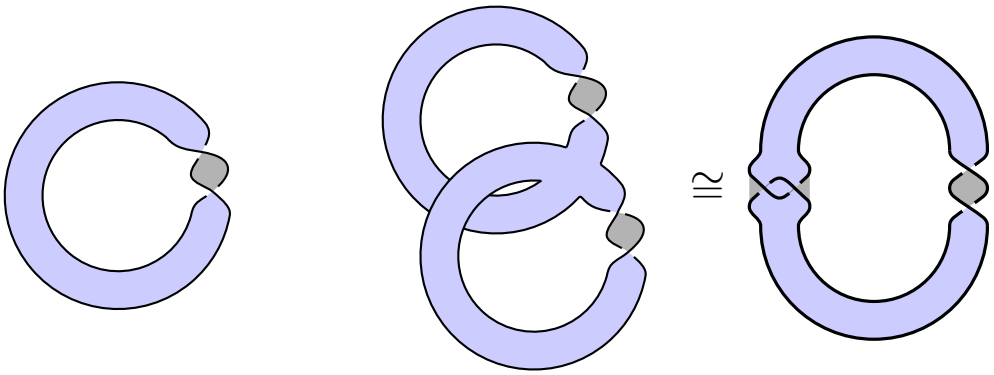

Figure 1.28: Left: a Hopf band. Right: a plumbing of two Hopf bands yielding a figure-eight knot.

To prove this theorem, we draw inspiration from structural graph theory by defining a minor relation, called link-minor, on this class of links and proving that it is a **well-quasi-order**. If a set is ordered by a well-quasi-order, every infinite sequence of this set will contain at least two comparable elements. This implies that no infinite **antichain** exists in the set, no infinite sequence of pairwise non comparable elements. Hence any set defined by a property stable for the well-quasi-order can be characterised by the **finite** antichain of minimum elements of its complement. This yields an algorithm for verifying the property on an element x when the relation can be decided: it is enough to decide that no minimum element is in relation with x.

Our proof technique follows this strategy and works by showing that the defect is stable for the link-minor relation and providing an algorithm to decide this relation. The link-minor relation relies on a precise definition of Hopf plumbings seen as a construction operation carried out on trees. From this definition, we can associate a tree and a specific Seifert surface to each Hopf arborescent link. First, we leverage the Kruskal tree Theorem [73] on these trees to prove that the link-minor is a well-quasi-order. Then, we provide an algorithm to decide the link-minor relation. Next, we prove the stability of the defect by link-minor by studying a second relation, called surface-minor, which is weaker than our link-minor relation but behaves well with respect to the defect.

The surface-minor is a containment relation between the specific Seifert surfaces associated with our links. Furthermore, since the link-minor is a well-quasi-order that is finer than the surface-minor, our proof techniques also yield Theorem B which states that the surface minor relation is a well-quasi-order on the Seifert surfaces associated to our links.

- We introduce a class of knots and links called Hopf arborescent links, which are obtained by iterated plumbings of Hopf bands.
- We prove the decidability of the defect on the class of Hopf arborescent links. See Theorem A.
- We prove that the surface-minor relation is a well-quasi-order on the class of Hopf arborescent links. See Theorem B.

Width invariant of links and spatial graphs from a structural graph theory approach. We develop further the context and concepts discussed in the last part of the previous section. The search for knot invariants that can be computed efficiently is a way to circumvent the difficulty of knot recognition. Another successful method to tackle hard problems is to develop algorithms, called *Fixed-Parameter-Tractable* (FPT [27]), whose design depends on additional information of the input: the parameter. Their main interest is that the complexity of these algorithms is small when the parameter is fixed. The treewidth is a primary parameter in that regard: algorithmic designs on graphs benefit greatly from the underlying tree-like structure of graphs of small treewidth (see for reference this survey of Bodlaender [14]). Applying this method to knot diagrams with low treewidth leads to the efficient computation of many knot invariants (see e.g. [86, 19, 87]), that are otherwise known to be hard to compute. As explained earlier, this situation led to the question [22, 86] of whether there exists a family of knots for which all diagrams have high treewidth, which was later answered positively [28]. Our work focuses on reproving and generalising this answer using inspiration from structural graph theory.

The treewidth is also a concept at the heart of the proof of the graph minor theorem from Robertson and Seymour [118]. This parameter led to the birth of many other related width invariants which can present both theoretically and practically relevant features for solving problems. One of them is the *branchwidth*, which is equivalent to treewidth up to a constant factor. Treewidth and branchwidth are both defined as min max of a measure taken over a set of decompositions. By essence, they are troublesome to lower bound since, in order to do so, one must prove that every decomposition has high width. Branchwidth is particularly interesting for us since it can be interpreted geometrically as tree-like sweepouts by circles of a sphere on which a planar graph is embedded (recall that the knot diagrams are planar graphs) [124]. Furthermore, it admits a tight *obstruction*, called a *tangle*, which presents topological aspects and whose existence grants a lower bound on branchwidth. To be more precise, a tangle has an *order* which represents its size, and a graph has a tangle of order k if and only if the branchwidth is larger than k. Hence providing an obstruction of size k also ensures that k is a lower bound of the branchwidth. These two facts led us to design our width invariant of knots, called **spherewidth**, inspired by branchwidth. It relies on **sphere**

decomposition of $\mathbb{S}^3$ which can be seen as a generalisation in $\mathbb{S}^3$ that uses spheres of the aforementioned sweepouts. Formally, a sphere decomposition is a continuous map $\mathbb{S}^3 \to T$ where T is a trivalent tree such that the preimage of each leaf is a point, the preimages of each point in each edge is a 2-sphere, and the preimages of each inner vertex form a double bubble. A **double bubble** is made of two spheres which intersect on a disc and represents intuitively the moment at which two spheres merge (see Figure 1.29 for a double bubble and a 2D representation of a sphere decomposition). The width of the decomposition is the maximal number of intersections between a sphere and the knot (one can complete the sphere decomposition pictured in Figure 1.29 so that it has width 4). Then, the spherewidth is the infimum of the width among all sphere decompositions.

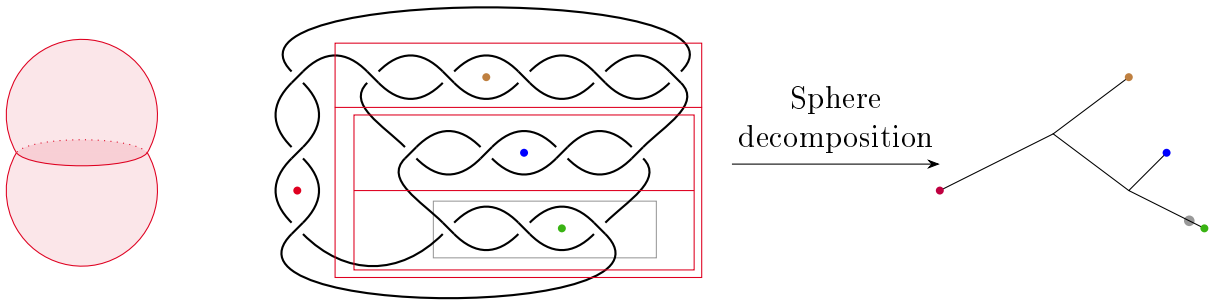

Figure 1.29: Left: a double bubble. Right: an example of a sphere decomposition where preimages of inner vertices are presented in red, and preimages of a point inside an edge are presented in grey.

Since sweepouts of knot diagrams by circles can be "lifted" to sphere decompositions, it follows that the spherewidth is a lower bound to the branchwidth of diagrams, which is itself a lower bound for the treewidth of diagrams. Similarly to graph theory, we design an obstruction, called bubble tangle, mimicking the one defined on graphs and providing lower bounds on spherewidth. We prove via Theorem C that this obstruction is also tight: for all k, either there exists a sphere decomposition of width k or a bubble tangle of order k, where the order also represents the size of our obstruction.

Furthermore, as a major feature of our approach, we provide tools to obtain such an obstruction. Theorem D states that a bubble tangle of order $\Omega(r)$ exists whenever the knot is embedded on a surface with *compression-representativity* r. The compression-representativity quantifies how well the knot represents the surface on which it is embedded. For example, the torus knot $T_{p,q}$, when embedded on the standard torus associated to its definition, has compression-representativity $\min(p, q)$. Intuitively, our obstruction designates a small side for every sphere with a small number of intersections with the knot (this is again inspired by the obstruction on graphs). When the knot is embedded on a surface with non zero genus, spheres with a small number of intersections with the knot will cut discs from the torus on one side, while the other will contain the topology of the torus. The small side is then the one that contains only the discs (see, for example, Figure 1.30).

Hence, our work reproves, using new techniques inspired by structural graph theory, the fact that any diagram of a torus knot $T_{p,q}$ has treewidth $\Omega(\min(p, q))$. On a more general

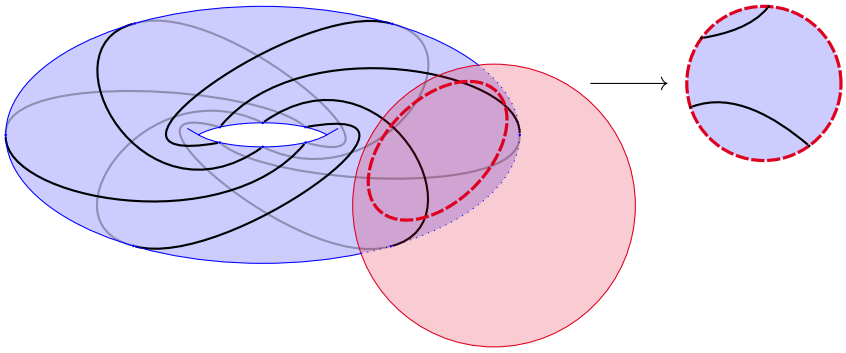

Figure 1.30: Intersection of small size between a torus knot $T_{5,6}$ embedded on a torus and a sphere. It follows that the topology of the torus lies on one side, while the other contains only discs of the torus. The latter is the small side.

level, we devise a systematic way to handle such questions by the study of surfaces on which we can embed knots, links, and, more generally, spatial graphs.

- We define **spherewidth**: a width measure of knots, and more generally spatial graphs inspired by structural graph theory. We show that it is a lower bound to the treewidth of any diagram.
- We define **bubble tangles**: an obstruction to spherewidth inspired by structural graph theory.
- We prove Theorem C stating that our bubble tangles are a tight obstruction to the spherewidth.
- We prove Theorem D stating that bubble tangles with size $\Omega(k)$ exist whenever there exists a surface on which a knot, link, or spatial graph is embedded with compression-representativity k.
- We reprove, using elementary techniques stemming from structural graph theory, the result of [28] stating the existence of knots for which all diagrams have high treewidth.

A superconstant lower bound on the complexity of splitting link diagrams. As we explained in the previous section, Reidemeister moves are a powerful method to study knots and links via their diagrams, thanks to the Reidemeister theorem. They appear as a very convenient and natural tool to study basic considerations on knots, like their crossing numbers or unknotting numbers. A primary example of a knot theory problem is recognising the trivial knot, which is the first instance of the major problem of knot theory: deciding whether two knots are equivalent or not. To do so, a natural strategy is to try to untangle a given knot by applying Reidemeister moves on its diagram in a brute force or random manner

until the diagram corresponds to a simple curve. However, some unknot diagrams [21], called **hard unknots**, exhibit an annoying behaviour for this algorithm: the maximum number of crossings of a diagram during the algorithm is larger than the initial one. We first *need* to add crossings before being able to reach the untangled diagram.

Formally, let us denote by $\mathbf{cr}(D)$ the number of crossings in the diagram D. Then, for two equivalent diagrams D_1, D_2 and a sequence of Reidemeister moves R transforming D_1 into D_2, we define $\mathbf{Top}(D_1, R)$ which is the maximum of $\mathrm{cr}(D^i) - \mathrm{cr}(D_1)$ throughout the sequence of Reidemeister moves where D^i is the diagram D_1 after performing the first i moves of the sequence. The quantity we are interested in is $\mathbf{Add}(D_1, D_2)$ which is the minimum of $\mathrm{Top}(D_1, R)$ taken among all the sequences of Reidemeister moves that transform D_1 into D_2. When we see D_2 as a goal diagram, $\mathrm{Add}(D_1, D_2)$ is a lower bound on the number of crossings to add during the running of the aforementioned algorithm that applies Reidemeister moves on D_1 to reach D_2.

Studying $\mathrm{Add}(D_1, D_2)$ turns out to be trickier than one might initially think. Apart from an exhaustive search of the possible Reidemeister moves, which quickly becomes intractable, no method is known to lower bound this quantity. In the context where D_1 is an unknot diagram and D_2 a diagram of a simple curve, D_1 is a hard unknot if $\mathrm{Add}(D_1, D_2)$ is positive. In fact, only diagrams on which $\mathrm{Add}(D_1, D_2) \leq 2$ are known [21] although it is conjectured that there exist unknot diagrams D for which $\mathrm{Add}(D, D_2)$ is arbitrarily large.

We will focus on this quantity in the splitting problem: deciding if a link L is **split** i.e., is there a sphere disjoint from L separating at least 2 link components of L? If such a sphere exists, there exists a link diagram in which two unlinked sublinks are disjoint: they are separated by a circle in the plane. Therefore, in terms of Reidemeister moves, we will study $\mathrm{Add}(D_1, D_2)$ where D_2 is a link diagram of a link L where a circle separating the link diagram can be drawn in the plane, and D_1 is any diagram of L. We will call a link diagram D_1 of a link L for which $\mathrm{Add}(D_1, D_2) > 0$ a **hard split link**. Finding a sphere in space separating two links is easier than the problem of finding a disc that has a knot as its boundary. Therefore, this problem, which is interesting in itself, has been studied several times as a useful and easier problem for understanding the unknot recognition problem [36, 78].

We exhibit a family of link diagrams $\mathcal{D}(p, q)$ with two unlinked sublinks: the first sublink is made of two tangled torus knots, and the second one is an unknot surrounding one of the torus knots (see $\mathcal{D}(7, 13)$ in Figure 1.31 for an example). For $\mathcal{D}'(p, q)$, any link diagram on which U is disjoint from the other link components, we prove Theorem E, implying that $\mathrm{Add}(\mathcal{D}(p, q), \mathcal{D}'(p, q)) = \Omega(\min(p, q))$. If we call **crossing-complexity** of D_1 the minimum of $\mathrm{Add}(D_1, D_2) > 0$ among all split link diagrams D_2 of L, we provide hard split links of arbitrarily large crossing-complexity. More precisely, Theorem E states that for each n, there exists a diagram D_n of a split link L_n of $\mathbb{S}^3$ with 3 components such that any sequence of Reidemeister moves converting it to a split diagram of L_n passes through a diagram with at least $2n^2 + \frac{2}{3}n$ crossings.

The method used here is to exploit the unknot U present in every one of our diagrams and which is separated in $\mathbb{S}^3$ from the remaining link components (see the blue link component of Figure 1.31). Our approach is to show that if there exists a sequence of Reidemeister moves where $\mathrm{Add}(\mathcal{D}(p, q), \mathcal{D}'(p, q))$ stays small, we can use the evolution of U throughout these moves to define a sweepout of the two linked torus knots with spheres, where each sphere has

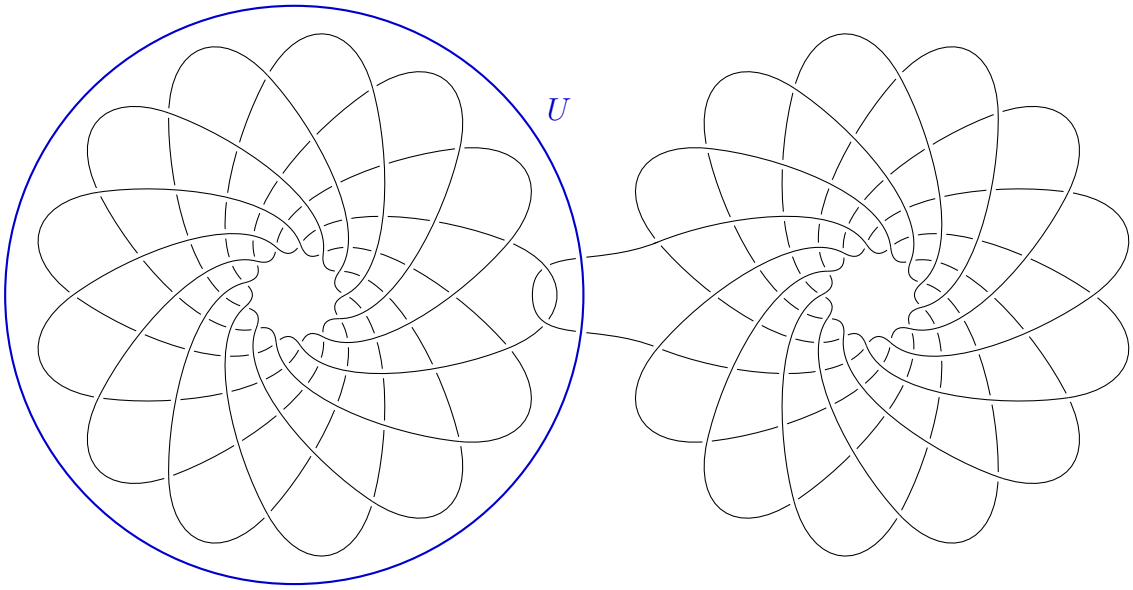

Figure 1.31: The link diagram $\mathcal{D}(7, 13)$: two linked $T_{7,13}$ and an unknot U.

a small number of intersections with that link. But our obstruction from Chapter 4 has been exactly designed to show that such a sweepout is impossible. Hence $\mathrm{Add}(\mathcal{D}(p, q), \mathcal{D}'(p, q))$ must be large enough.

- We define several families of link diagrams on which the number of crossings to add to be able to separate the links is arbitrarily large, as stated by Theorem E.
- We outline a method to provide lower bounds on the minimal number of crossings to add during the execution of algorithms relying on trying Reidemeister moves.

1.3 Organisation

In Chapter 2, we go through global preliminaries for our work. In particular, we define properly most of the concepts that were only loosely defined in this introduction.

Chapter 3 focuses on Hopf arborescent links and the decidability of the defect. We present there Theorems A and B from our article [B], written with Pierre Dehornoy and Arnaud de Mesmay, which appeared in the Proceedings of the 40th International Symposium on Computational Geometry and has been invited to a Discrete & Computational Geometry special issue of the Symposium on Computational Geometry 2024.

Chapter 4 concentrates on tree-like sweepouts of $\mathbb{S}^3$ and their obstruction. Our duality theorem, Theorem C, and existence theorem, Theorem D, are introduced and proved there. This chapter is mostly coming from our article [A], written with Arnaud de Mesmay, which appeared in the Proceedings of the 39th International Symposium on Computational Geometry.

Chapter 5 is the fruit of a project led with Arnaud de Mesmay and Jonathan Spreer. In it, we leverage the obstruction developed in Chapter 4 to prove our Theorem E.

Finally, Chapter 6 presents the remaining main conjectures and lines of research arising from our work.

In addition, Chapter 7 is a French translation of Chapter 1.

Chapter 2

Preliminaries

2.1 Topological background

In the following we will go through a brief introduction of the majority of common topological definitions and notions we will use in this thesis. Getting used to these objects or developing insights about them would require a deeper exposition. We refer again to the book of Schultens [130] for further presentations of the following notions, as well as the most fundamental objects of topology that the reader might not be familiar with.

2.1.1 Basic definitions and assumptions

Piecewise-linear setting. Following standard practice of low-dimensional topology, we will work in the Piecewise-Linear (PL) category [121]. It means that almost all the objects that we use in this thesis are assumed to be piecewise-linear, i.e., made of a finite number of linear pieces with respect to a triangulation of $\mathbb{S}^3$. This last statement will be discussed further within Chapter 3 since we will be led to consider smooth surfaces in $\mathbb{B}^4$ which are more restrictive than the PL ones.

Hence, in the following, we compactify $\mathbb{R}^3$ and work within $\mathbb{S}^3$ which is assumed to be triangulated (recall that such a compactification is presented in Chapter 1 illustrated by Figure 1.7). Furthermore, all embeddings will be PL: a **knot** will be a PL embedding of S^1 into $\mathbb{S}^3$, a **link** is a disjoint union of knots, and a **spatial graph** is a PL embedding of a graph G into $\mathbb{S}^3$. A **diagram** is a PL embedding of a decorated planar graph in the sphere that is obtained from one of the previous objects by a regular projection (we will come back to this point in Section 2.1.3).

More on manifolds. A natural extension to the n-manifolds presented in Chapter 2 are **n-manifolds with boundary** which are spaces everywhere locally homeomorphic to $\mathbb{R}^n$ or the closed half space $[0, +\infty[\times\mathbb{R}^{n-1}$ (see Figure 2.1 for instance). The set of points of an n-manifold with boundary M whose neighbourhood is homeomorphic to the closed half

space forms the **boundary** of M and will be denoted by ∂M. It is worth noticing that the boundary of an n-manifold with boundary naturally inherits from it a structure of an $n-1$-manifold. For example, the boundary of the closed ball $\mathbb{B}^{n+1}$ is the n-sphere $\mathbb{S}^n$.

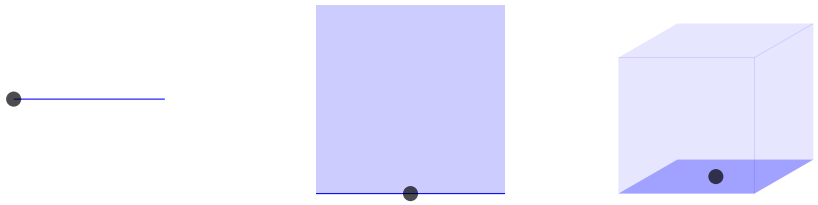

Figure 2.1: A neighbourhood of a point on the boundary of manifolds with boundary.

A subspace A of a connected space X is said to be **separating** if $X \setminus A$ is disconnected. Furthermore, such a space A is said to separate two subsets of X if the two aforementioned subsets lie in distinct connected components of $X \setminus A$.

Formal definition of some continuous operations and objects. If X is a topological space, a **curve of** X is a continuous map $\mathbb{S}^1 \to X$, while an **arc of** X is a continuous map $[0,1] \to X$. Each of these objects is said to be **simple** if it is an embedding: it never crosses itself.

A **homotopy** h between two continuous functions $f : X \to Y$ and $g : X \to Y$ is a continuous map $X \times [0,1] \to Y$ such that $h(\cdot,0) = f(\cdot)$ and $h(\cdot,1) = g(\cdot)$. Two functions are said to be **homotopic** if there exists a homotopy between them. In this thesis, many objects that we manipulate are intrinsically defined as continuous maps: they are said to be homotopic if there exists a homotopy between these maps. Among them are figure knots, that are embeddings of $\mathbb{S}^1$ into $\mathbb{S}^3$ or $\mathbb{R}^3$; and curves on surfaces that are continuous maps of $\mathbb{S}^1$ into the aforementioned surface. A usual way to picture a homotopy is to see it as a continuous transformation when the second parameter in $[0,1]$ is seen as the time (see for instance the homotopies of Figure 2.2).

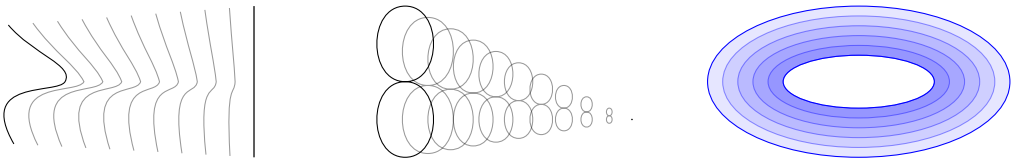

Figure 2.2: Three visualisations of homotopies, one on an arc (left), and one on a curve (middle) which is turned into a point, and an annulus (right) which retracts onto one of its boundaries.

A map is said to be **contractible** if it is homotopic to a constant function. For example, all curves on the sphere or the plane are contractible: they are all homotopic to a constant function whose image is a single point (see the middle of Figure 2.2). If this point is part of X, a homotopy contracting a map to a point is a special case of what is called a retraction. A **retraction** h from X to $A \subset X$ is a homotopy h from X to X such that $h(\cdot,1) \subset A$ and $h(A,1) = A$. When such a retraction exists, we say that X **retracts** onto A (see right of Figure 2.2).

While homotopies allow the objects to self-intersect, their strengthening, called isotopies, does not. An **isotopy** ϕ between two continuous maps $f : X \to Y$ and $g : X \to Y$ is a homotopy between f and g such that $\phi(\cdot, t)$ is an embedding for all $t \in [0, 1]$. In this case, f and g are said to be isotopic. In particular, it follows directly from these definitions that every curve isotopic to a simple curve is also simple.

Finally, we want to define homotopies and isotopies relative to a subspace, which are operations that keep elements of the subspace within it. Formally, for A a subspace of X and f, g two continuous maps, $X \to Y$, h is a **homotopy relative to** A (resp. **isotopy relative to** A) if it is a homotopy (resp. an isotopy) such that: $\forall a \in A$, $h(a, \cdot) = f(a) = g(a)$. This notion can be used to consider arcs on a surface with boundary Σ whose endpoints lie in $\partial\Sigma$. All such arcs are contractible by definition; however, different classes of arcs up to homotopy relative to the boundary can appear.

2.1.2 Surfaces

Orientability and compactness. Recall from Chapter 1 that surfaces are 2-manifolds. All surfaces shown in the previous chapter were orientable. It means that it is possible to define a consistent notion of right and left direction. That is not the case on the **Möbius band**, which is the surface obtained by glueing two opposite sides of a rectangle in the opposite direction as illustrated by Figure 2.3. Indeed, after going once through the red curve of the right side of the same figure, the point and arrows depicted will end up in the same place but with the arrows swapped. A surface is said to be **orientable** if it does not contain a sub-surface homeomorphic to a Möbius band. In this thesis, every surface that we consider will be orientable. Furthermore, all the surfaces will be compact, except the plane.

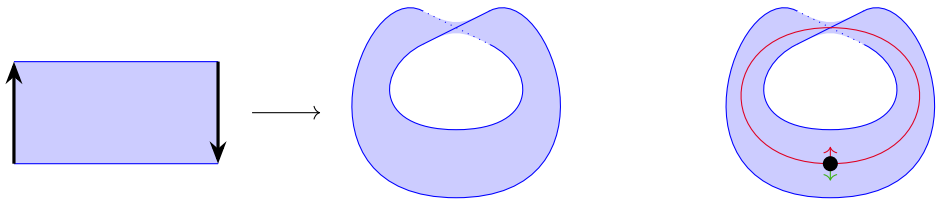

Figure 2.3: A Möbius band.

Genus. We have seen in Chapter 1 that some curves on surfaces may be non-contractible and can be interpreted as the existence of a hole in the surface. We call the **genus** of the surface the maximal number of disjoint simple closed curves that can be removed from a surface without rendering it disconnected (see Figure 2.4). That definition also holds for surfaces with boundaries. For example, 2-spheres, discs, and annuli have genus 0 while the torus has genus 1.

We will say that a surface with a boundary has $b \geq 0$ boundaries if its boundary is made of b connected components. The genus is a powerful surface invariant since compact connected surfaces with boundaries are characterised by their genus and number of boundaries as stated by Theorem 2.1.

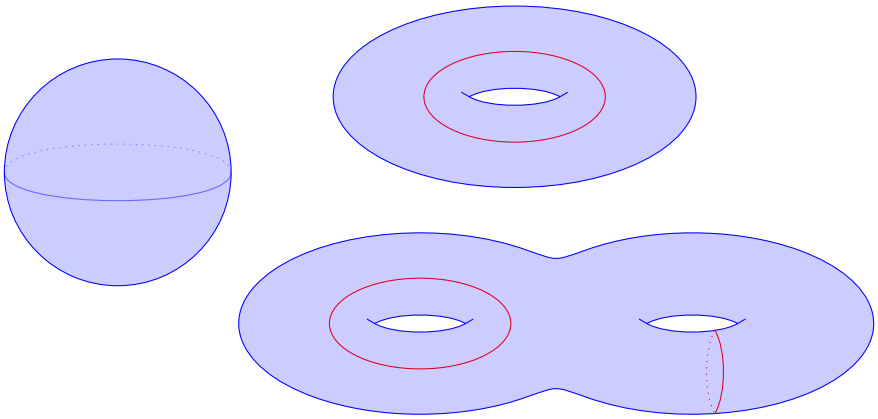

Figure 2.4: The orientable compact surfaces with genus less than 3. The maximal number of disjoint curves along which one can cut without disconnecting the surfaces is shown.

Theorem 2.1 (Classification of surfaces)**.** *Set $g \geq 0$ and $b \geq 0$, all compact connected orientable surfaces with b boundaries and genus g are homeomorphic.*[1]

Furthermore, the genus can easily be computed when the surface with boundary is triangulated. Indeed, when this is the case, the edges and vertices of the triangles used to triangulate the surface form a graph which is cellularly embedded on it. An embedded graph G is **cellularly embedded** on a surface with boundary Σ if $\Sigma \setminus G$ is a disjoint union of open discs where each disc is a face of the graph (such embeddings can be seen in Figure 1.6). In that case, the **Euler characteristic** $\chi(\Sigma)$ of the surface, which is still defined as $V - E + F$, is equal to $2 - 2g - b$.

In a surface with boundary Σ, a curve or an arc is said to be **boundary-parallel** if it is homotopic to the boundary relative to $\partial\Sigma$. Cutting along such a curve is quite irrelevant since it yields two surfaces with boundaries: one which is homeomorphic to Σ by Theorem 2.1 and a second which is a disc or an annulus. A curve or an arc is called **essential** if it is non-contractible and not boundary-parallel; intuitively, essential curves are the interesting curves on a surface. For example, the red curves of Figure 2.4 are all essential curves.

Finally, let us formalise the notion of **cutting**. Let γ be a curve of Σ. In our setting, it is possible to find a neighbourhood N of γ which is homeomorphic to $\gamma \times [-1, 1]$ and unique in that respect. Cutting Σ along γ yields $\Sigma \setminus \mathring{N}$, where $\mathring{N}$ refers to the interior of N. This intuitively represents slicing Σ along γ and letting a copy of A on both sides of the slice. For instance, cutting a torus along a non-contractible simple curve yields a closed annulus, and cutting a sphere along a simple curve yields two closed discs by the Jordan theorem (see right of Figure 2.5). This operation can be generalised to cutting an orientable 3-manifold along a surface embedded in it.

Fundamental group. Submanifolds within manifolds are relevant from a topological point of view. We have seen, for example, that curves within surfaces are significant with respect

[1] This theorem can be extended to compact surfaces by adding the orientability as a criterion.

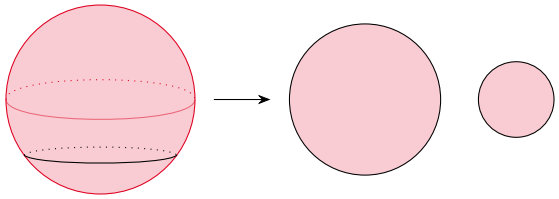

Figure 2.5: Cutting a sphere along the black simple curve.

to their topology: they exhibit "holes" in the surface. Higher dimensional holes can be defined in a similar way using non-contractible k-manifolds within n-manifolds where $k < n$. The study of these sub-spaces uncovered underlying algebraic structures that led to the definition of *homotopy groups* and *homology groups*. Their study is the subject of algebraic topology, which is a very rich mathematical field. Their precise definitions are complex, and understanding the scope of their results and applications exceeds by far the ambition of this section. Hence, we refer to Hatcher [55] for a deep introduction to these concepts; we will only broadly present here the **fundamental group** as it will be used in Chapter 4.

Let Σ be a surface. A **loop** of Σ is an arc α whose images of the endpoints coincide. It follows that $\alpha(0) = \alpha(1)$ is called the **basepoint** of the loop, and α is **based** at $\alpha(0) = p$. Loops based at p are considered up to homotopy that preserves the basepoint, i.e., homotopies h such that for all $t \in [0,1], h(0,t) = h(1,t) = p$. We can now remark that loops based at p can be added. If α and β are two loops based at p, $\alpha + \beta$ is defined as going through α and then through β. Formally, $\alpha + \beta(t) = \alpha(2t)$ if $t \leq \frac{1}{2}$ and $\beta(2t - \frac{1}{2})$ otherwise. It is straightforward to check that it is indeed a loop based at p (such an addition can be seen on Figure 2.6). We also define $-\alpha$ as $t \mapsto \alpha(1-t)$ and notice that $\alpha - \alpha = -\alpha + \alpha = 0$ (the $=$ is up to homotopy preserving the basepoint). Hence, we exhibited a group structure on σ, this group is called the **fundamental group based at** p and denoted $\pi_1(\Sigma, p)$.

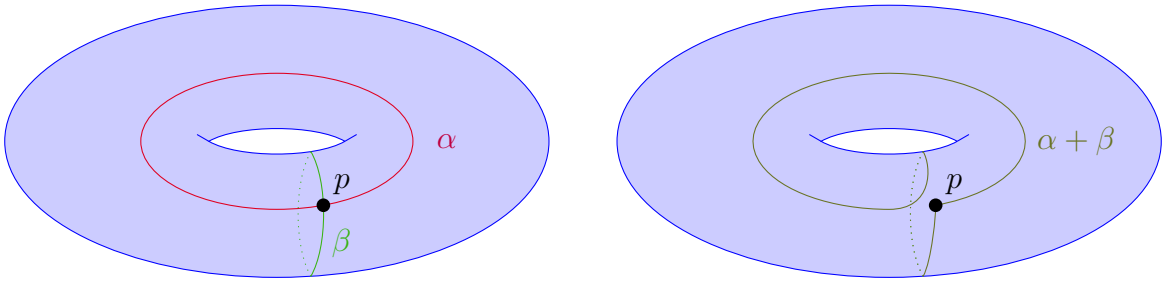

Figure 2.6: Addition of two loops based at p on a torus.

For connected surfaces, the basepoint of loops does not really matter: all fundamental groups based at any point are homomorphic. Hence, it makes sense to forget the basepoint and talk about the group of the homotopy classes of curves and denote it $\pi_1(\Sigma)$: this is the **fundamental group of** Σ. As a first application, we can study the case where $\Sigma = \mathbb{T}$, the torus. One can check in Figure 2.6 that $\alpha + \beta$ and $\beta + \alpha$ coincide (it can be visualised by sliding the "wrapping" part of $\alpha + \beta$ along α to place it on the other side of p). Hence, $\pi_1(\mathbb{T})$ is abelian. A further study shows that $\pi_1(\mathbb{T}) = \mathbb{Z}^2$, which admits $\{\alpha, \beta\}$ as a base. Hence,

all curves on the torus can be decomposed as a sum of p times $\alpha + q$ times β where α and β are the curves depicted in Figure 2.6.

2.1.3 Knot theory

We move on to knot theory and refer to the textbooks of Burde and Zieschang [18] or Rolfsen [120] for readers wishing to explore the topics covered here in greater depth or to find proofs for results that are only stated.

Framework. Recall that we work within a PL framework so that all knots and links that we consider are polygonal embeddings of the circle $\mathbb{S}^1$, or unions of $\mathbb{S}^1$, into $\mathbb{S}^3$. For a knot, admitting such an embedding is called being **tame**, which is to oppose **wild** knots. A wild knot and a polygonal embedding of the trefoil knot are pictured in Figure 2.7. In the following, pictures of knots, links, and their diagrams will often be depicted by rather smooth curves since it is more visually pleasant. This does not matter since smooth knots admit polygonal embeddings, so that one can always find an equivalent depiction with straight lines.

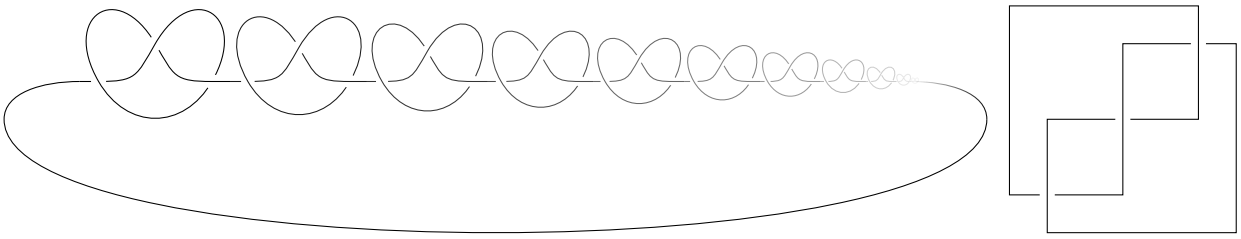

Figure 2.7: Left: a wild knot, right: a polygonal embedding of the trefoil knot.

We recall that neither homotopy nor homeomorphism appear as relevant notions to study embeddings in $\mathbb{S}^3$ according to the current definition of these operations. The former sees all knots as points (they are contractible in $\mathbb{S}^3$) and the latter sees all of them as $\mathbb{S}^1$. The relevant notion to consider is that of **ambient isotopy**, a physical interpretation of which was given in Chapter 1. An ambient isotopy between two embeddings f, g from a space X into a manifold M is a continuous map $i : M \times [0,1] \to M$ such that:

1. $i(\cdot, 0) = id_M$.
2. for all $t \in [0,1]$, $i(\cdot, t) : M \to M$ is a homeomorphism.
3. $i(\cdot, 1) \circ f = g$.

It is necessary that the equivalence relation put on knots takes into account not only knots, which are homeomorphic to circles, but also the space "outside" of them. In the case of knots, this space characterises them:

Theorem 2.2 (Gordon-Luecke Theorem [49]). *Two tame knots are equivalent if and only if their complements*[2] *are homeomorphic.*

Our definition of ambient isotopy readily applies to surfaces with boundaries embedded in space. Notice that such surfaces with boundaries have knots as their boundaries and ambient isotopy is the natural equivalence notion on these embeddings. When these surfaces with boundary are orientable, they are called **Seifert surfaces** of the knot, or more generally link, which constitutes their boundary. For example, Figure 2.8 presents the Seifert surface of minimal genus associated to the trefoil knot. The unknot is the only knot which admits a disc as a Seifert surface. Seifert surfaces are highly useful for knot theory since many invariants are defined by using Seifert surfaces of the knot, and in fact the set of Seifert surfaces of a knot is in itself an invariant. The class of knots and links that we study in Chapter 3 entirely stems from the definition of specific Seifert surfaces.

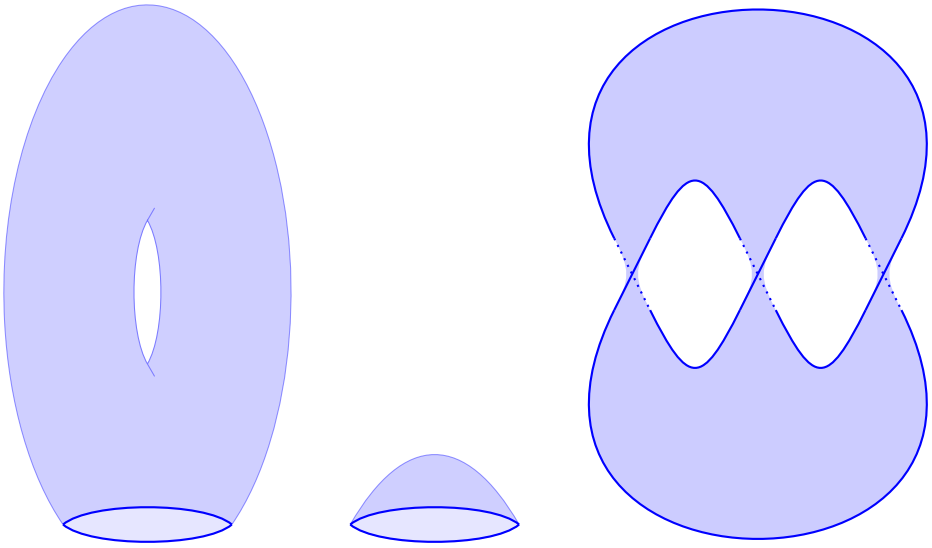

Figure 2.8: Three Seifert surfaces. Left: a torus with one unknot boundary component. Middle: a closed disc with one unknot boundary component. Right: a torus with one boundary component that is a trefoil knot.

Torus knots. In this thesis, we will be especially interested in torus knots. Providing a lower bound for the width of their diagrams is one of the main motivations of Chapter 4, and they are also crucial for us to design the diagrams used to prove Theorem E. Torus knots are knots that can be embedded on a **standard torus**.

A standard torus is depicted on the left side of Figure 2.9, its counterpart is a *knotted* surface on the right side of the same picture. Intuitively, every knot can be embedded on a torus because one can continuously attach a tube along a knot until the two ends of the tube meet again and merge (this process applied to a trefoil knot would yield the knotted

[2]To be more precise, the complement considered is $\mathbb{S}^3 \setminus N(K)$ where $N(K)$ is a regular neighbourhood of the knot. Furthermore, one should add that the homeomorphisms are orientation preserving, since our definition of ambient isotopy forbids orientation reversal.

surface of Figure 2.9). Formally, a standard torus is a torus which is ambient isotopic to a non self-intersecting torus of revolution. Such a torus of revolution can be obtained in $\mathbb{R}^3$ by rotating a translated unit circle $\mathbb{S}^1 = \{(x, y, z) \in \mathbb{R}^3 \mid (x-2)^2 + z^2 = 1\}$, which lies within the plane of equation $y = 0$, around the z-axis. Its core is $\{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 = 4\}$: the circle "inside" the torus.

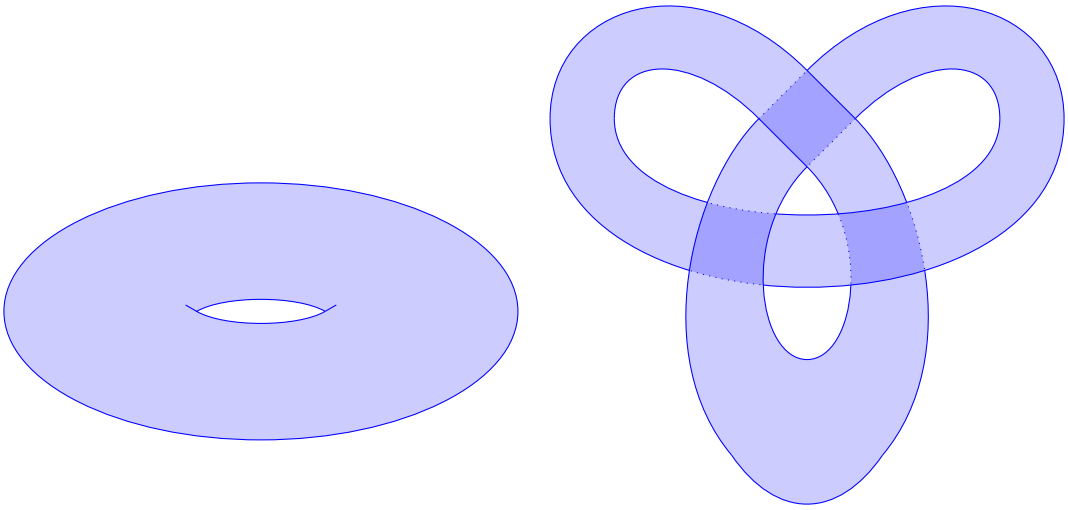

Figure 2.9: Left: a standard torus. Right: a knotted torus.

A **torus knot** is a curve embedded on a standard torus. We have seen in Section 2.1.2 that such a curve can be parametrised by two integers p, q. Therefore, we will call $T_{p,q}$, the torus knot which winds p times around the revolution axis and q times around the core of the torus. In other other words, it intersects q times α and p times β, where α and β are the curves depicted in Figure 2.6. We refer to Figure 2.10 for an illustration of $T_{5,6}$ on the standard torus and one of its diagrams.

Notice that the construction was done in $\mathbb{R}^3$ for clarity purposes. Furthermore, when the standard torus is put in $\mathbb{S}^3$ by compactification, the notion of core as the "inside" no longer exists. Indeed, cutting $\mathbb{S}^3$ along a torus yields two solid tori (3-manifold homeomorphic to $\mathbb{S}^1 \times \mathbb{B}^2$) so that the space is symmetric for the torus. Anyway, this is not an issue for p, q since the torus knot $T_{p,q}$ is always equivalent to $T_{q,p}$ in both $\mathbb{R}^3$ and $\mathbb{S}^3$.

Diagrams. One of the main features of tame knots is that they can be drawn unambiguously, as we explained in Chapter 1. Let us be more formal: tame knots, seen in $\mathbb{R}^3$, admit **regular projections**. A regular projection is a linear projection p from $\mathbb{R}^3$ to a plane P such that it is injective everywhere, except on a finite number of points, called **crossings**, where two preimages are allowed. Furthermore, the crossings must come from transverse parts of the knot. On Figure 2.11, the projection on the red plane presents two points with more than 2 preimages: one at the intersections of 3 strands of the knot and the other, which can be seen as the result of a segment colinear with the direction of the projection.

The existence of a regular projection for a polygonal embedding is fairly intuitive. There is a finite number of directions of projection that are forbidden: one for each segment. Then,

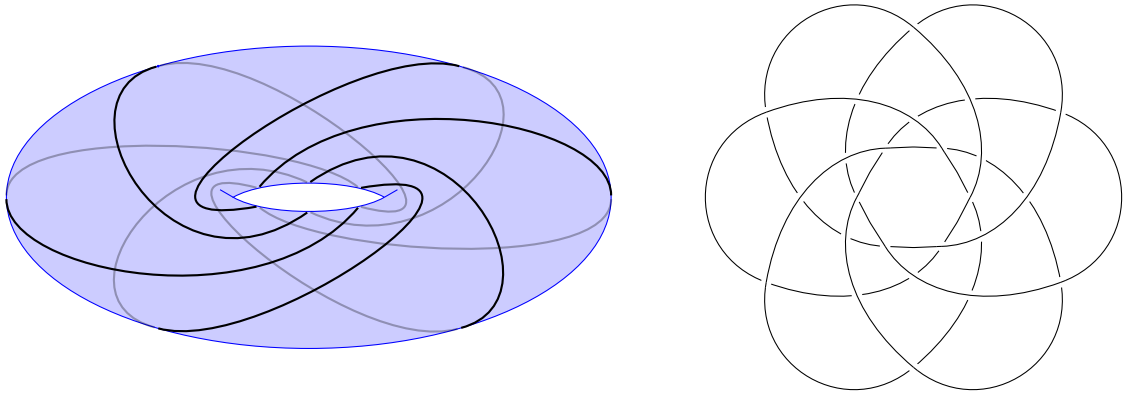

Figure 2.10: Left: the torus knot $T_{5,6}$ embedded on the standard torus. Right: a knot diagram of one $T_{5,6}$.

if for any direction of projection 3 or more strands cross on a diagram, one can slightly perturb the direction of projection to solve that issue. [3]

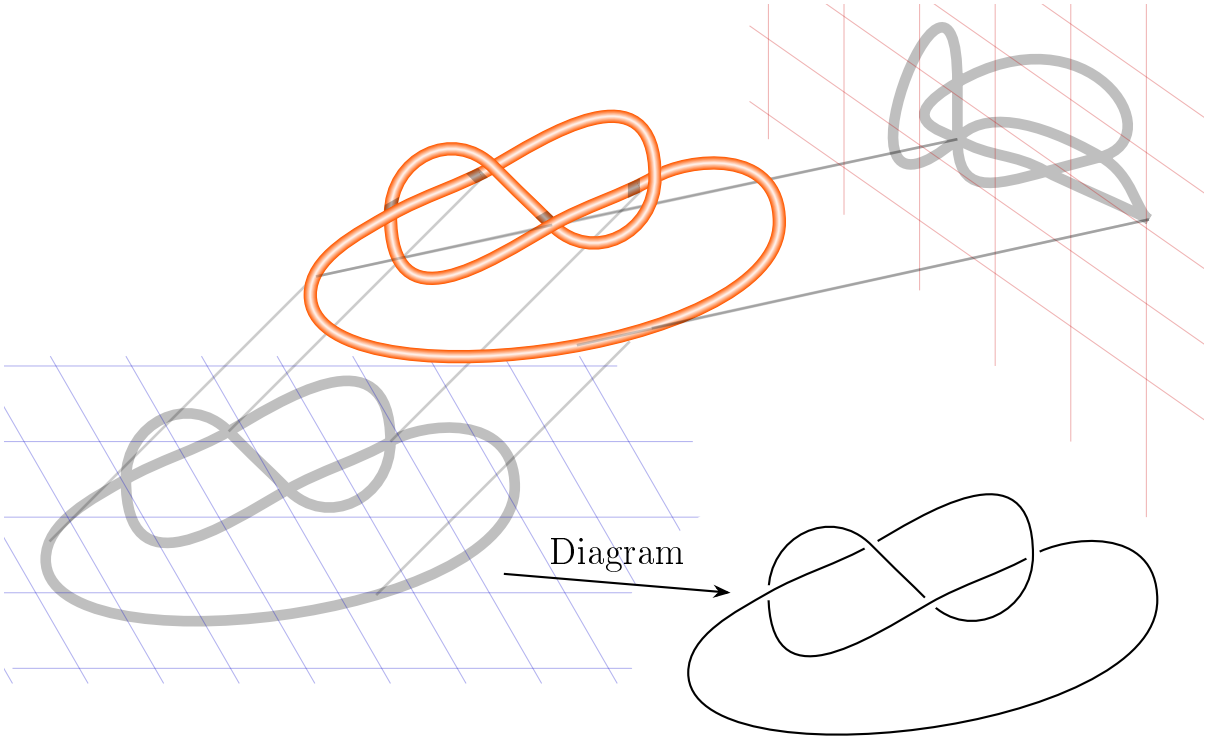

Figure 2.11: A regular projection of the figure eight knot on the blue plane, and a non regular one on the red plane: two points of the projection have more than 2 preimages. Bottom right: the 4-regular diagram arising from the regular projection.

[3] Formally, this last point can be proved by transversality arguments [74].

The image of a regular projection is a continuous map of $\mathbb{S}^1$ in $\mathbb{R}^2$ which induces an embedded decorated planar 4-regular graph in the plane. By compactification of $\mathbb{R}^2$, we obtain the **knot diagram** as an embedding in $\mathbb{S}^2$ with the same properties. This definition directly extends to links so that **link diagrams** are defined the same way. In Chapter 4, we will also be interested in diagrams of spatial graphs. They are also defined the same way, except that the 4-regular property is dropped around projections of vertices. These last points must have only one preimage. Furthermore, diagrams are considered up to ambient isotopies of $\mathbb{S}^2$.

Recall from Chapter 1 that link and knot diagrams encompass the equivalence between knots and links via Reidemeister moves depicted in Figure 2.12:

Theorem 2.3 (Reidemeister theorem [138])**.** *Two link diagrams represent the same ambient isotopy class of a link in* $\mathbb{S}^3$ *if and only if they are related by a finite number of Reidemeister moves*[4]*.*

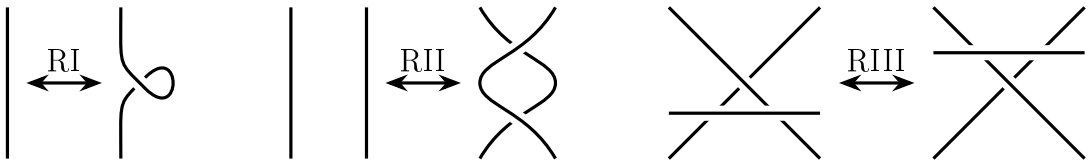

Figure 2.12: The three Reidemeister moves RI, RII, RIII.

Finally let us remark that a sequence of Reidemeister moves induces a homotopy on the projection of the link, which is a union of curves of $\mathbb{S}^2$.

Linking number. When links on the projection are given orientations, it is possible to give a sign to crossings which are illustrated on Figure 2.13. A way to remind oneself which crossing is of which sign is to remember the following rule: when it is possible to get the under-strand direction vector from the over-strand direction vector by an anticlockwise rotation[5] of angle less than π, the crossing is said to be positive, and negative otherwise.

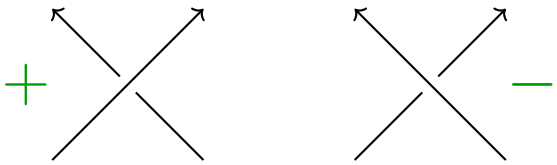

Figure 2.13: A positive crossing on the left, and a negative one on the right.

Half the number of positive crossings minus the number of negative crossings between two link components of a link diagram defines a link invariant called the **linking number**. It is one of the easiest link invariants to compute. Intuitively, it represents how many times a

[4]This theorem is often stated with the addition of "and an isotopy of the plane" but we consider diagrams up to ambient isotopy anyway, so that last part is already covered.

[5]The author particularly prefers this direction of rotation so that readers in the same case might find this rule useful.

link component winds around another one. When two link components have non zero linking number, one can readily conclude that they are **linked** i.e., cannot be separated by a sphere. The converse is not true, as can be observed in the middle of Figure 2.14 where a Whitehead link is depicted (it has linking number 0, but its components are linked). Deciding whether two links are linked is the root of the problem studied in Chapter 5.

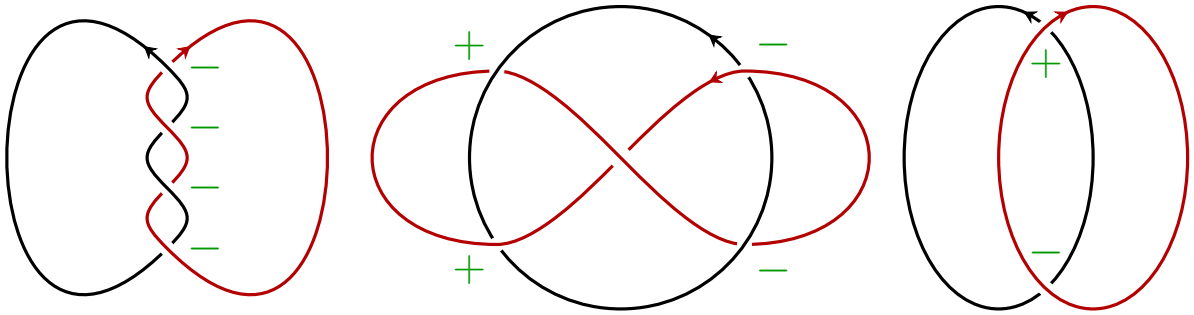

Figure 2.14: Three oriented link diagrams, the sign of crossings between different components is shown in green. Their linking numbers are from left to right: -2, 0, 0.

Computational background on knots and links. From the very definition of ambient isotopy comes the most fundamental algorithmic question in knot theory: the Knot Equivalence problem. We are given two knots K_1 and K_2 and we are tasked with deciding whether they are equivalent, that is, whether it is possible to deform one into the other by an ambient isotopy. The best algorithm for this problem, due to Kuperberg, is elementary recursive [75], yet the problem is not even known to be **NP**-hard (see for example [80, Conclusion]).

A first instance of this last problem is when K_2 is trivial: the problem then becomes the unknot recognition problem. A fundamental computational question is then to figure out the best algorithm for testing whether a given knot is the unknot. This was famously posed as an open problem by Turing [137]. The current state of the art on this problem is that it lies in **NP** [54] and co-**NP** [82]; a quasipolynomial time algorithm has been announced [83] but no polynomial-time algorithm is known.

More generally, algorithmic questions surrounding knots typically display a wide gap between the best known algorithms (which are almost never polynomial-time, and sometimes the complexity is a tower of exponentials) and the best known complexity lower bounds. We refer to the survey of Lackenby for a panorama of algorithms in knot theory [81].

2.2 Background on structural graph theory

The works presented in this thesis are inspired by structural graph theory. We want here to provide some background and insights into the main topics from which we draw our inspirations. These topics are graph minors, well-quasi-orders on graphs, and treewidth. The results and definitions of Section 2.2.1 and 2.2.2 will be used in Chapter 3, while the remaining sections are here for background and inspiration. For a further introduction to graph theory, and in particular its structural aspects, we refer to Diestel [30].

Our graphs may have **multi-edges** (several edges between two vertices) and **self-loops** (edges from a vertex to itself). We will respectively denote by $V(G)$, $E(G)$, and $L(G)$ the sets of vertices, edges, and leaves (degree one vertices) of a graph G. Furthermore, recall that a **face** of a graph G embedded on a surface Σ is a connected component of $\Sigma \setminus G$. Finally, in Chapter 4, we will consider that trees are canonically endowed with a continuous structure, i.e., that each edge is a continuous interval.

2.2.1 Well-quasi-order

The following contains only basic results and definitions relating to well-quasi-orders.

An order $\preccurlyeq$ on a set X is said to be a **well-quasi-order** if for every infinite sequence $(x_n)_{n\in\mathbb{N}}$ there exist $i, j \in \mathbb{N}$ such that $i < j$ and $x_i \preccurlyeq x_j$. Equivalently, $\preccurlyeq$ is a well-quasi-order if it is well-founded and has no infinite **antichain**, that is, no infinite sequence $(x_n)_{n\in\mathbb{N}}$ such that no two elements of (x_n) are comparable for $\preccurlyeq$. A property P is said to be stable for an order $\preccurlyeq$ if for any x satisfying P and $y \preccurlyeq x$, then y satisfies P.

A **minimal element** of X for an ordering $\preccurlyeq$ is an element x such that no $y \in X$ satisfies $y \preccurlyeq x$. If X is well-quasi-ordered by $\preccurlyeq$, it directly follows that the set of minimal elements of any subset A of X is finite. Indeed, the set of minimal elements of A forms an antichain.

A property P is said to be **stable** on a set X ordered by $\preccurlyeq$ if: for all $x, y \in X^2$, if x satisfies P and $y \preccurlyeq x$ then y satisfies P.

Lemma 2.4. *If $\preccurlyeq$ is a well-quasi-order on some set X and P is a property that is stable for $\preccurlyeq$, then there exists a finite family $\mathcal{F}$ of elements of X such that $x \in X$ satisfies P if and only if there is no f in $\mathcal{F}$ such that $f \preccurlyeq x$.*

Proof. Let S be the set of elements of X satisfying P. Then $X \setminus P$ is a subset of X whose set of minimal elements $\mathcal{F}$ is an antichain. Since $\preccurlyeq$ is a well-quasi-order, $\mathcal{F}$ is finite.

Let x be some element of X and f an element of $\mathcal{F}$. If $f \preccurlyeq x$, then x cannot satisfy P because this property is stable by $\preccurlyeq$ and $f \notin S$. For the other way around, notice that for every element y of $X \setminus S$ there exists $f \in \mathcal{F}$ such that $f \preccurlyeq y$ by definition of minimal element and reflexivity of $\preccurlyeq$. Thus, if no element $f \in \mathcal{F}$ satisfies $f \preccurlyeq x$, then $x \in S$. □

The family $\mathcal{F}$ of Lemma 2.4 is called a family of **excluded minors** for the property P. Notice that if a parameter $p : X \to \mathbb{N}$ is monotone with respect to $\preccurlyeq$ i.e., $x \preccurlyeq y$ implies $p(x) \leq p(y)$, then for each $k \in \mathbb{N}$, the property $p(\cdot) < k$ is stable for $\preccurlyeq$ and hence is characterised by such a finite family of excluded minors for X.

2.2.2 Kruskal tree theorem

A first example of a well-quasi-order for us is one formed on trees. Whenever one is given an infinite family of trees, it is always possible to find one tree as a subpart of another one. Before explaining the precise definition of subpart, let us introduce a slight refinement of the notion of trees, since this is the framework that we will use in Chapter 3.

Plane trees. We call a **plane tree** a rooted tree where each vertex v has a label $\ell(v)$ from an alphabet A, and the tree is provided with the combinatorial data of an embedding in the plane. The embedding is given by a permutation at each vertex recording the ordering of the edges to its children. The root induces an orientation on the tree: every edge $\{u, v\}$ is directed from u to v, written $u \to v$, when u is closer to the root of the tree than v (i.e., edges go toward the leaves).

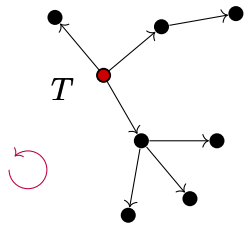

Figure 2.15: A plane tree embedded in the plane. The root is marked in red.

Homeomorphic embedding. We consider plane trees whose labels belong to a set which is well-quasi-ordered under a relation $\leq$. A plane tree T_1 has a **homeomorphic embedding into** T_2, written $T_1 \hookrightarrow T_2$, if T_1 can be obtained from T_2 by iteratively (i) removing a leaf and its adjacent edge, (ii) removing a root with a single child, its adjacent edge, and rerooting at that child, (iii) reducing labels (with respect to $\leq$) and (iv) contracting paths into edges while preserving the labels of the endpoints, where all these operations must be consistent with the plane embedding.

Kruskal proved in 1960 that plane trees labelled by elements that are well-quasi-ordered are also well-quasi-ordered [73]. In 1963, Nash-Williams provided a shorter proof for this theorem [101]:

Theorem 2.5 (Kruskal's tree theorem)**.** *The homeomorphic embedding on the set of labelled plane trees labelled by a well-quasi-order forms a well-quasi-order.*

2.2.3 Graph minor

Graph minors. A graph H is said to be a **minor** of a graph G if it can be obtained from G by contracting edges, deleting vertices, and deleting edges. Deleting an edge simply removes it from $E(G)$, while deleting a vertex also removes all its incident edges. An **edge contraction** is the result of deleting the edge and merging its endpoints (see Figure 2.16).

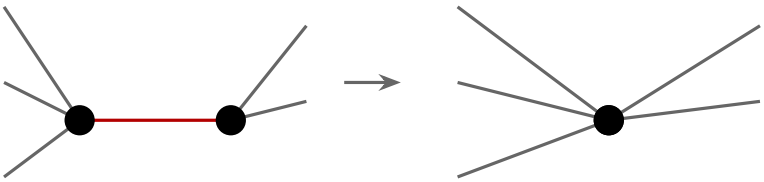

Figure 2.16: An edge contraction of the red edge.

A family of graphs is said to be **minor-closed** if it is stable by minor, i.e. if every minor of a graph of the family also belongs to the family. The family of graphs embeddable on a surface of genus g is an example of a minor-closed family since contracting an edge preserves an embedding. A very powerful result from Robertson and Seymour, whose proof spanned over 20 papers and took as many years, states that every minor-closed family of graphs is characterised by a finite family of excluded minors:

Theorem 2.6 (Robertson-Seymour theorem [118])**.** *The minor relation is a well-quasi-order on graphs.*

Their result can be seen as a powerful extension of Wagner's theorem on planar graphs and was in fact conjectured by him. For a given minor-closed family, finding the actual finite set of forbidden minors is generally a highly difficult problem that requires a deep understanding of the specific class. Linkless graphs are a success in that regard since their set of excluded minors is explicitly known. On the contrary, finding the set of excluded minors for graphs embeddable on the torus is still open, even if more than 17000 minimal elements have been found [99].

Computational applications. The Robertson and Seymour theorem is fundamental from a computational perspective. Indeed, they provided, along with the proof, an algorithm to show that when a graph H is fixed, it is possible to decide in cubic time if H is a minor of G. It follows that deciding any property that is stable under minors on graphs can be done in polynomial time: it is enough to run the aforementioned algorithm on each element of the finite set of forbidden minors. As a result, some algorithmic problems, like deciding if a graph is **knotless** i.e., can be embedded in $\mathbb{R}^3$ with no cycle being an unknot, are known to admit a cubic time algorithm. Strikingly, no explicit algorithm solving this problem is known, and only an exponential-time one is conjectured to be solving this problem [92].

Although the minor testing algorithm has been improved by Kawarabayashi, Kobayashi, and Reed [62] to a quadratic time algorithm, the practical running time of these algorithms turns out to be inapplicable. However, the proof of the Robertson-Seymour theorem also had a huge impact on algorithm design in a more indirect way: it introduced the notion of treewidth, which, as hinted at in the introduction, is very well-adapted to dynamic programming and, more generally, to understanding the underlying structure of graphs.

Treewidth. We gave in Chapter 2 an idea of what treewidth is, let us now give a formal definition. A **bag** is a subset of $V(G)$, and a **tree decomposition** of a graph G is a tree T with vertex set $X_1, \ldots, X_n$ where each X_i is a bag such that:

1. Each vertex of G is in at least one bag: $\bigcup_i X_i = V(G)$.
2. Every edge of G has its endpoints in a bag: $\forall \{u, v\} \in E(G), \exists X_i : \{u, v\} \subset X_i$.
3. For every vertex v of G, the bags containing v induce a subtree of T.

The width of this decomposition is the maximum size of a bag, minus 1. Such a tree decomposition is depicted in Figure 2.17. The **treewidth** is then the minimal width of a tree

decomposition of the graph. It follows clearly from this definition that trees have treewidth 0, they constitute in themselves a tree decomposition achieving this width. On the opposite, trying to have two nodes of size less than $n-1$ in a tree decomposition of K_n will quickly lead to a contradiction showing that the complete graph K_n has treewidth $n-1$. With a straightforward but lengthy analysis, one can prove that grids $p \times q$ have treewidth $\min(p, q)$ and that this treewidth can be achieved by a linear tree, i.e., a path.

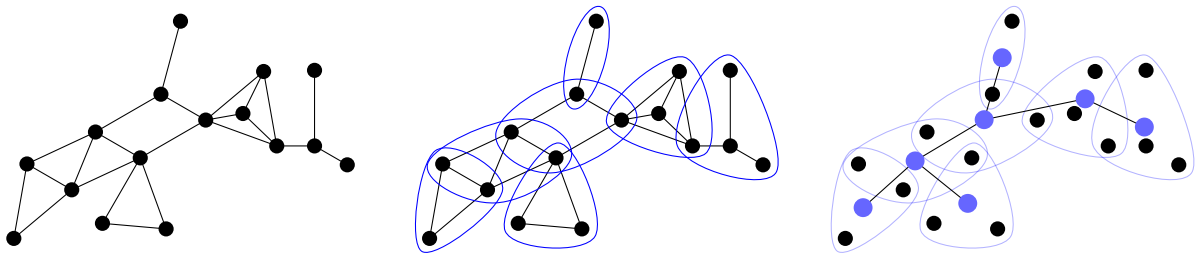

Figure 2.17: A tree decomposition of the graph of size 3.

When a problem belongs to the complexity class **NP**, there always exists a brute-forcing algorithm solving the problem in time exponential n where n is the size of the entry. When a graph is provided with a tree decomposition of width k, one can hope to solve it faster using the following general approach. First, we brute-force the solution in each bag, which can be done in time exponential in k. Then we attempt to recompose these solutions together by following the tree. When the recomposing step can be done in polynomial time, this yields an algorithm whose complexity is $\mathcal{O}(f(k)P(n))$ where P is a polynomial and f any computable function. This is a special case of **fixed-parameter tractable** algorithms, often shortened to **FPT**, which admit various parameters k. The design of FPT algorithms constitutes one of the primary methods to practically solve **NP**-hard problems since their running time is efficient in practice for small values of k. We refer to [27] for a general introduction to parameterized complexity.

2.2.4 Minors on planar graphs

It may not be obvious to the reader how this notion of treewidth is connected to the excluded minors of Theorem 2.6. While sketching the proof of Theorem 2.6 would be too intricate for these preliminaries, we would like to offer a glimpse of the connection to treewidth by explaining the idea of its proof in the easier case of planar graphs.

By adapting the proof of Nash-Williams of the Kruskal tree theorem [101] to graphs of treewidth at most k, Robertson and Seymour proved that planar graphs of treewidth at most k are well-quasi-ordered by the minor relation:

Proposition 2.7 ([112]). *For any integer k, the class of planar graphs of treewidth at most k is well-quasi-ordered by the minor relation.*

We have seen that planar graphs admit graphs of arbitrarily large treewidth: the grids. It turns out that any planar graph can be obtained from a large enough grid:

Proposition 2.8 ([119]). *Let G be a planar graph; there exists k such that G can be obtained as a minor of the $k \times k$ grid.*

Furthermore, grids are intimately connected to high-treewidth due to the following theorem:

Proposition 2.9 ([117]). *There exists a function f such that each graph of treewidth larger than $f(k)$ admits a grid minor of size $k \times k$.*

Hence, if we consider an infinite family of planar graphs $\mathcal{G}$, we can make the following case disjunction. Either the treewidth of the graphs in $\mathcal{G}$ is not bounded. In that case, arbitrarily large grids can be found as minors of graphs in this family by Proposition 2.9. Furthermore, any planar graph can be obtained as a minor of a grid by Proposition 2.8. It follows that any planar graph of the family can be obtained as a minor of another graph of $\mathcal{G}$, and thus $\mathcal{G}$ is well-quasi-ordered under the minor relation. The other case is if the treewidth of elements of $\mathcal{G}$ is uniformly bounded by some constant: this case is handled by Proposition 2.7. Altogether, this proves Theorem 2.10.

Theorem 2.10 ([112]). *Planar graphs are well-quasi-ordered by the minor relation.*

From a more general perspective, if a family of planar graphs is minor-closed, either it has unbounded treewidth, and this family is the family of planar graphs as a whole. Or it has bounded treewidth, and algorithm design on this class can exploit tree decompositions of bounded size.

Chapter 3

Hopf arborescent links, and decidability of the defect

In this chapter, we investigate a class of knots and links called Hopf arborescent links, which are obtained as the boundaries of tree-like iterated plumbings of Hopf bands. We show that for such links, computing the genus defects, which measure how much the four-dimensional genera differ from the classical genus, is decidable. We also show that Hopf arborescent links and some of their Seifert surfaces form a well-quasi-order under various containment relations.

The results of this chapter stem from our article [B], written with Pierre Dehornoy and Arnaud de Mesmay, which appeared in the Proceedings of the 40th International Symposium on Computational Geometry and has been invited to a Discrete & Computational Geometry special issue of Symposium on Computational Geometry 2024. In this chapter, all the 2-manifolds that we consider will be surfaces with boundary. For the sake of simplicity, in this chapter, we will write surface instead of surface with boundary.

3.1 Introduction

Given how seemingly hard testing the equivalence of knots is (see Section 2.1.3), a huge body of research has been devoted to designing and studying knot invariants in order to tell them apart. A classical invariant of a knot is its **genus**: this is the smallest possible genus of one of its **Seifert surfaces**. Computing the genus of a knot turns out to be significantly more tractable: celebrated works of Hass, Lagarias, and Pippenger [54] and Agol, Hass, and Thurston [4], building on the normal surface theory of Haken [53], have shown that deciding if a knot has genus at most g is in **NP**. Later, Lackenby has proved that it is also in co-**NP** [82]. These algorithms run also well in practice within the software Regina [20].

There are, however, different notions of genus that are much less understood: considering $\mathbb{S}^3$ as the boundary of the 4-dimensional ball $\mathbb{B}^4$, the 4-genus of a knot $g_4(K)$ is roughly the smallest possible genus of a surface with boundary in $\mathbb{B}^4$ having the knot as its boundary. This comes in two flavours that are known to not be equivalent: the **topologically locally flat 4-genus** and the **smooth 4-genus**, depending on the regularity of the surface with boundary. We refer to Section 3.2 for precise definitions. A knot is (topologically or smoothly) **slice** if it bounds a disc in $\mathbb{B}^4$, i.e., has 4-genus zero. One of the motivations for the study of such 4-dimensional invariants comes from algebraic geometry, as such surfaces arise naturally around singularities of algebraic curves in $\mathbb{C}^2$ [72, 122]. Another motivation is the slice-ribbon conjecture [41], which states that a knot is smoothly slice if and only if it is ribbon, i.e., it bounds an immersed disc in $\mathbb{S}^3$ with only ribbon-type singularities. From a more elementary point of view, the **unknotting number**, which is the minimal number of crossings to change in a knot diagram to transform the associated knot into the trivial one, is a basic measure for the complexity of a knot. However, this remains, up to now, an extremely hard quantity to compute and study: for example, the unknotting number of all 11 crossing prime knots is still unknown. Interestingly, the smooth 4-genus of knots is a lower bound for it [102]. The idea is that modifying k crossings can be seen as a homotopy in $\mathbb{S}^3 \times [0,1]$ which induces a surface with k "singularities". Replacing these singularities with annuli yields a surface with genus k in $\mathbb{B}^4$ that has the knot as its boundary.

Unfortunately, no algorithmic framework at all is known to attack topological problems in 4-dimensional topology, and indeed many of these problems are known to be undecidable, e.g., the homeomorphism of 4-manifolds [89]. For some other problems, the decidability is a well-known open problem: this is the case for 4-sphere recognition [141] or embeddability of 2-dimensional complexes in $\mathbb{R}^4$ [90]. Similarly, no algorithm is known to decide the 4-genus of a knot or even to decide whether it is slice. To illustrate how hard these problems are, it is only in a recent breakthrough of Picirillo [109] that it was proved that the Conway knot is not smoothly slice, although it only has 11 crossings. From the perspective of lower bounds, recent work of de Mesmay, Rieck, Sedgwick, and Tancer [29] has proved that an analogue of the 4-genus for links, the 4-ball Euler characteristic, is **NP**-hard to compute, but it is also not known to be decidable.

Our results. The goal of this chapter is to investigate the structure of a particular class of links, which we call **Hopf arborescent links**, in order to prove the decidability of some of their 4-dimensional invariants. This family of links is informally defined as follows (we refer to Section 3.2 for more precise definitions).

A **Hopf band** is the surface pictured in Figure 3.1 (top left) that has a Hopf link as its boundary, it can be either **positive** or **negative** depending on how it twists. If they are unlinked, i.e., if there exists a sphere separating them, two Hopf bands can be **plumbed** together by identifying a square in one to a square in the other, as pictured in Figure 3.1 (top). The class of **Hopf arborescent links** is the class of links arising as the boundary of some iterated tree-like sequence of plumbings of Hopf bands. Hopf arborescent links can naturally be described using labelled trees, see Figure 3.1 (bottom), and are a subfamily of the more general arborescent links [16, 45]. The (topological or smooth) **genus defect** of a knot is defined as $\Delta g(K) = g(K) - g_4(K)$. It measures how much its 4-genus differs from its

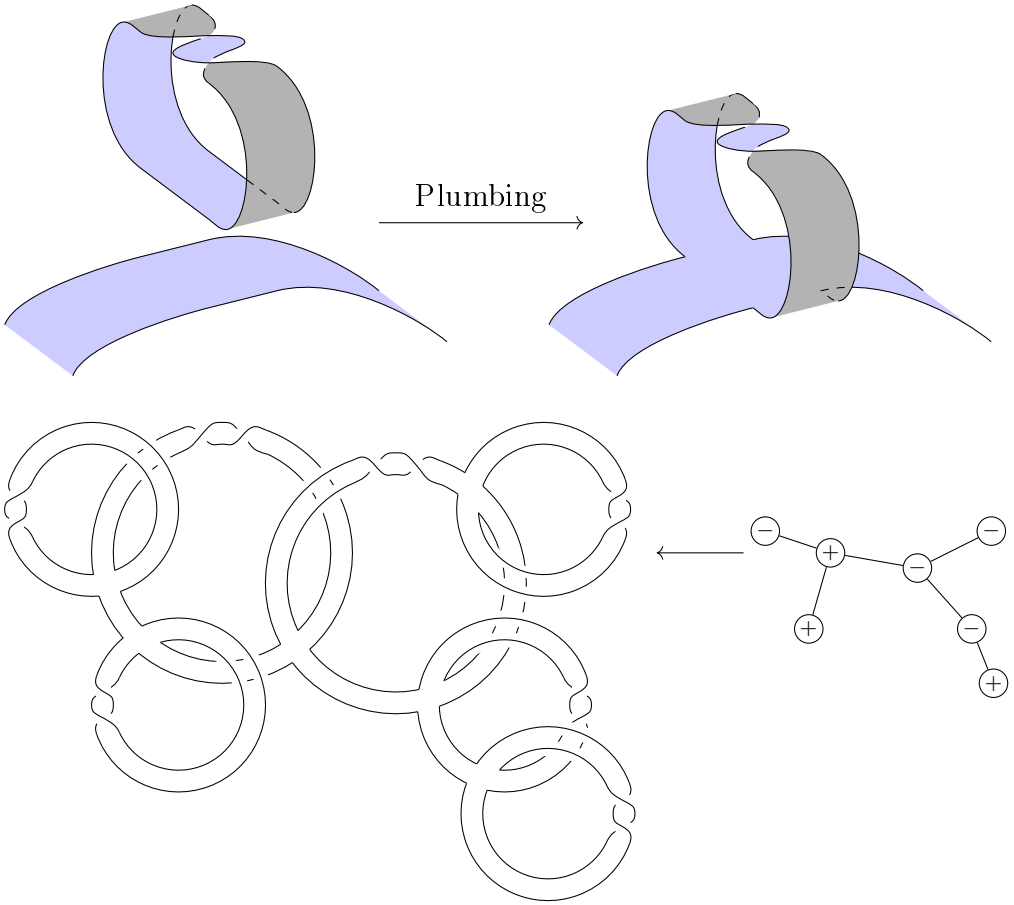

Figure 3.1: Top: A positive Hopf band and a plumbing. Bottom: An Hopf arborescent link and an associated planar tree.

classical genus, and this definition can be extended to oriented links by considering surfaces having that oriented link as their boundary. Our main result is the following:

Theorem A. *For any fixed k, deciding whether a Hopf arborescent link L has genus defect at most k is decidable. This holds both in the topological and smooth categories.*

The proof of Theorem A is not constructive. It is obtained as a corollary of another result, which establishes a well-behaved minor theory for Hopf arborescent links. A subsurface $\Sigma' \subseteq \Sigma$ of a Seifert surface Σ is **incompressible** if the complement $\Sigma \setminus \Sigma'$ has no open disc component. Given two Seifert surfaces Σ_1 and Σ_2 in $\mathbb{R}^3$, we say that Σ_1 is a **surface-minor** of Σ_2 (or **minor** for short), denoted by $\Sigma_1 \preccurlyeq \Sigma_2$, if Σ_1 is isotopic to an incompressible subsurface of Σ_2. This minor relation was introduced by Baader [7] (see also [8, 9]) with the goal of characterising those links that are closures of positive braids and whose signature is equal to twice their genus. The underlying question, still open today, is whether the canonical Seifert surfaces associated with positive braid closures form a well-quasi-order. A specificity of Hopf arborescent links is that they are **fibred** (see the definition in Section 3.2). This implies that to each link is associated a canonical Seifert surface (see Theorem 3.2), that we

call a **Hopf arborescent surface**. The notion of surface-minor naturally implies a minor relation for Hopf arborescent links. Our second result proves that Hopf arborescent surfaces are well-quasi-ordered under surface-minors.

Theorem B. *The minor relation $\preccurlyeq$ is a well-quasi-order for the set of Hopf arborescent surfaces, that is, for any infinite sequence $(\Sigma_n)_{n\in\mathbb{N}}$ of Hopf arborescent surfaces, there exists $i < j$ in $\mathbb{N}$ such that $\Sigma_i \preccurlyeq \Sigma_j$.*

The idea behind the proof of Theorem B is to study a specific subset of the possible surface-minors that interacts nicely with an encoding of Hopf arborescent surfaces via labelled plane trees. We can then leverage the celebrated Kruskal tree theorem [73], presented in Section 2.2 in the form of Theorem 2.5, to prove that the minor relation is a well-quasi-order. The connection from Theorem B to Theorem A follows from the fact that the genus defect is minor-monotone, i.e., if $\Sigma_1 \preccurlyeq \Sigma_2$ are Seifert surfaces of minimal genus for links K_1 and K_2, then $\Delta g(K_1) \leq \Delta g(K_2)$. This is not a new observation (see [10, Lemma 6]), we provide a proof in Proposition 3.11 for completeness. Therefore, for Hopf arborescent links, having genus defect at most k is characterised by a finite number of forbidden minors, and the algorithm of Theorem A proceeds by checking those. However, testing whether a surface is a minor of another one seems to be a very hard problem: even testing whether two tori are isotopic is already as hard as knot equivalence, and an algorithm for genus two surfaces was only very recently found [11]. This problem is circumvented thanks to our restriction of the minor relation to one that is well-tailored to the arborescent structure of our links, which allows us to work entirely at the level of trees. In particular, Theorem A does not strictly follow from Theorem B but rather from its proof, relying on Proposition 3.9.

While the algorithms behind Theorem A are not explicit, we would like to offer three reasons to motivate our results. First, Theorem A proves that the corresponding problems are *not* undecidable, which is a significant result in the landscape of 4-dimensional topology. Second, this kind of existential algorithmic result has been a strong guiding light in algorithm design in the past decades. Indeed, for a vast family of graph problems, the fact that an algorithm merely exists follows from Robertson-Seymour theory, and this has provided a strong impetus to actually look for explicit algorithms and optimise their complexity. This has been particularly influential in parameterized algorithms: we refer, for example, to the discussion in Chapter 6.3 in the book on parameterized algorithms [27], where it is conjectured that a result like our Theorem A precludes W[1]-hardness. Similarly, we are hopeful that our results can inspire future work aiming at developing explicit algorithms in 4-dimensional topology. Additionally, our framework directly proves that any property that is stable with respect to our link-minor relation (see Section 3.3.4) is decidable on the class of Hopf arborescent links. Finally, while it is certainly not the case that minor-based approaches can encompass the entirety of knot theory, it is fruitful to delineate exactly the classes which they can illuminate. In that respect, we find it interesting that our proof of Theorem B strongly relies on the structure of Hopf bands and does not seem to generalise to the wider family of arborescent knots, even when one bounds the number of twists in each band (see Remark 3.10).

Related work. It was observed by Baader and Dehornoy [8, 84] that the natural Seifert surfaces for another class of knots, the positive braid knots with bounded braid index (we refer to the papers for the relevant definitions) also form a well-quasi-order. Furthermore, Liechti [84] proved that even without bounding the braid index, the set of positive braid knots of bounded genus defect is characterised by a finite number of forbidden minors. Since the minor relation in that setting simply amounts to removing letters in the braid presentation, this readily yields decidability as in our Theorem A. While the two results are incomparable, we emphasise that our result also applies to links and also features negative crossings (coming from negative Hopf bands). This last point extends the impact of our result to the smooth category, while for strongly quasipositive knots (and thus positive braid knots), the smooth defect is zero since the smooth 4-genus and the classical genus coincide [122].

All the knots and links we consider, as well as those considered by various authors in the context of surface-minor theory, are fibred (see again the definition in Section 3.2). This property is important as it brings control to the classical genus of the links. Also, it is easy to construct infinite families of incomparable surfaces when dropping this assumption: the set $(A_n)_{n\in\mathbb{Z}}$ of those unknotted annuli in $\mathbb{S}^3$ with n twists forms an infinite antichain. In this direction, an optimistic conjecture would be that the collection of all fibred surfaces in $\mathbb{S}^3$ is a well-quasi-ordered set. If true, that would provide a strong generalisation of Theorem B. However, no strategy of proof is known to the authors for such a statement.

Also, it follows from a result of Giroux and Goodman [48] that *any* fibred link can be obtained from the unknot from a sequence of plumbings and deplumbings (a natural reverse operation to plumbing) of Hopf bands. While these (de)plumbings might not have the arborescent structure that characterises ours, this shows that Hopf bands can be considered as basic building blocks for a wide class of three-dimensional objects.

Organisation of this chapter. After providing background and going through the specifics of this chapter in Section 3.2, we focus on Hopf arborescent links in Section 3.3. There, we define a precise construction of these links from plane trees, investigate this class of links, and prove Theorem B. In Section 3.4, we prove Theorem A. Finally, in Section 3.5, we first provide examples of Hopf arborescent links with non zero defect, and then explain how to combine them to obtain examples with arbitrarily large defect.

3.2 Specific preliminaries

Knot theory. We only recall the definitions that are critical to this chapter and refer to Teichner [136] for a beginner-friendly introduction to the 4-dimensional aspects of knot theory.

In this chapter, every knot and link component is endowed with an orientation inherited from the orientation of $\mathbb{S}^1$. For algorithmic purposes, we assume that an input link is given as a link diagram, encoded by a directed 4-valent planar graph with decorations at vertices indicating which strands are going over and under. Since this chapter focuses on decidability problems, switching to a different input, e.g., polygonal curves in $\mathbb{R}^3$, makes no difference. Furthermore, Seifert surfaces need to be consistent with orientation: for an oriented link L a **Seifert surface** is a compact connected oriented surface Σ embedded in $\mathbb{S}^3$ such that the

oriented boundary of Σ is L (see, for example, the oriented Hopf links of Figure 3.2). The **genus** of a link L, denoted by $g(L)$, is the smallest possible genus of a Seifert surface for L.

The 3-dimensional sphere can be seen as the boundary of the 4-dimensional ball $\mathbb{B}^4$. Being embedded in $\mathbb{S}^3$, a knot or a link can also be obtained as the boundary of surfaces embedded in $\mathbb{B}^4$. However, any knot K in $\mathbb{S}^3$ can be used as a base that tapers to a point, the apex, inside B^4 to define a cone that bounds K. Hence, any knot in $\mathbb{S}^3$ bounds a topological disc in $\mathbb{B}^4$. This motivates the following definition: a surface Σ embedded in $\mathbb{B}^4$ is **locally flat** if for each point $x \in \Sigma$, there is a neighbourhood U in Σ and a neighbourhood V in $\mathbb{B}^4$ such that the pair (U, V) is homeomorphic to the standard $(\mathring{\mathbb{B}}^2, \mathring{\mathbb{B}}^4)$. In the coning construction above, the latter condition is not satisfied at the apex, where the boundary of a disc is the knot instead of a standard $\mathbb{S}^1$. The **topological** (respectively **smooth**) **4-dimensional genus**, or simply **4-genus** of a link L, denoted by[1] $g_4(L)$ is the smallest possible genus of a compact connected oriented surface that is locally flat (respectively smoothly) embedded in $\mathbb{B}^4$, and that has L as its boundary.

A knot is **topologically** (resp. **smoothly**) **slice** if it bounds a locally flat (resp. a smooth) disc in $\mathbb{B}^4$. The **(topological or smooth) defect** of a link L is the quantity $\Delta g(L) = g(L) - g_4(L)$, where $g_4(L)$ denotes the topological or smooth 4-genus of L.

A **positive**, resp. **negative, Hopf band** is an unknotted annulus as pictured in Figure 3.2 with a positive, resp. negative, full twist. A **Hopf link** is the boundary of a Hopf band. Note that the two components of a positive Hopf link have linking number $+1$, while the two components of a negative Hopf link have linking number -1. A Hopf band naturally retracts to a trivial knot, which we call its **core**.

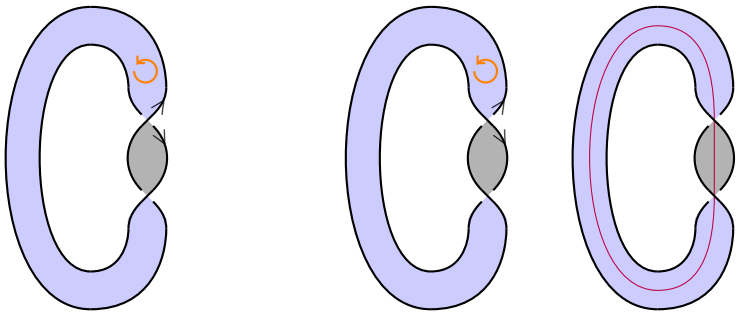

Figure 3.2: A negative Hopf band on the left and a positive one on the right with its core in red.

A link is **fibred** if the complement $\mathbb{S}^3 \setminus L$ **fibres** over $\mathbb{S}^1$, that is, if there exists a one-parameter continuous family of Seifert surfaces $(\Sigma_t)_{t\in\mathbb{S}^1}$ for L which are disjoint except for their boundaries and whose interiors together foliate $\mathbb{S}^3 \setminus L$. These are called **fibres** or **fibre surfaces**. Note that they are all ambient isotopic by definition.

A positive, resp. negative, Hopf link is fibred, with fibre the positive, resp. negative, Hopf band. Indeed, seeing $\mathbb{S}^3$ as the unit sphere $\{(z_1, z_2) \,|\, |z_1|^2 + |z_2|^2 = 1\}$ in $\mathbb{C}^2$, a positive Hopf link is given by the equation $z_1 z_2 = 0$. For every argument $\theta \in \mathbb{S}^1$, the equation

[1] Our notation is intentionally ambiguous with respect to smooth or topological genus, because all our arguments will apply equally well in both categories.

$\arg(z_1 z_2) = \theta$ describes a Seifert surface bounded by the Hopf link, and the collection of these surfaces describes the desired fibration.

It is a folklore result that goes back at least to Stallings [134] that a fibre surface of a fibred link is a Seifert surface of minimal genus, and moreover this surface is unique up to isotopy (see next paragraph). So, for fibred links - and all links in this chapter will be fibred - it makes sense to speak of the **canonical** Seifert surface, by which we mean the unique fibre surface for that link up to ambient isotopies.

Fibre surface, minimal genus, and uniqueness. Here, we comment on the following statements, which are considered folklore in the knot theoretical community. The first one is fundamental for this chapter since it brings control over the genus of Hopf arborescent links. The second one justifies the notion of *canonical* Seifert surface for fibred links.

Theorem 3.1. *Fibred surfaces minimise the genus over surfaces with the same oriented boundary.*

Theorem 3.2. *Two fibred surfaces with the same oriented boundary are isotopic, relatively to the boundary.*

Theorem 3.1 is often presented as a consequence of Stallings' Theorem [134]. Here we present a semi-elementary proof (for which we claim no novelty), in the sense that it only uses basic notions and results from homology theory, we refer to [55] for an introduction to this theory.

Proof of Theorem 3.1. Let $L = (L_1 \sqcup \cdots \sqcup L_k)$ be an oriented link in $\mathbb{S}^3$. We consider the 3-manifold N_L which is the complement of an open tubular neighbourhood of L in $\mathbb{S}^3$. It has k boundary components, which are all tori.

Note the following isomorphisms from elementary algebraic topology: by excision [55, p119], one has $\mathrm{H}_2(\mathbb{S}^3, L; \mathbb{Z}) \simeq \mathrm{H}_2(N_L, \partial N_L; \mathbb{Z})$; in the long exact sequence $\cdots \to \mathrm{H}_2(\mathbb{S}^3; \mathbb{Z}) \to \mathrm{H}_2(\mathbb{S}^3, L; \mathbb{Z}) \to \mathrm{H}_1(L; \mathbb{Z}) \to \mathrm{H}_1(\mathbb{S}^3; \mathbb{Z}) \to \dots$ the first and last term are trivial so the boundary map gives an isomorphism $\partial : \mathrm{H}_2(\mathbb{S}^3, L; \mathbb{Z}) \to \mathrm{H}_1(L; \mathbb{Z})$; and by Alexander duality [55, p254], $\mathrm{H}^1(\mathbb{S}^3 \setminus L; \mathbb{Z})$ is isomorphic to $\mathrm{H}_1(L; \mathbb{Z}) \simeq \mathbb{Z}^k$.

The second isomorphism states that a class in $\mathrm{H}_2(\mathbb{S}^3, L; \mathbb{Z})$ is described by its boundary. That means that when we restrict our attention to oriented Seifert surfaces for L, i.e., surfaces S so that $\partial S = L$, all such surfaces lie in the class $\partial^{-1}(1, \dots, 1)$. In particular, they are homologous. Given two such surfaces, one can consider their restrictions S_1, S_2 to N_L, where they are also homologous. Their boundaries are curves on the tori ∂N_L, which therefore have the same homology class. With an isotopy it is thus possible to make S_1, S_2 parallel and disjoint in a neighbourhood of ∂N_L.

The class of the Seifert surfaces bounded by L is dual to the class $\mathrm{Lk}(\cdot, L_1) + \cdots + \mathrm{Lk}(\cdot, L_k) \in \mathrm{H}^1(\mathbb{S}^3 \setminus L; \mathbb{Z})$. We denote the latter by ℓ_L. We then consider the infinite cyclic covering $\hat{N}_L \to N_L$ associated to ℓ_L. This is the covering associated to the morphism $\ell_L : \pi_1(\mathbb{S}^3 \setminus L) \to \mathbb{Z}$, so that $\pi_1(\hat{N}_L) = \ker(\ell_L)$.

One way to construct $\hat{N}_L$ is to consider an arbitrary Seifert surface S for L, cut $\hat{N}_L$ along S, thus obtaining a 3-manifold with boundary $N_{L,S}$. This boundary consists of three parts: a "horizontal part" composed of two copies of S that we denote by S^+ and S^-, and

a "vertical part" coming from ∂N_L which consists of k annuli (or *sutures*) connecting S^+ to S^-. Now consider $\mathbb{Z}$ copies $(N^n_{L,S})_{n\in\mathbb{Z}}$ of $N_{L,S}$. Call S^+_n and S^-_n the horizontal part of the boundaries, and for every $n \in \mathbb{N}$ glue S^+_n to S^-_{n+1}. The resulting 3-manifold $\cup_{n\in\mathbb{Z}} N^n_{L,S}$ has a natural projection to N_L, and the loops in N_L that lift to closed loops in $\hat{N}_L$ are exactly those in the kernel of ℓ_L. Therefore we indeed constructed $\hat{N}_L$.

Now suppose that $\mathbb{S}^3 \setminus L$ fibres over the circle with fibres having L as oriented boundary. Then the previous construction can be made by choosing for the surface S a fibre. In this case $N_{L,S}$ is the product manifold $S \times [0,1]$, and $\hat{N}_L$ is homeomorphic to $S \times \mathbb{R}$.

Consider now an arbitrary Seifert surface S' for L with oriented boundary L. Then S' is homologous to S, and in particular it is also dual to ℓ_L. So S' lifts in $\hat{N}_L$ as $\mathbb{Z}$ parallel copies $(S'_n)_{n\in\mathbb{Z}}$. Take any such component S'_0, since it lies in $\hat{N}_L \simeq S \times \mathbb{R}$, the projection on the first coordinate induces a map $f : S'_0 \to S$. Looking at the neighbourhood of the common boundary $\partial S = L = \partial S'_0$ we see that f has degree 1. Degree 1 maps induce surjections in homology, which implies the desired inequality. More concretely, one can argue as follows: if there was a class $a \in \mathrm{H}^1(S;\mathbb{Z})$ in the kernel of f^*, then there would be a class b such that the cup-product $a \smile b$ is the fundamental class $[S] \in \mathrm{H}^2(S;\mathbb{Z})$, and therefore one would have $f^*([S]) = f^*(a) \smile f^*(b) = 0 \neq [S'_0]$, a contradiction. Therefore there is a linear injection from $\mathrm{H}^1(S;\mathbb{Z}) \simeq \mathbb{Z}^{2g(S)+k-1}$ into $\mathrm{H}_1(S'_0;\mathbb{Z}) \simeq \mathbb{Z}^{2g(S'_0)+k-1}$. This implies $g(S) \leq g(S'_0)$, and proves the statement. □

Theorem 3.2 is sometimes also attributed to Stallings, but we could not find a corresponding statement in Stallings' articles. However, a proof is proposed by Whitten [142], which relies on the fundamental work of Waldhausen [140]. Indeed, if S_1 is a fibre surface for a link L, then the infinite cyclic covering $\hat{N}_L$ constructed in the proof of Theorem 3.1 is homeomorphic to $S_1 \times \mathbb{R}$. If S_2 is a Seifert surface for L that is also a fibred surface, then Theorem 3.1 implies that it is of minimal genus. Therefore any of its lifts in $\hat{N}_L$ is incompressible (for otherwise it would not minimise the genus). One can pick a lift S'_2 that is disjoint of $S_1 \times \{0\}$ and assume it lives in $S_1 \times (0,N)$ for some large enough N. Proposition 3.1 of Waldhausen [140] then implies that S'_2 is isotopic to $S_1 \times \{0\}$. Projecting back in N_L yields an isotopy from S_2 to S_1. The proof of [140, Proposition 3.1] seems too long to be detailed here. We only mention that it works by an induction on the complexity of the surface S_1 and by cutting it along curves until it is a disc.

Plane trees. We recall and precise plane tree and homeomorphic embedding in this chapter.

A **plane tree** is a rooted tree where each vertex v has a label $\ell(v)$ from the alphabet $\{+,-\}$ endowed with the empty ordering $\leq$. Furthermore, the tree is provided with the combinatorial data of an embedding in the plane i.e., each vertex is additionally given a permutation recording the ordering of the edges to its children. In addition, the root induces an orientation on the tree: every edge $\{u,v\}$ is directed from u to v, written $u \to v$, when u is closer to the root of the tree than v, i.e., edges go toward the leaves, we refer to the trees of Figure 3.5 for examples.

A plane tree T_1 has a **homeomorphic embedding into** T_2, written $T_1 \hookrightarrow T_2$, if T_1 can be obtained from T_2 by iteratively (i) removing a leaf and its adjacent edge, (ii) removing a

root with a single child, its adjacent edge and rerooting at that child, and (iv) contracting paths into edges while preserving the labels of the endpoints, where all these operations must be consistent with the plane embedding. Throughout this chapter, the relation $\leq$ will be trivial, so that (iii) of Section 2.2 will never apply. We recall that this notion it is more restrictive than the notion of minor on graphs, which allows us to contract any edge to a point: in our case we can only contract paths to at least one edge (see Figure 3.3). This property turns out to be critical in order to make our proofs work. Theorem 2.5 states that this order is a well-quasi-order on the set of labelled plane trees.

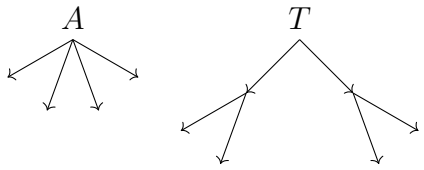

Figure 3.3: The tree A is a graph minor of T but does not embeds homeomorphically on it.

3.3 Hopf arborescent links

Arborescent links are a class of knots and links originally defined and studied by Conway[2]. This class has received much attention from knot theorists [16, 45, 126]. In this chapter we study a subclass that we call Hopf arborescent links.

3.3.1 Hopf plumbing

The links that we investigate here are boundaries of surfaces which are defined iteratively from Hopf bands using an operation called **plumbing**.

Let H be a Hopf band and Σ be an oriented surface with boundary, and let us assume that they are unlinked, i.e., that there exists a sphere S in $\mathbb{S}^3$ separating them. To plumb H on Σ, pick an arc α on Σ whose endpoints lie on $\partial\Sigma$ and which is not boundary parallel (i.e., α is not isotopic relatively to its endpoints to an arc in $\partial\Sigma$). Let D be a small neighbourhood of α in Σ that we see as a rectangle with two sides on $\partial\Sigma$ and two sides in the interior of Σ. Isotope Σ within $\mathbb{S}^3 \setminus S$ so that it intersects S exactly on D, see Figure 3.4, left. Then, define similarly D' a neighbourhood of the unique (up to isotopy) non boundary parallel arc in H with endpoints in ∂H. The orientations of Σ and H induce an orientation of the normal direction to the surface (so that concatenating the orientation of the surface with the positive normal direction gives a positive basis in $\mathbb{S}^3$). Finally, isotope H within its component of $\mathbb{S}^3 \setminus S$, so that D and D' are identified on S in a way that the sides of D that are on ∂S are matched with the sides of D' that are not on ∂H and the orientations of both rectangles match. The resulting surface is said to be obtained from Σ by **Hopf plumbing H on top**

[2]Conway called them algebraic links, but this denomination is now more used for the links that come from algebraic curves in $\mathbb{C}^2$.

of Σ along α, see Figure 3.4. For the sake of simplicity, and when we do not need to specify α, we will denote this operation by **Hopf plumbing**.

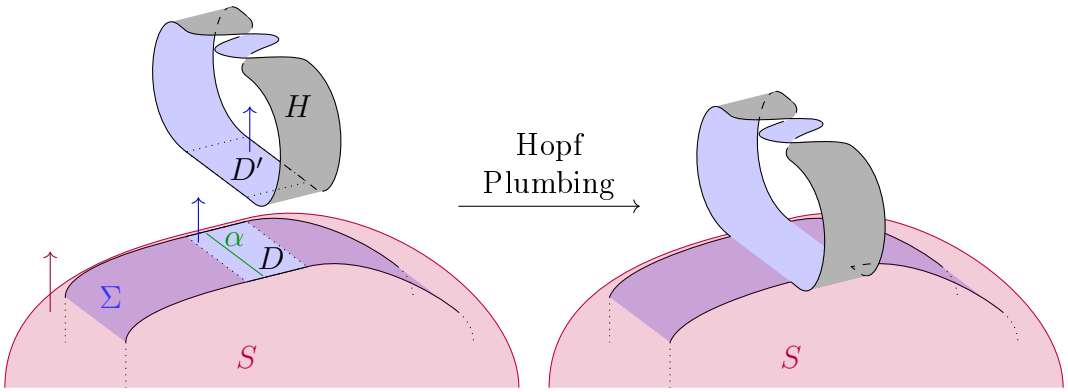

Figure 3.4: A Hopf plumbing of a Hopf band H on top of a Seifert surface Σ along α.

Hopf plumbing is a special case of a more general operation called a Murasugi sum, see [97, 104]. A key property of Murasugi sums, proved by Gabai [44], is that it preserves fibredness. In the above setting, since Hopf bands and Σ are fibred, then the surface obtained from Σ by Hopf plumbing H on Σ along any arc is also fibred.

3.3.2 From plane trees to Hopf arborescent surfaces and links

Let T be a plane tree. The associated surface $\Sigma(T)$ we construct is an oriented surface with boundary that retracts on the union of a finite set of oriented simple curves $\mathcal{C}_T$ parametrised by the vertices of T. Furthermore, every $\alpha \in \mathcal{C}_T$ is the core of a Hopf band embedded on Σ whose sign is the label of the corresponding vertex in T. For a vertex v in T, the curve $\alpha(v)$ intersects another curve $\alpha(v')$ if and only vv' is an edge of T, and the two curves intersect exactly once. Moreover, following $\alpha(v)$ with its given orientation, the cyclic ordering of the intersection points with the curves $\alpha(v')$ coincides with the cyclic orderings of the neighbours of v in the plane tree T.

We now describe the construction inductively, see Figure 3.5 for an illustration.

Construction of Hopf arborescent surfaces:

1. Start from a Hopf band $H(v_r)$ where v_r is the root of T, and whose sign is the label $\ell(v_r)$.

2. For the induction step, assume that the tree T' is obtained from T by adding at a leaf v a finite number of leaves $v_1, \dots, v_k$ appearing in the plane in this order around v, and that the surface $\Sigma(T)$ is already constructed with its set of core curves $\mathcal{C}_T$.

   a) Since v is a leaf in T, the curve $\alpha(v)$ intersects only one curve $\alpha(v')$: the curve associated to v', the parent of v in T.

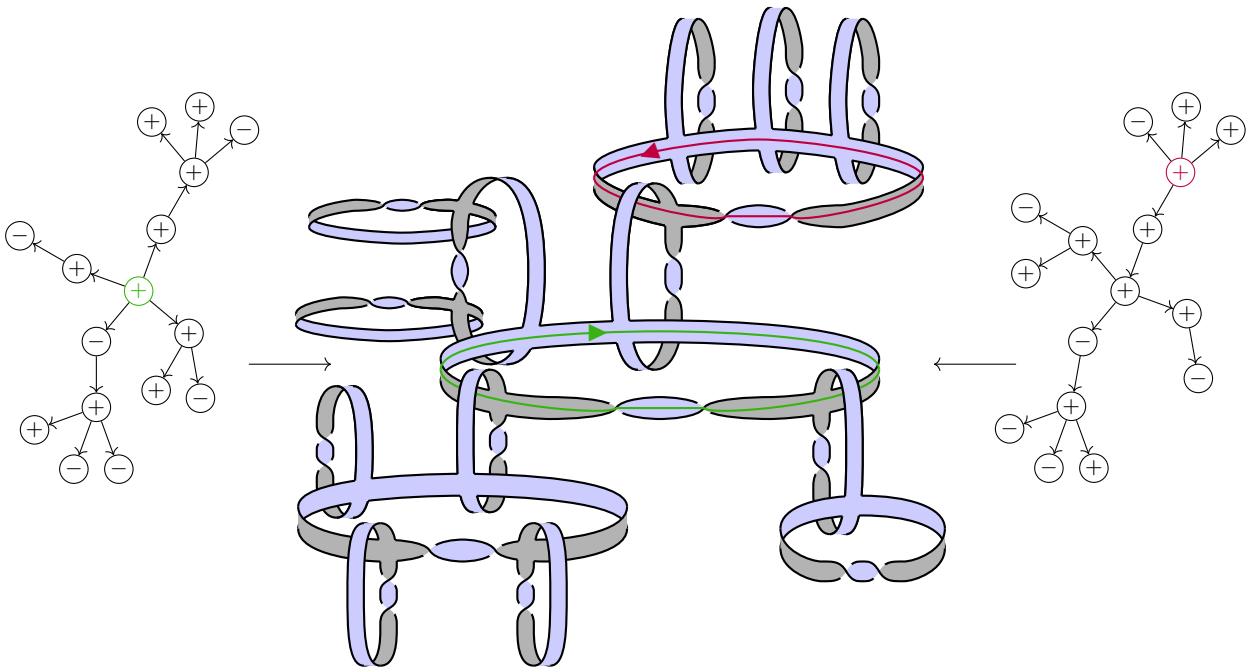

Figure 3.5: A 3D-view of a Hopf arborescent link and its construction from two different plane trees. The chosen orientation of the root of each tree is indicated on the coloured core of the matching Hopf band. The orientation of the plane is counter-clockwise.

b) Starting from this intersection point, we place k points $p_1, \dots, p_k$ on $\alpha(v)$ in this order. Then we draw on $\Sigma(T)$ a family of k arcs $\beta_1, \dots, \beta_k$ from ∂S to itself that correspond to those arcs that retract on $p_1, \dots, p_k$. Each such arc β_i intersects the collection $\mathcal{C}_T$ exactly at the point p_i.

c) For $i = 1, \dots, k$, perform the Hopf plumbing of a Hopf band $H(v_i)$ of sign $\ell(v_i)$ on top of $\Sigma(T)$ along the arc β_i. The resulting surface is $\Sigma(T')$.

d) Finally, for every i orient the core of $H(v_i)$ so that when going from $\alpha(v)$ to $\alpha(v_i)$, we follow this rule: if $\ell(v)$ is positive, one turns to the left (with respect to the orientation of $\Sigma(T)$), and if $\ell(v)$ is negative, one turns to the right (once again with respect to the orientation of $\Sigma(T)$), see Figure 3.6. The set $\mathcal{C}_{T'}$ is the union of $\mathcal{C}_T$ with $\alpha(v_1), \dots, \alpha(v_k)$.

Definition 3.3. *A **Hopf arborescent surface** is a surface $\Sigma(T)$ obtained from a plane tree T by this construction (see Figure 3.5 for an example). A **Hopf arborescent link** is the boundary of a Hopf arborescent surface (see the boundary of Figure 3.5 or the bottom left of Figure 3.1 for instance).*

Since Hopf bands are fibred and this property is preserved under plumbing, Hopf arborescent surfaces are fibres for their boundaries, and are thus of minimal genus. The arbitrary-looking rule that we use to orient the cores of the Hopf band in Step 2d is new and will turn out to be key for our proofs of Theorem B and Proposition 3.4.

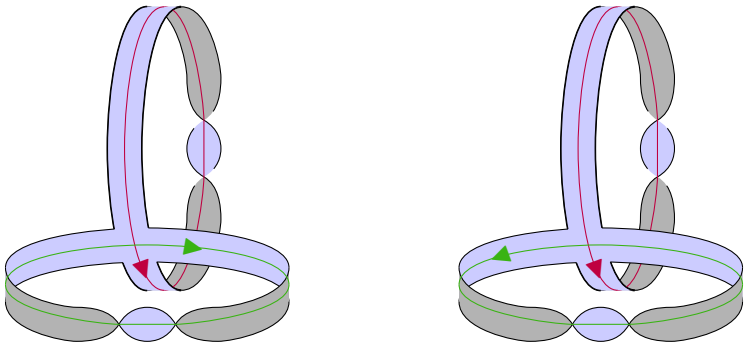

Figure 3.6: Orientation of the green core when its associated Hopf band (unsigned as it does not matter for the rule) is plumbed on top of the Hopf band with the red core.

3.3.3 Connections to other classes of knots

Here, we provide some additional background on Hopf arborescent links, their plumbing structure and their relations to other classes (arborescent links and fibred links). In particular, one may think that considering rooted trees and always plumbing the new Hopf bands on top of the surface is a strong restriction. We will show that this is not the case, i.e., that the family of surfaces and links obtained with less restriction on the sides on which one plumbs is the same as the family considered here.

Arborescent knots. Our definition differs from the classical definition(s) of arborescent links, e.g., in [16, 45] and their encoding via plane trees in three ways. Firstly, we restrict our attention to Hopf bands and not general unknotted annuli, hence the name Hopf arborescent links. This restriction is essential to our approach, and the minor theory we develop does not extend when one drops this assumption, as Remark 3.10 shows.

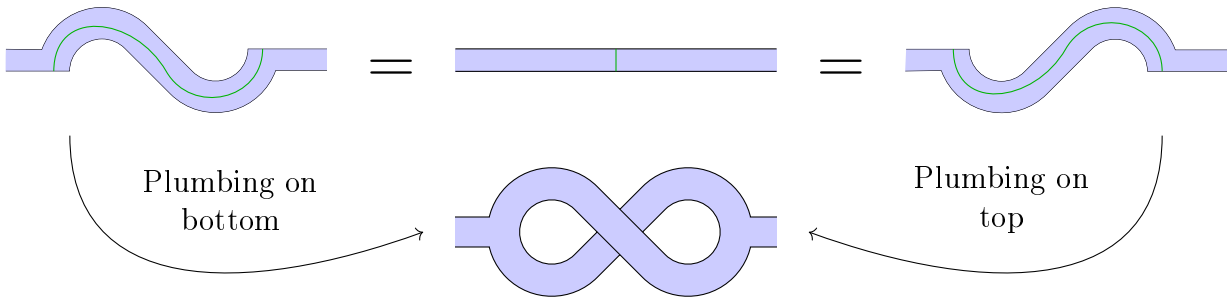

Figure 3.7: Surfaces obtained by plumbing a Hopf band inside or outside a Seifert surface are isotopic.

Secondly, an arborescent surface is, in general, defined as a surface that retracts on a collection of core curves whose intersection pattern is a tree. One can show that in this case, the surface can be reconstructed using an algorithm similar to the one we propose in

Section 3.3.2. The main difference is that when gluing a new band to the surface, it may be glued on top or below the surface. This difference could let one think that our construction is more peculiar than the standard one. However, it was remarked by Misev [94] that gluing a Hopf band on top or below an arc yields isotopic surfaces, see Figure 3.7. The only difference is that the orientation of the core curves of the glued band is then reversed. This change of orientation then implies that all vertices of the subtree subsequently glued to this band should be reversed to have the same surface. This argument shows that the set of surfaces we construct actually coincides with the set of surfaces obtained by the classical construction of arborescent links using Hopf bands only. The reason why we chose our presentation, i.e., always gluing on top when going away from the root of the tree, is that it is more practical with respect to our minor theory, and in particular it simplifies the proof of Lemma 3.8.

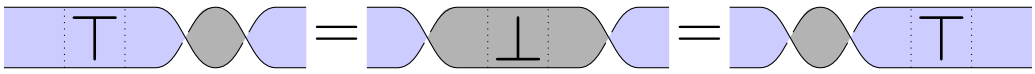

Figure 3.8: The surface does not depend on where the twists are.

Thirdly, the usual encoding of arborescent knots with plane trees [16, Chapter 12] is more detailed in that it puts the number of twists between each pair of adjacent edges around each vertex. Here, since bands have either 2 positive crossings or 2 negative crossings, Figure 3.8 shows that we do not need to specify where the twists are on the band, and thus indicating the sign of the band on each vertex is sufficient.

As noticed in Section 3.4.2 with the set $\mathcal{T}(L)$, it may happen that several trees yield the same isotopy class of oriented link and the same surface. In other words, our construction induces a well-defined map from plane trees to isotopy classes of surfaces in $\mathbb{S}^3$, but this map may not be injective, see Figure 3.5 for two different plane trees yielding the same link. This lack of injectivity is not an issue for the problems we address in this chapter, for the preimage of a given link or surface is finite. However, it is carefully studied by Bonahon and Siebenmann [16] who give a recipe to detect, in the more general case of arborescent links, when two $\mathbb{Z}$-labelled plane trees yield isotopic links.

More general plumbing structures. By construction, the tree T involved in defining a Hopf arborescent surface $\Sigma(T)$ is the intersection graph of the set of core curves $\mathcal{C}_T$. It may well happen that a surface retracts on different collections of Hopf curves [95, 16] whose intersection pattern is very different from T, and in particular is not a tree. In that case, the surface can still be obtained by a sequence of Hopf plumbings, but some plumbings are made on arcs that intersect the core of more than one Hopf band. On Figure 3.9 for example, the green curve on the right is indeed the core of a Hopf band, see the top of Figure 3.10 to picture how a tubular neighbourhood of the core is a band with two positive crossings.

3.3.4 Minors on surfaces, links, and plane trees

Since we focus our investigation on Hopf arborescent links, we define a stronger notion of minor that is well-tailored to these links. We say that a Hopf arborescent link L_1 is a **link-minor** of L_2 if there exist T_1 and T_2, two labelled plane trees such that $\Sigma(T_1)$ and $\Sigma(T_2)$ are

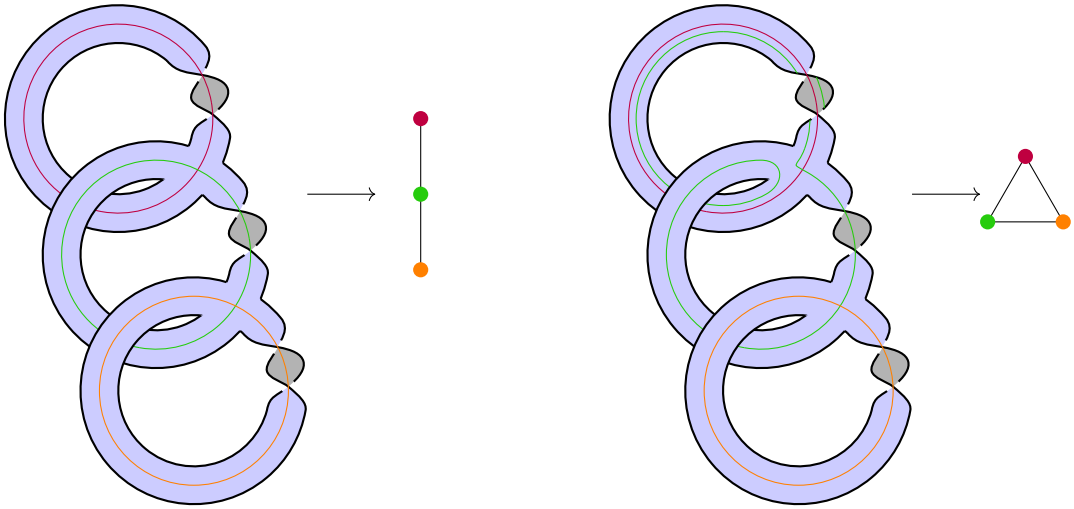

Figure 3.9: A Hopf arborescent plumbing with a cyclic plumbing graph.

canonical Seifert surfaces of L_1 and L_2 respectively, and $T_1 \hookrightarrow T_2$. The main result of this section is the following one: establishing that if L_1 is a link-minor of L_2, then the Seifert surface of L_1 is a surface-minor of the Seifert surface of L_2.

Proposition 3.4. *Let T_1 and T_2 be two plane trees such that T_1 admits a homeomorphic embedding into T_2. Then the Hopf arborescent surface $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

The proof relies on Lemmas 3.6, 3.7, and 3.8, which correspond respectively to the operations (i), (ii), and (iv) defining homeomorphic embeddings of trees. We first prove:

Lemma 3.5. *Let Σ be a surface and γ be an arc that is not boundary-parallel in Σ with both extremities in $\partial\Sigma$. Then cutting Σ along γ yields a surface $\Sigma' = \overline{\Sigma \setminus \gamma}$ such that $\Sigma' \preccurlyeq \Sigma$.*

Proof. By definition of Σ', there is a natural map $h : \Sigma' \to \Sigma$ that is injective except on $h^{-1}(\gamma) = \gamma_1 \cup \gamma_2$. Let T be a tubular neighbourhood of γ in Σ. Its boundary can be decomposed into t_1, t_2, two arcs isotopic to γ in Σ and two open arcs of $\partial\Sigma$. Isotope h within Σ so that $h(\gamma_1) = t_1$, $h(\gamma_2) = t_2$, and $h(\Sigma') \cap T = t_1 \cup t_2$. It follows that $h(\Sigma')$ is a subsurface of Σ such that $\Sigma \setminus \Sigma'$ is not an open disc since γ is not boundary-parallel. □

Lemma 3.5 essentially states that our surfaces behave well with respect to the surface-minor relation when cut along any essential arc. An important point is that cutting along an arc that is the diagonal of a plumbing rectangle merges two bands into one new band with two extra crossings that are either negative or positive depending on the diagonal, see Figure 3.10. So, cutting the plumbing of two positive Hopf bands along the diagonal that produces two negative crossings yields a positive Hopf band. Symmetrically, one can merge two negative Hopf bands into one negative by cutting along the other diagonal. Furthermore, when having a plumbing of two bands with opposite signs, one can merge them into a band with either sign depending on which cut one chooses.

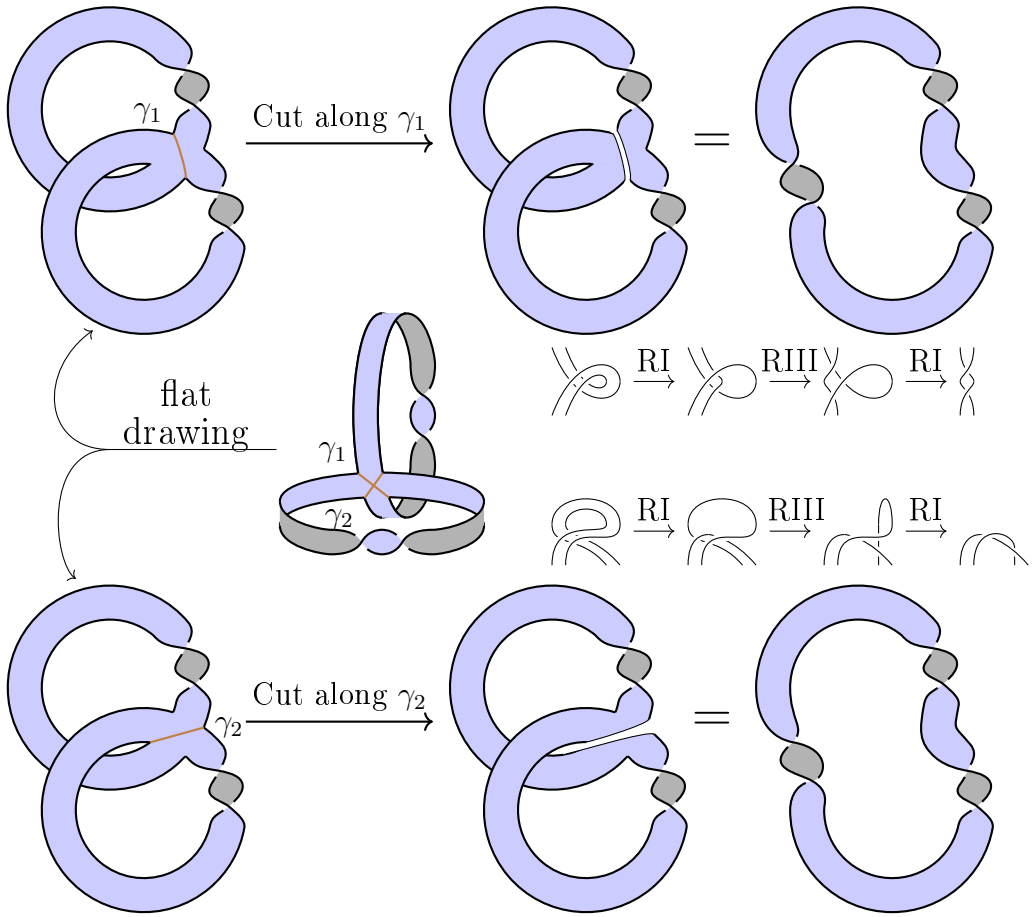

Figure 3.10: Cutting a Hopf plumbing of two positive Hopf bands along the diagonal γ_2 (resp. γ_1) induces 2 positive (resp. negative) crossings. The Reidemeister moves are shown to help understand which crossings are obtained. The top right surface is a Hopf band, while the bottom right one is a band with 3 positive full twists.

Lemma 3.6. *Assume that T_1 is obtained from T_2 by deleting a leaf. Then $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

Proof. Let D be the plumbing rectangle of the Hopf band H associated to the additional leaf v of T_2 compared to T_1. By definition, D has two sides γ_1, γ_2 in $\partial\Sigma(T_1)$. Thus γ_1 is also an arc of $\Sigma(T_2)$ with its extremities in $\partial\Sigma(T_2)$. By Lemma 3.5, $\Sigma(T_2)$ cut along γ_1 is an incompressible subsurface Σ' of $\Sigma(T_2)$. Furthermore, the remaining of H is a disc that can be isotoped into a neighbourhood of γ_2 so that $\Sigma(T_1) = \Sigma' \preccurlyeq \Sigma(T_2)$, see Figure 3.11. □

A very similar proof yields the following lemma.

Lemma 3.7. *Assume that T_1 is a plane tree whose root has only one child, and T_2 is the subtree rooted at that child. Then $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

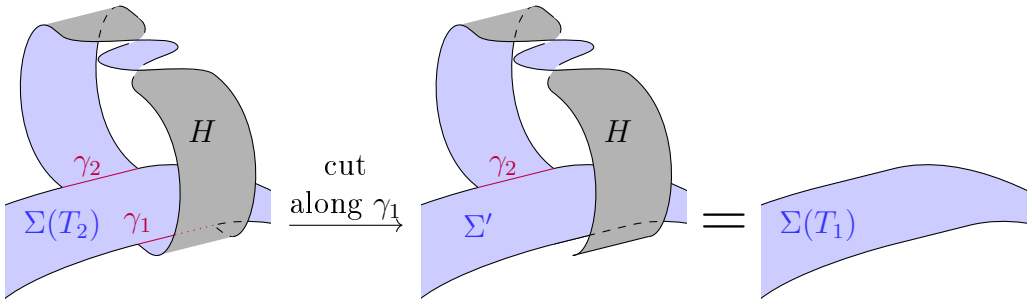

Figure 3.11: Deleting a leaf yields a surface-minor.

Proof. The proof is identical to the previous one: in that situation, $\Sigma(T_1)$ is obtained from $\Sigma(T_2)$ by plumbing a Hopf band, and cutting along one of the boundaries of the plumbing disc provides the needed incompressible subsurface, as in Figure 3.11. □

Lemma 3.8. *Assume that T_2 is a plane tree in which $u \to v \to w$ are three consecutive vertices where v has degree 2, and that T_1 is obtained from T_2 by contracting $u \to v \to w$ into a single edge $u \to w$, while preserving the labels of the endpoints. Then $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

Proof. By the construction of Hopf arborescent surfaces, the edge between u and v in T_2 corresponds to a plumbing rectangle D. It is important to recall here our orientation convention: if u is labelled positively, the cores $\alpha(u)$ and $\alpha(v)$ are oriented so that one turns to the left when going from u to v at the rectangle D, while if v is labelled negatively, one turns to the right, see Figure 3.6. Now, we consider two diagonal arcs γ_1 and γ_2 on the plumbing rectangle D as pictured in Figure 3.10. When cutting along such a diagonal arc, we obtain a new surface in which the cores $\alpha(u)$ and $\alpha(v)$ merge into a single core. However, their orientations might mismatch, depending on whether we cut along γ_1 or γ_2. We take the convention that γ_1 is the arc that preserves the orientations, while γ_2 induces an orientation mismatch, see Figures 3.10 and 3.12.

Now, let us first consider the case where the labels of u and v are the same. In this case, we consider the subsurface Σ' of $\Sigma(T_2)$ obtained by cutting along γ_1. This has the effect of merging the core curves $\alpha(u)$ and $\alpha(v)$ in a way that respects their orientations. However, it might seem that since each curve $\alpha(u)$ and $\alpha(v)$ corresponds to a Hopf band, merging them like that yields a band that twists too much. But a key observation is that cutting along γ_1 adds a twist between these two bands, as pictured in Figure 3.10, and this twist is negative when the bands are positive, while it is positive when the bands are negative (indeed, this is the reason for our orientation convention). Therefore, the resulting surface Σ' is exactly the same as the one corresponding to the tree T_1, and therefore $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$ by Lemma 3.5. See the top and bottom pictures of Figure 3.12 for an illustration.

Now, let us consider the case where the label of u is $+$ while the label of v is $-$. In that case, we consider the surface Σ' of $\Sigma(T_2)$ obtained by cutting along γ_2. This has the effect of merging the core curves $\alpha(u)$ and $\alpha(v)$ but with an orientation mismatch. We take the

convention that the resulting core curve α' is oriented by $\alpha(u)$, and therefore disagrees with the orientation of $\alpha(v)$ while it follows it. Since u and v are labelled with opposite signs, the two twists on their Hopf bands cancel out, but cutting along γ_2 adds a new positive twist, therefore we can consider α' as being the core curve of a positive Hopf band. Now, let us consider the plumbing rectangle D' corresponding to the edge between v and w. Due to the orientation mismatch, arriving at this rectangle from α', we are oriented in the direction opposed to the one we would arrive with if we were arriving from $\alpha(v)$. But due to the orientation convention, when going from $\alpha(v)$ to $\alpha(w)$ in $\Sigma(T_2)$ we turn to the right since v is negative, while when going from α' to $\alpha(w)$ in Σ' we turn to the left since α' is a positive band. Therefore, this effect cancels out the orientation mismatch, and Σ' coincides exactly with the surface $\Sigma(T_1)$ corresponding to the tree T_1. Therefore $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$ by Lemma 3.5. See the third row of Figure 3.12 for an illustration.

The same cancellation effect happens when the label of u is $-$ and the label of v is $+$: when cutting along γ_2 we have an orientation mismatch which is cancelled out by the fact that the new band is negative, and thus the orientation convention makes it turn in the opposite direction in the plumbing rectangle between v and w. Therefore, in that case $\Sigma(T_1)$ is also an incompressible subsurface of $\Sigma(T_2)$ thanks to Lemma 3.5. This is illustrated in the second row of Figure 3.12. □

As a corollary, contracting any path of T_1 into an edge whose labels match with the labels of the extremities of the path produces a tree T_2 such that $\Sigma(T_1) \preccurlyeq \Sigma(T_2)$.

Proof of Proposition 3.4. By definition, if T_1 admits a homeomorphic embedding into T_2, it can be obtained iteratively from T_2 by (i) removing a child leaf, (ii) removing a parent leaf, (iii) reducing a label, or (iv) contracting a path while preserving the labels of the endpoints. Since no two elements on the alphabet $\{+, -\}$ are comparable, case (iii) cannot happen. Then the cases (i), (ii) and (iv) are handled respectively by Lemma 3.6, Lemma 3.7 and Lemma 3.8. □

On the other hand, the Kruskal tree theorem, Theorem 2.5, directly yields the following proposition.

Proposition 3.9. *Hopf arborescent links are well-quasi-ordered under the link-minor relation.*

Proof. Take an infinite sequence $(L_n)_{n\in\mathbb{N}}$ of Hopf arborescent links, and let $(T_n)_{n\in\mathbb{N}}$ a sequence of plane trees such that for all $n \in \mathbb{N}$, $\Sigma(T_n)$ is a Seifert surface of L_n. Then, by Theorem 2.5, there exists $i < j$ such that T_i admits a homeomorphic embedding into T_2. Hence L_i is a link-minor of L_j. □

We can deduce Theorem B as a direct corollary of Proposition 3.4 and Proposition 3.9.

Proof of Theorem B. Take an infinite sequence $(\Sigma_n)_{n\in\mathbb{N}}$ of canonical Seifert surfaces of Hopf arborescent links. Then by Proposition 3.9, $\partial\Sigma_i \hookrightarrow \partial\Sigma_j$ for some $i < j$. By Proposition 3.4 we have $\Sigma_i \preccurlyeq \Sigma_j$ i.e., the surface-minor order is a well-quasi-order on Hopf arborescent surfaces. □

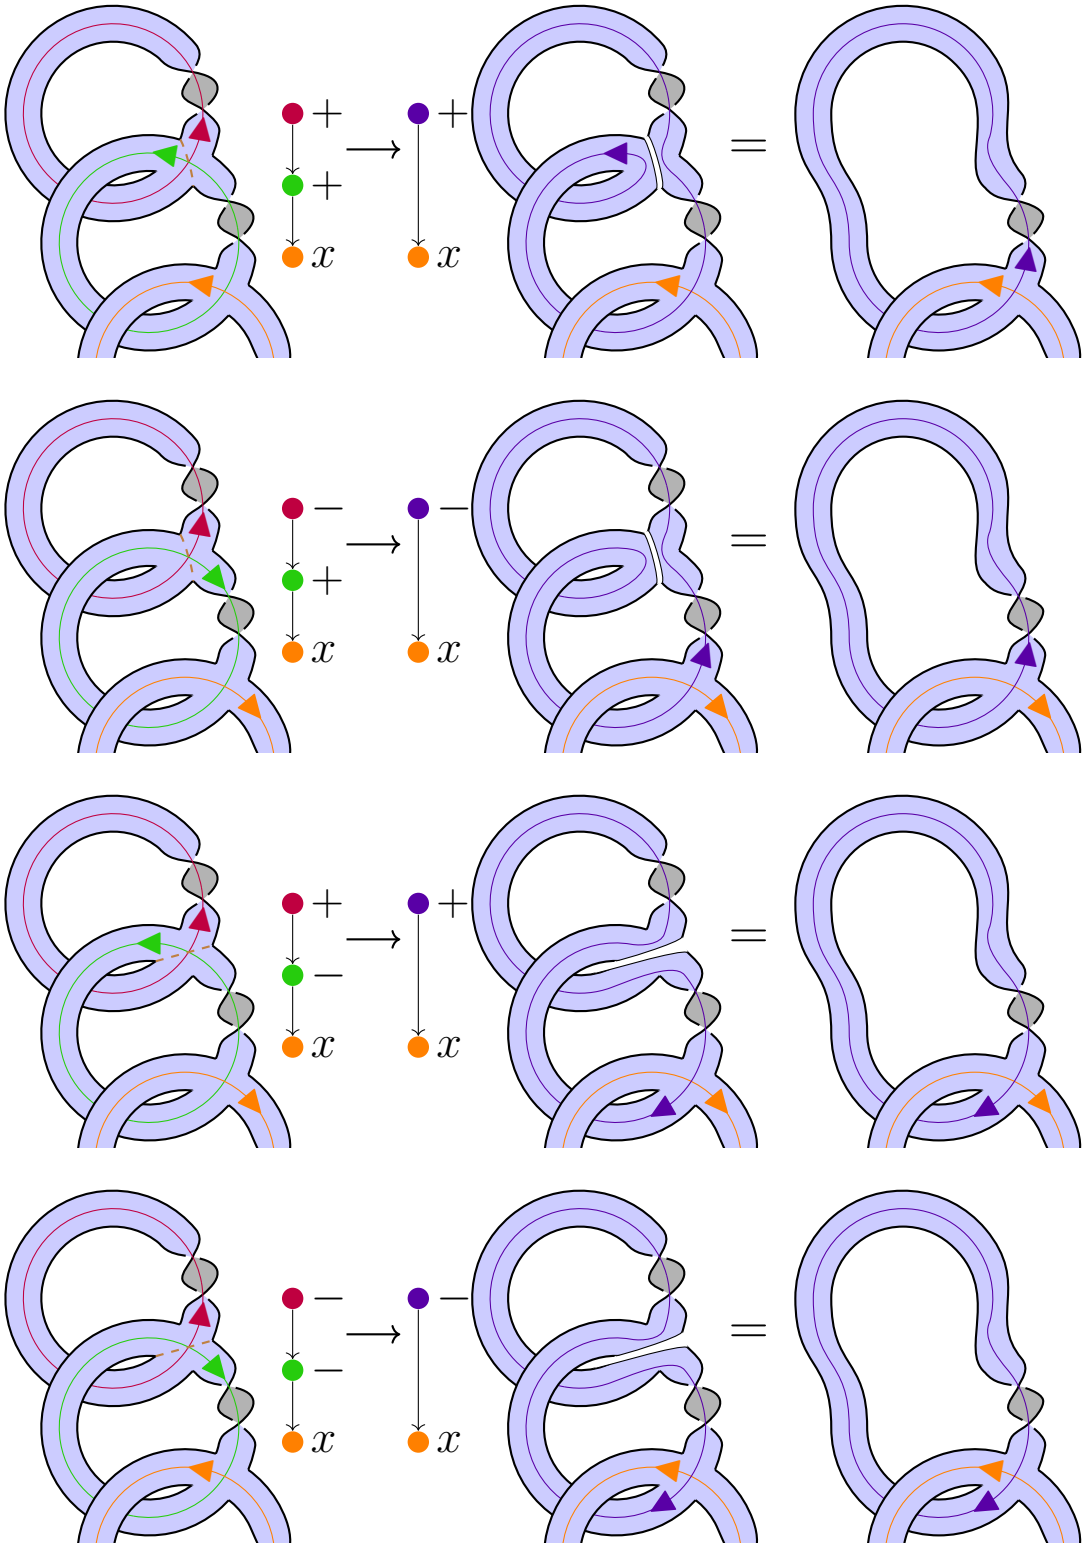

Figure 3.12: All cases of contraction of a 3-path to an edge preserving the labels of the endpoints.

Remark 3.10. *The proof of Proposition 3.4 highlights that the minor relation on the set of Hopf arborescent surfaces is more subtle and fragile than one might expect. Indeed, the cuts involved when taking an incompressible subsurface in the proof of Lemma 3.8 inevitably merge Hopf bands and thus one needs to be careful in order to control the number of resulting twists. In particular, the proof does not seem to generalise to the more general classes of surfaces obtained by plumbing bands with a bounded number of twists (even though everything works well at the level of trees).*

3.4 Decidability of the defect for Hopf arborescent links

3.4.1 Monotonicity of the genus defect

Now that we proved that link-minor is a well-quasi-order on the set of Hopf arborescent links, we want to highlight a property that is stable for this minor relation. Recall that the genus defect $\Delta_g(L)$ of an oriented link L is the difference $g(L) - g_4(L)$ between its classical genus and its 4-dimensional genus. The latter can be either in the topologically locally flat or in the smooth category. All statements in this section (and in particular Theorem A) hold in both categories. We reprove Lemma 6 of [10] in the form of Proposition 3.11 using the fact that link-minor implies that the associated Seifert are surface-minors.

Proposition 3.11. *The genus defect Δ_g is monotone on the family of Hopf arborescent links with respect to the link-minor relation, i.e., if L_1 is a link-minor of L_2, then $\Delta g(L_1) \leq \Delta g(L_2)$.*

We rely on the following lemma that highlights how the 4-genus behaves with respect to surface-minors. It uses a cut-and-paste construction and an Euler characteristic argument.

Lemma 3.12. *Let Σ be an oriented surface of $\mathbb{S}^3$ and Σ' be a surface-minor of Σ. If we write $L = \partial\Sigma$ and $L' = \partial\Sigma'$, then we have $g(\Sigma) - g_4(L) \geq g(\Sigma') - g_4(L')$.*

Proof. Seeing $\mathbb{S}^3$ as the boundary of the 4-ball $\mathbb{B}^4$, consider a surface S' in $\mathbb{B}^4$ such that $\partial S' = L'$ and $S' \cap \mathbb{S}^3 = L'$. gluing the remaining pieces of $\Sigma \setminus \Sigma'$ to S' along L' yields a surface S in B^4 such that $\partial S = L$, see Figure 3.13.

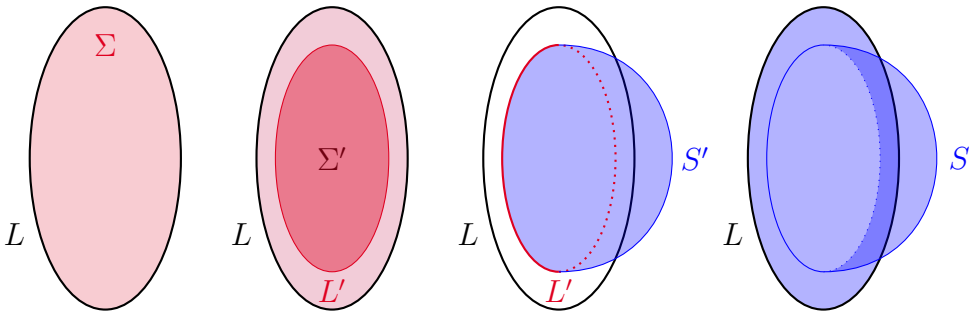

Figure 3.13: Illustration of the construction of Σ, Σ', L', S, and S'. Red surfaces are in $\mathbb{S}^3$ while blue one are considered to be in B^4.

By the definition of $g_4(L)$, we have $g_4(L) \leq g(S)$, and thus $g(\Sigma) - g_4(L) \geq g(\Sigma) - g(S)$. Furthermore, the genus of S is given by $g(S) = g(\Sigma) - g(\Sigma') + g(S')$. Thus one has $g(\Sigma) - g_4(L) \geq g(\Sigma') - g(S')$. Now if we assume S' to minimise the 4-genus over surfaces bounded by L', we conclude: $g(\Sigma) - g_4(L) \geq g(\Sigma') - g_4(L')$. □

As Hopf arborescent surfaces are of minimal genus for Hopf arborescent links, Lemma 3.12 can be used to prove Proposition 3.11.

Proof of Proposition 3.11. Let L and L' be two Hopf arborescent links such that L' is a link-minor of L. Consider Σ and Σ' the corresponding canonical Seifert surfaces. Since Hopf arborescent links are fibred, Σ is a Seifert surface of L of minimal genus (see Theorem 3.1), i.e. $g(L) = g(\Sigma)$ and similarly, $g(\Sigma') = g(L')$. By Proposition 3.4 one has $\Sigma \preccurlyeq \Sigma'$. Hence, by Lemma 3.12, one gets $\Delta_g(L) \geq \Delta_g(L')$. □

3.4.2 Proof of Theorem A

We first show that the link-minor relation can be decided using the decidability of link equivalence. For knots, equivalence can be tested as a combination of an algorithm that allows one to decide whether two 3-manifolds with boundary are homeomorphic [75, 91] and the Gordon-Luecke Theorem [49] that states that two knots are equivalent if their complements, which are 3-manifolds with boundaries, are equivalent. In the case of links, we additionally need to keep track of a longitude of each torus boundary component in the complement of the link. We refer to the survey of Lackenby [79, Section 2] for a summary of the techniques that allow us to prove the following theorem:

Theorem 3.13 (Link equivalence)**.** *Given two links L_1 and L_2, the problem of testing whether L_1 is ambient isotopic to L_2 is decidable.*

Given a Hopf arborescent link L, denote by $\mathcal{T}(L)$ the set of plane trees T whose associated Hopf arborescent surface $\Sigma(T)$ has L as oriented boundary. As a corollary, we obtain:

Lemma 3.14. *Given a Hopf arborescent link L, the set $\mathcal{T}(L)$ is computable.*

Proof. For increasing $k \in \mathbb{N}$, we enumerate all plane trees T_i with k vertices labelled by $\{-, +\}$ and store the trees T_i such that $\partial\Sigma(T_i)$ is isotopic to L, where we test isotopy using Theorem 3.13. If we find a k for which such a tree exists, we finish the enumeration for this value and return the stored trees. Indeed, plumbing n Hopf bands produces a surface with Betti number n. By Theorem 3.1, all the trees T_i such that $\partial\Sigma(T_i) = L$ produce surfaces $\Sigma(T_i)$ with the same genus, hence have the same number of vertices. As the entry is a Hopf arborescent link, there exists a tree T such that $\partial\Sigma(T) = L$. So the algorithm terminates. □

Alternatively, if one wants some control over the complexity of that algorithm, one can avoid blindly testing for increasing k by first computing the genus of the link [54, 91], or just computing an upper bound to it using, e.g., Seifert's algorithm, and then enumerating only the trees that produce surfaces up to that genus. From Lemma 3.14, we obtain:

Lemma 3.15. *Given two Hopf arborescent links L_1 and L_2, testing if L_1 is a link-minor of L_2 is decidable.*

Proof. Using Lemma 3.14, we compute $\mathcal{T}(L_1)$ and $\mathcal{T}(L_2)$. The trees in $\mathcal{T}(L_1)$ (resp. $\mathcal{T}(L_2)$) all have the same number k_1 (resp. k_2) of vertices. Then we exhaustively try out every possible path contraction to an edge and iterated leaf deletion on trees of $\mathcal{T}(L_2)$ such that the result is a tree with k_1 vertices and test whether it is equal to a tree of $\mathcal{T}(L_1)$. If such a test succeeds, we output yes, otherwise we return no. There is a finite number of trees in both $\mathcal{T}(L_1)$ and $\mathcal{T}(L_2)$ and a finite number of trees with k_1 vertices that homeomorphically embed into a tree of $\mathcal{T}(L_2)$. Hence that algorithm eventually terminates. Its correctness follows directly from the definition of link-minor. □

Finally we prove Theorem A by using the stability of the genus-defect by link-minor, the previous algorithms, and the well-quasi-order properties.

Proof of Theorem A. By Proposition 3.9, the order defined by link-minors is a well-quasi-order on the set of Hopf arborescent links. Hence, by Lemma 2.4, the set of Hopf arborescent links $\mathcal{H}_k$ that have defect at most k is characterised by a finite family $\mathcal{F}_k$ of forbidden minors. It follows, by Proposition 3.11 that $\Delta_g(L) \leq k$ if and only if for all f in $\mathcal{F}_k$, f is not a link-minor of L. Using Lemma 3.15 we test for each $f \in \mathcal{F}_k$ if f is a link-minor of L. If such a test succeeds, output no; otherwise, the input link has $\Delta_g(L) \leq k$. □

As said in the introduction, our proof is not constructive as it relies at its core on the existence of a set of forbidden minors for having defect at most k. This set of forbidden minors is not explicit and hard-coded in the algorithm. Furthermore, the sets of excluded minors will be different for the two different notions of defect (smooth and locally flat). It is likely that computing them is a topological challenge requiring arguments of different nature.

Theorem B provides the existence of a set of forbidden minors for having defect at most k but for a different and stronger definition of minors on links that relies only on the surface-minor relation on the Seifert surface and not the trees. However, deciding this relation, even by a brute force argument, seems challenging: in addition to the fact that no algorithm seems to be known for testing isotopy of surfaces, one would also need to control the complexity of the cutting arcs. Even with positive Hopf arborescent links only, this seems delicate [93].

3.5 Examples: Hopf arborescent links with non-trivial defect

It is not clear a priori that the topological and/or smooth defects of Hopf arborescent links are nonzero. For instance, our building block, the Hopf band, has both defects equal to 0 since it bounds an annulus in $\mathbb{S}^3$; which has genus 0. Furthermore, as mentioned in the introduction, when Hopf arborescent links are only made with positive Hopf bands, they belong to a class of links called *positive* links. This implies that they are *strongly quasi-positive* [123], and this implies in turn that their smooth 4-genus equals their 3-genus [122]. So for this class of links, the smooth defect is always zero. In contrast, in this section, we provide an example of a Hopf arborescent knot for which both the topological and the smooth defects are nonzero, and an argument explaining how to use this knot to provide examples with arbitrarily large defects.

Example 1. A Hopf arborescent link with non trivial topological defect is given by Baader, Feller, Lewark and Liechti [10, Example 4]. It consists in plumbing 6 positive Hopf bands in a latin cross-like pattern. Since it is positive, its smooth defect is zero. However, one can find a torus sub-surface of S_X such that the restricted Alexander polynomial vanishes. By Freedman's Disc Theorem [43], this implies that one can remove this subsurface and glue along its boundary a locally flat disc in the 4-ball. This means that the topological defect is at least 1. We refer to the article [10] for more details.

Example 2. By the aforementioned result of Rudolph, in order to obtain a Hopf arborescent link with nonzero smooth defect, we need to use both positive and negative Hopf bands. Here is such an example: the Hopf arborescent link obtained from the tree pictured in Figure 3.14 is in fact a knot, named 8_{10} in the Rolfsen table [120]. One can readily check from knot censuses (for example, from the Knot Atlas [1]) the smooth and topological genera of this knot, and we have $\Delta_g^{smooth}(8_{10}) = \Delta_g^{top}(8_{10}) = 2$. For completeness, we provide here another argument to explain why the defects are nonzero.

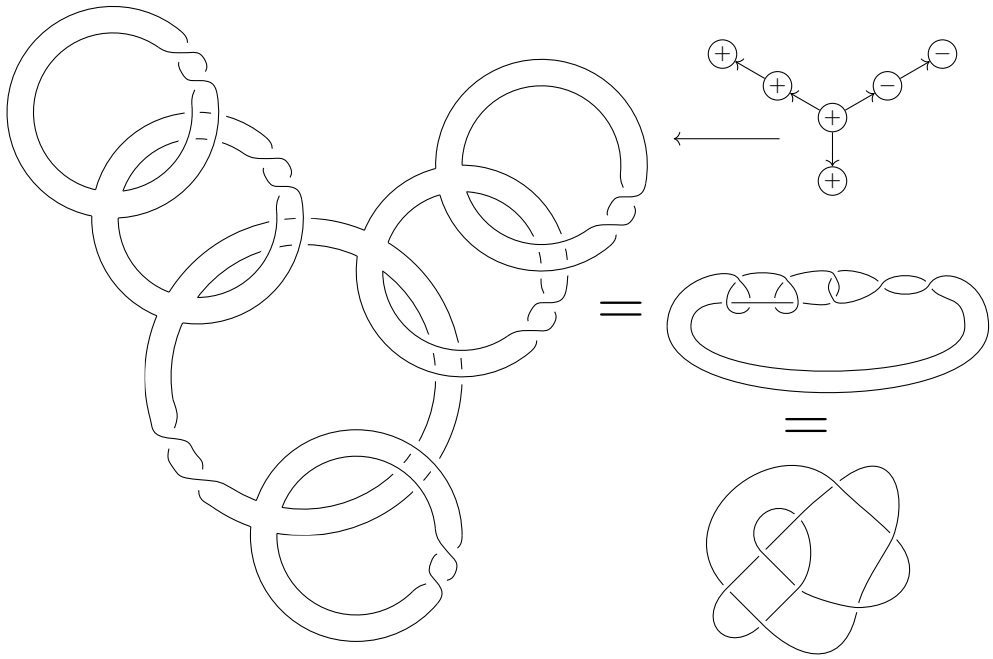

Figure 3.14: The Hopf arborescent knot 8_{10} (from the Rolfsen table).

The genus-minimising surface is the Hopf arborescent surface pictured in Figure 3.15, whose genus is 3, and with 1 boundary component. The red curve on the left of Figure 3.15 is the core of an embedded annulus, and its two boundaries are unlinked. Replacing this annulus by two discs glued on its two boundary components yields an immersed surface with genus 1 less, still bounded by $K_{8_{10}}$. The singularities of this immersion are of a particular type called ribbon singularities [120, p.225] (right side of Figure 3.15). Such singularities can be resolved in the 4-ball, and thus this implies that the Hopf arborescent surface can be pushed into a smoothly embedded surface in the 4-ball so that its smooth defect is at least 1.

Furthermore, for any link L, we have that $g_4^{smooth}(L) \geq g_4^{top}(L)$ so that its topological defect is also at least 1.

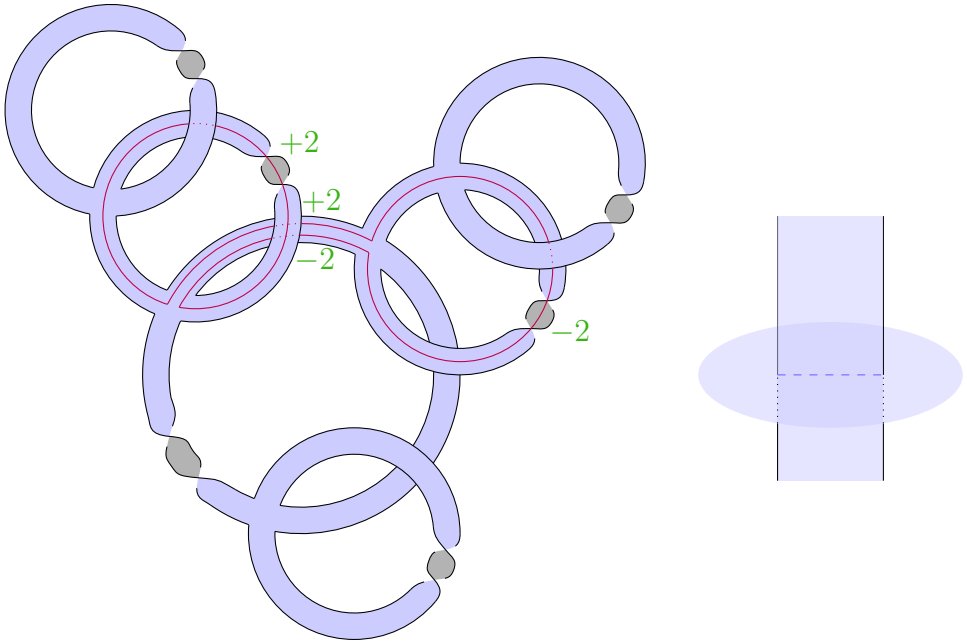

Figure 3.15: Left: A simple red curve on the surface associated to $\Sigma(8_{10})$. The intersection number in green show that the boundaries of the corresponding annulus are unlinked. Right: A ribbon intersection.

One can also check in the censuses that these two examples are hyperbolic.

Arbitrarily large defect. We can use the previous examples to create Hopf arborescent links with arbitrarily large (smooth or topological) defects. The following proposition proves that the defects coming from tree patterns add up if they are isolated enough.

Proposition 3.16. *Let L be an arborescent link associated to a plane tree T. Let $T_1, \dots, T_n$ be disjoint subtrees of T that are pairwise at distance at least two within T. Denote $L_i = \partial(\Sigma(T_i))$, then $\sum_{1 \leq i \leq n} \Delta_g(L_i) \leq \Delta_g(L)$.*

Proof. For each i, identifying the Hopf bands of the vertices of T_i within $\Sigma(T)$ provides a family of embeddings of the $\Sigma(T_i)$. These embeddings are disjoint by the distance assumption.

We can now repeat the proof of Lemma 3.12 and Proposition 3.11 with a family of links to obtain the desired inequality. Since the $\Sigma(T_i)$ are disjoint subsurfaces of $\Sigma(L)$, $\bigcup \Sigma(T_i)$ is a surface-minor of $\Sigma(L)$. As in the proof of Lemma 3.12, we define for each L_i, a surface S_i' in $\mathbb{B}^4$ that has L_i as its boundary. gluing each S_i' to the pieces of $\Sigma(T) \setminus \bigcup \Sigma(T_i)$ along its associated L_i yields a surface S in $\mathbb{B}^4$ such that $\partial S = L$.

Now we have $g(S) = g(\Sigma(T)) - \sum_i g(\Sigma(T_i)) + \sum_i g(S_i')$. It follows, by definition of g_4 that $g(\Sigma(T)) - g_4(L) \geq g(\Sigma(T)) - g(S) = \sum_i g(\Sigma(T_i)) - g(S_i')$. By Theorem 3.1, and taking each S_i' to be of minimal genus, $\Delta_g(L) = g(L) - g_4(L) \geq \sum_i g(L_i) - g_4(L_i) = \sum_i \Delta_g(L_i)$. □

Consider a Hopf arborescent link L associated to some tree T such that L has non zero defect. By identifying the roots of n copies of T with the leaves of a n-star graph (one vertex of degree n, n vertices of degree 1), we create a tree T' such that the link $L' = \partial\Sigma(T')$ has defect at least $n\Delta_g(L)$ by Proposition 3.16, therefore obtaining links with arbitrarily large topological and smooth defects.

Chapter 4

Tree-like decompositions of knots and spatial graphs

In this chapter, we initiate a thorough investigation of tree decompositions of knot diagrams, or more generally diagrams of spatial graphs, using ideas from structural graph theory. We define an obstruction on spatial embeddings that forbids low treewidth diagrams, and we prove that it is optimal with respect to a related width invariant that we develop. We then show the existence of this obstruction for knots of high compression-representativity, which include, for example, torus knots, providing a new and self-contained proof that those do not admit diagrams of low treewidth.

The results of this chapter were obtained with Arnaud de Mesmay and appeared in the Proceedings of the 39th International Symposium on Computational Geometry [A]. Additional remarks on this work are presented in Section 4.6.

4.1 Introduction

In recent years, many attempts have been made to attack seemingly hard problems, like the unknot recognition problem (see last part of Section 2.1.3), via the route of parameterized algorithms. In particular, recall that the **treewidth** of a graph is a parameter quantifying how close a graph is to a tree, and thus algorithmic problems on graphs of low treewidth can often be solved very efficiently using dynamic programming techniques on the underlying tree structure of the instance. The concept of **branchwidth**, which we will also use, is a somewhat equivalent width parameter on graphs and is always within a constant factor of treewidth [116]. It offers a similar insight as treewidth on graphs from a computational perspective, in addition to features that will interest us.

When one is provided with a knot diagram of low treewidth or branchwidth, one can therefore use dynamic programming techniques to solve seemingly hard problems very efficiently. While this approach has not yet been successful for unknot recognition beyond treewidth 2 [15], it has proved effective for the computation of many knot invariants, including: Jones and Kauffman polynomials [86] (which are known to be $\#P$-hard to compute in general [66]), HOMFLY-PT polynomials [19], and quantum invariants [87, 23]. Since any knot admits infinitely many diagrams, these algorithms naturally lead to the following question raised by Burton [22, p.2694], and Makowsky and Mariño [86, p.755]: do all knots admit diagrams of constant treewidth, or conversely, does there exist a family of knots for which *all* the diagrams have treewidth going to infinity. This question was answered recently by de Mesmay, Purcell, Schleimer, and Sedgwick [28] who proved that, among other examples, torus knots $T_{p,q}$ are such a family. The proof relies at its core on an intricate result of Hayashi and Shimokawa [57] on *thin position* of *multiple Heegaard splittings*.

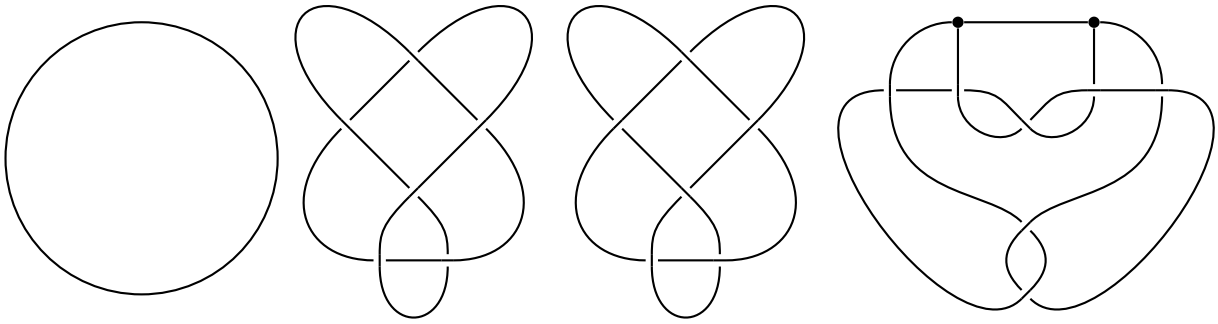

Figure 4.1: Diagrams of two trivial knots on the left, a bowline knot and a knotted spatial graph.

Our results. The main purpose of this work is to provide new techniques to characterise which knots, or more generally which **spatial graphs** (see for example Figure 4.1 for a knotted spatial graph), do not admit diagrams of low treewidth. Our starting point is similar to the one in [28]: we first observe that if a knot or a spatial graph admits a diagram of low treewidth, then there is a way to sweep $\mathbb{R}^3$ using spheres arranged in a tree-like fashion which intersect the knot a small number of times (Proposition 4.3). This corresponds roughly to a map $f : \mathbb{R}^3 \to T$ where T is a trivalent tree, where the preimage of each point interior to an edge is a sphere with a small number of intersections with the knot (we refer to Section 4.2 for the precise technical definitions of all the concepts discussed in this introduction). We call this a **sphere decomposition**[1], and the resulting measure, which is roughly the maximal number of intersections with the knot, is called the **spherewidth** of the knot.

Thus, in order to lower bound the treewidth of all the diagrams of a knot, it is enough to lower bound its spherewidth. We provide a systematic technique to do so using a perspective taken from structural graph theory. In the proof of the celebrated Graph Minor Theorem of Robertson and Seymour [118], handling families of graphs with bounded treewidth turns out

[1] Our sphere decompositions are different from the ones in [28] but functionally equivalent for knots.

not to be too hard [115], and in contrast, a large part is devoted to analysing the structure shared by graphs of large treewidth. There, a fundamental contribution is the concept of **tangle**[2]. We refer to Diestel [33] or Grohe [50] for nice introductions to tangles and their applications. Informally, a tangle of order k in a graph G is a choice, for each separation of size at most k, of a "big side" of that separation, where the highly connected part of the graph lies. In addition, there are some compatibility properties: in particular, no three "small sides" should cover the whole graph. Such a tangle turns out to be exactly the structure dual to branchwidth, in the sense that, for any graph G, the maximal possible order of a tangle is exactly equal to its branchwidth [116]. We develop a similar concept dual to sphere decompositions, which we call a **bubble tangle**. Informally, a bubble tangle of order k for a knot or spatial graph K is a map that, for each sphere intersecting K at most k times, chooses a "big side" indicating where the complicated part of K lies. There are again some compatibility conditions which add topological information to the collection of "small sides". Then our first result is the following.

Theorem C. *For any knot or spatial graph K, the maximum order of a bubble tangle for K is equal to the spherewidth of K.*

This theorem provides a convenient and systematic pathway to prove lower bounds on the spherewidth, and thus on the treewidth of all possible diagrams: it suffices to prove the existence of a bubble tangle of high order. However, making choices for the uncountable family of spheres with a small number of intersections with K, and then verifying the needed compatibility conditions is very unwieldy. Our second contribution is to provide a way to define such a bubble tangle in the case of knots (or spatial graphs) which are embedded on some surface Σ in $\mathbb{R}^3$. Given a surface Σ in $\mathbb{R}^3$, a **compression disc** is a disc properly embedded in $\mathbb{R}^3 \setminus \Sigma$ whose boundary is a non-contractible curve on Σ. The **compression-representativity of an embedding** of a knot or spatial graph K on a surface Σ in $\mathbb{R}^3$ is the smallest number of intersections between K and a cycle on Σ that bounds a compression disc (this was originally defined by Ozawa [103] under the name of representativity). The **compression-representativity** of a knot or spatial graph is the supremum of that quantity over all embeddings on surfaces . Our second theorem is the following.

Theorem D. *For any knot or spatial graph K embedded on a surface Σ in $\mathbb{R}^3$, there exists a bubble tangle of order $2/3$ of the compression-representativity of the embedding. Therefore, for any knot or spatial graph K, there exists a bubble tangle of order $2/3$ of the compression-representativity of K.*

Combining together Theorems C and D with Proposition 4.3 provides a large class of knots of high spherewidth, and our tools are versatile enough to apply to spatial graphs, while previous ones did not. In particular, observing that torus knots $T_{p,q}$ have high compression-representativity, we obtain the following corollary, which improves the lower bound obtained by [28], without relying on deep knot-theoretical tools.

Corollary 4.1. *A torus knot $T(p,q)$ has spherewidth at least $2/3\min(p,q)$, and thus any diagram of $T(p,q)$ has treewidth at least $1/3\min(p,q)$.*

[2]It turns out that the word *tangle* holds a completely different meaning in knot theory, and, to avoid confusion, in this thesis we will always use it with the graph theory meaning.

Related work and proof techniques. While there have been some recent works aiming at building bridges between combinatorial width parameters and topological quantities (for example, the aforementioned [28], but also [60, 61, 88] for related problems in 3-manifold theory), the main contribution in this chapter is that we dive deeper into the structural graph theory perspective via the concept of a tangle. The latter has now proved to be a fundamental tool in graph theory and beyond (see for example Diestel [30, Preface to the 5th edition]).

The duality theorem of Robertson and Seymour between branchwidth and tangles in [116] has been generalised many times since its inception. For example, in order to encompass other notions of decompositions and their obstructions [6, 85], to apply more generally to matroids [47], and to the wide-ranging setting of abstract separations systems [31, 32]. The key difference in our work, and why it does not fit into these generalisations, is that our notions of sphere decomposition and bubble tangles inherently feature the *topological* constraint of working with 2-spheres. This is a crucial constraint, as it would be easy to sweep any knot with width at most 2 if one were allowed to use arbitrary surfaces during the sweeping process. It should also be noted that in planar graphs, it was shown by Seymour and Thomas [133] that the separations involved in an optimal branch decomposition can always be assumed to take the form of 1-spheres, i.e., Jordan curves. This property led to the celebrated ratcatcher algorithm to compute the branchwidth of planar graphs in polynomial time [133] as well as to sphere-cut decompositions and their algorithmic applications (see for example [34]). Our sphere decompositions are the generalisation one dimension higher of these sphere-cut decompositions, and Theorem C identifies bubble tangles as a correct notion of dual obstruction for those. We believe that these notions could be of further interest beyond knots: in the study of graphs embedded in $\mathbb{R}^3$ with some topological constraints, e.g., linkless graphs [125].

The **representativity** (also called **facewidth**) of a graph embedded on a surface S is the smallest number of intersections of a non-contractible curve with that graph. Theorem D will not come as a surprise for readers accustomed to graph minor theory, as Robertson and Seymour proved a very similar-looking theorem in Graph Minors XI [117, Theorem 4.1]. They showed by exhibiting a tangle that the branchwidth of a graph embedded on a surface is lower bounded by its representativity. The key difference is that our notion of compression-representativity only takes into account the length of cycles bounding compression discs, instead of all the non-contractible cycles. Here again, this topological distinction is crucial to give a meaningful concept for knots, as for example the graph-theoretical representativity of a torus knot is zero. Due to this difference, the proof technique of Robertson and Seymour does not readily apply to prove Theorem D; instead we have to rely on more topological arguments.

From the knot theory side, there is a long history in the study of the "best" way to sweep a knot while trying to minimise the number of intersections in this sweepout. One of the oldest knot invariants, the bridge number, can be seen through this lens (see for example [131]). A key concept in modern knot theory, introduced by Gabai in his proof of the Property R conjecture [46], is the notion of **thin position** that quantifies more precisely the best way to place a knot to minimise its width. It is at the core of many advances in modern knot theory (see, for example, Scharlemann [127]). Recent developments in thin

position have highlighted that in order to obtain the best topological properties, it can be helpful to sweep the knot in a tree-like fashion compared to the classical linear one. This approach leads to definitions bearing close similarities to our sphere decompositions (this is one of the ideas behind generalised Heegaard splittings [129, 128], see also [57, 58, 135]). The concept of compression-representativity of a knot or a spatial graph finds its roots in the works of Ozawa [103], and Blair and Ozawa [13] who defined it under the simple name of representativity, taking inspiration from graph theory. They proved that it provides a lower bound on the bridge number and on more general linear width quantities. Our Theorem D strengthens their results by showing that it also lower bounds the width of tree-like decompositions. Furthermore, while specific tools have been developed to lower bound various notions of width of knots or 3-manifolds, we are not aware of duality theorems like our Theorem C. It shows that our bubble tangles constitute an obstruction that is, in a precise sense, the optimal tool for the purpose of lower bounding spherewidth.

Finally, an important inspiration for our proof of Theorem D comes from a seemingly unrelated breakthrough of Pardon [105], who solved a famous open problem of Gromov [51] by proving the existence of knots with arbitrarily high **distortion**. The distortion for two points on an embedded curve in $\mathbb{R}^3$ is the ratio between the intrinsic and Euclidean distance between the points. The distortion of the entire curve is the supremum over all pairs of points. The distortion of a knot is the minimal distortion over all embeddings of the knot. While this metric quantity seems to have nothing to do with tree decompositions, it turns out that the technique developed by Pardon can be reinterpreted in our framework. With our terminology, his proof amounts to first lower bounding the distortion by a constant factor of the spherewidth, and then defining a bubble tangle for knots of high representativity. The lower bound is nicely explained by Gromov and Guth [52, Lemma 4.2], where the simplicial map is similar to our sphere decompositions, up to a constant factor. Then our proof of Theorem D is inspired by the second part of Pardon's argument, with a quantitative strengthening to obtain the 2/3 factor, whereas his argument would only yield 1/2.

Organisation of this chapter. After providing background and defining the specifics for this chapter in Section 4.2, we prove Theorem C in Section 4.3, and Theorem D in Section 4.4. We cover some examples in Section 4.5. Section 4.6 contains an additional result that will be crucial in Chapter 5 and remarks about computational aspects of compression-representativity.

4.2 Specific preliminaries

Low-dimensional topology. We denote by $C(A)$ the connected components of a subset A of $\mathbb{S}^3$, and thus by $|C(A)|$ its number of connected components. Recall that we work in the PL setting. This allows us to avoid pathologies such as wild knots or pathological surfaces like the Alexander horned sphere [5]. An embedding of a compact topological space X into another one Y is said to be **proper** if it maps the boundary ∂X within the boundary ∂Y. A 3-dimensional version of the Schoenflies theorem guarantees that for any PL 2-sphere S embedded in $\mathbb{S}^3$, both components of $\mathbb{S}^3 \setminus S$ are balls (see for example [12, Theorem XIV.1]). Again, all these objects are considered equivalent when they are ambient isotopic. Knots and

links are a special instance of spatial graphs, and henceforth we will mostly focus on spatial graphs, generally denoted by the letter G. For technical reasons, it is convenient to thicken our embedded graphs as follows. A **thickened embedding** φ of a graph G is an embedding of G in $\mathbb{S}^3$ where each vertex is thickened to a small ball, two balls are connected by a polygonal edge if and only if they are adjacent in the graph G, and pairs of edges are disjoint. We emphasise that we do not thicken edges, which might be considered non-standard. We will also work with graphs embedded on surfaces which are themselves embedded in $\mathbb{S}^3$: such embeddings will also always be thickened, that is, vertices on the surface are thickened into small discs. *From now on, all the graph embeddings will be thickened, and thus to ease notations we will omit the word thickened.*

As mentioned in this introduction, for Σ a surface embedded in $\mathbb{S}^3$, a **compression disc** is a properly embedded disc D in $\mathbb{S}^3 \setminus \Sigma$ such that the boundary ∂D is a non-contractible curve on Σ. A **compressible curve** γ of Σ is the boundary of a compressing disc of Σ (the two compressible of the torus are pictured in Figure 4.2). For a spatial graph G embedded on a surface Σ in $\mathbb{S}^3$, the **compression-representativity** of G on Σ, written c-rep(G, Σ) is $\min\{|C(\alpha \cap G)| \mid \alpha \text{ compressible curve of } \Sigma\}$. We count connected components to correctly handle thickened vertices. The compression-representativity c-rep(G) of G is the supremum of c-rep(G, Σ) over all nested embeddings $G \hookrightarrow \Sigma \hookrightarrow \mathbb{S}^3$.

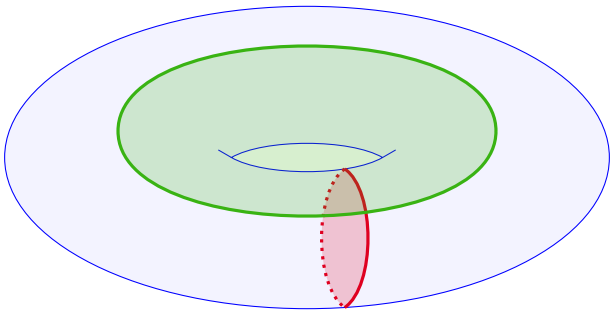

Figure 4.2: Two compression discs on the torus, whose boundaries are the two compressible curves of the torus.

In order to define spherewidth and bubble tangles, we require a precise control of the event when two spheres merge together to yield a third one, which is mainly encapsulated in the concept of double bubble. A **double bubble** is a triplet of closed discs (D_1, D_2, D_3) in $\mathbb{S}^3$, disjoint except on their boundaries, that they share: $D_1 \cap D_2 = D_1 \cap D_3 = D_2 \cap D_3 = D_1 \cap D_2 \cap D_3 = \partial D_1 = \partial D_2 = \partial D_3$, see Figure 4.3. Such a double bubble defines three spheres, which, by the PL Schoenflies theorem, bound three balls.

Two surfaces (resp. a knot and a surface) embedded in $\mathbb{S}^3$ are **transverse** if they intersect in a finite number of connected components, where the intersection is locally homeomorphic to the intersection of two orthogonal planes (resp. to the intersection of a plane and an orthogonal line). Likewise, we say that a surface is transverse to a ball if it is transverse to

its boundary. A surface is transverse to a graph if it is transverse to all the thickened vertices and edges it intersects. A double bubble is transverse to a graph or a surface if each of its three spheres is and if the vertices of the graph do not intersect the spheres on their shared circle ∂D_i. Intersections are **tangent** when they are not transverse, and a sphere S is said **finitely tangent** to a graph G embedded in $\mathbb{S}^3$ if they do not intersect transversely but the number of intersections $|E(G) \cap S|$ is still finite.

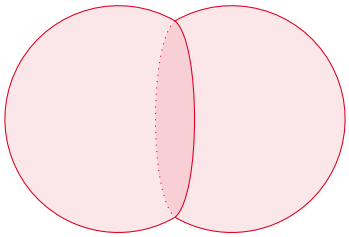

Figure 4.3: A double bubble: two spheres that intersect in a single disc.

4.2.1 Spherewidth

We now move on to the definition of sphere decompositions, which are the main way that we use to sweep $\mathbb{S}^3$, and thus knots and spatial graphs embedded within it using spheres.

Definition 4.2 (Sphere decomposition)**.** *Let G be a graph embedded in $\mathbb{S}^3$. A **sphere decomposition** of G is a continuous map $f : \mathbb{S}^3 \to T$ where T is a trivalent tree (vertices are either of degree 3 or 1) with at least one edge:*

- *For all $x \in L(T)$, $f^{-1}(x)$ is a point disjoint from G.*
- *For all $x \in V(T) \setminus L(T)$, $f^{-1}(x)$ is a double bubble transverse to G.*
- *For all x interior to an edge, $f^{-1}(x)$ is a sphere transverse or finitely tangent to G.*

The **weight** of a sphere S (with respect to G) is the number of connected components in its intersection with G. The **width** of a sphere decomposition f is the supremum of the weight of $f^{-1}(x)$ over all points x interior to edges of the tree T. The **spherewidth** of the graph G, denoted by $\mathrm{sw}(G)$, is the infimum, over all sphere decompositions f, of the width of f: $\mathrm{sw}(G) = \inf_{\mathbb{S}^3 \mapsto T} \sup_{x \in \mathring{e} \in E(T)} |C(f^{-1}(x) \cap G)|$. Therefore, a sphere decomposition is a way to continuously sweep $\mathbb{S}^3$ using spheres, which will occasionally merge or split in the form of double bubbles. The spherewidth is a measure of how well we can sweep a graph G using sphere decompositions. This is similar to the level sets given by a Morse function, but note that our double-bubble singularities are not of Morse type, and those are key to the proof of Theorem C. Figure 4.4 pictures a 2D representation of a sphere decomposition: some spheres of the sweepout, each one being preimages of an inner point in an edge, are pictured by grey circles (represented as rectangles here), and double bubbles are represented by rectangular θ shaped subsets of the plane.

The point of using thickened embeddings instead of usual ones is that this allows disjoint spheres of a sphere decomposition to intersect the same vertex of a graph embedding. This

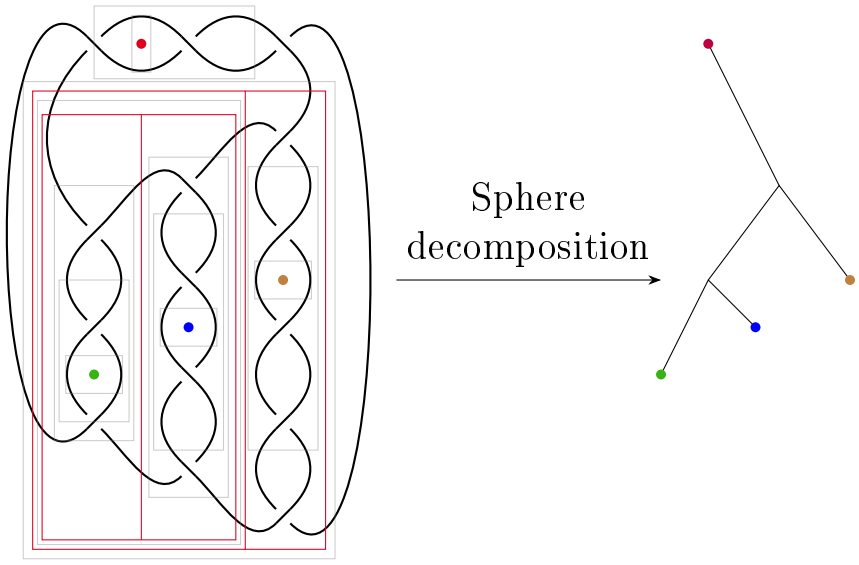

Figure 4.4: A representation in the plane of sphere decomposition of width 4 of a knot.

is motivated by the following proposition, which provides a bridge between sphere decompositions and tree decompositions of diagrams of knots and spatial graphs.

Proposition 4.3. *Let G be a knot or a graph embedded in $\mathbb{S}^3$, and D be a diagram of G. Then the spherewidth of G is at most twice the treewidth of D.*

The proof is very similar to that of Lemma 3.4 in [28].

Proof of Proposition 4.3. If D is a tree, then the spherewidth of G is one, and the proposition is trivial. So we assume henceforth that G contains at least one cycle. We will refrain from providing a precise definition of treewidth (we refer to Diestel [30]) as we will rely on a variant with more structure, more adapted to planar graphs: sphere-cut decompositions [34], which we first introduce. Given a planar graph D, a **sphere-cut decomposition** is a trivalent tree T, whose leaves are in bijection with the edges of D, and a family of Jordan curves γ_e parametrised by the edges of T. An edge e of T partitions the leaves of T and thus the edges of D into two subsets E_1^e and E_2^e. Each Jordan curve γ_e is required to intersect D only at its vertices (such a curve is often called a **noose**), and thus partitions the edges of D into two subsets: we require that this partition matches the partition E_1^e and E_2^e. The **weight of a Jordan curve** $w(\gamma_e)$ is the number of vertices that it intersects. The width of a sphere-cut decomposition $w(T, (\gamma_e))$ is the maximum weight $w(\gamma_e)$ over all the edges of T, and the sphere-cut width is the minimum width over all sphere-cut decompositions. Given a graph D that is not a tree, if we denote by k its treewidth, it is known that the sphere-cut width of D is at most k. This follows from the fact that the branchwidth of D is at most its treewidth [116] and the results of Seymour and Thomas on bond carving

decompositions [133] (see for example [65, Lemma 2.2] for the exact statement that we use connecting branchwidth and sphere-cut width).

We now work with a sphere-cut decomposition $(T, (\gamma_e))$ achieving the optimal width. There are two types of vertices in D: those corresponding to vertices of G and those corresponding to crossings of G in the projection. Since G is a thickened embedding in $\mathbb{S}^3$, it makes sense to also thicken D, i.e., we thicken a bit the vertices of the planar graph D so that they are small discs, and so that the Jordan curves γ_e intersect the interior of these discs. Then, by perturbing them a little and, if necessary, removing bigons, we can assume that the Jordan curves are all pairwise disjoint while still intersecting the graph D only at the vertex discs. Now, since D is a planar projection of G, each Jordan curve can be lifted to a sphere in $\mathbb{S}^3$ by capping it off with one disc above and one disc below the projection. The resulting spheres S_e are pairwise disjoint and intersect the graph G in two different ways: whenever γ_e traverses one of the vertices of G, S_e intersects the corresponding thickened vertex; and whenever γ_e traverses one of the crossings of G, S_e intersects transversely one or two of the edges of G (see Figure 4.5).

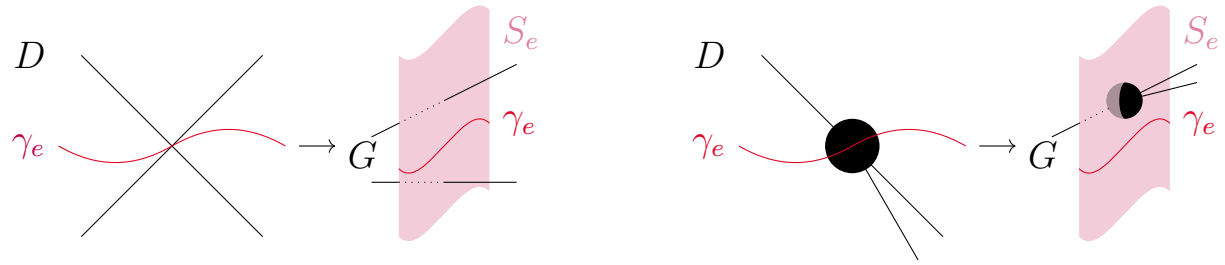

Figure 4.5: The two cases for the intersection between γ_e and D and their lift in $\mathbb{S}^3$.

Therefore, the weight of a sphere S_e is upper bounded by $c_1^e + 2c_2^e$, where c_1^e is the number of vertices of G traversed by γ_e, and c_2^e is the number of crossings of G traversed by γ_e. The spheres S_e are used as the backbone of a sphere decomposition, there just remains to interpolate between them:

- A sphere S_e for e adjacent to a leaf of T encloses exactly one edge of D, therefore it bounds a G-trivial ball. It is straightforward to define a continuous family of nested spheres of width 2 and converging to a point disjoint from G, which together sweep such a G-trivial ball.

- A vertex v of T is adjacent to three edges, corresponding to three spheres S_1, S_2 and S_3. The space inbetween these three spheres is homeomorphic to a solid pair of pants P, i.e., $\mathbb{S}^3$ with three balls removed. We look at $P \cap G$ and observe that since every edge of D is in bijection with one of the leaves of T, it is enclosed by a ball in the previous item and thus does not appear in $P \cap G$. Therefore, the only components that we see in $P \cap G$ are preimages under the projection map of thickened vertices or crossings. By construction, both these preimages are topologically trivial (they do not involve any knotting). The thickened vertices connect one, two, or the three boundaries of P. So we can define a double bubble DB in the middle of P and transverse to G, and three families of nested spheres interpolating between the three spheres of DB and S_1, S_2,

and S_3, with the following property: each of these interpolating spheres intersects at most twice each thickened vertex connecting the three boundaries, and at most once any of the other preimages (see Figure 4.6).

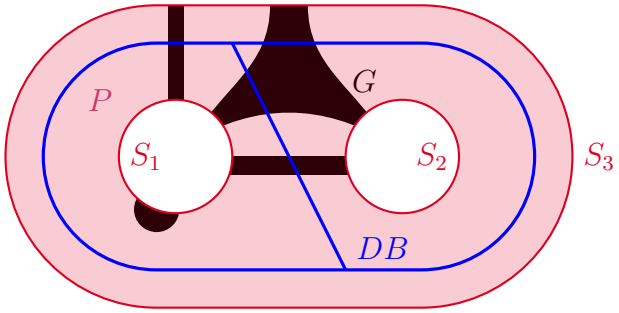

Figure 4.6: Definition of a double bubble DB within a solid pant P.

Each sphere involved in the previous two items has weight at most $2c_1^e + 2c_2^e$ for some Jordan curve γ_e, and thus the width of the sphere decomposition is upper bounded by twice the sphere-cut width. We obtain a sphere decomposition of weight at most twice the treewidth of D, which concludes the proof. □

4.2.2 Bubble tangle

Bubble tangles are our second main concept in this chapter. They will constitute an obstruction to spherewidth, by designating, for each sphere in $\mathbb{S}^3$ not intersecting the graph too many times, the side of the sphere that is easy to sweep. We first observe that some balls have to be easy to sweep: intuitively, this will be the case of any unknotted segment or empty ball (see Figure 4.7). Let G be a graph embedded in $\mathbb{S}^3$. A closed ball B in $\mathbb{S}^3$ is said to be G**-trivial** if its boundary is transverse to G and one of the following holds (where $B(0,1)$ is the unit ball of $\mathbb{R}^3$):

- $B \cap G = \varnothing$.
- $B \setminus G$ is homeomorphic to $B(0,1) \setminus [-1,1] \times \{(0,0)\} \subset \mathbb{R}^3$.
- $B \setminus G$ is homeomorphic to $B(0,1) \setminus [-1,0] \times \{(0,0)\} \subset \mathbb{R}^3$.

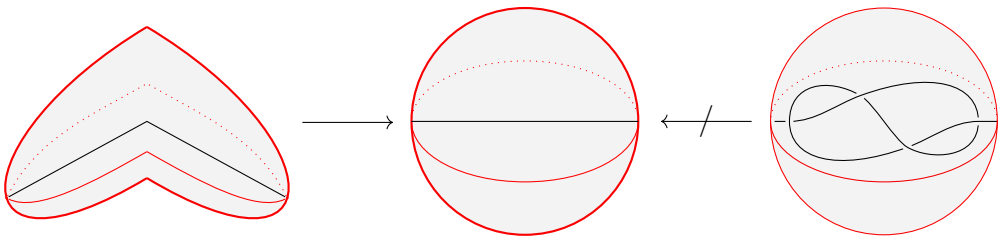

Figure 4.7: Representation of a G-trivial ball and a non G-trivial ball.

We can now introduce bubble tangles.

Definition 4.4. *Let G be an embedding of a graph in $\mathbb{S}^3$ and $n \in \mathbb{N}$. A **bubble tangle** $\mathcal{T}$ **of order** $n \geq 2$, is a collection of closed balls in $\mathbb{S}^3$ such that:*

(T1) For every closed ball B in $\mathcal{T}$, $|C(\partial B \cap G)| < n$.

(T2) For every sphere S in $\mathbb{S}^3$ transverse to G, if $|C(S \cap G)| < n$ then exactly one of the two closed balls $\bar{B}_1$ is in $\mathcal{T}$ or $\bar{B}_2$ is in $\mathcal{T}$, where $\mathbb{S}^3 \setminus S = \{B_1, B_2\}$.

(T3) For every triple of balls B_1, B_2 and B_3 induced by a double bubble transverse to G, $\{B_1, B_2, B_3\} \not\subset \mathcal{T}$.

(T4) For every closed ball B in $\mathbb{S}^3$, if B is G-trivial and $|C(\partial B \cap G)| < n$, then $B \in \mathcal{T}$.

For every transverse sphere S such that $|C(S \cap G)| < n$, a bubble tangle chooses one of the two balls having S as the boundary. We think of the ball in $\mathcal{T}$ as being a "small side", since (T4) stipulates that balls containing trivial parts of G are in $\mathcal{T}$, while the other one is the "big side". Then the key property (T3) enforces that no three small sides forming a double bubble should cover the entire $\mathbb{S}^3$ ((T2) and (T3) are illustrated in Figure 4.8).

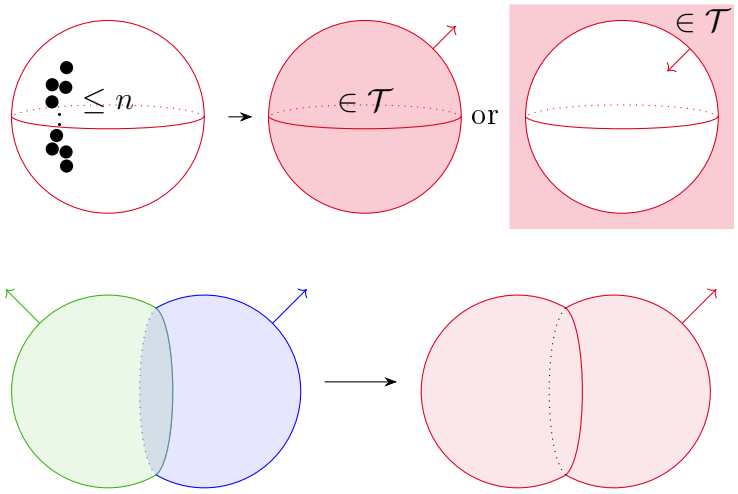

Figure 4.8: Illustration of (T2) (top) and (T3) (bottom), small sides are indicated by arrows pointing outward.

Remark 4.5. *Tangles in graph theory are often endowed with an additional axiom, specifying that small sides should be stable under inclusion (see e.g., [47, Axiom (T3A)]). Our bubble tangles are weaker in the sense that we do not enforce this axiom, but still strong enough to guarantee duality (Theorem C) and the connection to compression-representativity (Theorem D). Whether such an axiom can be additionally enforced in our definition of bubble tangle while preserving these properties is left as an open problem. As explained by a discussion of Section 4.6 and Proposition 4.20, this property is satisfied by bubble tangles provided by Theorem D.*

4.3 Obstruction and duality

In this section, we prove Theorem C: given a graph G embedded in $\mathbb{S}^3$, the highest possible order of a bubble tangle is equal to the spherewidth of G. In the following, G is an embedding of a graph in $\mathbb{S}^3$ and the order of all bubble tangles that we consider is at least 3, the theorem being trivial otherwise. The proof is split into two inequalities: Proposition 4.6 and Proposition 4.9 which together immediately imply Theorem C.

4.3.1 Bubble tangles as obstruction.

We first show that a bubble tangle of order k and a sphere decomposition of width less than k cannot both exist at the same time.

Proposition 4.6. *Let G be an embedding of a graph in $\mathbb{S}^3$. If G admits a bubble tangle $\mathcal{T}$ of order k, then $sw(G) \geq k$.*

The proof of this proposition is similar to its graph-theoretical counterparts showing that tangles are an obstruction to branchwidth (see, e.g., [116]). The main difference with these proofs lies in the continuous aspects of our sphere decomposition, which we control using Lemmas 4.7 and 4.8.

Let S and S' be two disjoint spheres in $\mathbb{S}^3$. Then $\mathbb{S}^3 \setminus (S \cup S')$ has three connected components: two balls and a space I homeomorphic to $\mathbb{S}^2 \times [0,1]$. The spheres S and S' are said to be **braid-equivalent** if $(I \cup S \cup S') \setminus G$ is homeomorphic to $S_k \times [0,1]$ where S_k is the 2-sphere with k holes (see Figure 4.9). The intuition behind this definition is that it means that G forms a braid between S and S'. The following lemma explains how braid-equivalent spheres interact with a bubble tangle.

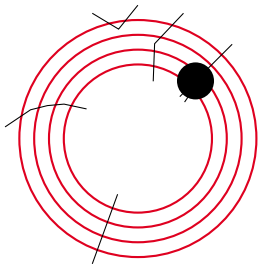

Figure 4.9: The three innermost spheres are braid-equivalent, not the fourth one.

Lemma 4.7. *Let $\mathcal{T}$ be a bubble tangle and S, S' be two braid-equivalent spheres. Let us write $\mathbb{S}^3 \setminus S = \{B_1, B_2\}$ and $\mathbb{S}^3 \setminus S' = \{B'_1, B'_2\}$ such that $B_1 \subset B'_1$. If $B_1 \in \mathcal{T}$ then $B'_1 \in \mathcal{T}$.*

Proof of Lemma 4.7. The idea of the proof is to show that if S and S' are braid-equivalent, we can cover the space $I = B'_1 \setminus \mathring{B}_1$ by G-trivial balls in order to "grow" S into S', which will yield the result by invoking (T3) inductively. Notice that braid equivalence implies that S and S' are transverse to G, and that $|C(S \cap G)| = |C(S' \cap G)| = k$ where k is less than the order of $\mathcal{T}$ as $S \in \mathcal{T}$; hence, by (T2) either $B'_1 \in \mathcal{T}$ or $B'_2 \in \mathcal{T}$. Let $h : S_k \times [0,1] \to I \setminus G$

be a homeomorphism induced by the braid equivalence where S_k is the 2-sphere with k holes $t_1, \dots, t_k$, and let us work on $S_k \times [0,1]$.

Let $D_1, \dots, D_k$ be k disjoint closed discs on $\mathbb{S}^2$, each one covering a hole of S_k: for all $i \in [\![1,k]\!]$, $\mathring{D}_i \smallsetminus t_i$ is an open annulus. For $i \in [\![1,k]\!]$ notice that $h((D_i \smallsetminus t_i) \times [0,1])$ is homeomorphic to $B(0,1) \smallsetminus [-1,1] \times \{(0,0)\}$ (see Figure 4.10).

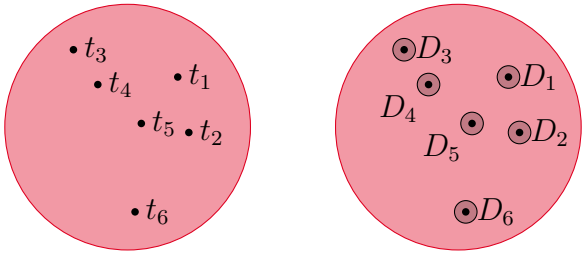

Figure 4.10: discs $D_1, \dots, D_k$ over holes of $\mathbb{S}^2$.

We first extend h^{-1} to a continuous function $\varphi : I \to \mathbb{S}^2 \times [0,1]$ such that $\varphi(D_i \times [0,1])$ is a closed ball of S^3 that is G-trivial, and belongs to $\mathcal{T}$ by $(T4)$. It suffices to extend h^{-1} on G. Notice that G is a union of balls and edges with endpoints on these balls. On I, the closure of each $h(S_k \times \{x\})$ for $x \in [0,1]$ naturally extends h^{-1} to $\partial G \cap I$ by continuity: send the boundary components of $h(S_k \times \{x\})$ to $\{t_i\} \times \{x\}$ with φ when the boundary component extends the image of a neighbourhood of $\{t_i\} \times \{x\}$. All that remains is to deal with the balls of the thickened embedding. Let B be a vertex ball of G such that $I \cap B \neq \varnothing$ and let X be one of the connected components of $B \cap I$. We layer X continuously with discs such that each disc is bounded by a circle induced by one of the boundary components of $h(S_k \times \{x\})$. Then we send each disc to $\{t_i\} \times \{x\}$ with φ.

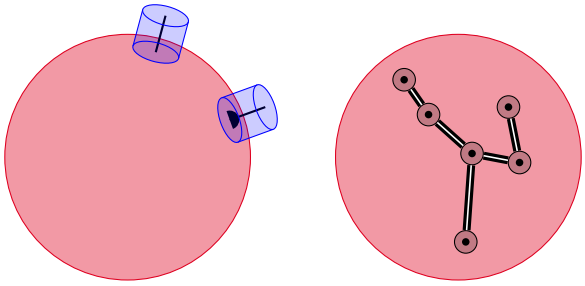

Figure 4.11: A G-trivial ball covering a strand and an example of tree connecting the discs D_i.

Now $\varphi^{-1}(D_i \times [0,1])$ is a closed ball β_i in $\mathbb{S}^3$ by construction and is G-trivial as $\varphi^{-1}(D_i \times [0,1]) \smallsetminus G = h((D_i \smallsetminus t_i) \times [0,1])$. Hence β_i belongs to $\mathcal{T}$. Notice that B_1 intersects β_i on a disc D_i, so that these balls induce a double bubble. It follows by (T3) that $B_1 \cup \beta_1$ belongs to $\mathcal{T}$. As the balls D_i are disjoint, $(B_1 \cup \bigcup_{j \in [\![1,i]\!]} \beta_j) \cap \beta_{i+1} = D_{i+1}$. Hence, these balls induce a double bubble. By induction on i, we obtain that $B_1 \cup \bigcup_{j \in [\![1,k]\!]} \beta_j$ is in $\mathcal{T}$.
At that point, we simplify the notations, and indexes i, j will span the integer interval $[\![1,k]\!]$ in $\bigcup$. We have that $(B_1 \cup \bigcup_j \beta_j) \cap S' = h(\bigcup_j D_j \times \{1\})$. Now, we consider a tree T connecting the discs D_i as in the Figure 4.11. Then, we thicken the edges of that tree to replace each

edge e joining two discs with a band r_e. It follows that each $h(r_e \times [0,1])$ is a closed ball in $\mathbb{S}^3$ that intersects $B_1 \cup \bigcup_j \beta_j$ on a U shaped disc and is disjoint from G (thus it is G-trivial). Using both (T3) and (T4) on each of the balls induced by the bands, we add balls to $B_1 \cup \bigcup_j \beta_j$ to get $B \in \mathcal{T}$ such that B intersects S' on a single connected component: $h((\bigcup_{e\in T} r_e \cup \bigcup_i D_i) \times \{1\})$. Finally, we add to B the ball $\overline{B'_1 \setminus B}$ disjoint from G which intersects S' on a disc : the complement of $h((\bigcup_{e\in T} r_e \cup \bigcup_i D_i) \times \{1\}$. Since its intersection with B is a disc, they induce a double bubble and $B'_1 \in \mathcal{T}$ by (T3). □

In the following, we will assume that there exists a bubble tangle $\mathcal{T}$ of order k and a sphere decomposition $f : \mathbb{S}^3 \to T$ of G of width less than k in order to reach a contradiction. Let $e = (u, v) \in E(T)$ be an edge and x be a point of e so that $f^{-1}(x)$ is transverse to G. Notice that x cuts T in two trees : $T_u(x)$ and $T_v(x)$, where $T_u(x)$ is the tree containing the endpoint u. By definition, $f^{-1}(x) = S$ is a sphere in $\mathbb{S}^3$ such that $|C(G \cap S)| < k$. It follows by (T2) that exactly one of $f^{-1}(T_u(x))$ or $f^{-1}(T_v(x))$ belongs to $\mathcal{T}$. We define an **orientation** $o : T \to V(T)$ induced by $\mathcal{T}$ as follows: if $f^{-1}(x)$ is transverse to G, $o(x) := v$ if $f^{-1}(T_u(x)) \in \mathcal{T}$, or $o(x) := u$ if $f^{-1}(T_v(x)) \in \mathcal{T}$. In other words, at a point x where $f^{-1}(x)$ is transverse to G, the orientation o orients x outwards, toward the "big side". If $f^{-1}(x)$ has a tangency with G, note that for any close enough neighbour y of x, $f^{-1}(y)$ is transverse to G, and we define $o(x) := o(y)$, making an arbitrary choice if needed. As we consider edges of the tree T to be intervals, we will use interval notations: we write $[u, v]$ for the edge (u, v), and more generally $[x, y]$ to describe all the points on the edge between x and y. We say that an orientation o is **consistent** if for any x on some edge such that $f^{-1}(x)$ is transverse to G, o is constant on $[x, o(x)]$. The following lemma shows that the orientation o can be assumed to be consistent on all the edges of the tree T.

Lemma 4.8. *Let us assume that there exists a bubble tangle $\mathcal{T}$ of order k and a sphere decomposition $f : \mathbb{S}^3 \to T$ of G of width less than k. Then there exists a sphere decomposition of the same tree such that o is consistent on T.*

Proof of Lemma 4.8. Lemma 4.7 shows that the orientations of spheres are consistent between braid-equivalent spheres. Therefore, in order to prove Lemma 4.8, we want to control what happens outside of these braid-equivalences, that is, when the quantity $|C(f^{-1}(x) \cap G)|$ varies on an edge of T. These variations, local maxima and minima, occur only when spheres have a tangency with G, since two spheres transverse to G and close enough are braid-equivalent. By applying a small perturbation to f if needed, we can assume that the spheres of f tangent to G have at most one tangency.

Let $e = [u, v]$ be an edge of T and let x be some point interior to it such that $f^{-1}(x)$ has no tangency with G. Up to switching u and v, we can assume that $o(x) = v$. Then, we let $t_1, \ldots, t_n \in [x, v]^n$ be all the points in $[x, v]$ where $f^{-1}(t_j)$ has a tangency with G. By Lemma 4.7, for all $y \in [x, t_1)$, we have that $o(x) = o(y) = v$ since $f^{-1}(y)$ is braid-equivalent to $f^{-1}(x)$ (since all the intersections between the spheres and G are transverse between x and y, one can glue the local homeomorphisms to get one between $f^{-1}([x, y])$ and $S_j \times [0, 1]$ for some j). There only remains to show that $o(x) = o(y)$ for all $y \in (t_1, t_2)$ and we will be done by induction, since this will imply by our definition of o that $o(t_1)$ is also equal to $o(x)$. We set $S_t = f^{-1}(t_1)$ and $S' = f^{-1}(y)$ for a fixed $y \in (t_1, t_2)$.

- We first consider the case where there is a local maximum on t_1, i.e., locally around the tangency τ the embedding of G has either the "V" shape described in Figure 4.12 (recall that both the spheres and the embedding of the graph are piecewise-linear) or S_t is tangent to a vertex. The idea is to deal first with this tangency and then with every other strand in a similar manner as in the proof of Lemma 4.7. As we will use an argument similar to the proof of Lemma 4.7, let us introduce similar notations: $f^{-1}(x) = S$, $B_1 = f^{-1}(T_u(x))$, $f^{-1}(t_1) = S_t$, and I the space between S and S': $I = \mathbb{S}^3 \setminus (\mathring{B}_1 \cup \mathring{f}^{-1}(T_v(y)))$.

Figure 4.12: The tangency τ on S_1 and a cover by a G-trivial ball.

Let g be the connected component of $G \cap I$ containing τ. Because S and G are transverse, $g \cap S$ is either a closed disc if τ belongs to a ball, a point if τ is a degree one vertex, or two points otherwise. Let us choose D a closed disc on S containing $g \cap S$ on its interior and with no other intersection with G: $D \cap G = g \cap S \subset \mathring{D}$. We can now define B, a closed ball contained in I intersecting S on D, S' on a disc, and containing g in its interior (see Figure 4.12). Then B is G-trivial. Indeed, either τ is the point of a "V" shape and $B \setminus G$ is homeomorphic to $B(0,1) \setminus [-1,1] \times \{(0,0)\}$, or $B \setminus G$ is homeomorphic to $B(0,1) \setminus [-1,0] \times \{(0,0)\}$ in the other cases (up to homeomorphism of $B \setminus G$ the thickness of g here does not matter, it can be a point or a segment or a closed ball). Furthermore, recall that we assume that $k \geq 3$ in this section and B is by definition transverse to G, so B belongs to $\mathcal{T}$ and thus by (T3), the ball $B \cup B_1$ also belongs to $\mathcal{T}$.

At this stage, the space $I \setminus (B \cup G)$ is homeomorphic to $S_k \times [0,1]$ so that we can use a similar technique as in the proof of Lemma 4.7 to deal with the other elements of G in I and cover what remains with balls disjoint from B, we apply (T3) inductively and conclude that $f^{-1}(y)$ belongs to $\mathcal{T}$.

- If there is a local minimum on S_t, we let τ denote the point of tangency on S_t. Either $o(y) = o(x) = v$ and the proof is over, or $o(y) = u$. In the latter case, the setup is exactly the same as in the previous case: τ is now a local maximum from S' to S so that $o(x) = o(y) = u$ which is in contradiction with (T2) on $f^{-1}(x)$ by definition of o.

□

Lemma 4.8 ensures that for any edge $e = (u, v)$ of T, there exists a point x_e so that all the points in (x_e, v) are oriented towards v, while all the points in (u, x_e) are oriented towards u. Hence, by subdividing each edge e of T at this x_e, we can think of o as assigning a direction to each edge. This directed tree is the main tool that we use in the proof of Proposition 4.6.

Proof of Proposition 4.6. Let us assume that there exists both a bubble tangle of order k and a sphere decomposition $f : \mathbb{S}^3 \to T$ of width less than k. By Lemma 4.8, there exists

a sphere decomposition of width less than k so that the orientation o as defined above is consistent. Denoting by T' the tree T where each edge has been subdivided once, this orientation corresponds to a choice of direction for each edge of T'. Every directed acyclic graph, and thus in particular the tree T' contains at least one sink, see Figure 4.13.

This sink cannot be a leaf of the tree. Indeed, let $e = [\ell, u]$ be an edge of T incident to a leaf ℓ. By definition, $f^{-1}(\ell)$ is a point disjoint from G, and thus for any y in (ℓ, u) close enough to ℓ, $f^{-1}(y)$ is a sphere disjoint from G. Hence $f^{-1}(T_\ell(y))$ is a G-trivial ball and belongs to $\mathcal{T}$. It follows that all edges incident to leaves of T' are oriented inward. This sink cannot be a degree-two vertex either, as the tree T' was defined in such a way that the two edges adjacent to a degree-two vertex are always oriented outwards. Finally, this sink cannot be a degree-three vertex, as this would mean that the three balls induced by a double bubble are in $\mathcal{T}$, which would violate (T3). We have thus reached a contradiction. □

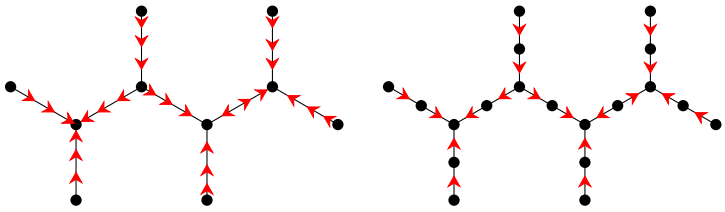

Figure 4.13: An example of T' from T leading to at least one sink.

4.3.2 Tightness of the obstruction.

In this part, we want to show that bubble tangles form a tight obstruction to sphere decompositions, in the sense that a bubble tangle of order k exists whenever a sphere decomposition of width less than k does not exist.

Proposition 4.9. *Let G be an embedding of a graph in $\mathbb{S}^3$ and k be an integer at least three. If G does not admit a sphere decomposition of width less than k, then there exists a bubble tangle of order k.*

The idea of the proof is to show that, given a collection of closed balls satisfying the axioms (T1) and (T4) of bubble tangles, then either we can extend this collection to a bubble tangle, or there exists a **partial sphere decomposition** of width k which sweeps the space "between" the balls of the collection. We first introduce the relevant definition.

Let G be a graph embedded in $\mathbb{S}^3$. A **partial sphere decomposition** of G is a continuous map $f : \mathbb{S}^3 \to T$ where T is a trivalent tree with at least one edge such that:

- For all $x \in L(T)$, $f^{-1}(x)$ is a point disjoint from G or a closed ball B.
- For all $x \in V(T) \setminus L(T)$, $f^{-1}(x)$ is a double bubble transverse to G.
- For all x interior to an edge, $f^{-1}(x)$ is a sphere transverse or finitely tangent to G.

The leaves of T having preimages by f that are not points are called **non-trivial leaves**. Let G be a graph embedding in $\mathbb{S}^3$ and $\mathcal{A}$ be a collection of closed balls in $\mathbb{S}^3$. A partial sphere decomposition f **conforms** to $\mathcal{A}$ if, for all $x \in L(T)$, $f^{-1}(x)$ is either a point disjoint from G, or a closed ball B such that there exists $A \in \mathcal{A}$ such that ∂B and ∂A are braid equivalent and $B \subset A$. In the latter case we say that x conforms to A. The width of a partial sphere decomposition is defined like the width of standard sphere decompositions: it is the supremal weight of spheres that are preimages of points in the interiors of edges of T.

Now, the proof of Proposition 4.9 hinges on the following key lemma. Its proof is similar to branchwidth-tangle duality proofs [116] in that it builds a bubble tangle inductively, but the continuous nature of our objects makes us rely on transfinite induction in the form of Zorn's lemma.

Lemma 4.10. *Let G be an embedding of a graph in $\mathbb{S}^3$, k be an integer at least 3, and $\mathcal{A}$ be a collection of closed balls in $\mathbb{S}^3$ satisfying (T1) and (T4). Then one of the following is true :*

- *$\mathcal{A}$ extends to a bubble tangle of order k.*
- *there is a partial sphere decomposition of width less than k that conforms to $\mathcal{A}$.*

The proof of Lemma 4.10 relies on the following preliminary lemma, which allows us to slightly move the non-trivial leaves of a partial sphere decomposition.

Lemma 4.11. *Let $f : \mathbb{S}^3 \to T$ be a partial sphere decomposition that conforms to a set of closed balls $\mathcal{T}$. Let $\ell \in L(T)$ be a non trivial leaf such that $\partial f^{-1}(\ell)$ conforms to $A \in \mathcal{T}$. Let B be a closed ball such that $B \subset A$ and ∂B is braid-equivalent to $f^{-1}(\ell)$. Then there exists a partial sphere decomposition $f' : \mathbb{S}^3 \to T$ that conforms to $\mathcal{T}$ such that $w(f) = w(f')$ and $f'^{-1}(\ell) = B$.*

Proof of Lemma 4.11. Let us denote by I the subset of $\mathbb{S}^3$ that is swept by f, that is, the union $\bigcup f^{-1}(x)$ where x ranges over all the points of T except the non-trivial leaves. Now, $(I \setminus B) \setminus G$ is homeomorphic to $I \setminus G$, since the subset of $\mathbb{S}^3 \setminus G$ in $B \setminus f^{-1}(\ell)$ is homeomorphic to the product of a sphere with holes with an interval, by definition of braid-equivalence. We can easily extend this homeomorphism to a homeomorphism between $I \setminus B$ and I as is done in the proof of Lemma 4.7, and then we compose f with this homeomorphism to obtain a new partial sphere decomposition f' so that $w(f) = w(f')$ and $f'^{-1}(\ell) = B$. □

Proof of Lemma 4.10. The proof relies on Zorn's lemma. Assume that there is no partial sphere decomposition of G of width less than k that conforms to $\mathcal{A}$, for ease of notation, we denote this assumption by (NPD). Let X be the set of collections of closed balls in $\mathbb{S}^3$ satisfying (T1), (T4) and (NPD) and whose boundaries are transverse to G. Then, we order X by inclusion. Our aim is to apply Zorn's lemma on X in order to get a maximal set of closed balls containing A and still satisfying (T1), (T4), and (NPD).

We first show that every chain of X admits an upper bound in X. Let C be a chain of X and $\mathcal{T} = \bigcup_{c \in C} c$. If C is empty, then the set of G-trivial balls is an upper bound by definition of X. Otherwise, there exists c in C, and as c satisfies (T4) and $c \subset \mathcal{T}$, we also have that $\mathcal{T}$ satisfies (T4). Let B be a closed ball in $\mathcal{T}$. Then there exists c in C such that B belongs to C. Since c satisfies (T1), we have $|C(\partial B \cap G)| < k$, and thus $\mathcal{T}$ satisfies (T1) as well.

Finally, we establish (NPD) for $\mathcal{T}$: we assume by contradiction that there exists a partial sphere decomposition f of G of width less than k that conforms to $\mathcal{T}$. For each non-trivial leaf ℓ of T, by definition of conformity, there exists a ball $B_\ell \in c_\ell \in C$ such that $\partial f^{-1}(\ell)$ is braid-equivalent to ∂B_ℓ. As C is totally ordered and $L(T)$ is finite, there is an element c of C that contains all the sets c_ℓ. This implies that f conforms to c which is absurd because c satisfies (NPD). We conclude that $\mathcal{T}$ satisfies (NPD). Hence $\mathcal{T}$ belongs to X and is an upper bound of C.

By Zorn's lemma, since every chain of X admits an upper bound, it admits a maximal element: there exists $\mathcal{T}$ such that $\mathcal{A} \subset \mathcal{T}$, and $\mathcal{T}$ is maximal with respect to (T1), (T4), and (NPD).

Notice that $\mathcal{T}$ satisfies (T3). Indeed, if there exist $B_1, B_2, B_3 \in \mathcal{T}^3$ such that B_1, B_2, B_3 induce a double bubble $\mathcal{B}$ transverse to G and $\mathbb{S}^3 = B_1 \cup B_2 \cup B_3$, then each of the spheres S_1, S_2, S_3 induced by the double are transverse to G. Each of them admits a braid-equivalent sphere in its neighbourhood, we choose for each of them a sphere S'_i, braid-equivalent to S_i and disjoint from $\mathcal{B}$ (such a sphere exists because each S_i has braid-equivalent spheres on both of its sides).

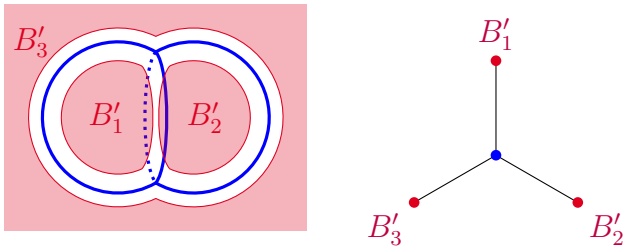

Figure 4.14: A partial sphere decomposition that conforms to $\mathcal{B}$.

Let B'_i be the closed ball disjoint from $\mathcal{B}$ such that $\partial B'_i = S'_i$. As S_i and S'_i are braid-equivalent, there is a homeomorphism $h_i : S_k \times [0,1] \to I_i \setminus G$; where $I_i = B_i \setminus \mathring{B}'_i$. As seen in the proof of Lemma 4.7, we can extend h_i^{-1} to a continuous function $\phi_i : I_i \to \mathbb{S}^2 \times [0,1]$. We can now construct a partial sphere decomposition that conforms to B_1, B_2, B_3 (see Figure 4.14) on the trivalent tree made of 3 edges $(\{v_0, v_1, v_2, v_3\}, \{[v_0, v_1], [v_0, v_2], [v_0, v_3]\})$.

$$f(x) = \begin{cases} v_0 \text{ if } x \in \mathcal{B} \\ v_i \text{ if } x \in B'_i \\ y \in [v_0, v_i] \approx [0,1] \text{ if } x \in I_i \text{ and } h_i(x) = (\cdot, y) \end{cases}$$

This partial sphere decomposition is in contradiction with $\mathcal{T}$ satisfying (NPD), thus $\mathcal{T}$ satisfies (T3).

Now, if $\mathcal{T}$ satisfies (T2), it is a bubble tangle, and we are done. Otherwise $\mathcal{T}$ does not satisfy (T2): there exists some sphere of weight less than k so that none of the two balls that it bounds are in $\mathcal{T}$. In that case, we can define a partial sphere decomposition that conforms to $\mathcal{T}$. Let S be such a sphere in $\mathbb{S}^3$: $|C(S \cap G)| < k$ and for B_1, B_2, its sides, neither $B_1 \in \mathcal{T}$ nor $B_2 \in \mathcal{T}$. Notice that $\mathcal{T} \cup \{B_1\}$ is a collection of closed balls satisfying both (T1) and (T4). By maximality of $\mathcal{T}$ under (T1), (T4), and (NPD), there exists a partial sphere decomposition $f_1 : \mathbb{S}^3 \to T_1$ of width at most k that conforms to $\mathcal{T} \cup \{B_1\}$. As f_1 does not conform to $\mathcal{T}$,

it necessarily admits a non-trivial leaf ℓ_1 such that $\partial f_1^{-1}(\ell_1)$ is braid-equivalent to $\partial B_1 = S$ and $f_1^{-1}(\ell_1) \subset B_1$. Similarly, there exists a partial sphere decomposition $f_2 : \mathbb{S}^3 \to T_2$ of width at most k that conforms to $\mathcal{T} \cup \{B_2\}$ with a non-trivial leaf ℓ_2 such that $\partial f_1^{-1}(\ell_2)$ is braid-equivalent to $\partial B_2 = S$ and $\partial f_2^{-1}(\ell_2) \subset B_2$. Now, by Lemma 4.11, f_1 and f_2 can be modified so that $f_1^{-1}(\ell_1)$ is B_1, and $f_2^{-1}(\ell_2)$ is B_2. Now, f_1 and f_2 can be pasted together at ℓ_1 and ℓ_2 to yield a single partial sphere decomposition of width k which conforms to $\mathcal{T}$. We reach a final contradiction that concludes the proof. □

Remark 4.12. *We use the axiom of choice in this proof for convenience, but it seems likely that one can just rely on the countable axiom of choice since, while there are uncountably many spheres with a low number of intersections with G, there are only countably many isotopy classes of those.*

We can now prove Proposition 4.9.

Proof of Proposition 4.9. We denote by $\mathcal{A}$ the collection of G-trivial balls. By definition, $\mathcal{A}$ satisfies (T4), and since G-trivial balls have weight at most 2, it also satisfies (T1) for k at least 3. Therefore, by Lemma 4.10, either $\mathcal{A}$ extends to a bubble tangle of order k, or there exists a partial sphere decomposition of width less than k conforming to it. In the first case, we are done. In the second case, we are also done, since, given a partial sphere decomposition of width less than k conforming to G-trivial balls, it is straightforward to sweep within the G-trivial balls so as to obtain a sphere decomposition of width less than k. □

4.4 From compression-representativity to bubble tangles

The goal of this section is to show Theorem D: when a graph G is embedded on a compact, orientable, and non-zero genus surface Σ, there exists a bubble tangle naturally arising from the compression-representativity of G on Σ. *In the following, we assume Σ is compact, orientable, and not a sphere.*

4.4.1 Compression bubble tangle

Under the hypotheses above, the idea of the proof is to show that there exists a natural choice of a small side for every sphere with fewer intersections with G than the compression-representativity. Intuitively, such a sphere will only cut discs or "trivial parts" of Σ on one of its sides, which we will designate as the small one. That is justified by the following lemma.

Lemma 4.13. *Let Σ be a surface embedded in $\mathbb{S}^3$ and S be a sphere in $\mathbb{S}^3$ that intersects Σ transversely such that there is at least one non-contractible curve in the intersection. Then one of the non-contractible curves is compressible.*

Proof. As Σ and S are transverse, the intersection of S and Σ consists of a disjoint union of simple closed curves. Each one of these curves bounds two discs on S. Let α be a curve of $S \cap \Sigma$ that is innermost in S, i.e., it bounds a disc D in S that does not contain any other curve of $S \cap \Sigma$. If α is non-contractible, then the disc D is a compression disc for α, and thus α is compressible. Otherwise, α bounds a disc D_Σ in Σ (see for example Epstein [37,

Theorem 1.7]). We deform S continuously by "pushing" D through D_Σ while keeping S embedded (see Figure 4.15) until α disappears from $\Sigma \cap S$.

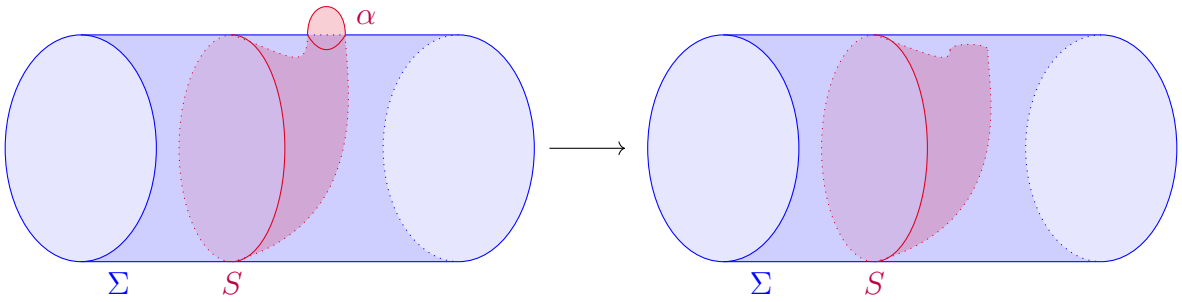

Figure 4.15: Removing a trivial curve from $S \cap \Sigma$.

Repeating this process on a new innermost curve of S will eventually yield a non-contractible compressible curve. Indeed, the number of curves in the intersection is finite (recall that both surfaces are piecewise linear), decreases at each step, and one of the curves in $\Sigma \cap S$ is non-contractible. □

A direct consequence of this lemma is that if G is embedded on a surface Σ intersected by a sphere S and the intersection has weight less than c-rep(G, Σ), then all the simple closed curves in the intersection are contractible. Therefore, one of the two balls bounded by S contains the meaningful topology of Σ, while the other one only contains spheres with holes (see Figure 4.16). In order to formalise this, we will rely on fundamental groups. The inclusion of a subsurface X on Σ induces a morphism $i_* : \pi_1(X) \to \pi_1(\Sigma)$. If this morphism is trivial, we say that X is π_1**-trivial** with respect to Σ.

Definition 4.14 (Compression bubble tangle on an embedded surface)**.** *Let G be a graph embedded on Σ, a surface embedded in $\mathbb{S}^3$ such that c-rep$(G, \Sigma) \geq 3$ and set $k = \frac{2}{3}$c-rep(G, Σ). The **compression bubble tangle** c-$\mathcal{T}$, is the collection of balls in $\mathbb{S}^3$ defined as follows: for any sphere S in $\mathbb{S}^3$ transverse to G such that $|C(S \cap G)| < k$, by Lemma 4.13, there is exactly one connected component A of $\Sigma \setminus S$ that is not π_1-trivial. Exactly one of the open balls B of $\mathbb{S}^3 \setminus S$ does not contain A, i.e., contains only $\pi_1 -$ trivial components of Σ, we put this side in c-$\mathcal{T}$: $\bar{B} \in$ c-$\mathcal{T}$.*

The main step in the proof of Theorem D is to prove that a compression bubble tangle on the torus is indeed a bubble tangle.

Proposition 4.15. *A compression bubble tangle is a bubble tangle.*

Note that Proposition 4.15 immediately implies Theorem D by Proposition 4.6 (the theorem is trivial if c-rep$(G, \Sigma) < 3$). Therefore, the remainder of this section is devoted to proving Proposition 4.15.

By definition, a compression bubble tangle satisfies (T1) and (T2). We then notice that (T4) is verified whenever the compression-representativity of G on Σ is greater than 2.

Lemma 4.16. *If c-rep$(G, \Sigma) \geq 3$ then for all G-trivial balls B, $B \cap \Sigma$ is π_1-trivial.*

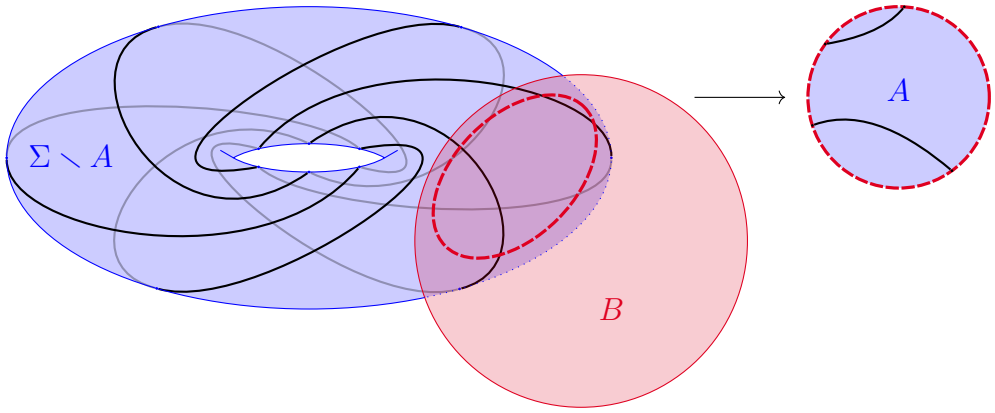

Figure 4.16: Intersection between a torus knot $T_{5,6}$ embedded on a torus and a sphere. Here the ball B containing the disc on the right is in the compression bubble tangle.

Proof of Lemma 4.16. Let B be a G-trivial ball. Then its boundary S has weight at most two. By Lemma 4.13 and the definition of compression-representativity, this implies that S intersects Σ only on contractible curves. Therefore, S bounds two balls, exactly one of which, denoted by B', is such that $B' \cap \Sigma$ is π_1-trivial. This implies that B' is G-trivial. Now either $B = B'$ and we are done, or they are different. In the latter case, this would directly imply that G is made of two trivial pieces and is thus an unknot, plus possibly some trees attached to it. Since an unknot has compression-representativity one (see for example [103, Example 3.2]), this would contradict the assumption that the compression-representativity is at least three. □

The hard part of the proof is to show that (T3) is satisfied. This is more delicate than it seems at first glance, since any surface can be obtained by gluing three discs, and these three discs can even come from a double bubble: we provide an example in Appendix 4.6.3 and Figures 4.23, 4.25, and 4.26 in the case of the torus.

Henceforth, we will proceed by contradiction and assume that we can cover $\mathbb{S}^3$ by three closed balls B_1, B_2, B_3 of c-$\mathcal{T}$ that induce a double bubble DB transverse to Σ and G. Thus Σ is covered by three surfaces with boundary: $\Sigma \cap B_1, \Sigma \cap B_2$ and $\Sigma \cap B_3$ which are π_1-trivial by definition of c-$\mathcal{T}$. In the following, we write $S_i = \partial B_i$. We first show that we can furthermore assume that these surfaces are a disjoint union of closed discs on Σ.

Lemma 4.17. *Let G be a graph embedded on Σ, a surface embedded in $\mathbb{S}^3$. Let c-$\mathcal{T}$ be the compression bubble tangle associated to G and Σ. If there is a double bubble DB transverse to Σ, inducing three balls $B_1, B_2, B_3 \in$ c-$\mathcal{T}^3$ such that $B_1 \cup B_2 \cup B_3 = \mathbb{S}^3$, then we can isotope the double bubble so that we additionally have that $B_i \cap \Sigma$ is a union of closed discs.*

Proof of Lemma 4.17. By Proposition 4.13, each simple closed curve c of $\bigcup C(\partial B_i \cap \Sigma)$ is a contractible curve of Σ. Hence it bounds a unique closed disc D_c on Σ (the other connected component being a surface with non-zero genus and a puncture), and we can define a **slope**[3]

[3]This terminology mirrors the one of Robertson and Seymour in [117].

s that associates the disc D_c to each c:

$$\begin{array}{rccl} s: & \bigcup\limits_{i\in\{1,2,3\}} C(\partial B_i \cap \Sigma) & \to & \mathcal{P}(\Sigma) \\ & c & \mapsto & \text{the closed disc } D_c \text{ of } \Sigma \text{ such that } \partial D_c = c \end{array}$$

Let c be a simple closed curve of $S_i \cap \Sigma$ such that $s(c) \not\subset B_i$ and that is innermost in S_i with respect to that property, i.e., c bounds a disc D_c in S_i such that all $c' \in C(\mathring{D}_c \cap \Sigma)$ satisfy $s(c') \subset B_i$ (see Figure 4.17). Denote by $D_{c'}$ the disc associated to each c' on S_i.

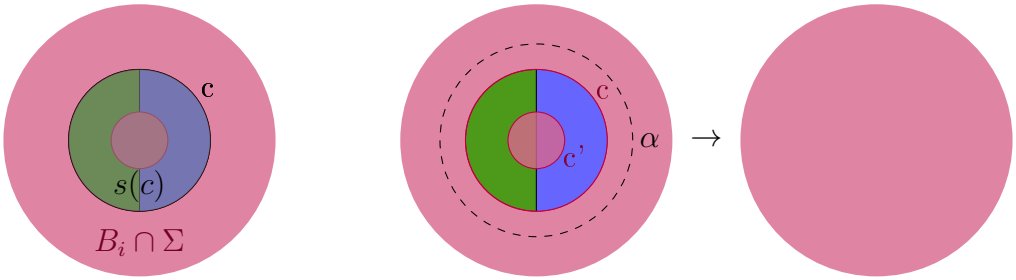

Figure 4.17: An example of c, innermost with the property that $s(c) \not\subset B_i$. The green and blue parts depict the intersection between Σ and the two other balls induced by the double bubble.

The idea is now to push $s(c)$, a disc of Σ, through D_c until $s(c) \subset B_i$ (see Figure 4.18). In order to remove both the interior of D_c and its boundary, we define α, a simple closed curve homotopic to c in $\Sigma \cap B_i$ lying in a tubular neighbourhood of c and disjoint from S_i. Then $D_\alpha = s(\alpha)$ satisfies $D_c \subset D_\alpha$.

To be more precise, let Δ be the disc of $S_i \setminus c$ not disjoint from $s(c)$. As c is innermost on S_i for the property that $s(c) \not\subset B_i$, $\Delta \setminus \Sigma$ has one connected component W which contains c in its boundary, and possibly other connected components which are open discs (any other surface would contradict the fact that c is innermost). These discs may cobound with D_α some closed balls of B_i. But in any case, it suffices for D_α to cross through W during the push (see Figure 4.18). As c is innermost on S_i with respect to the property above, $\mathbb{S}^3 \setminus (W \cup D_\alpha)$ has two connected components. Let A be the one disjoint from B_i, it is necessarily disjoint from Σ (again because c is innermost). Hence the push described here makes D_α sweep A until it crosses W to end in B_i.

This transformation can be made a piecewise-linear isotopy since $A \cup \mathring{W}$ is disjoint from Σ and both W and Σ are piecewise-linear. We apply this transformation on Σ to obtain a new embedding Σ'. The number of elements in $C(S_i \cap \Sigma')$ is smaller than $C(S_i \cap \Sigma)$ so that repeating this process will eventually end.

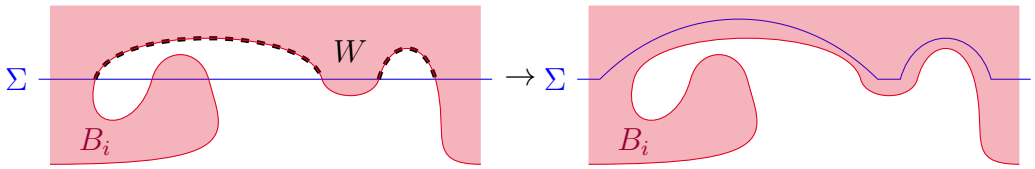

Figure 4.18: Pushing the surface Σ until c is removed from it intersection with B_i.

During this transformation, for all $j \in \{1,2,3\}$ the only modifications of $\Sigma \cap B_j$ were made on $s(c)$ such that $DB \cap \Sigma' = DB \cap \Sigma \setminus s(c)$ and the transversality is preserved. Furthermore, $|S_i \cap G|$ did not increase so that the balls B_j still verify (T1) and are still in c-$\mathcal{T}$ since they are either not intersecting G (hence G-trivial) or they contain the π_1-trivial subparts of Σ'.

We repeat this transformation on Σ until there remains no simple closed curve of $S_i \cap \Sigma$ such that $s(c) \not\subset B_i \cap \Sigma$ for all $i \in \{1,2,3\}$. We obtain that $B_i \cap \Sigma$ is a disjoint union of closed discs. Finally, we have described the isotopy on Σ for convenience, but up to applying a homeomorphism, we can instead keep Σ fixed and apply an isotopy on DB and obtain the same properties. This concludes the proof. □

Then we define Γ **induced by the double bubble** DB to be the intersection of the double bubble with Σ: where vertices are the intersection of the common boundary of the three discs with Σ and edges are the intersections of Σ with the discs. By Lemma 4.17, we can assume that this graph is trivalent and cellularly embedded. It is naturally weighted by endowing each edge with its weight, i.e., the number of connected components in its intersection with G. Let us now state the lemma that we will use for the sake of contradiction.

Lemma 4.18. *The total weight of* Γ *is less than c-rep*(G,Σ)*:*

$$\sum_{e \in E(\Gamma)} |C(e \cap G)| < \textit{c-rep}(G,\Sigma).$$

Proof. Since each edge of Γ bounds exactly two faces of Γ, i.e, discs of Σ; and $\Gamma = DB \cap \Sigma$ we get the following equality:

$$|C(S_1 \cap G)| + |C(S_2 \cap G)| + |C(S_3 \cap G)| = 2 \sum_{e \in E(\Gamma)} |C(e \cap G)| \tag{4.1}$$

By definition of c-$\mathcal{T}$, each ball B_i satisfies $|C(S_i \cap G)| < \frac{2}{3}\text{c-rep}(G,\Sigma)$ so that:

$$|C(S_1 \cap G)| + |C(S_2 \cap G)| + |C(S_3 \cap G)| < 3 \cdot \frac{2}{3}\text{c-rep}(G,\Sigma) = 2\text{c-rep}(G,\Sigma). \tag{4.2}$$

Combining (4.1) and (4.2) concludes the proof: $2 \sum_{e \in E(\Gamma)} |C(e \cap G)| < 2\text{c-rep}(G,\Sigma)$. □

Hence, if Γ contained a simple closed curve that is compressible, we would obtain the contradiction that we are looking for. The rest of the proof almost consists of finding such a compressible curve, leading to the following proposition.

Proposition 4.19. *There exists a set of edges* X *on* Γ *such that:*

$$\sum_{e \in X} |C(e \cap G)| \geq \textit{c-rep}(G,\Sigma).$$

4.4.2 Merging process

The proof of Proposition 4.19 is the technical crux of Theorem D. It consists in defining a merging process, which gradually merges two balls of a double bubble, and proving that at some point in this merging process, one ball will intersect Σ in a non-trivial way. Thus, this yields a compressible curve via Lemma 4.13. An additional difficulty is that this curve might be non-simple in Γ; we circumvent this issue by finding a fractional version of such a curve instead, which will be strong enough to prove Proposition 4.19.

Proof of Proposition 4.19. In order to prove Proposition 4.19, we consider three balls B_1, B_2, B_3 inducing a double bubble DB on their boundaries and define a merging process that gradually merges two of these balls in a controlled way. This is illustrated in Figure 4.20 with a double bubble intersecting a torus.[4] The point of this merging process is to yield a family of balls, one of which will have a non-trivial intersection with Σ and thus allow us to (almost) find a compressible curve.

The merging process depends on the shape of one specific disc D of the double bubble, say the one between B_1 and B_2, which we will call a **membrane** to avoid confusion with the many discs that we deal with. Each connected component of $D \cap \Sigma$ is a separating curve of D with both ends on ∂D since intersections between D and Σ are edges of Γ. Considering these arcs as the embeddings of edges of a graph on D, we consider the dual of this graph, which we call the **membrane tree** and denote $MT(D)$. Its set of vertices are connected components of $D \setminus \Sigma$, $V(MT(D)) = C(D \setminus \Sigma)$, and there is an edge between $f, f' \in V(MT(D))^2$ if $\bar{f} \cap \bar{f}' \neq \varnothing$. In the following we will consider $E(MT(D)) = C(D \cap MT(D))$ as there is a natural one-to-one correspondence between the sets (see Figure 4.19). Note that $MT(D)$ is indeed a tree since it is the (weak) dual of an outerplanar graph.

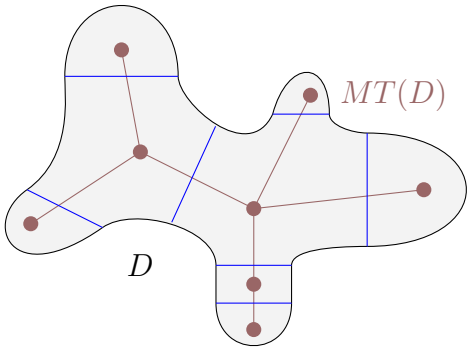

Figure 4.19: The membrane tree (brown) of a membrane D whose edges of Γ stem from $D \cap \Sigma$ (blue).

The merging process starts by shrinking a bit B_2 and surrounding what remains by B_3 so that B_2 is now a closed ball within what was previously B_2 while B_3 is homeomorphic to $\mathbb{S}^2 \times \mathbb{S}^1$. On Σ, the edges that came from $D = B_1 \cap B_2$ are thus now replaced by thin bands

[4] In this figure, the double bubble does not induce a cellularly embedded graph for clarity purposes, since even on a torus the resulting picture would be too intricate to describe the merging process (compare with pictures in Appendix 4.6.3).

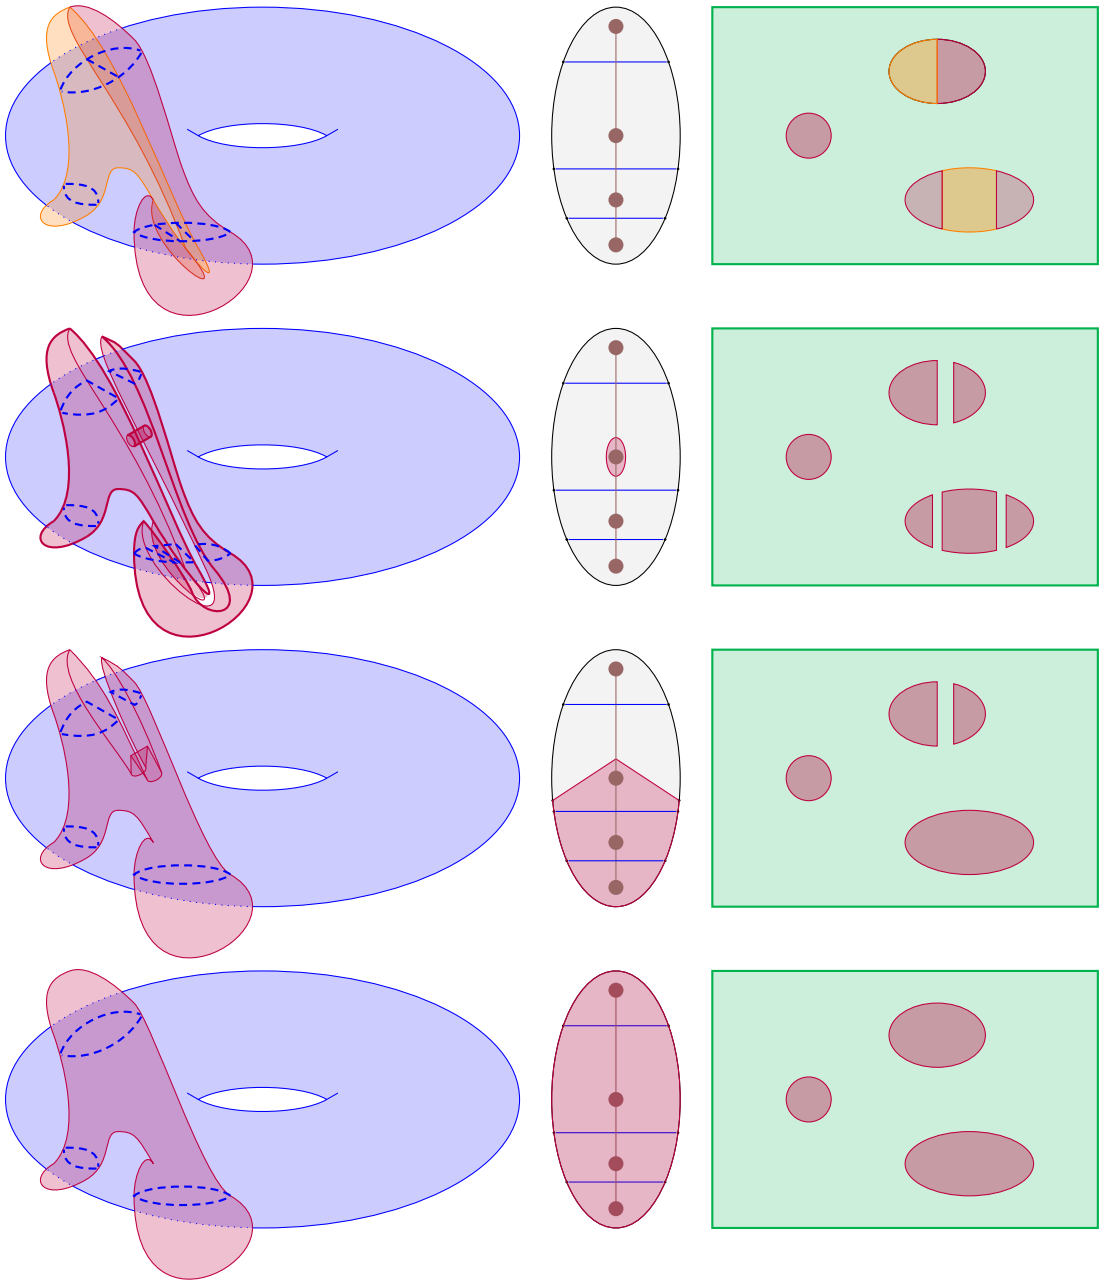

Figure 4.20: Three views of the merging process. Left: a 3D view of a double bubble (purple and orange) intersecting Σ (blue). Middle: the membrane tree, which is progressively covered by a subtree. Right: The unfolded torus and the graph Γ evolving throughout the merging process. It is first covered by 3 balls (orange, purple, green), then 2 (purple and green) after the cylinder merges two of them. The green ball, which covers the "remaining" space, is not shown on the left view for clarity purposes.

belonging to the transformed B_3 with two copies of each edge as boundaries of the bands. We assume the shrinking to be small enough so that the copied edges intersect G the same way the original edges do: this is always possible because the double bubble is transverse to G and the vertices of G have been thickened. It follows that one face of Γ is now $B_3 \cap \Sigma$, which is Σ with as many boundaries as there are connected components in $(B_2 \cup B_1) \cap \Sigma$ (see the unfolded torus on the right of the second row compared to the one on the right of the first row in Figure 4.20). All the other faces of Γ are now discs whose boundaries consist of edges with B_3.

We now pick an arbitrary vertex v of $MT(D)$ and connect B_1 with B_2 by a small solid cylinder disjoint from Σ whose ends are glued on B_1 and B_2 in the face indicated by v. This new ball is denoted by $B_{1,2}$. This step can be seen on the second row of Figure 4.20.

For e an edge of $MT(D)$ adjacent to v, we then grow this cylinder so that it entirely covers e. This amounts to connecting the surfaces of $\Sigma \cap B_{1,2}$ (intuitively, the bands between the copies of edges disappear, see the right column of Figure 4.20). Indeed, when the edge e of $MT(D)$ is covered by the growing cylinder, the two surfaces containing the initial discs of B_1 and B_2 incident to e are glued along e. This process is then iterated by adding more edges of $MT(D)$, thus growing a subtree of $MT(D)$, until that subtree has fully grown and is $MT(D)$. This is illustrated in the third and fourth rows of Figure 4.20. Eventually, the whole membrane is covered so that $B_{1,2} = B_1 \cup B_2$

We track the evolution of the ball $B_{1,2}$ during the merging process, which is parametrised by the subtree T of $MT(D)$ that has been merged. This leads to the following **merging function** m:

$$\begin{array}{rccc} m: & MT(D)_{\subset} & \to & \mathcal{P}(\Sigma) \\ & T & \mapsto & B_{1,2}(T) \cap \Sigma \end{array}$$

where $MT(D)_{\subset}$ is the set of subtrees of $MT(D)$ and $B_{1,2}(T)$ is the closed ball of the merging process parametrised by T. By definition, we have that $B_{1,2}(T)$ is contained in $B_1 \cup B_2$, transverse to Σ, and $B_{1,2}(T) \cap \Sigma = E(T)$.

It will be convenient in the following to define what kind of topology of Σ is captured by a closed connected surface W embedded on Σ. For example, an annulus can be embedded on the torus in two ways: either it contains a non-contractible curve of the torus, or it is π_1-trivial and contains irrelevant holes. The following definition aims to "fill" these irrelevant holes. Each contractible boundary component b of W bounds a disc D_b on Σ that is disjoint from $\mathring{W}$. We define $\bullet(W) = W \cup \bigcup_{\substack{b \in C(\partial W) \\ b \text{ contractible}}} D_b$, the **filled surface** induced by W.

It directly follows from the definition that $\partial\bullet(W) \subset \partial W$. It is also worth noticing that if W is π_1-trivial, then $\bullet(W) = D_b$ for one of the b.

We can now finally use all these objects to find a fractional packing of curves in Γ with enough weight. Let T be a minimal tree of $MT(D)_{\subset}$ such that $\bullet(m(T))$ is not a union of discs. Note that such a tree exists, since for $T = MT(D)$, we have $B_{1,2}(T) = B_1 \cup B_2$ and thus $\bullet(m(T)) = \Sigma$. Necessarily, the surface $\bullet(m(T))$ contains exactly one component that is not a disc. Indeed, let e be an edge incident to a leaf of T. By minimality of T, $\bullet(m(T \smallsetminus e))$ is a union of disjoint discs. By definition of the merging process, $m(T)$ is $m(T \smallsetminus e)$, where two copies of an edge have been merged back. If these copies were on two different π_1-trivial

surfaces, their merging would have been π_1-trivial too. Hence the two copies of an edge were on the same disc of $\bullet(m(T \smallsetminus e))$, so that the merging of e glues together two distinct segments on the boundary of that disc. This operation yields either a Möbius band, which is impossible since Σ is orientable, or an annulus a.

By definition of $\bullet(m(T))$, the boundaries of a are non-contractible curves of Σ. Lemma 4.13 ensures that one of these two curves is compressible. Since they bound an annulus on Σ, they are homotopic on Σ and thus are both compressible. Let us denote them by b and b'.

At any stage of the merging process, some edges that were originally in Γ are now duplicated. Therefore, while the curves b and b' are simple and disjoint, they might be using some of those duplicated edges of Γ, but each such edge is used at most twice. Therefore we have, denoting by X the set of edges of Γ used by b and b'.

$$\sum_{e \in X} |C(e \cap \Gamma)| \geq \frac{1}{2} \sum_{e \in b} |C(e \cap \Gamma)| + \frac{1}{2} \sum_{e \in b'} |C(e \cap \Gamma)| \geq 2 \cdot \frac{1}{2} \text{c-rep}(G, \Sigma),$$

where the last inequality comes from the fact that b and b' are compressible and the definition of c-rep(G, Σ). This last inequality concludes the proof. □

This proposition directly implies Proposition 4.15, and thus Theorem D:

Proof of Proposition 4.15. A compression bubble tangle immediately satisfies the bubble tangle axioms (T1) and (T2) by definition, and (T4) by Lemma 4.16. For the axiom (T3), assume by contradiction that there exist three closed balls $B_1, B_2, B_3 \in \text{c-}\mathcal{T}$ covering $\mathbb{S}^3$ and inducing a double bubble transverse to Σ. By Lemma 4.17, we can assume the graph Γ induced by the intersection of the double bubble with Σ is cellularly embedded. Then by Proposition 4.19, the total weight of Γ is at least c-rep(G, Σ). This is a contradiction with Lemma 4.18. □

4.5 Examples

We now combine our results to lower bound the treewidth of any diagram of a torus knot, thus proving Corollary 4.1.

Proof of Corollary 4.1. Combining Proposition 4.3, Theorem C and Theorem D, it suffices to lower bound the compression-representativity of $T_{p,q}$ by $\min(p, q)$. The torus $\mathbb{T}$ has two curves which obviously bound compression discs, we prove that there are no other compressible curves. In $\mathbb{S}^3$, the torus $\mathbb{T}$ bounds one solid torus $\mathbb{D}^2 \times \mathbb{S}^1$ on each side, and in a solid torus S, exactly one of the homotopy classes of simple closed curves on the boundary is compressible, namely the one in the kernel of the inclusion map $i_* : \pi_1(\mathbb{T}) = \mathbb{Z}^2 \to \pi_1(S) = \mathbb{Z}$. The result follows by observing that $T_{p,q}$ intersects p times the first of these homotopy classes and q times the other one, and thus has compression-representativity $\min(p, q)$. □

More generally, the same argument can be applied to lower bound the treewidth of the (p, q)-cabling [3, Section 5.2] of any nontrivial knot. We refer to Ozawa [103, Theorem 6] for examples of spatial embeddings of any graph with high compression-representativity, and thus high spherewidth.

4.6 Additional results and remarks

4.6.1 Computability of compression-representativity

A natural question arising from our Theorem D is whether the compression-representativity can be computed. A first reasonable step to do so would be to recognise compressible curves. An algorithm from Matveev [91, Theorem 4.1.10] is able to do[5] so by leveraging the normal surface theory [53]. The setting, in our words, is the following: if α is a non-contractible curve embedded on an orientable surface Σ embedded in $\mathbb{S}^3$, we first cut $\mathbb{S}^3$ along Σ to obtain two 3-manifolds M_1, M_2 with boundary which both have a copy of α on their boundary. The aforementioned algorithm will find a compressing disc for α within M_1, or M_2, if such a disc exists. In that case, α is compressible on the corresponding side of Σ.

However, as soon as the surface on which G is embedded has genus 2, one can construct infinitely many non-contractible compressible curves. The construction starts by taking two disjoint curves compressible "on the same side" and joining them by a path. Merging the curves along the path yields a curve which is compressible if the path was complex enough: the initial compression discs of the two curves can be linked by a rectangle with a "green-house" shape that follows the aforementioned path on the same side of the surface as the discs. This construction is illustrated by Figure 4.21.

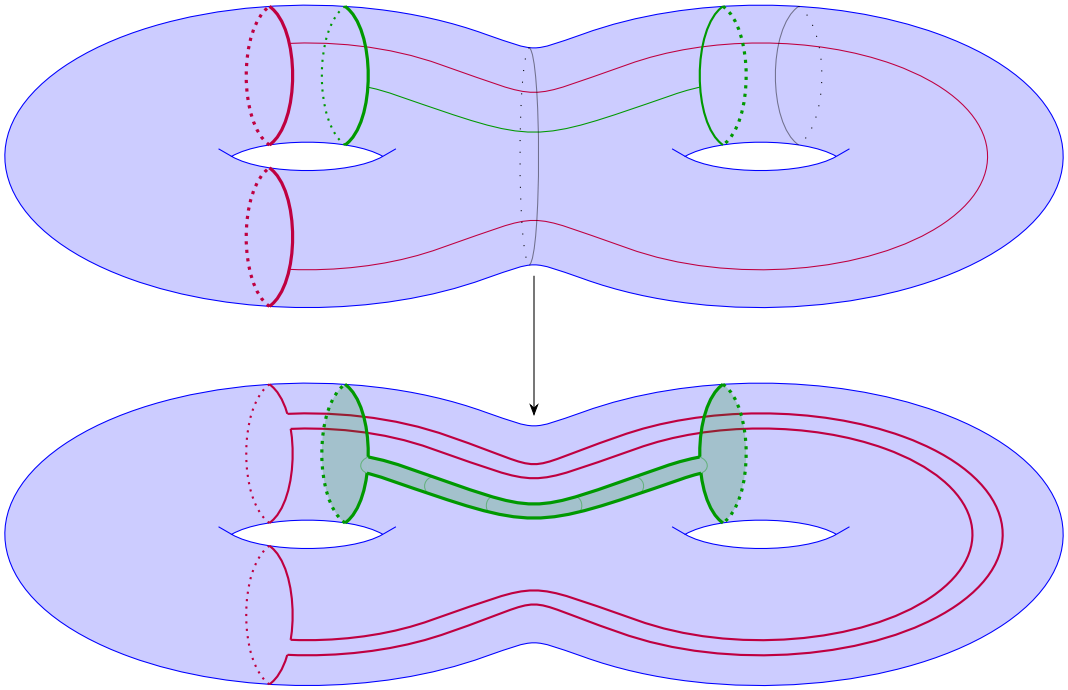

Figure 4.21: Construction of two compressible non-contractible curves from two pairs of such curves and two paths. Dehn twists along the grey curves on top yield other such curves. A compression disc is shown for the curve obtained from the green pair.

[5]The algorithm is expressed in terms of knot genus, which is equivalent to what is expressed here since a knot is of genus 0, i.e., an unknot, if it is the boundary of a properly embedded disc in the 3-manifold with boundary considered.

A **Dehn twist**[6] is a type of self-homeomorphism of a surface which can be thought of as a local "twist" of the surface around a given curve as pictured in Figure 4.22. By applying Dehn twists along a compressible curve that intersects only the aforementioned path, we obtain a new compressible curve which can be non-contractible depending on the path (such valid curves for Dehn twists are pictured in grey on top of Figure 4.21). This behaviour is not observed on the torus, on which there are only two non-contractible compressible curves which only produce contractible curves.

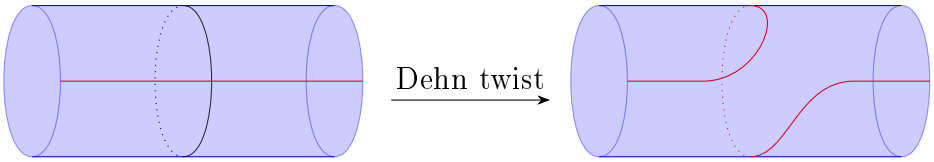

Figure 4.22: A Dehn twist of the red curve along the black curve.

Hence, to compute the compression-representativity, one should find a way to deal with this infinity of curves, for instance, by limiting the number of curves considered. A way to do so could be to prove that all compressible curves can be constructed using the aforementioned operation from a finite number of initial compressible curves; and then limit the complexity of paths used for the construction by their number of intersections with G. That way, one could expect to deal with a finite number of compressible curves, among which one will achieve the compression-representativity.

4.6.2 Monotonicity of compression bubble tangles

The reader familiar with branchwidth theory and tangles from [117] may notice that our axiom (T3) stands out from the others. Indeed, while the three other axioms mimic closely the axioms of tangles in graph theory, (T3) only applies to 3 small balls which induce a double bubble. That behaviour differs from graph theory, where (T3) would be expressed by: "the whole space cannot be covered by 3 elements of the tangle". In fact, it is easier to manipulate tangles when the latter axiom is swapped with a combination of "the whole space cannot be covered by 3 *disjoint* elements of the tangle" and, "if A is in the tangle, and $B \subset A$ is small, then B is in the tangle". By comparison, (T3) is weaker: we lose the stability by inclusion. From their definition, one can see that compression bubble tangles satisfy the stronger property. Intuitively, the small side of a sphere is the one that only cuts discs from the surface, if we have an inclusion, the sphere will only cut smaller discs.

Proposition 4.20. *Let G be a graph embedded on Σ such that c-$\mathcal{T}$ is defined. Let k be the order of c-$\mathcal{T}$, then c-$\mathcal{T}$ is stable by inclusion up to (T1): for every closed balls A, B of $\mathbb{S}^3$, if $A \in$ c-$\mathcal{T}$, $B \subset A$, and $|C(\partial B \cap G)| < k$, then $B \in$ c-$\mathcal{T}$.*

Proof. Let A and B be two closed balls of $\mathbb{S}^3$ such that $|C(\partial A \cap G)| < k$, $|C(\partial B \cap G)| < k$, and $A \in$ c-$\mathcal{T}$. By definition, $A \cap \Sigma$ is π_1-trivial. Since $B \subset A$, we have $B \cap \Sigma \subset A \cap \Sigma$ so that the canonical inclusion morphisms i_* satisfy $\pi_1(B \cap \Sigma) \to \pi_1(A \cap \Sigma) \to 0$. Hence, $B \cap \Sigma$ is π_1-trivial and $B \in$ c-$\mathcal{T}$. □

[6] We refer to [130] for a formal definition.

4.6.3 Covering a torus with three discs from a double bubble

In this part, we describe a standard torus that can be covered by 3-balls induced by a double bubble which intersects the torus on three discs, as pictured in Figure 4.23. This double bubble is described in Figures 4.24, 4.25, 4.26, 4.27 and one can prove that this is, up to homeomorphism, the only double bubble realising such a cover. We will now try to describe these intricate depictions of this double bubble. Furthermore, it is worth noticing that this construction can be generalised to standard surfaces of higher genus.

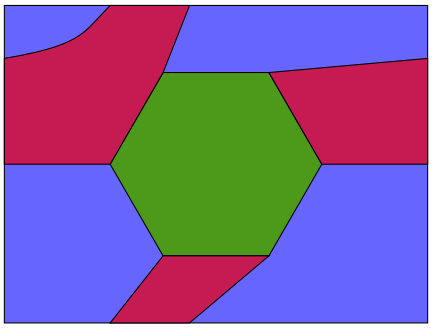

Figure 4.23: Covering a torus with three discs from a double bubble.

In the following, the remaining space of each figure is covered by a green ball that we do not represent for the purpose of preserving the clarity of our embeddings in space. The intersection of this ball with the torus is a green disc that we picture in Figure 4.23 and 4.26. First, we describe the two balls in Figure 4.24: the first one, B_1 in blue, is a rectangular cuboid that will contain our rectangular torus shown in Figure 4.26 and Figure 1.3, except on a rectangular notch deep enough to let an annulus of the torus uncovered. The second ball B_2, in red, fills this notch except for a hole, so that the second ball looks like a thick rectangular U. We then pull a disc out of B_2 and wrap it around B_1, as illustrated on the right of Figure 4.25. By thickening this disc to dig inside B_1 enough to reach the torus, we obtain the final form of B_2 whose boundary is depicted in the right part of Figure 4.24.

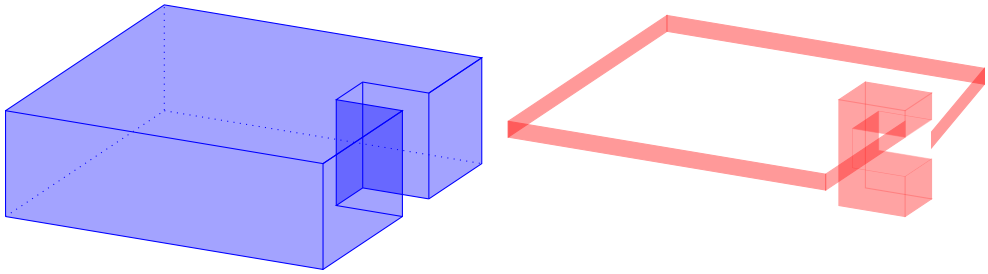

Figure 4.24: The spheres S_1 (blue) and S_2 (red) up to thickening the ribbon part.

The union of B_1 and B_2 is the initial rectangular cuboid except for the hole left by the construction of B_2 as pictured on the left of Figure 4.25. At this point, their intersection with the torus is the union of the blue disc and purple disc shown in Figure 4.26 on the unfolded torus. The remaining part of the torus is the rectangular green rectangle.

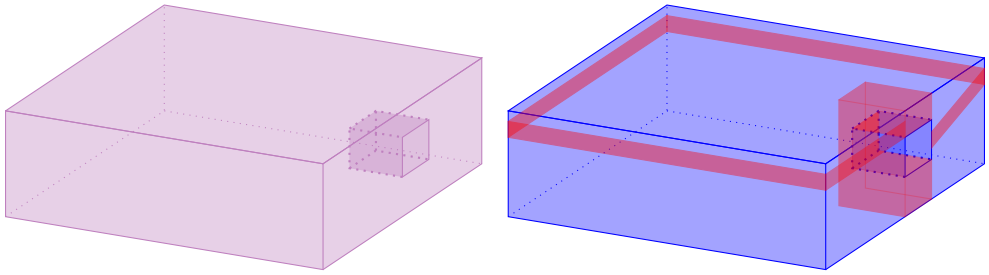

Figure 4.25: The ball $B_1 \cup B_2$ in purple, B_3 is the complement.

To be sure that this construction is a double bubble, it is enough to check that the intersections $B_1 \cap B_2$ and $B_2 \cap B_3$ are discs. These discs are pictured in Figure 4.27. We will not attempt to describe their shape further than the U shape from which the disc wrapping around B_1 is pulled. However, as a sanity check, one can verify that the boundaries of the discs indeed are the same.

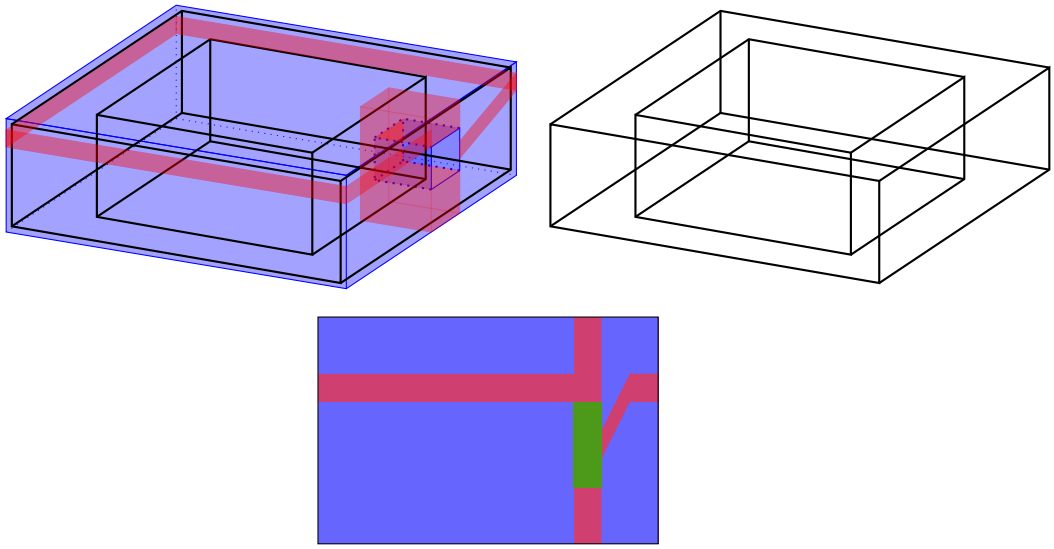

Figure 4.26: The double bubble and a torus inside, covered by three discs (up to thickening the ribbon part of S_2 by pushing it into S_1).

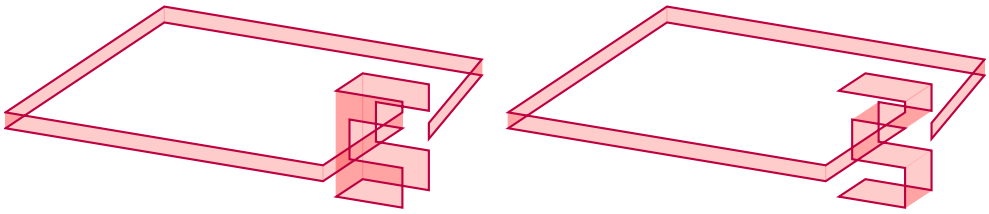

Figure 4.27: The discs $S_1 \cap S_2$ and $S_3 \cap S_2$.

Chapter 5

A lower bound on the complexity of splitting link diagrams

In this chapter, we study the minimal number of crossings one might need to add in order to split a link diagram using Reidemeister moves. We show that link diagrams requiring an arbitrarily large number of such additional crossings exist.

This chapter is the fruit of a project with Arnaud de Mesmay and Jonathan Spreer. We leverage the obstructions developed in Chapter 4 to prove our results.

5.1 Introduction

The Reidemeister theorem [138] is a fundamental and powerful result of low-dimensional topology stating that any two diagrams of the same knot, or more generally link, can be related by a sequence of Reidemeister moves, pictured in Figure 2.12. This theorem paved the way for both theoretical results and computational applications. Indeed, many knot invariants can be shown to be invariant by showing that they are not modified by Reidemeister moves, like tricolorability (see Section 1.1) or the Jones polynomial [69]. From a computational point of view, this theorem allows for a discretisation of the space of possible transformations to consider in order to study knot equivalence: it is enough to focus on sequences of Reidemeister moves. This fact is at the root of a straightforward algorithm to study many knot problems: at the level of diagrams, apply Reidemeister moves in a random or brute force manner until a desired property is verified.

Recall that a primary example of such a knot theory problem is recognising the trivial knot, which is a first instance of the major problem of knot theory: deciding whether two knots are equivalent or not. It turns out that some unknot diagrams [21], called "hard" unknots or *culprits*, exhibit an annoying behaviour for this algorithm. The maximum number

of crossings of a diagram during the algorithm is larger than in the initial diagram: one first needs to *add* crossings before being able to reach the untangled diagram. The existence of such diagrams implies that it is not possible to untangle an unknot by applying only Reidemeister moves that will not increase the number of crossings of a diagram, which represents intuitively the complexity of the diagram. Such a culprit, called the Goeritz culprit, is presented in Figure 5.1, and proved to be hard in [21] where it is shown that at least one crossing is required to untangle this unknot. We do not know of systematic techniques or methods to prove easily that an unknot is a hard unknot. All the known results seem to resort to an exhaustive search in the Reidemeister graph, which very quickly becomes unfeasible.

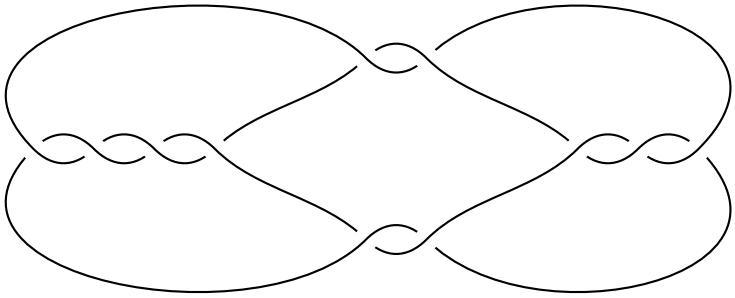

Figure 5.1: The Goeritz culprit: using Reidemeister moves, one must add at least 1 crossing to untangle this unknot.

Following the notations of [70, 21] we will denote by $\mathbf{cr}(D)$ the number of crossings in the diagram D. Then, for two equivalent diagrams D_1, D_2 and a sequence of Reidemeister moves R transforming D_1 into D_2, we define $\mathbf{Top}(D_1, R)$ which is the maximum of $\mathrm{cr}(D^i) - \mathrm{cr}(D_1)$ throughout the sequence of Reidemeister moves R where D^i is the diagram D_1 after performing the first i moves of the sequence. The minimal number of extra crossings to pass from D_1 to D_2 is denoted $\mathbf{Add}(D_1, D_2)$ which is formally the minimum of $\mathrm{Top}(D_1, R)$ taken among all the sequences of Reidemeister moves that transform D_1 into D_2. When we see D_2 as a goal diagram, $\mathrm{Add}(D_1, D_2)$ is a lower bound on the number of crossings to add during the running of the aforementioned algorithm that applies Reidemeister moves on D_1 to reach D_2.

In this last context, where D_1 is an unknot diagram and D_2 is a diagram of a simple curve, D_1 is a **hard unknot** if $\mathrm{Add}(D_1, D_2)$ is positive. This measure of complexity is called m in [21] and has a ratio version called *recalcitrance* in [70]. Studying these complexity measures and hard unknots turns out to be trickier than one might initially think. In fact, one of the purposes of [21] is to confirm or invalidate many claims about hard unknot diagrams (using a computer search of the Reidemeister graph). Currently, only diagrams on which $\mathrm{Add}(D_1, D_2) \leq 2$ are known, although it is conjectured that there exist unknot diagrams D for which $\mathrm{Add}(D, D_2)$ is arbitrarily large. Let us note that a proof of this conjecture, formulated in terms of recalcitrance, is claimed in [70], but concerns about this proof are raised in [21].

In this chapter, we will focus on the same measure of complexity, but the problem considered is the **splitting problem**: deciding whether a link L is split, i.e., whether there exists a sphere disjoint from L separating at least 2 link components of L. If such a sphere exists,

there exists a link diagram in which two unlinked sublinks are disjoint: they are separated by a circle in the plane. Such a diagram will be called a **split diagram**. By capping off the aforementioned circle with one disc above and one disc below the projection, we verify that the converse of the last assertion is true. In other words, the splitting problem is equivalent to the search of a split diagram. Therefore, in terms of Reidemeister moves, we will study $\mathrm{Add}(D_1, D_2)$ where D_2 is a split diagram of a link L and D_1 is any diagram of L. In this setting, if the minimum of $\mathrm{Add}(D_1, D_2)$, denoted by $\mathbf{Unl}(D_1)$, is positive among all split diagrams D_2 of L, we will call D_1 a **hard split link** and its **crossing-complexity** will be $\mathrm{Unl}(D_1)$.

Our results. We exhibit a family of link diagrams $\mathcal{D}(p, q)$ of a split link $L(p, q)$ with exactly two unlinked sublinks. The first sublink K is made of two linked torus knots $T_{p,q}$ and the second one is an unknot U surrounding one of the torus knots (see $\mathcal{D}(7, 13)$ in Figure 1.31 for an example).

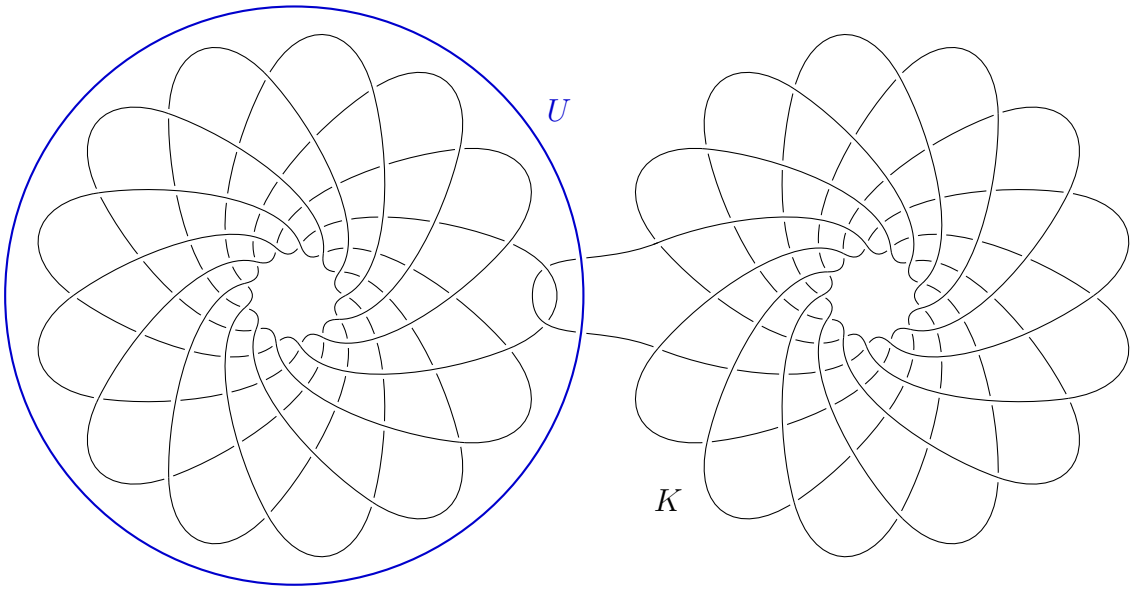

Figure 5.2: The link diagram $\mathcal{D}(7, 13)$ made of two split links: two linked $T_{7,13}$ and an unknot U.

For $\mathcal{D}'(p, q)$, any link diagram of $L(p, q)$ in which U is disjoint from the other link components, we prove Theorem E implying that $\mathrm{Add}(\mathcal{D}(p, q), \mathcal{D}'(p, q)) = \Omega(\min(p, q))$.

Theorem E. *For each $n \geq 2$, there exists a diagram D_n with $2n^2 + 2$ crossings that is the diagram of a split link L of $\mathbb{S}^3$ with 3 components such that any sequence of Reidemeister moves converting it to a split diagram of L_n passes through a diagram with at least $2n^2 + \frac{2}{3}n$ crossings.*

In other words, we exhibit diagrams D_n for which $\mathrm{Unl}(D_n) = \frac{\sqrt{2}}{3}\sqrt{n-2} - 2$, so that:

Corollary 5.1. *There exist **hard split links** with arbitrarily large crossing-complexity.*

To prove Theorem E, we exploit the unknot U present in every one of our diagrams and which is separated in $\mathbb{S}^3$ from the remaining link components (one can picture a sphere around K in Figure 1.31 which lies in between K and U). Our approach is to show that if there exists a sequence R of Reidemeister moves where $\mathrm{Top}(\mathcal{D}(p,q),R)$ stays small, we can use the evolution of U throughout these moves to define a sweepout of K with spheres. By construction, each of these spheres has a small number of intersections with K. This sweepout presents two notable differences from the sphere decompositions of Chapter 4. On the one hand, it is simpler: it is linear and features no double bubbles. On the other hand, it is not monotone: a sphere involved in this sweepout may go back-and-forth, this behaviour never happens in sphere decompositions. Despite this last difference, the obstructions that we developed in Chapter 4 are versatile enough to oppose the thinness of this sweepout.

Building the sweepout is not straightforward. Intuitively one would like to lift each unknot U to a sphere by capping it off above and below the diagram. However, the unknot U may intersect itself during the sequence of Reidemeister moves, which complicates the process. We alleviate this problem by leveraging results and methods from the article [24], which provides a way to transform homotopies of curves on a Riemannian surface into isotopies of similar length. This can be applied to our problem as follows: any sequence of Reidemeister moves induces a homotopy of U in the plane, and we think of the projection of the link K as a discrete metric for the curves in this homotopy: the length of such a curve will be its number of intersections with K. In this framework, our assumptions imply that there exists such a homotopy where the intermediate curves have a small length. Now, the techniques of [24] show that this implies that there also exists an isotopy with the same bounds on the length. Since an isotopy consists of simple curves, it is then easy to lift this planar isotopy into a sweepout of $\mathbb{S}^3$ with 2-spheres, which all have a controlled number of intersections with K. A subtlety is that we cannot exactly use the results of [24] as a black-box. Indeed, the metric which is defined by the link K is not fixed since it might also be evolving with Reidemeister moves. We explain in Section 5.3 why the proof techniques of [24] also deal seamlessly with this issue.

The way we leverage our obstructions directly stems from the construction of our split links. They consist of two torus knots $T_{p,q}$ that are each embedded on a torus with high compression-representativity. Hence, each one of them yields an obstruction: a bubble tangle of order $\frac{2}{3}\min(p,q)$ by Theorem D. Each of these obstructions defines a small side for each sphere of the sweepout since we assume that the number of intersections between the spheres and K is small enough. We then obtain a contradiction by proving that under these assumptions, the two small sides will always agree. However, by definition, they disagree on the initial diagram. This work is done in Section 5.4.

Related work. Finding a sphere in space separating two links is easier than the problem of finding a disc that has a knot as its boundary. Therefore, the splitting problem has been studied several times as a useful and easier problem for understanding the unknot recognition problem [36, 78]. The problem is interesting in itself and several of its aspects are known.

It was shown by Haken that this problem is decidable in 1961 [53] using the theory of normal surfaces. It was later shown by Hass, Lagarias, and Pippenger that this problem is in **NP** [54]. And several decision problems that stem from the splitting problem, like deciding

whether changing at most k crossings can transform a link diagram into the diagram of a split link, are even **NP**-hard [71].

Another natural consideration in relation to the iterated application of Reidemeister moves is how many of those are needed to split a diagram. In 2005, Hayashi showed an exponential bound [56] for this number. This bound was later greatly improved in [78], where Lackenby provided a polynomial bound by refining the arguments of [54]. This proof leveraged the work of Dynnikov on arc representations and grid diagrams, which are a way to represent knots in grids [36]. Lackenby [78] showed that there exists a polynomial sequence of moves on grid diagrams which leads to a split grid diagram while not increasing the number of lines in the grid diagram. Since these diagrams with n lines naturally have at most n^2 crossings, and moves on grid diagrams translate to Reidemeister moves, this result yields a quadratic upper bound on our quantity of interest: $\mathrm{Add}(D_1, D_2)$.

Organisation of this chapter. After going through the specifics of this chapter in Section 5.2, we will explain how to use the results of [24] in Section 5.3. This step is crucial for our definition of sweepouts. Then, we will exploit our obstructions of Chapter 4 to prove Theorem E in Section 5.4. Finally we will present some other link diagrams which our method covers in Section 5.5.

5.2 Specific preliminaries

In the following, the foundation of our reasoning will be a link diagram $\mathcal{D}$ of a split link L and the existence of a sequence of Reidemeister moves R on this diagram such that $\mathrm{Top}(\mathcal{D}, R) \leq k$ for some k that we will specify later. Our goal will be to prove that such a sequence cannot exist if k is too small.

Furthermore, we will be led to swap regularly our point of view between one in $\mathbb{S}^3$ and one in $\mathbb{R}^3$, recall that they are equivalent by compactification (see Chapter 2). We will make use of the point of view in $\mathbb{S}^3$ to define our obstruction while the one in $\mathbb{R}^3$ is more intuitive to define diagrams and projections.

Link diagrams. For p, q coprime integers, the link diagram $\mathcal{D}(p, q)$, is a link diagram which consists of two sublinks: the first one K is made of two torus knots $T_{p,q}$ linked together such that they cross twice, as pictured in Figure 5.2 and 5.3, the second one is an unknot U which circles one of the torus knots and intersects K twice. In the following proof, we will be mostly considering one such diagram at a time so that $\mathcal{D}(p, q)$ will be shortened to $\mathcal{D}$ whenever there is no ambiguity.

A first interesting remark is that these link diagrams are initially in minimal position, except for the two intersections between U and K. Indeed, Murasugi computed the crossing number of torus knots using bridge number [98, Proposition 7.5]:

Proposition 5.2. *The crossing number of the torus knot $T_{p,q}$ is* $\min(p(q-1), q(p-1))$.

As a direct corollary, which can be visualised on Figure 5.3:

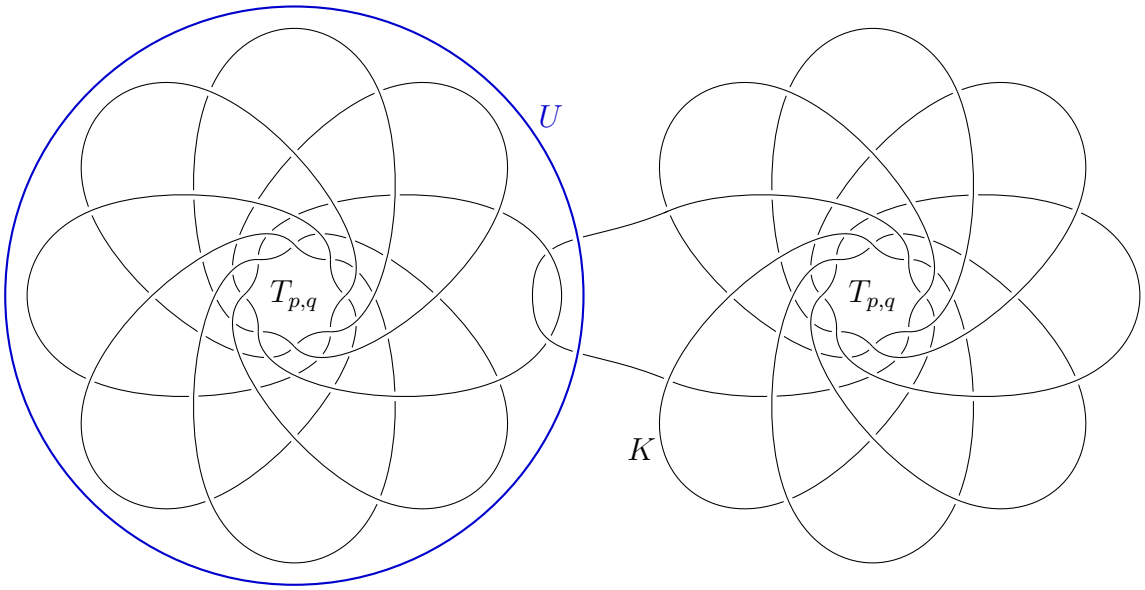

Figure 5.3: The link diagram $\mathcal{D}(7,8)$ made of two split links: two linked $T_{7,8}$ and an unknot U.

Lemma 5.3. *Let D be a link diagram equivalent to $\mathcal{D}(n,n+1)$ by Reidemeister moves. Then, D has at least $2n^2$ crossings: $cr(D) \geq 2n^2$.*

Proof. Let L be the link associated to $\mathcal{D}(n,n+1)$ and D a diagram of L. By Proposition 5.2, both torus knot components $T_{n,n+1}$ of L have at least $(n+1)\times(n-1) = n^2-1$ crossings. Since these two knots are linked, they share at least 2 crossings in each diagram equivalent to $\mathcal{D}$ so that $\mathrm{cr}(D) \geq 2n^2$. □

In Chapter 4, our framework was designed to handle both spatial graphs and links. Since we consider no vertices here, we will no longer make use of thickened balls and only work with embeddings of links into $\mathbb{S}^3$. In this context, and in order to simplify notation, we will be writing $|K \cap S|$ instead of $|C(K \cap S)|$ since these numbers coincide whenever K and S are transverse or finitely tangent.

From a sequence of Reidemeister moves to continuous operations. Throughout this chapter, we will be working with a sequence of Reidemeister moves R converting the link diagram $\mathcal{D}$ to a split diagram. We will consider that all of the intermediate diagrams are obtained from the same projection $p : \mathbb{R}^3 \to \mathbb{R}^2$, and that the sequence of Reidemeister moves corresponds to an ambient isotopy Φ_R of $\mathbb{R}^3$, so that the diagrams correspond to $p(\Phi_R(L,t))$ for a finite number of times t. The projection p is regular except at the **critical times** of R, which are times where the projection $p \circ \Phi_R(L)$ displays a tangency or a triple point (these critical times are pictured in Figure 5.4). For any time t that is not critical, we denote by L_t the link $\Phi_R(L,t)$ and by $\mathcal{D}_t$ the diagram $p(L_t) = p(\Phi_R(L,t))$.

Figure 5.4: The critical times of the two Reidemeister moves RII and RIII.

Notice that the definition of $\mathrm{Top}(\mathcal{D}, R)$ naturally coincide with $\sup_t \mathrm{cr}(\mathcal{D}_t) - \mathrm{cr}(\mathcal{D}_0)$. Indeed, the diagram $\mathcal{D}_t$ at critical times has fewer intersections than one of $\mathcal{D}_{t+\epsilon}$,$\mathcal{D}_{t-\epsilon}$ for ϵ small enough. Furthermore, $\mathrm{cr}(\mathcal{D}_t)$ is constant between 2 critical times.

In the next Section 5.3, it will be useful for us to consider separately the movements of U and of K under the Reidemeister moves. We will use K_t as a shorthand for the diagram of K at time t: $p(\Phi_R(K,t))$. Furthermore, for technical reasons, for U we will focus our attention on the homotopy $\phi_U : \mathbb{S}^1 \times [0,1] \to \mathbb{R}^2$ induced in the plane by the projection $p(\Phi_R(U))$. We denote the corresponding curves by $U_t = \phi_U(\mathbb{S}^1, t)$ and emphasise that we consider these as closed curves in the plane, i.e., we forget about the information of which strand is over which at each crossing.

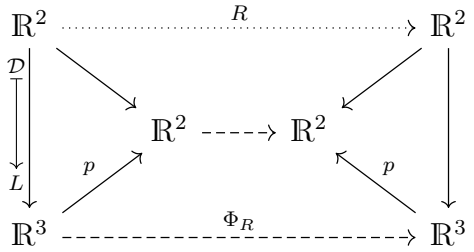

Figure 5.5: Definition of our homotopies from the sequence of Reidemeister moves R.

5.3 From homotopies to isotopies

The aim for this section is to leverage results of [24] to show that we can assume that our unknot U remains a simple curve in the projection throughout the sequence of Reidemeister moves:

Proposition 5.4. *Let R be a sequence of Reidemeister moves bringing $\mathcal{D}$ to a split diagram $\mathcal{D}_2$ and such that for all $t \in [0,1]$, $cr(\mathcal{D}_t) \leq k$ for some integer k. Then there exists an ambient isotopy Φ' of $\mathbb{R}^3$ and an isotopy $h : \mathbb{S}^1 \times [0,1] \to \mathbb{R}^2$ such that:*

1. *$h(\mathbb{S}^1, 0) = U_0$ and $h(\mathbb{S}^1, 1) = U_1$.*
2. *For any $t \in [0,1]$, the total number of crossings in the overlay of $p(\Phi'(K,t))$ and $h(\mathbb{S}^1, t)$ in $\mathbb{R}^2$ is at most k.*

We emphasise that in the second item of the proposition, we only consider the projection of the sublink K in $\Phi'(K,t)$, and not the entire link. Therefore, the proposition provides an

ambient isotopy for K in $\mathbb{R}^3$ and an isotopy for $p(U)$ in $\mathbb{R}^2$, while preserving a bound on the total number of intersections when projecting K via p.

This proposition follows from the techniques of Chambers and Liokumovich in [24]. We first state one of their main results.

Definition 5.5 (Chambers, Liokumovich [24, Definition 1.3]). *For γ a curve on $\mathbb{S}^2$, we let $-\gamma$ denote the curve γ with reversed orientation. Two curves α and β are **ϵ-image equivalent**, $\alpha \sim_\epsilon \beta$, if there exists a finite collection of disjoint intervals $\bigsqcup_{i=1}^n I_i \subset \mathbb{S}^1$ such that $|\alpha(\mathbb{S}^1 \setminus \bigsqcup I_i)| + |\beta(\mathbb{S}^1 \setminus \bigsqcup I_i)| < \epsilon$. We also require that there exists a permutation σ of $\{1, \dots, n\}$ and a map $f : \{1, \dots, n\} \to \{0, 1\}$, such that $\alpha|_{I_i} = (-1)^{f(i)}\beta|_{I_{\sigma_i}}$ for all i.*

Theorem 5.6 (Chambers, Liokumovich [24, Theorem 1.1']). *Suppose that γ is a smooth homotopy of closed curves on a 2-manifold M and that γ_0 is a simple closed curve. Then, for every $\epsilon > 0$, there exists an isotopy $\overline{\gamma}$ such that $\overline{\gamma_0} = \gamma_0$ and $\overline{\gamma_1}$ is ϵ-image equivalent to a small perturbation of γ_1. Additionally, for every t, there exists a t' such that $\overline{\gamma_t}$ is ϵ-image equivalent to a small perturbation of $\gamma_{t'}$. If γ_1 is simple or is a point, then this homotopy also ends at γ_1, up to a change in orientation.*

We think of this definition and theorem in more informal terms as follows. When a curve α self-intersects in the plane, each crossing point can be *resolved* in two ways by reconnecting the endpoints in a small ball around the crossing point (see the curve on the right side of Figure 5.6 compared to the one in the middle). When all the crossings have been resolved in some way and we obtain a simple closed curve α', we say that α' is a *resolution* of α. The theorem states that if we have a homotopy γ on a surface between two simple curves γ_0 and γ_1, one can obtain an *isotopy* $\overline{\gamma}$ between γ_0 and γ_1 (or its reverse) where each intermediate curve $\overline{\gamma_t}$ is a resolution of some intermediate curve $\gamma_{t'}$ (the times t and t' need not coincide).

In particular, if M is endowed with a metric (for example a Riemannian one), the lengths of γ_t and of $\overline{\gamma_{t'}}$ will differ by an arbitrarily small quantity. Therefore, Theorem 5.6 immediately implies that for any $\varepsilon > 0$, if there exists a homotopy between two simple curves γ_0 and γ_1 where each intermediate curve has length at most L, then there also exists an isotopy between γ_0 and γ_1 (or its reverse) where each intermediate curve has length at most $L + \varepsilon$.

Theorem 5.6 can be applied as a black-box to prove Proposition 5.4 in the particular case where the sublink K stays invariant throughout the sequence of Reidemeister moves R, i.e., $\Phi_R(K, t) = \Phi_R(K, 0)$ for all $t \in [0, 1]$. Indeed, in that case, we can think of K_t as a discrete metric which measures the length of a curve U_t by its number of intersections with K_t. More formally, we can take γ to be the homotopy ϕ_U between U_0 and U_1, which are both simple closed curves by the definition of the link diagram $\mathcal{D}$. Then applying Theorem 5.6 provides us with an isotopy $\overline{U}$ between U_0 and U_1 where all intermediate curves $\overline{U}_t$ are obtained from resolving intersections of some $U_{t'}$. In particular, for any $t \in [0, 1]$, the number of intersections between $\overline{U}_t$ and K_t is at most the number of intersections between $U_{t'}$ and K_t, which is at most k because $K_t = K_0$ and thus does not depend on t.

A careful reading of [24] shows that the general case of Proposition 5.4, where the diagram K_t of the components K also evolves during the sequence of Reidemeister moves, can also be obtained using the exact same proof techniques. The basic idea of the proof of Theorem 5.6 is to first decompose the homotopy in a sequence of local moves, and then replace each curve γ_t by one of its resolutions.

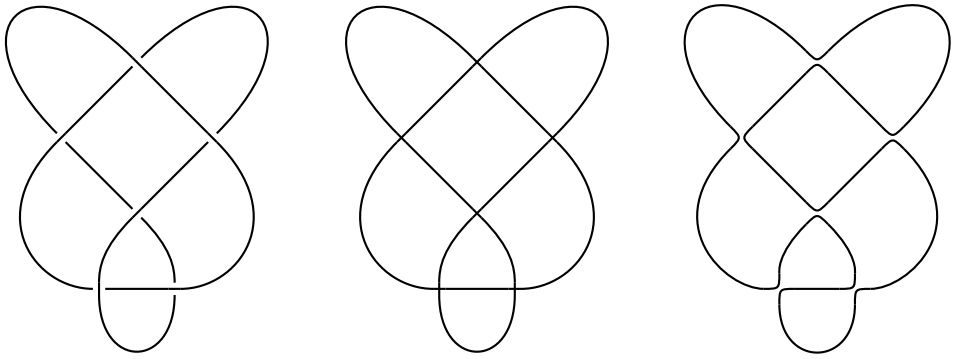

Figure 5.6: Left: A diagram of the unknot. Middle: A projection U_t of this unknot, where the crossing information has been forgotten. Right: A resolution of U_t.

For the first step, they track discrete times where the self-intersection pattern (i.e., the homeomorphism type) of γ_t changes. By an argument similar to the proof of the Reidemeister theorem, one can assume that this only happens at *critical events*, when a curve γ_t undergoes a homotopy move, which is a transformation analogue to a Reidemeister move but without any crossing information, this step is formalised by their Proposition 2.1 and Lemma 2.2.

Inbetween the critical events, any homotopy of a curve γ_t can be applied similarly on any of its resolutions, yielding an isotopy. Hence, the idea is to replace each of these Reidemeister moves by a resolution of the move, as explained by their Figure 2. However, doing so in a naive way runs into discontinuity issues, a basic example of which is detailed by Example 2 of [24] which is associated to their Figures 3 and 4. Therefore, the authors provide a more intricate workaround: the crux of the proof of Theorem 5.6 is to show how to choose the correct resolutions and connect their isotopies together. This is achieved by defining an auxiliary graph of resolutions (see their Figure 7 for instance), synthesising how they are connected by local isotopies: the precise definition for this graph follows their Figure 8. The proof is achieved by finding a path through it by using the handshaking lemma [39].

Now, in our case, the critical events are exactly the times when U undergoes a Reidemeister move. In-between these critical events, there are other Reidemeister moves involving either (i) both U and K, or (ii) only K. In case (i), the Reidemeister move only changes the relative position of U and K, and the diagram of K stays unchanged. Therefore this is a homotopy of U which can be applied to give an isotopy of any of its resolutions. When case (ii) happens, the diagram $p(K)$ changes, but U, considered up to isotopy, does not. Therefore, by applying the same Reidemeister moves on $p(K)$, any motion of U between critical events can be applied to any of its resolutions while preserving the number of intersections with K. These motions can then be connected using the same handshaking argument as in the proof of Theorem 5.6. Summarising, the proof technique directly adapts to the case of an evolving metric, as long as these evolutions are applied appropriately throughout the new sequence of isotopies.

5.4 Leveraging bubble tangles

From now on, we assume thanks to Proposition 5.4 that the homotopy ϕ_U of $\mathbb{R}^2$ rendering U disjoint from K is an isotopy. In other words, for all t in $[0,1]$, U_t is a simple curve. Furthermore, the sequence of Reidemeister moves R, as well as the ambient isotopy Φ_R have been modified in accordance.

5.4.1 Definition of the sweepout

For each t, we will associate a sphere to the simple curve U_t. We glue one of the boundaries of two infinite annuli on U_t in the direction of the diagram projection, one on top of it and one on the bottom. Now seen in $\mathbb{S}^3$, these annuli form a torus pinched at ∞. Cutting the surface at ∞ provides our sphere S_t (see Figure 5.7) such that S_t intersects K_t only on the preimages of $U_t \cap K_t$ by the projection p. We now have a continuous family of spheres S_t sweeping a continuous family of links K_t. Equivalently, by applying the ambient isotopy Φ_R^{-1} to S_t and K_t, we can assume that K_t is fixed. We will slightly abuse notation and make this assumption while still denoting by S_t the family of spheres. We will denote by Θ an ambient isotopy so that $\Theta(S_0, t) = S_t$.

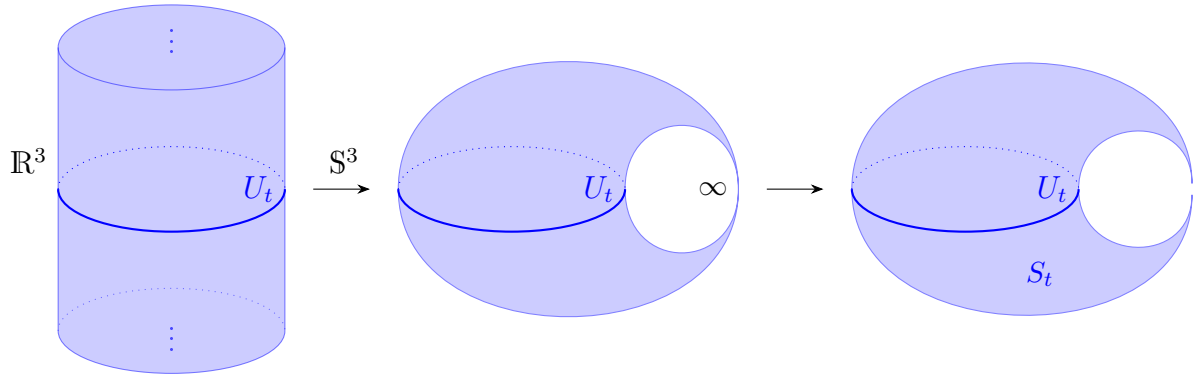

Figure 5.7: Gluing two infinite annuli on U and cutting the resulting pinched torus at ∞ in $\mathbb{S}^3$.

We now put in relation $\mathrm{Top}(\mathcal{D}(n, n+1), R)$ and S_t:

Lemma 5.7. *Our objects verify:* $\sup_{t\in[0,1]} |S_t \cap K_t| - 2 \leq Top(\mathcal{D}(n, n+1), R)$.

Proof. Let us write $\mathrm{cr}(K, K')$ for the number of intersections between K and K', where K and K are two link components of a link diagram. If $K = K'$, then $\mathrm{cr}(K, K')$ is the number of self crossings of K.

Since U_t is simple, we have that $\mathrm{cr}(\mathcal{D}_t) = \mathrm{cr}(K_t, U_t) + \mathrm{cr}(K_t, K_t)$. By definition, $\mathcal{D}_t$ is the diagram $\mathcal{D}(n, n+1)$ at time t of the sequence of Reidemeister moves. Hence, it is equivalent to $\mathcal{D}(n, n+1)$, so that adapting the proof of Lemma 5.3 yields $\mathrm{cr}(K_t, K_t) \geq 2n^2$. Furthermore, $\mathrm{cr}(K_t, U_t) = |S_t \cap K_t|$ by construction of S_t. By definition of $\mathcal{D}(n, n+1) = \mathcal{D}_0$, $\mathrm{cr}(\mathcal{D}_0) = 2n^2 + 2$. Hence, $\mathrm{cr}(\mathcal{D}_t) - \mathrm{cr}(\mathcal{D}_0) \geq 2n^2 + |K_t \cap S_t| - 2n^2 - 2$.

Thus, $\sup_{t\in[0,1]} |S_t \cap K_t| - 2 \leq \sup_{t\in[0,1]} \mathrm{cr}(\mathcal{D}_t) - \mathrm{cr}(\mathcal{D}_0) = \mathrm{Top}(\mathcal{D}(n, n+1), R)$. □

Hence, we will focus on the number of intersections between S_t and K_t in $\mathbb{S}^3$, in order to lower bound $\mathrm{Top}(\mathcal{D}, R)$. Changes in $|K_t \cap S_t|$ can happen only at critical times of R where the number of crossings involving U_t in $\mathcal{D}_t$ increases or decreases. Therefore, we do not consider RIII moves between U_t and K_t. Since RI cannot happen to U_t which stays simple, only RII moves are relevant for our study. We define the $(d_j)_j$ to be the critical times of RII moves involving both U_t and K_t. According to our definitions, the times d_j, are the only times where S_t is not transverse to K_t: at these times K_t and S_t are finitely tangent.

Remark that compared to sphere decompositions of Chapter 4, the sweep-out defined by the spheres S_t is most likely to do many back-and-forths and does not branch off. However, as long as the number of intersections of S_t with K_t is low enough, we will be able to define small sides using our obstructions from Chapter 4.

5.4.2 Obstruction to the sweepout

The link K_t is made of two linked torus knots K^1, K^2. Each of them can be embedded on a standard torus as pictured in Figure 5.8. The compression-representativity of each embedding is n. Notice that these tori intersect, and this will be of no consequence. By Theorem D, there exist two compression bubble tangles $\mathcal{T}^1, \mathcal{T}^2$ of order $\frac{2}{3}n$, where $\mathcal{T}^i$ is the compression bubble tangle induced by the torus on which K^i is embedded.

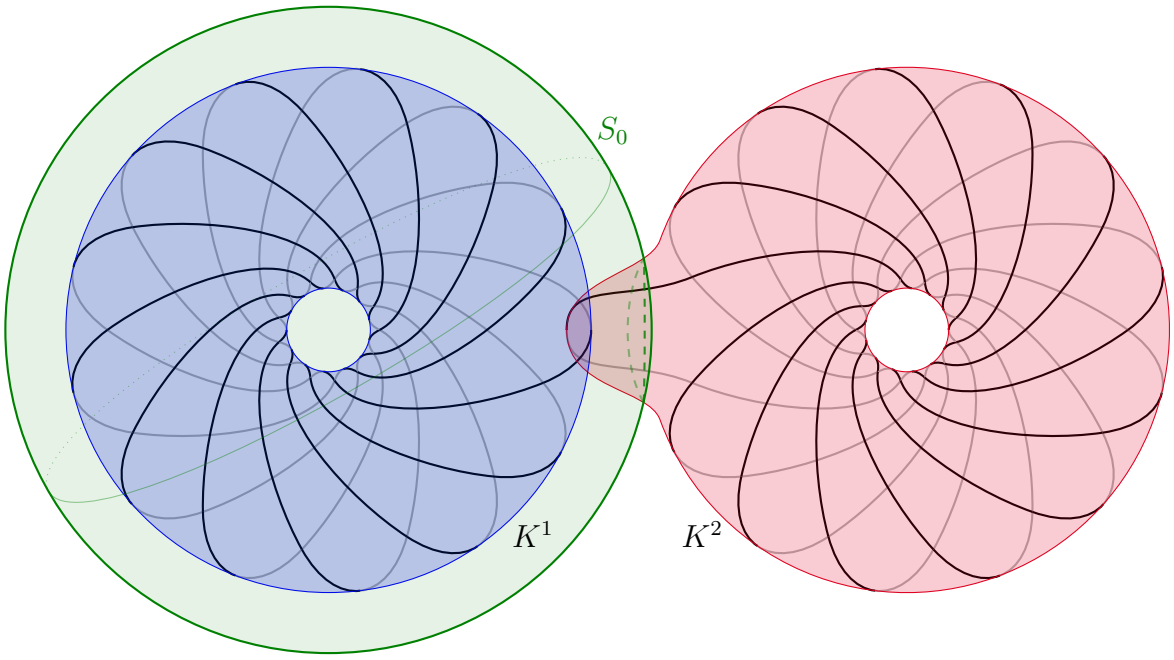

Figure 5.8: Each link component embedded on a torus, and S_0.

From now on, we assume that $\mathrm{Top}(\mathcal{D}(n, n+1), R) < k = \frac{2}{3}n - 2$ so that $|K_t^i \cap S_t| \leq |K_t \cap S_t| < \frac{2}{3}n$ for $i \in \{1, 2\}$ by Lemma 5.7. Hence, for all $t \in [0, 1] \setminus \{d_j\}_j$, S_t has a small side B_t^i for each $\mathcal{T}^i$. We will study how these small sides evolve throughout the sweep-out. In particular, we will be careful when using the notion of braid-equivalence introduced in Chapter 4: the back and forths make it so that two spheres of the sweep-out, even when

there is no S_{d_j} between them, may intersect. Also, in a similar way to Chapter 4, we will be careful with tangencies, where small sides are not defined.

5.4.3 Proof of Theorem E

We define the **agreement**, which is a map $a : [0,1] \smallsetminus \{d_j\}_j \to \{0,1\}$, such that $a(t) = 1$ if $B^1_t = B^2_t$ and 0 otherwise. In order to show Theorem E, we will prove by contradiction that $\mathrm{Top}(\mathcal{D}(n, n+1), R) \geq \frac{2}{3}n - 2$. For the sake of contradiction, recall that we have assumed that $\mathrm{Top}(\mathcal{D}(n, n+1), R) < k$. We will first show that $a(0) = 0$ and $a(1) = 1$.

First, it is easy to see that $a(1) = 1$. Indeed, U_1 is disjoint from K_1 so that S_1 is also disjoint from K_1. It follows that one side of S_1 is an empty ball with respect to both K^1 and K^2 so that $a(1) = 1$ by (T4).

Then, notice that the side of S_0 in $\mathcal{T}^1$ is, by (T4), the ball not containing K^1: it is disjoint from the torus. For the other compression tangle, as it is illustrated by Figure 5.8, S_0 contains a K_2-trivial ball and K^1 on one side, and the rest of K^2 on the other side. Hence, $a(0) = 0$, the small sides disagree at time 0.

In a similar way to Chapter 4, in particular to the proofs of Section 4.3, we now show that a stays constant on its domain of definition. To handle the back and forths, we introduce a new equivalence relation on spheres. Two intersecting spheres S, S' of $\mathbb{S}^3$ are said to be **intersection-equivalent** if there exists an isotopy between them which stays constant on their intersection $S \cap S'$, see Figure 5.9.

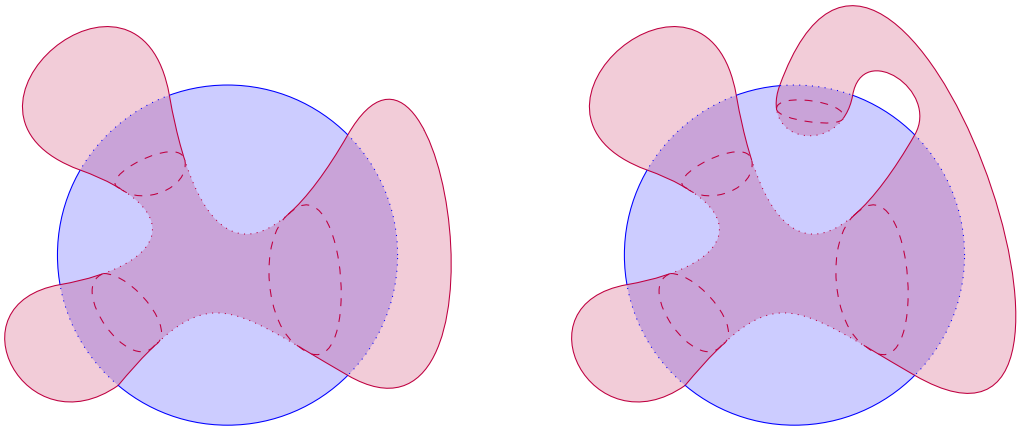

Figure 5.9: Left: two intersection-equivalent spheres. Right: two spheres that are not intersection-equivalent: the red sphere has an annulus component that cannot be mapped to a component of the blue one. Indeed, the components of the blue sphere are discs and a sphere with 4 punctures.

Two spheres S_t, $S_{t'}$, if t and t' are close enough, are either disjoint or intersection-equivalent.

Lemma 5.8. *Let t be a time of $[0,1] \smallsetminus \{d_j\}_j$. There exists a neighbourhood V of t such that for all spheres S_v where $v \in V \smallsetminus \{t\}$, S_v and S_t are either disjoint or intersection-equivalent.*

Proof. Since we work in a PL setting, surfaces have a naturally polygonal structure. It follows that on a small enough neighbourhood, these polygons are locally pushed in one direction or stay in place. Hence, in this neighbourhood, every sphere is either disjoint or intersection-equivalent to S_t. □

We can now show that under our assumption, a has the same value for two intersection-equivalent spheres that are close enough for t.

Lemma 5.9. *Let $t \in [0,1]$ be a time. There exists V, a neighbourhood of t in $[0,1]$ such that a is constant on $V \setminus \{t\}$.*

Proof. We distinguish two cases depending on whether t is a critical time or not.

-Case 1: let us assume that t is not a critical time and let V be an open connected neighbourhood of t in $[0,1] \setminus \{d_j\}_j$. Up to intersection with a neighbourhood provided by Lemma 5.8, we assume that spheres S_v for $v \in V$ are either disjoint or intersection-equivalent to S_t. For i in $\{1,2\}$:

•If $S_v \subset B_t^i$, and if we denote by B_v^i the component of $\mathbb{S}^3 \setminus S_v$ included in B_t^i, by Proposition 4.20, $B_v^i \in \mathcal{T}^i$. If both i were in this case, then $a(t) = 1$ and $a(v) = 1$.

•If $S_v \subset (B_t^i)^c$, and if we denote by B_v^i the component of $\mathbb{S}^3 \setminus S_v$ containing B_t^i, S_v and S_t are braid equivalent, and by Lemma 4.7 we conclude that $B_v \in \mathcal{T}^i$. If $a(t) = 1$ then $a(v) = 1$. Otherwise, if $a(t) = 0$, this argument on one of the i, and the previous one on the other yields $a(v) = 0$.

•If S_v and S_t intersect, they are intersection-equivalent by assumption on V. The idea is to manage pieces of S_v in the small side of S_t using the monotony of $\mathcal{T}^i$, and the one in the big side using braid equivalence relations.

Let us denote by $\mathcal{I} = C(\mathring{B}_t^i \cap S_v)$: the connected components of $S_v \setminus S_t$ within B_t^i. Since S_t and S_v are intersection-equivalent, there is a natural injection $\psi : \mathcal{I} \hookrightarrow C(S_t \setminus S_v)$ which associates to each $I \in \mathcal{I}$, a connected component of $C(S_t \setminus S_v)$ to which I is isotopic while keeping $S_t \cap S_v$ fixed. There is no d_j between t and v. It follows that $\mathcal{I}$ and $\psi(\mathcal{I})$ have the same number of intersections with K: Θ sweeps a braid. Hence, $S_v' = (S_v \setminus \psi(\mathcal{I})) \cup \mathcal{I}$ is a sphere (S_t and S_v' are isotopic via an isotopy that keeps $(B_t^i)^c \cap S_v$ fixed) such that $|S_v' \cap K^i| < k$ and $S_v' \subset B_t^i$. Hence, by Proposition 4.20, $A^i \in \mathcal{T}^i$ where A^i is the side of S_v' included in B_t^i.

Now, there only remains to deal with $C(S_v \setminus B_t^i)$, the connected components of $S_v \setminus S_t$ outside of B_t^i. By pushing S_v' a bit into $\mathring{A}^i$ so that S_v' is disjoint from S_v, we get that S_v' and S_v are braid equivalent. Let B^i be the side of S_v containing A^i, by Lemma 4.7, $B^i \in \mathcal{T}^i$.

In this construction, if $a(t) = 1$, all the small sides coincide so that $a(v) = 1$. Otherwise, during the construction, the closed ball $A^i \in \mathcal{T}^i$ lie on different sides of S_t. Then, each closed ball $B_t^i \in \mathcal{T}^i$ contains A^i, therefore $B_v^1 \neq B_v^2$ so that $a(v) = 0$ too.

-Case 2: let us assume that $t = d_j$ for some j. Let V be an open connected neighbourhood of t in $[0,1]$ disjoint from the other d_j. Up to intersection with a neighbourhood provided by Lemma 5.8, we assume that the spheres S_v for $v \in V$ are either disjoint or intersection-equivalent to S_t. We want to show that $a(u) = a(v)$ where u and v are in disjoint components of $V \setminus t$. To do so, we use the point at the tangency, and show that it switches sides during the sweepout.

Let p be the point of K^1 tangent to S_t at time t (we assume that it is K^1 since the knots play symmetrical roles). Let x be the point of S_u such that $\Theta(x,t) = p$. Denote by P the path followed by x during the sweepout by the spheres between u and v: $P = \Theta(x,[u,v])$. Let us cover P by a closed ball B that intersects S_u and S_v on a single disc each. This ball is K^1-trivial, and we assume that it covers the strand of K^1 around p (similarly to the proof of Lemma 4.8).

Notice that the knot K^2 presents no tangency with S_t in between S_u and S_v. So the method of case 1 applies for $\mathcal{T}^2$ and we remark that p must switch sides for K^2, i.e., $p \in B^2_u$ if and only if $p \notin B^2_v$. Let us show that the same is true for $\mathcal{T}^1$.

•If $p \notin B^1_u$, we notice that $|\partial(B^1_u \cup B) \cap K^1| < k$ by our construction of B and assumptions on S_t. By (T3), $B^1_u \cup B \in \mathcal{T}^1$. Furthermore, by construction, $\partial(B^1_u \cup B)$ and S_u are intersection-equivalent. The methods of case 1 apply here and tell us that the side of S_v containing p is in $\mathcal{T}^1$.

•If $p \in B^1_u$, we notice that $|\partial(B^1_u \smallsetminus B) \cap K^1| < k$ by our construction of B and assumptions on S_t. By Proposition 4.20, $B^1_u \smallsetminus B \in \mathcal{T}^1$. By construction, $\partial(B^1_u \cup B)$ and S_u are intersection-equivalent. The methods of case 1 apply here and tell us that the side of S_v not containing p is in $\mathcal{T}^1$.

Hence, for $i \in \{1,2\}, p \in B^i_u \Leftrightarrow p \notin B^i_v$. This implies that $a(u) = a(v)$, and concludes our proof. □

Proposition 5.10. *The agreement a is constant on $[0,1] \smallsetminus \{d_j\}_j$.*

Proof. We can cover $[0,1]$ by open discs from Lemma 5.9 on which a is constant. Since $[0,1]$ is compact, only finitely many of them are enough to cover it. Each connected component $[0,1] \smallsetminus \{d_j\}_j$ is connected, the agreement function is then constant on these components by continuity of a (a is locally constant). Around each d_j, a is constant by Lemma 5.9. Hence, a is constant on $[0,1] \smallsetminus \{d_j\}_j$. □

Proof of Theorem E. The initial discussion of this subsection states that $a(0) = 0$ and $a(1) = 1$. This contradicts Proposition 5.10. Thus, our assumption that $\mathrm{Top}(\mathcal{D}(n,n+1),R) < \frac{2}{3}n-2$ does not hold. Therefore, during the sequence, at least a diagram has at least $2n^2 + \frac{2}{3}n$ crossings. □

As corollary:

Corollary 5.11. *For each $n \geq 2$ there exists a hard split link D_m with $m \geq n$ crossings and crossing complexity $Unl(D_m) \geq \frac{2}{3}\sqrt{\frac{1}{2}m - 1} - 2$.*

Proof. By Theorem E, $\mathrm{Top}(\mathcal{D}(n,n+1),R) \geq \frac{2}{3}n - 2$ for every Reidemeister move sequence converting $\mathcal{D}(n,n+1)$ to a split diagram. Then, by definition, $\mathrm{Unl}(\mathcal{D}(n,n+1)) \geq \frac{2}{3}n-2$. The link diagram $D_{2n^2+2} = \mathcal{D}(n,n+1)$ has $2n^2+2$ crossings, and we have $\mathrm{Unl}(D_{2n^2+2}) \geq \frac{2}{3}n-2$.

Hence for $m = 2n^2+2 \geq n$ and the link diagram D_m with m crossings satisfies $\mathrm{Unl}(D_m) \geq \frac{2}{3}\sqrt{\frac{1}{2}m-1} - 2$. □

To sum up, we showed that the small sides designated by $\mathcal{T}^1$ and $\mathcal{T}^2$ were different at the start of the sweepout, but that they had to match at the end of it due to the properties of

our compression bubble tangles. Then we showed that locally these small sides are consistent with the sweepout, i.e., that points not swept stay on the same side. Doing so we proved that the small sides of the sphere at $t = 1$ are both different and the same, which is absurd.

The agreement function allows us to circumvent the impossibility of defining an orientation for each sphere similarly to what was done in Chapter 4. This natural orientation that we would like to define is an arrow pointing outward of the small side with respect to bubble tangles. However, doing so could lead to a situation where two spheres S_u and S_v are the same, but the associated arrows point in opposite directions. A $2D$ equivalent of this behaviour can be pictured by the following sweepout of $\mathbb{S}^2$: first, start with a circle on a sphere around the north pole. Then sweep the sphere by making the circle pass through the latitudes to finish close to the south pole. And finally, make this small circle go back to its initial position by following a meridian. Tracking the small side throughout this operation should illustrate the aforementioned phenomenon.

Furthermore, from an intuitive point of view, by defining two bubble tangles, we made sure that the sphere and U would sweep at least one of the torus knots, which is enough to get a contradiction. The agreement function turned out to be a way to formalise this approach.

5.5 Other settings

We specified our diagrams $\mathcal{D}_{p,q}$ for $p = n$ and $q = n + 1$, but our proof can easily be adapted to handle any coprime p, q and prove that $\mathcal{D}_{p,q}$ is a hard split link. However, our proof will provide a lower bound for $\mathrm{Unl}(\mathcal{D}_{p,q})$ that only depends on $\min(p, q)$ while the number of crossings of the diagram is larger than $pq - \min(p, q)$. Hence, the lower bound on crossing-complexity, which depends on the number of initial crossings, is the highest possible on the diagrams $\mathcal{D}(n, n + 1)$.

Additionally, Figure 5.10 shows a similar setting for which our arguments apply equally well, so that these diagrams are also hard split links with arbitrarily large crossing complexity.

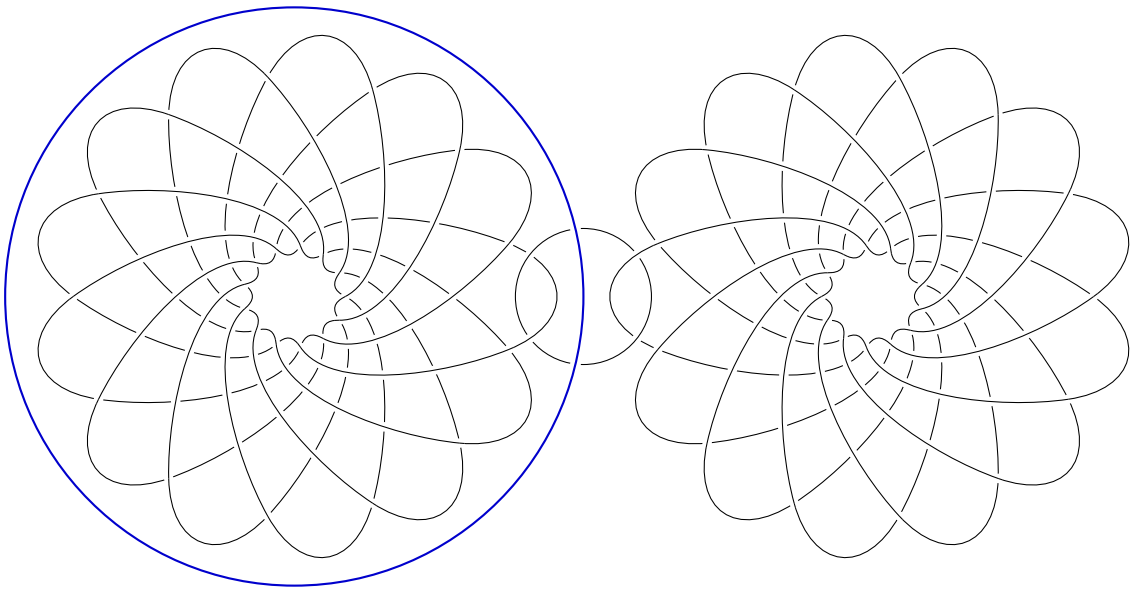

Figure 5.10: A set up with 4 link components. A first linked sublink is made of two unlinked torus knots $T_{p,q}$, each linked to the same unknot. The second sublink, is an unknot (blue), that can be separated from the other sublink by a sphere.

Chapter 6

Perspectives

In this thesis, we shed light on the computational properties of knots and links by drawing inspiration from techniques and methods of structural graph theory. In the following, we present some research lines opened up by our results.

6.1 Perspectives on the genus defects

Direct extensions. Our proof of Theorem A allows us to show that any stable property, or any monotonous invariant for our surface-minor relation, is decidable on Hopf arborescent links. Identifying any such property or invariant, in particular if it is connected to the 4-dimensional behaviour of knots, is then significant from a computational point of view.

Computing the 4-genus. We are interested in showing the decidability of the smooth and the topological 4-genus on vast classes of knots. Let us note that our proof could suggest that we settle this problem on the class of arborescent Hopf links. However, that is not the case for Hopf arborescent links of unbounded genus. Indeed, a natural approach is to try the link minor relation for a set of excluded minors characterising the property $\Delta_g(L) \leq k$ and do that iteratively for each k up to the genus of L. However, since our algorithm relies on hardcoded excluded families of minors, which we do not know, we can not apply it to test whether $\Delta_g(L) \leq k$ for non-constant values of k. Hence, this method could only work on classes of arborescent Hopf links with bounded genus. But there are finitely many diagrams of arborescent Hopf links, and thus this method proves nothing in that regard.

However, the relations that we exhibit remain relevant to study the 4-genus, and they could be applied successfully to set the decidability of 4-genus on other classes of links.

Relaxing the tree structure. Furthermore, the results of Chapter 3 exploit another key property of the surfaces associated to the class of knots that we study: their fibredness. We are interested in relaxing the tree structure of the plumbings to show that the defects

can be computed on knots obtained as boundaries of fibred surfaces. Such surfaces can be obtained through plumbings and deplumbings of Hopf bands [94], where the latter operation corresponds to cutting a band along an arc. This seemingly innocuous operation is in fact more challenging than it appears. In particular, allowing cutting along any arc makes it difficult to exhibit an underlying graph structure for the surfaces obtained and does not preserve their fibredness in the general case [93].

A line of work close to this question is an open conjecture about whether positive *braid links* are well-quasi-ordered in some sense [8]. Similarly to our knots, these links can be obtained as boundaries of surfaces obtained by plumbing of Hopf bands, except that their intersection patterns yield planar graphs that are not restricted to trees. However, they are also in some way more restrictive than our links since only positive Hopf bands are allowed. The direction of proof that we would like to explore is one inspired by what was presented in Section 2.2.4. That is to say, when considering a family of positive braid links, either their graphs admit arbitrarily large grids as minors, or they have bounded treewidth. The latter case can be handled by the Kruskal tree theorem. The former case requires establishing a connection between the relation of minors in graphs and the surface-minors.

Computational tractability. Finally, our methods provide the existence of an algorithm for each value of each defect but no explicit algorithms or upper bound on the complexity of the problem. We are interested in studying the complexity of computing the defects and finding explicit algorithms to compute them.

6.2 Perspectives on spherewidth and bubble tangles

Relation to treewidth. A natural question following this work is whether spherewidth is equivalent to the minimal treewidth over all possible diagrams of a knot? In other words, is the treewidth of a knot, up to a constant factor, a lower bound for its spherewidth? This question could be answered by studying how to recompose the knot and a diagram of the smallest possible treewidth from the information provided by a sphere decomposition.

The treewidth of the dual graph of the triangulation of a 3-manifold can be a relevant parameter for the design of FPT algorithms [87]. Hence, the question, similar to the one we answered in Chapter 4, is whether there exist 3-manifolds which do not admit triangulations of low treewidth [61]. A promising approach we want to explore is then to try to extend our obstruction to sweepouts of general 3-manifolds. Since sphere decompositions use only spheres, they are unfitted for 3-manifolds other than $\mathbb{S}^3$. Hence, we need to extend our objects to surfaces of higher genus.

Algorithm design. Our spherewidth relies at its core on tree-like sweep-outs of $\mathbb{S}^3$ highlighting an underlying tree structure of the knot. We want to investigate the possibility of designing parameterized algorithms that exploit our spherewidth. For example, the structure of the knot inbetween double bubbles of a sphere decomposition is close to a braid with a bounded number of strands. Algorithms could exploit this braid-like structure to solve a partial solution to a problem before recomposing solutions along double bubbles.

Furthermore, the sweep-outs of the 3-sphere could behave particularly well on linkless graphs that were presented in Chapter 1. An approach to connect branch decompositions and sphere decompositions when a graph is embedded in $\mathbb{S}^3$ is to encapsulate by spheres the parts of a spatial graph given by a branch decomposition of the abstract graph. It is difficult to achieve encapsulating only a given part in general because, for example, the branch decomposition could indicate only a cycle while some other cycle is linked with the first one in the embedding. This last condition can be avoided when the graph considered is linkless. Hence, we want to explore further the connections between branch decomposition of a linkless graph with sphere decompositions of one of its linkless embeddings. A possible application could be to generalise the celebrated ratcatcher algorithm [133] to the setting of linkless graphs, thereby showing that the spherewidth, branchwidth, or carvingwidth can be computed on these graphs in polynomial time.

Compression-representativity. We discussed the challenges of trying to compute the compression-representativity of an embedding of a spatial graph on a surface in Section 4.6. We are interested in facing this challenge, as well as the more difficult version consisting of computing the compression-representativity of the graph: that is to say, the largest compression-representativity of a surface on which the graph can be embedded.

As we discussed in Section 4.6, high genus surfaces raise specific problems for the computation of compression-representativity. In order to circumvent these, a promising line of work is to define an obstruction on subsurfaces only. When a family of spheres sweeps efficiently a surface of high genus, in particular, it sweeps efficiently each of the handles of the surface (which can be seen as the neighbourhood of a pair of once intersecting curves). Hence we want to prove that an obstruction, similar to ours, can be defined on handles, and not the surface as a whole. Since a handle is a torus with one boundary component, its compression-representativity is easier to compute. Hence, we would provide a lower bound on the spherewidth, which is a supremum of the compression-representativity of a handle, taken over all possible handles. The issue first discussed here is then partially solved, since exhibiting any handle with high compression-representativity would yield an obstruction to spherewidth.

Torus knots. The Theorem D provides a lower bound on the spherewidth of $\frac{2}{3}$ of the compression-representativity. A natural question is whether we can do better than $\frac{2}{3}$. We conjecture that it is possible to achieve $\frac{4}{3}$ in the case of torus knots.

Indeed, for surfaces of genus higher than 1, it is difficult to designate a small side for spheres with a number of intersections higher than the compression-representativity: such spheres can intersect the surface on non-contractible separating curves. In that situation, both sides of the sphere contain a non trivial part of the topology of the surface. Hence, it is something that our compression bubble tangles do not handle. However, if a sphere intersects a torus along a compressible curve, there must be a second one in the intersection. Hence, in that case, it seems easy to define our compression bubble tangle as long as its order is less than twice the compression-representativity. However, our proof that it is indeed a bubble tangle would require finding a fractional packing of two compressible curves instead of only

one. The announced bound of $\frac{4}{3}$ would be tight: we know a sphere decomposition of torus knots of such width.

6.3 Perspectives on the hardness of diagrams

Hard split links. Among the lines of research presented in the previous section, one offers to define an obstruction using a handle instead of the whole surface. Using an approach similar to the one we presented in Chapter 5, this stronger notion of obstruction could possibly be applied to prove the hardness of other families of diagrams, for example, the one pictured in Figure 6.1.

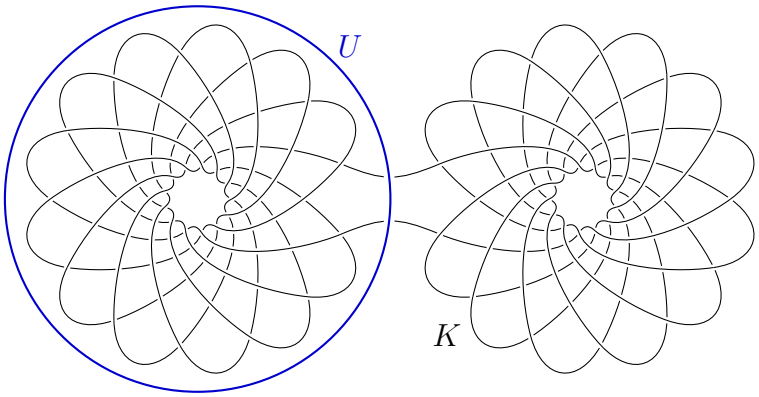

Figure 6.1: A link diagram made of two knots: K, a connected sum of two torus knots $T_{7,13}$ (black), and an unknot U (blue).

Hard unknots. Figure 6.2 pictures a diagrammatic construction transforming any connected sum of torus knots into an unknot (see [42]).

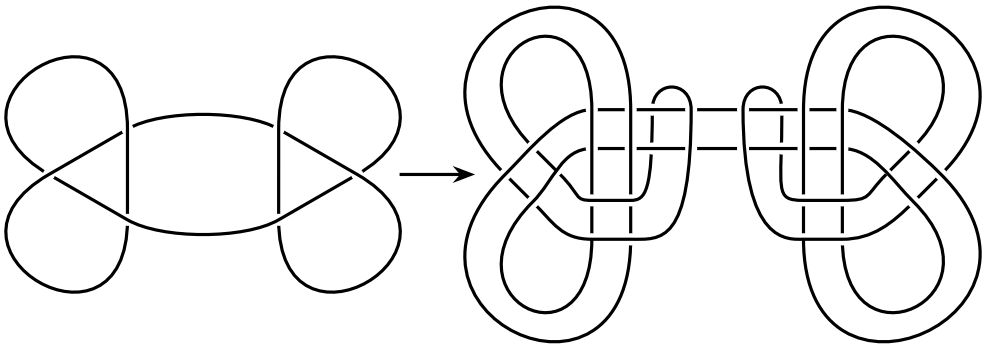

Figure 6.2: Obtaining an unknot from torus knots.

These diagrams feature characteristics similar to our examples of hard split links, and therefore we conjecture that these are hard unknots: when the torus knots are taken to be increasingly complicated, the number of extra crossings required to untangle the corresponding unknot tends to infinity.

List of publications

[A] Corentin Lunel and Arnaud de Mesmay, *A Structural Approach to Tree Decompositions of Knots and Spatial Graphs*, published in the Proceedings of the 39th International Symposium on Computational Geometry (SoCG 2023), Schloss Dagstuhl – Leibniz-Zentrum für Informatik. Arxiv: 2303.07982

[B] Pierre Dehornoy, Corentin Lunel, and Arnaud de Mesmay, *Hopf Arborescent Links, Minor Theory, and Decidability of the Genus Defect*, published in the Proceedings of the 40th International Symposium on Computational Geometry (SoCG 2024), Schloss Dagstuhl – Leibniz-Zentrum für Informatik. Invited to a Discrete & Computational Geometry special issue of Symposium on Computational Geometry 2024. Arxiv: 2312.09094

Bibliography

[1] The knot atlas: 8_{10}. `http://katlas.org/wiki/8_10`. Accessed: 2024-13-03.

[2] Wikipedia webpage of the Petersen family. `https://en.wikipedia.org/wiki/Petersen_family`. Accessed: 2024-05-02.

[3] Colin C. Adams. *The knot book*. American Mathematical Society, 1994.

[4] Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Transactions of the American Mathematical Society*, 358(9):3821–3850, 2006.

[5] James W. Alexander. An example of a simply connected surface bounding a region which is not simply connected. *Proceedings of the National Academy of Sciences of the United States of America*, 10 1:8–10.

[6] Omid Amini, Frédéric Mazoit, Nicolas Nisse, and Stéphan Thomassé. Submodular partition functions. *Discrete Mathematics*, 309(20):6000–6008, 2009.

[7] Sebastian Baader. Positive braids of maximal signature. *L'Enseignement Mathématique*, 59(3):351–358, 2014.

[8] Sebastian Baader and Pierre Dehornoy. Minor theory for surfaces and divides of maximal signature. *arXiv preprint arXiv:1211.7348*, 2012.

[9] Sebastian Baader, Pierre Dehornoy, and Livio Liechti. Minor theory for quasipositive surfaces. In Athanase Papadopoulos, editor, *Essays in geometry, dedicated to Norbert A'Campo*, pages 351–358. EMS Pres, 2023.

[10] Sebastian Baader, Peter Feller, Lukas Lewark, and Livio Liechti. On the topological 4-genus of torus knots. *Transactions of the American Mathematical Society*, 370(4):2639–2656, 2018.

[11] Filippo Baroni. Classification of genus-two surfaces in $\mathbb{S}^3$. *arXiv preprint arXiv:2309.05387*, 2023.

[12] R.H. Bing. *The geometric topology of 3-manifolds*, volume 40. American Mathematical Society, 1983.

[13] Ryan Blair and Makoto Ozawa. Height, trunk and representativity of knots. *Journal of the Mathematical Society of Japan*, 71(4):1105–1121, 2019.

[14] Hans L. Bodlaender. Treewidth: Characterizations, applications, and computations. *Graph-Theoretic Concepts in Computer Science*, pages 1–14, 2006.

[15] Hans L. Bodlaender, Benjamin A. Burton, Fedor V. Fomin, and Alexander Grigoriev. Knot diagrams of treewidth two. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 80–91. Springer, 2020.

[16] Francis Bonahon and Laurence C Siebenmann. New geometric splittings of classical knots and the classification and symmetries of arborescent knots. Preprint available on the authors' webpage, 1979.

[17] Nicolas Bourbaki. *Elements of the History of Mathematics*. Springer-Verlag Berlin Heidelberg, 1 edition, 1994.

[18] Gerhard Burde and Heiner Zieschang. *Knots*. Walter de gruyter, 2002.

[19] Benjamin A. Burton. The HOMFLY-PT polynomial is fixed-parameter tractable. In *Proceedings of the 34th International Symposium on Computational Geometry (SoCG 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[20] Benjamin A. Burton, Ryan Budney, William Pettersson, et al. Regina: Software for low-dimensional topology. http://regina-normal.github.io/, 1999–2023.

[21] Benjamin A. Burton, Hsien-Chih Chang, Maarten Löffler, Clément Maria, Arnaud de Mesmay, Saul Schleimer, Eric Sedgwick, and Jonathan Spreer. Hard diagrams of the unknot. *Experimental Mathematics*, 0(0):1–19, 2023.

[22] Benjamin A. Burton, Herbert Edelsbrunner, Jeff Erickson, and Stephan Tillmann. Computational geometric and algebraic topology. *Oberwolfach Reports*, 12(4):2637–2699, 2016.

[23] Benjamin A. Burton, Clément Maria, and Jonathan Spreer. Algorithms and complexity for Turaev–Viro invariants. *Journal of Applied and Computational Topology*, 2(1-2):33–53, 2018.

[24] Gregory Chambers and Yevgeny Liokumovich. Converting homotopies to isotopies and dividing homotopies in half in an effective way. *Geometric and Functional Analysis*, 24, 11 2013.

[25] J. H. Conway and C. McA. Gordon. Knots and links in spatial graphs. *Journal of Graph Theory*, 7(4), 1983.

[26] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.

[27] Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 4. Springer, 2015.

[28] Arnaud de Mesmay, Jessica Purcell, Saul Schleimer, and Eric Sedgwick. On the tree-width of knot diagrams. *Journal of Computational Geometry*, 10(1):164–180, 2019.

[29] Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. The unbearable hardness of unknotting. *Advances in Mathematics*, 381:107648, 2021.

[30] Reinhard Diestel. *Graph theory*. Number 173 in Graduate texts in mathematics. Springer, New York, 5th edition, 2016.

[31] Reinhard Diestel and Sang-il Oum. Tangle-Tree Duality: In Graphs, Matroids and Beyond. *Combinatorica*, 39(4):879–910, August 2019.

[32] Reinhard Diestel and Sang-il Oum. Tangle-tree duality in abstract separation systems. *Advances in Mathematics*, 377, 2021.

[33] Reinhard Diestel and Geoff Whittle. Tangles and the Mona Lisa. *arXiv preprint arXiv:1603.06652*, 2016.

[34] Frederic Dorn, Eelko Penninkx, Hans L Bodlaender, and Fedor V Fomin. Efficient exact algorithms on planar graphs: Exploiting sphere cut branch decompositions. In *European Symposium on Algorithms*, pages 95–106. Springer, 2005.

[35] Henry Dudeney. Perplexities, with some easy puzzles for beginners, 1913.

[36] I. A. Dynnikov. Arc-presentations of links: Monotonic simplification. *Fundamenta Mathematicae*, 190(1):29–76, 2006.

[37] David B.A. Epstein. Curves on 2-manifolds and isotopies. *Acta Mathematica*, 115:83–107, 1966.

[38] David Eppstein et al. Twenty-one proofs of euler's formula: V - e + f = 2. `https://ics.uci.edu/~eppstein/junkyard/euler/`, 2013. Accessed: 2024-05-02.

[39] Leonhard Euler. Solutio problematis ad geometriam situs pertinentis (in latin). *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 8:128–140, 1736.

[40] Leonhard Euler. Elementa doctrinae solidorum (in latin). *Novi Commentarii Academiae Scientiarum Petropolitanae*, pages 109–140, 1758.

[41] Ralph Fox. Some problems in knot theory. In *Topology of 3-manifolds and related topics (Proc. The Univ. of Georgia Institute)*, pages 168–176, Englewood Cliffs, N.J, 1961. Prentice-Hall.

[42] Michael H. Freedman, Zheng-Xu He, and Zhenghan Wang. Mobius energy of knots and unknots. *Annals of Mathematics*, 139(1):1–50, 1994.

[43] Michael Hartley Freedman. The topology of four-dimensional manifolds. *Journal of Differential Geometry*, 17(3):357–453, 1982.

[44] David Gabai. The Murasugi sum is a natural geometric operation. *Contemp. Math*, 20:131–143, 1983.

[45] David Gabai. *Genera of the Arborescent Links*. Memoirs of the American Mathematical Society. American Mathematical Society, 1986.

[46] David Gabai. Foliations and the topology of 3-manifolds. ii. *Journal of Differential Geometry*, 26(3):461–478, 1987.

[47] James Geelen, Bert Gerards, Neil Robertson, and Geoff Whittle. Obstructions to branch-decomposition of matroids. *Journal of Combinatorial Theory*, page 11, 2006.

[48] Emmanuel Giroux and Noah Goodman. On the stable equivalence of open books in three-manifolds. *Geometry & Topology*, 10(1):97–114, 2006.

[49] Cameron McA. Gordon and John S Luecke. Knots are determined by their complements. *Journal of the American Mathematical Society*, 2:371–415, 1989.

[50] Martin Grohe. Tangles and connectivity in graphs. In *Language and Automata Theory and Applications: 10th International Conference, LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings 10*, pages 24–41. Springer, 2016.

[51] Mikhael Gromov. Filling Riemannian manifolds. *Journal of Differential Geometry*, 18(1):1–147, 1983.

[52] Misha Gromov and Larry Guth. Generalizations of the Kolmogorov–Barzdin embedding estimates. *Duke Mathematical Journal*, 161(13):2549–2603, 2012.

[53] Wolfgang Haken. Theorie der normalflächen. *Acta Mathematica*, 105(3):245–375, 1961.

[54] Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *Journal of the ACM (JACM)*, 46(2):185–211, 1999.

[55] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge ; New York, 2002.

[56] Chuichiro Hayashi. The number of reidemeister moves for splitting a link. *Mathematische Annalen*, 332:239–252, 2005.

[57] Chuichiro Hayashi and Koya Shimokawa. Thin position of a pair (3-manifold, 1-submanifold). *Pacific Journal of Mathematics*, 197(2):301–324, February 2001.

[58] Qidong He and Scott A Taylor. Links, bridge number, and width trees. *Journal of the Mathematical Society of Japan*, 1(1):1–39, 2022.

[59] H. Heesch. *Untersuchungen zum Vierfarbenproblem (in german)*. B.I.-Hochschulskripten. Bibliographisches Institut, 1969.

[60] Kristóf Huszár and Jonathan Spreer. 3-manifold triangulations with small treewidth. In *Proceedings of the 35th International Symposium on Computational Geometry (SoCG 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.

[61] Kristóf Huszár, Jonathan Spreer, and Uli Wagner. On the treewidth of triangulated 3-manifolds. In *Proceedings of the 33th International Symposium on Computational Geometry (SoCG 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[62] Ken ichi Kawarabayashi, Yusuke Kobayashi, and Bruce Reed. The disjoint paths problem in quadratic time. *Journal of Combinatorial Theory, Series B*, 102(2):424–435, 2012.

[63] K. Ishihara. Enzyme action for topological entanglement in dna and knot theory. *Reactive and Functional Polymers*, 132:74–80, 2018.

[64] Sophie E Jackson, Antonio Suma, and Cristian Micheletti. How to fold intricately: using theory and experiments to unravel the properties of knotted proteins. *Current Opinion in Structural Biology*, 42:6–14, 2017.

[65] Hugo Jacob and Marcin Pilipczuk. Bounding twin-width for bounded-treewidth graphs, planar graphs, and bipartite graphs. *Graph-Theoretic Concepts in Computer Science*, pages 287–299, 2022.

[66] François Jaeger, Dirk L. Vertigan, and Dominic J.A. Welsh. On the computational complexity of the Jones and Tutte polynomials. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 108, pages 35–53. Cambridge University Press, 1990.

[67] Camille Jordan. *Cours d'analyse de l'École Polytechnique (in french)*. 1887.

[68] L.H. Kauffman. *Knots And Physics (Third Edition)*. Series On Knots And Everything. World Scientific Publishing Company, 2001.

[69] Louis H. Kauffman. The jones polynomial, knots, diagrams, and categories. *Bulletin of the American Mathematical Society*, 2023.

[70] Louis H. Kauffman and Sofia Lambropoulou. Hard unknots and collapsing tangles. *Introductory Lectures on Knot Theory*, pages 187–247, 2011.

[71] Dale Koenig and Anastasiia Tsvietkova. Unlinking, splitting, and some other np-hard problems in knot theory. *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1496–1507, 2021.

[72] Peter B Kronheimer and Tomasz S Mrowka. The genus of embedded surfaces in the projective plane. *Mathematical Research Letters*, 1(6):797–808, 1994.

[73] Joseph B. Kruskal. Well-quasi-ordering, the tree theorem, and Vazsonyi's conjecture. *Transactions of the American Mathematical Society*, 95:210–225, 1960.

[74] Greg Kuperberg. Quadrisecants of knots and links. *Journal of Knot Theory and its Ramifications*, 3(01):41–50, 1994.

[75] Greg Kuperberg. Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization. *Pacific Journal of Mathematics*, 301(1):189–241, September 2019.

[76] Greg Kuperberg and Eric Samperton. Coloring invariants of knots and links are often intractable. *Algebraic & Geometric Topology*, 21:1479–1510, 08 2021.

[77] Casimir Kuratowski. Sur le problème des courbes gauches en topologie (in french). *Fundamenta Mathematicae*, 15:271–283.

[78] Marc Lackenby. A polynomial upper bound on reidemeister moves. *Annals of Mathematics*, pages 491–564, 2015.

[79] Marc Lackenby. Elementary Knot Theory. In *Lectures on Geometry*. Oxford University Press, 01 2017.

[80] Marc Lackenby. Some conditionally hard problems on links and 3-manifolds. *Discrete & Computational Geometry*, 58:580–595, 2017.

[81] Marc Lackenby. Algorithms in 3-manifold theory. *Surveys in Differential Geometry*, 2020.

[82] Marc Lackenby. The efficient certification of knottedness and Thurston norm. *Advances in Mathematics*, 387:107796, 2021.

[83] Marc Lackenby. Unknot recognition in quasi-polynomial time, 2021. Talk with slides available on the author's webpage : http://people.maths.ox.ac.uk/lackenby/quasipolynomial-talk.pdf.

[84] Livio Liechti. On the genus defect of positive braid knots. *Algebraic & Geometric Topology*, 20(1):403–428, 2020.

[85] Laurent Lyaudet, Frédéric Mazoit, and Stéphan Thomassé. Partitions versus sets: a case of duality. *European journal of Combinatorics*, 31(3):681–687, 2010.

[86] J.A. Makowsky and J.P. Mariño. The parametrized complexity of knot polynomials. *Journal of Computer and System Sciences*, 67(4):742–756, December 2003.

[87] Clément Maria. Parameterized Complexity of Quantum Knot Invariants. In *Proceedings of the 37th International Symposium on Computational Geometry (SoCG 2021)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2021.

[88] Clément Maria and Jessica Purcell. Treewidth, crushing and hyperbolic volume. *Algebraic & Geometric Topology*, 19(5):2625–2652, 2019.

[89] Andrei Andreevich Markov. The insolubility of the problem of homeomorphy. In *Doklady Akademii Nauk*, volume 121, pages 218–220. Russian Academy of Sciences, 1958.

[90] Jiří Matoušek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *Journal of the European Mathematical Society*, 13(2):259–295, 2010.

[91] Sergei Vladimirovich Matveev. *Algorithmic topology and classification of 3-manifolds*. Springer, 2007.

[92] Jonathan Miller and Ramin Naimi. An algorithm for detecting intrinsically knotted graphs. *Experimental Mathematics*, 23:12 – 6, 2011.

[93] Filip Misev. Cutting arcs for torus links and trees. *Bulletin de la Société Mathématique de France*, 145:575–602, 2014.

[94] Filip Misev. *On the plumbing structure of fibre surfaces*. PhD thesis, Universität Bern, 2016.

[95] Filip Misev. Hopf bands in arborescent Hopf plumbings. *Osaka Journal of Mathematics*, 56(2):375 – 389, 2019.

[96] Edwin E. Moise. *Geometric Topology in Dimensions 2 and 3*. 1 edition, 1977.

[97] Kunio Murasugi. On a certain subgroup of the group of an alternating link. *American Journal of Mathematics*, 85(4):544–550, 1963.

[98] Kunio Murasugi. On the braid index of alternating links. *Transactions of the American Mathematical Society*, 326(1):237–260, 1991.

[99] Wendy J. Myrvold and Jennifer Woodcock. A large set of torus obstructions and how they were discovered. *Electron. J. Comb.*, 25:1, 2018.

[100] Jennifer Jones Thomas W. Mattman Nancy Ho, James Godzik and Dan Sours. Invisible knots and rainbow rings: Knots not determined by their determinants. *Mathematics Magazine*, 93(1):4–18, 2020.

[101] C. St. J. A. Nash-Williams. On well-quasi-ordering finite trees. *Mathematical Proceedings of the Cambridge Philosophical Society*, 59(4):833–835, 1963.

[102] Brendan Owens. On slicing invariants of knots. *Transactions of the American Mathematical Society*, 362(6):3095–3106, 2010.

[103] Makoto Ozawa. Bridge position and the representativity of spatial graphs. *Topology and its Applications*, 159(4):936–947, 2012.

[104] Burak Ozbagci and Patrick Popescu-Pampu. Generalized plumbings and murasugi sums. *Arnold Mathematical Journal*, 2(1):69–119, dec 2015.

[105] John Pardon. On the distortion of knots on embedded surfaces. *Annals of Mathematics*, 174(1):637–646, July 2011.

[106] Grisha Perelman. The entropy formula for the ricci flow and its geometric applications, 2002. `arXiv:math/0211159`.

[107] Grisha Perelman. Finite extinction time for the solutions to the ricci flow on certain three-manifolds, 2003. `arXiv:math/0307245`.

[108] Grisha Perelman. Ricci flow with surgery on three-manifolds, 2003. `arXiv:math/0303109`.

[109] Lisa Piccirillo. The Conway knot is not slice. *Annals of Mathematics*, 191(2):581–591, 2020.

[110] Serge A. Plotkin, Satish Rao, and Warren D. Smith. Shallow excluded minors and improved graph decompositions. *ACM-SIAM Symposium on Discrete Algorithms*, 1994.

[111] Candice Reneé Price. Applications of knot theory: Using knot theory to unravel biochemistry mysteries. In *Advances in the Mathematical Sciences*, pages 173–186. Springer International Publishing, 2016.

[112] Neil Robertson and P. D. Seymour. Graph minors. iv. tree-width and well-quasi-ordering. *Journal of Combinatorial Theory, Series B*, 48(2):227–254, 1990.

[113] Neil Robertson, P. D. Seymour, and Robin Thomas. Linkless embeddings of graphs in 3-space. *Bulletin of the American Mathematical Society*, 28(1):84–89, january 1993.

[114] Neil Robertson, P. D. Seymour, and Robin Thomas. Sachs' linkless embedding conjecture. *Journal of Combinatorial Theory, Series B*, 64(2):185–227, 1995.

[115] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, August 1986.

[116] Neil Robertson and Paul D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, July 1991.

[117] Neil Robertson and Paul D. Seymour. Graph minors. XI. Circuits on a Surface. *Journal of Combinatorial Theory. Series B*, 60(1):72–106, January 1994.

[118] Neil Robertson and Paul D. Seymour. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.

[119] Neil Robertson and P.D Seymour. Graph minors. vii. disjoint paths on a surface. *Journal of Combinatorial Theory, Series B*, 45(2):212–254, 1988.

[120] Dale Rolfsen. *Knots and links*. AMS Chelsea Pub, Providence, R.I, 2003. OCLC: ocm52901393.

[121] Colin Rourke and Brian Sanderson. *Introduction to Piecewise-Linear Topology*. Springer Berlin, Heidelberg, 1972.

[122] Lee Rudolph. Quasipositivity as an obstruction to sliceness. *Bulletin of the American Mathematical Society*, 29(1):51–59, 1993.

[123] Lee Rudolph. Positive links are strongly quasipositive. *Geometry & Topology Monographs, Volume 2: Proceedings of the Kirbyfest*, pages 555–562, 1998.

[124] Juanjo Rué, Ignasi Sau, and Dimitrios M. Thilikos. Dynamic programming for graphs on surfaces. *ACM Trans. Algorithms*, 10(2), feb 2014.

[125] Horst Sachs. On a spatial analogue of Kuratowski's theorem on planar graphs—an open problem. In *Graph theory*, pages 230–241. Springer, 1983.

[126] Makoto Sakuma. Minimal genus seifert surfaces for special arborescent links. *Osaka Journal of Mathematics*, 31:861–905, 1994.

[127] Martin Scharlemann. Thin position in the theory of classical knots. In *Handbook of knot theory*, pages 429–459. Elsevier, 2005.

[128] Martin Scharlemann, Jennifer Schultens, and Toshio Saito. *Lecture notes on generalized Heegaard splittings*. World Scientific, 2016.

[129] Martin Scharlemann and Abigail Thompson. Thin position for 3-manifolds. In *Geometric Topology: Joint US-Israel Workshop on Geometric Topology, June 10-16, 1992, Technion, Haifa, Israel*, volume 164, page 231. American Mathematical Society, 1994.

[130] Jennifer Schultens. *Introduction to 3-manifolds*, volume 151. American Mathematical Society, 2014.

[131] Jennifer Schultens. The bridge number of a knot. In *Encyclopedia of Knot Theory*, pages 229–242. Chapman and Hall/CRC, 2021.

[132] Paul D. Seymour. Packing circuits in eulerian digraphs. *Combinatorica*, 16:223–231, 1996.

[133] Paul D. Seymour and Robin Thomas. Call routing and the ratcatcher. *Combinatorica*, 14(2):217–241, 1994.

[134] John R. Stallings. Constructions of fibred knots and links. *Proc. Symp. Pure Math.*, AMS 27:315–319, 1975.

[135] Scott Taylor and Maggy Tomova. Additive invariants for knots, links and graphs in 3–manifolds. *Geometry & Topology*, 22(6):3235–3286, 2018.

[136] Peter Teichner. Slice knots: Knot theory in the 4th dimension, 2011.

[137] Alan Mathison Turing. *Solvable and unsolvable problems*. Penguin Books London, 1954.

[138] Kurt van Reidemeister. Elementare Begründung der Knotentheorie (in german). *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 5:24–32, 1927.

[139] Klaus Von Wagner. Über eine Eigenschaft der ebenen Komplexe (in german). *Mathematische Annalen*, 114:570–590, 1937.

[140] Friedhelm Waldhausen. On irreducible 3-manifolds which are sufficiently large. *Annals of Mathematics*, 87:56–88, 1968.

[141] Shmuel Weinberger. Homology manifolds. *Handbook of geometric topology*, pages 1085–1102, 2002.

[142] Wilbur Whitten. Isotopy types of knot spanning surface. *Topology*, 12:373–380, 1973.

Chapitre 7

Introduction en français

Cette thèse se place au sein du domaine de la topologie algorithmique : un domaine mathématique dont le but est de formaliser et traiter les questions topologiques d'un point de vue informatique. En particulier, nous nous intéressons aux riches connexions qui apparaissent entre la théorie des graphes et la théorie des nœuds. Dans ce qui suit, nous présenterons de manière générale ces domaines en soulignant les liens qui apparaissent entre eux. Ensuite, nous présenterons plus précisément les contributions et l'organisation de cette thèse.

Ce premier chapitre est destiné à être lisible sans connaissances mathématiques trop poussées, bien que les concepts et les notions présentés seront plus difficiles à comprendre lorsque nous passerons aux contributions.

7.1 Présentation Générale

Topologie. La topologie peut être définie comme la branche des mathématiques qui étudie les *formes*. Plus précisément, elle se concentre sur les propriétés qui sont préservées par les opérations continues, ces propriétés étant appelées propriétés topologiques. Par exemple, étirer, tordre ou courber sont des opérations continues, alors que découper ou coller ne le sont pas. Un exemple de propriété préservée par ces opérations continues est la capacité d'atteindre n'importe quel point d'un objet donné à partir de n'importe quel autre, un objet satisfaisant cette propriété est appelé **connexe**. Si un objet est composé d'exactement deux parties disjointes, aucune opération de pliage, de torsion ou d'expansion ne nous permettra d'atteindre un point de la seconde partie à partir d'un point de la première partie. Ce n'est que lorsque les deux parties fusionneront, c'est-à-dire seront collées (ce qui n'est pas une opération continue), que l'on pourra joindre les deux points susmentionnés et que l'objet sera connexe. Réciproquement, la flexion, la torsion ou l'expansion ne changeront jamais le fait qu'un objet est connecté.

Historiquement, la naissance de la topologie peut être attribuée [17, Chapter 10] à Riemann, dont les travaux dans les années 1850 ont conduit à la définition des espaces topologiques, qui sont les blocs élémentaires des théories topologiques. Cependant, des résultats

fondamentaux comme la formule d'Euler remontent à Descartes au XVIIe siècle (cette formule sera discutée plus loin dans cette section dans la preuve de la Proposition 7.1), et la notion clé de continuité avait déjà été imaginée par des philosophes grecs tels qu'Aristote. Aujourd'hui, c'est un domaine de recherche très actif dont les résultats sont répandus dans de nombreuses branches des mathématiques. Parmi les problèmes du prix du millénaire, qui sont sept problèmes mathématiques célèbres et très difficiles sélectionnés par le Clay Mathematics Institute en 2000, un seul a été résolu [106, 107, 108], il s'agit de la conjecture de Poincaré qui est de nature topologique. Récemment, avec l'essor de l'informatique théorique, un besoin de calculer les propriétés topologiques de modèles informatiques est apparu. Ceci a favorisé le développement de la topologie algorithmique telle que présentée ci-dessus, et a alimenté des échanges florissants d'outils, de méthodes et de problèmes à étudier entre la topologie classique et l'informatique théorique.

Comme présenté plus haut, la topologie se focalise sur les propriétés préservées par opérations continues. Il est donc naturel que la topologie considère que deux objets sont équivalents lorsqu'ils ne diffèrent que par une opération continue. Selon l'opération en question, la relation d'équivalence sera différente. L'une de ces relations d'équivalence, des plus fondamentales, découle des homéomorphismes : un **homéomorphisme** est une application bijective et continue entre deux espaces dont la bijection réciproque est elle aussi continue. Toutes les propriétés topologiques intrinsèques à deux objets homéomorphes sont les mêmes. Par exemple, $\mathbb{R}$ et $\mathbb{R}_+^*$ sont homéomorphes par la fonction $\exp : \mathbb{R} \to \mathbb{R}_+^*$. Cependant, il n'existe pas d'homéomorphisme entre $[0,1]$ et $\mathbb{R}$ car ils sont topologiquement différents. Une différence topologique entre $[0,1]$ et $\mathbb{R}$ est le fait que partout dans $\mathbb{R}$ il est possible de se déplacer vers l'avant et l'arrière, ce qui n'est pas le cas aux extrémités de $[0,1]$ où seul un de ces mouvements est possible (voir Figure 7.1).

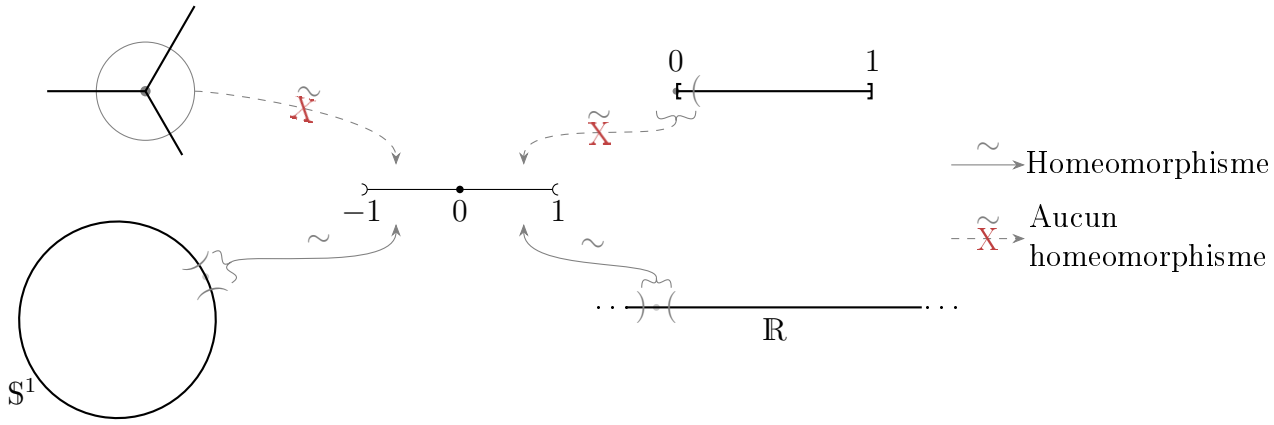

FIGURE 7.1 : Deux objets non homéomorphes à $]-1,1[$ au voisinage des points présentés (3 ou seulement 1 direction de déplacement sont possibles) tandis que $\mathbb{S}^1$ (un cercle) et $\mathbb{R}$ sont localement homéomorphes à $]-1,1[$.

Intuitivement, cette notion de nombre de directions de déplacement, ou nombre de degrés de liberté autour d'un point est ce que l'on appelle la **dimension** d'un objet. Pour l'instant, cette définition est vague, car elle peut dépendre du point considéré. Soyons plus formels et définissons plus précisément un objet commun étudié en topologie appelé **n-variété**. Un

ensemble M est une n-variété si localement il ressemble à l'espace a n dimension $\mathbb{R}^n$, c'est-à-dire si pour chaque point de M il existe un voisinage homéomorphe à la n-boule ouverte unité de $\mathbb{R}^n$: $\{x \in \mathbb{R}^n \mid \|x\| < 1\} = \mathbb{B}^n$. Cette dernière propriété est naturellement préservée par les homéomorphismes. Par conséquent, tout ensemble homéomorphe à une n-variété est aussi une n-variété. Ainsi, $\mathbb{R}$ est une 1-variété alors que $[0, 1]$ ne l'est pas.

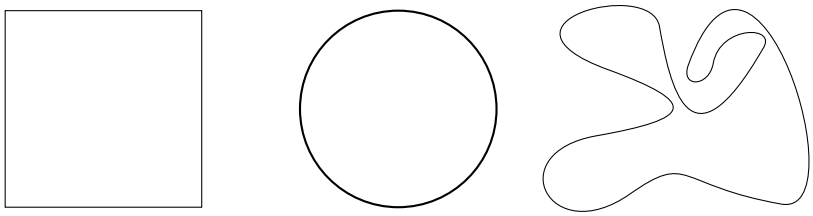

FIGURE 7.2 : Trois $\mathbb{S}^1$, c'est-à-dire trois objets homéomorphes à un cercle.

Nous pouvons d'ores et déjà classer les 1-variétés connexes. S'il est possible, en partant d'un point et en gardant une direction, d'atteindre à nouveau ce point de départ, alors la 1-variété est en fait homéomorphe à un cercle noté $\mathbb{S}^1$. Un tel cercle est communément décrit par $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$, et nous dirons que la 1-variété est un $\mathbb{S}^1$ (voir Figure 7.2). Sinon, la marche précédente ne se terminera jamais et l'ensemble est une ligne, c'est-à-dire qu'il est homéomorphe à $\mathbb{R}$. Plus généralement, comme il est commun pour les topologues, nous abrégerons l'expression « est homéomorphe à » par « est ». Ainsi, une n-sphère notée $\mathbb{S}^n$, est toute n-variété homéomorphe à la sphère euclidienne unité de $\mathbb{R}^{n+1}$: $\{x \in \mathbb{R}^{n+1} \mid \|x\| = 1\}$. En particulier, nous choisissons la sphère unité comme représentant, mais n'importe quel rayon donnerait une définition équivalente puisque augmenter ou diminuer le rayon d'une sphère sont des opérations continues.

Avant de passer aux surfaces, qui sont les 2-variétés, introduisons une deuxième propriété topologique qui sera un concept clé dans le reste de cette section : la plongeabilité. En effet, une opération topologique fondamentale est celle de plongement. Un **plongement** $j : X \to Y$ d'un objet X dans un objet Y est un homéomorphisme f de X sur son image $f(X)$. Si une telle application existe, on dit que X est **plongeable** dans Y. Puisque la composition de deux homéomorphismes est un homéomorphisme, il s'ensuit que la notion de plongeabilité est une propriété préservée par homéomorphismes.

Comme $[0, 1]$ peut être plongé dans $\mathbb{R}$ par l'inclusion naturelle, et $\mathbb{R}$ dans $[0, 1]$ par $\frac{1}{2} + \frac{1}{\pi} \arctan$, nous ne pouvons pas distinguer $[0, 1]$ et $\mathbb{R}$ par leurs plongements. En effet, les objets dans lesquels ils peuvent être plongés, ou les objets qui peuvent être plongés en eux sont les mêmes. Cependant, une telle notion nous permet de distinguer les deux 1-variétés que sont $\mathbb{R}$ et $\mathbb{S}^1$: $\mathbb{R}$ se plonge dans $\mathbb{S}^1 = \left\{e^{i\theta} \mid \theta \in [0, 2\pi)\right\}$[1] $\subset \mathbb{C}$ via $x \mapsto e^{i \arctan(x)}$ mais il n'y a pas de plongement de $\mathbb{S}^1$ dans $\mathbb{R}$ puisqu'un tel plongement nécessiterait deux chemins disjoints dans $\mathbb{R}$ joignant les images de deux points disjoints de $\mathbb{S}^1$.

[1]Notez que cette description de $\mathbb{S}^1$ par $\phi : [0, 2\pi) \to \mathbb{S}^1$ telle que $\phi(\theta) = e^{i\theta}$ est un exemple de bijection continue qui n'est pas un plongement : la bijection réciproque ϕ^{-1} n'est pas continue en $\phi(0)$. De fait, la limite de ϕ^{-1} en ce point, par continuité, est 0 d'un côté et 2π de l'autre.

Surfaces. Gagnons une dimension et concentrons-nous maintenant sur les 2-variétés appelées **surfaces**. D'abord, considérons quelques exemples. Le plan $\mathbb{R}^2$ et la 2-boule unité ouverte sont tous deux des surfaces : autour de chaque point de ces deux objets, il existe un voisinage homéomorphe à une 2-boule ouverte (prendre n'importe quel disque assez petit convient). Comme il est illustré sur la Figure 7.3, et ce à une petite courbure près, la propriété précédente est aussi vérifiée sur le tube infini $\mathbb{R} \times \mathbb{S}^1$ et le tore $\mathbb{T} = \mathbb{S}^1 \times \mathbb{S}^1$ qui peut être décrit comme la surface d'un donut.

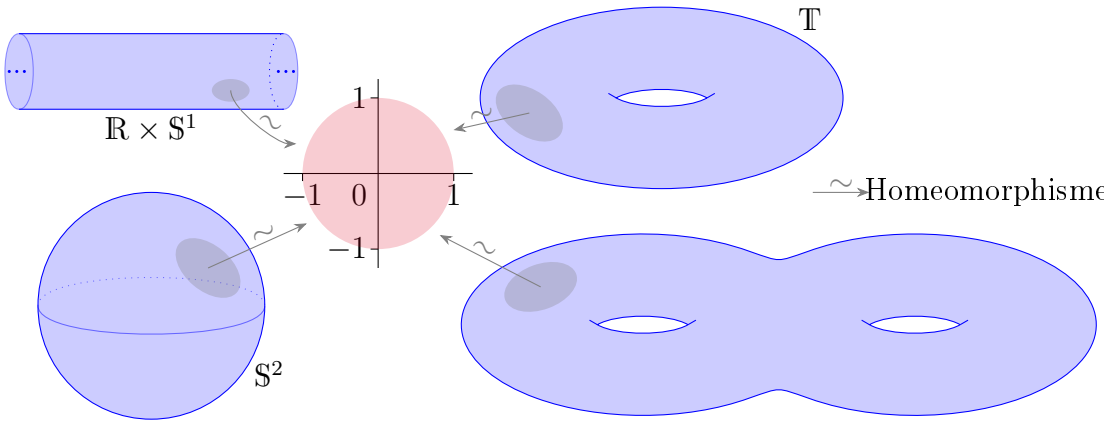

FIGURE 7.3 : Quelques exemples de surfaces, un voisinage de chacun des points est homéomorphe à une boule unité ouverte représentée au milieu.

Soulignons à nouveau que nous considérons généralement les objets à homéomorphisme près. Par conséquent, nous ne faisons pas de différence entre un cube et une 2-sphère ou entre un disque ouvert et un triangle, comme le montre la Figure 7.4. Une description commune de ce phénomène consiste à se représenter chaque objet comme s'il était élastique.

Nous allons maintenant expliquer pourquoi, parmi les exemples de surfaces précédents, le tore est topologiquement plus complexe que les autres surfaces. Ceci sera illustré par les transformations continues qui peuvent être appliquées à une **courbe** : qui est un dessin continu d'un cercle $\mathbb{S}^1$ sur la surface considérée (il n'est pas nécessaire qu'il s'agisse d'un plongement). Si une telle courbe peut être continûment contractée en un point, on dit qu'elle est **contractile**, sinon elle est non contractile. Sur une 2-sphère S, toute courbe est contractile (voir Figure 7.5). En effet, quitte à les imaginer comme des élastiques se déplaçant continûment sur S, nous pouvons déplacer ces courbes de façon à ce qu'elles restent sur une calotte de la sphère qui finira par se rétrécir jusqu'à un point (imaginez l'équateur glissant vers le pôle nord de la sphère de la Figure 7.5).

Sur le tore, il existe des courbes non contractiles, dont deux sont visibles sur la Figure 7.5, en rouge et en vert. De notre point de vue qui considère tout à déformation continue près sur la surface, ces courbes non contractiles ne sont pas triviales : elles ne sont pas équivalentes à des points alors que les courbes contractiles le sont. De telles courbes non contractiles témoignent de l'existence d'un trou dans le tore qui n'existe pas dans les 2-sphères. Parmi les surfaces de la Figure 7.3, seules la surface en bas à droite et $\mathbb{T}$, le tore, admettent des courbes non contractiles. Il est intéressant de noter que la surface en bas à droite possède

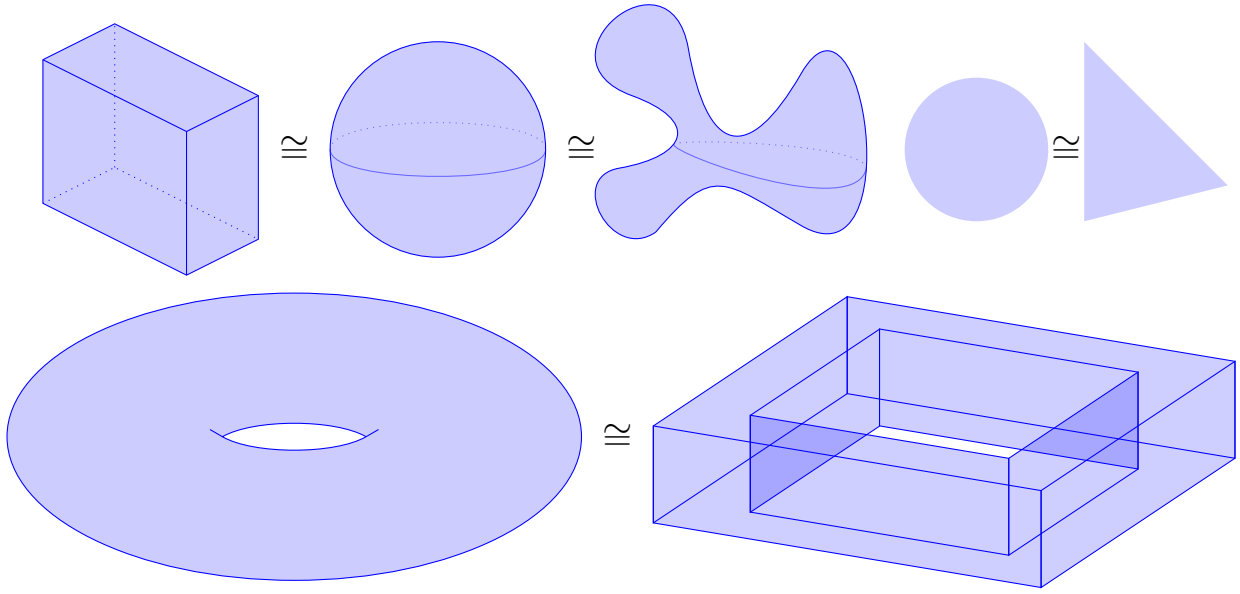

FIGURE 7.4 : Quelques surfaces homéomorphes à : la 2-sphère (en haut à gauche), la 2-boule unité ouverte (en haut à droite) et le tore (en bas). La relation d'homéomorphisme est notée par $\cong$.

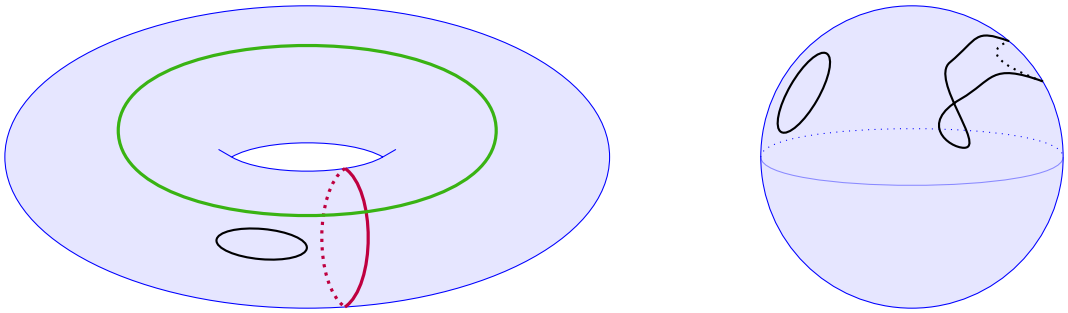

FIGURE 7.5 : Quelques courbes non contractiles du tore, un méridien en rouge et une latitude en vert. Les autres courbes, en noir, peuvent être continûment contractées en un point.

intuitivement deux trous (deux courbes non contractiles distinctes et disjointes) alors que le tore n'en possède qu'un seul.

Les exemples précédents de surfaces présentent une autre différence topologique : certaines d'entre elles sont *compactes*, comme la 2-sphère et le tore, alors que $\mathbb{R}^2$, la 2-boule ouverte $\mathbb{B}^2$ ou le tube infini $\mathbb{S}^1 \times \mathbb{R}$ ne le sont pas. Ces dernières surfaces présentent une sorte de comportement infini que les premières n'ont pas. La définition précise d'un ensemble compact nécessiterait un niveau de détail dépassant les ambitions de cette section. Nous nous contenterons donc de ce qui suit : une surface n'est pas compacte si l'on peut trouver une suite de points de l'ensemble qui converge « en dehors » de l'ensemble, comme la suite $(1 - \frac{1}{n}, 0)_{n \in \mathbb{N}}$ dans le 2-disque ouvert unité qui converge vers $(1, 0)$, qui est « en dehors » du disque. Du point de vue informatique, la propriété de compacité est particulièrement utile.

En effet, les surfaces compactes peuvent être décrites par des triangles fermés collés sur leurs bords : elles peuvent être *triangulées* par un nombre fini de triangles [130]. Par exemple, les tores sont souvent représentés par un rectangle dont les côtés opposés sont identifiés (voir le bas de la Figure 7.6) : en collant deux de ces côtés, on obtient un tube. La dernière étape consiste à coller ses extrémités pour obtenir le tore.

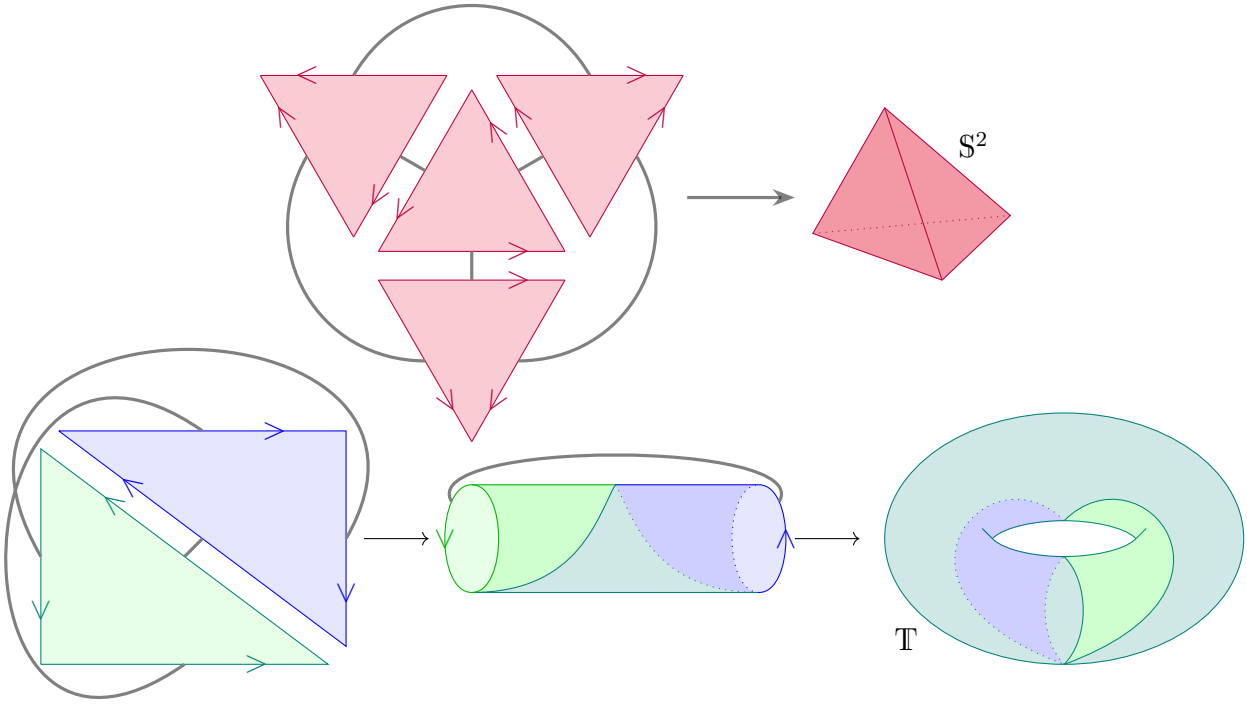

FIGURE 7.6 : Représentation d'une 2-sphère et d'un tore par identification de paires de segments sur les bords de triangles en accord avec les liens gris et les orientations.

Les surfaces compactes peuvent être données comme entrées à des ordinateurs, ou leurs équivalents abstraits que sont les **machines de Turing**, grâce à leurs descriptions par un nombre fini de triangles et la façon dont ils sont collés sur leurs bords. Une telle représentation discrète peut être surprenante, car la topologie se focalise sur des structures et opérations continues qui semblent être en opposition avec une description discrète. Cependant, ce n'est pas le cas : ces représentations discrètes sont équivalentes aux continues. Elles englobent donc l'ensemble des propriétés topologiques des surfaces tout en étant beaucoup plus commodes à manipuler algorithmiquement. C'est pourquoi les pratiquants de la topologie algorithmique ont tendance à préférer les variétés compactes à celles non compactes.

Nous avons vu précédemment que $\mathbb{S}^2$ et $\mathbb{R}^2$ n'ont pas de courbes non contractiles. En fait, ils ont une topologie très similaire, à tel point qu'ils sont identiques à un point près : il existe un homéomorphisme que nous désignerons par ϕ_n entre chaque $\mathbb{R}^n$ et $\mathbb{S}^n \setminus \{N\}$ où N est le « pôle nord » de $\mathbb{S}^n$. Détaillons la construction de cet homéomorphisme, appelé **projection stéréographique** (voir Figure 7.7). Nous commençons par placer $\mathbb{S}^n$ sur $\mathbb{R}^n$ vu comme un hyperplan de $\mathbb{R}^{n+1}$. Ensuite, pour chaque point x de $\mathbb{R}^n$, nous définissons la demi-droite ℓ_x commençant à $N = (0, \ldots, 0, 1)$ et passant par x. Il y a exactement un point

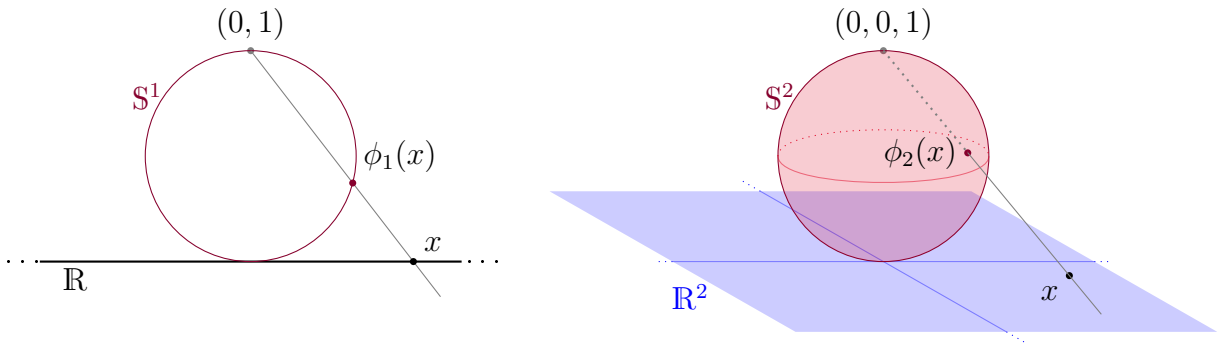

FIGURE 7.7 : Gauche : la ligne réelle $\mathbb{R}$ compactifiée par l'application ϕ_1 en $\mathbb{S}^1$, qui est représentée ici par la sphère unité de $\mathbb{R}^2$. Droite : le plan $\mathbb{R}^2$ compactifié par l'application ϕ_2 en $\mathbb{S}^2$, qui est représenté ici par la sphère unité de $\mathbb{R}^3$.

d'intersection entre ℓ_x et $\mathbb{S}^n$: il s'agit de $\phi_n(x)$. Remarquons qu'intuitivement, « l'infinité » de $\mathbb{R}^n$ est envoyée sur N par continuité. Puisque $\mathbb{S}^n$ est compact alors que $\mathbb{R}^n$ ne l'est pas, nous dirons que $\mathbb{S}^n$ est la *compactification d'Alexandrov* de $\mathbb{R}^n$. La projection stéréographique nous permet de réduire les problèmes dans $\mathbb{R}^n$ à des problèmes dans $\mathbb{S}^n$, qui sont des variétés que nous pouvons manipuler plus facilement d'un point de vue algorithmique.

Enfin, le fait que toute 2-variété puisse être triangulée est également vérifié par les 3-variétés [96, 130]. Dans ce dernier cas, les éléments constitutifs ne sont pas des triangles mais des tétraèdres (la 2-sphère, représentée en haut à droite de la Figure 7.6, si elle était remplie à l'intérieur, serait un tétraèdre) qui sont collés sur leurs faces triangulaires. Cela signifie que les 3-variétés peuvent également être données en entrée aux machines de Turing et étudiées par l'informatique théorique. Cette propriété essentielle n'est pas vérifiée par les n-variétés lorsque $n \geq 4$ et ceci est révélateur du fait que les espaces de basse dimension se prêtent particulièrement bien aux questions algorithmiques. Au contraire, en haute dimension, de nombreuses questions ne peuvent même pas être posées dans un cadre algorithmique.

Les triangulations de surfaces peuvent en pratique être observées sur de vielles images de synthèse en 3D, où les triangles qui composent les surfaces représentant des objets peuvent être distingués. Aujourd'hui, ces triangles sont plus difficiles à délimiter, ils sont suffisamment petits pour créer l'illusion d'une surface continue et lisse. Cependant, les triangulations sous-jacentes rendent les calculs et simulations possibles. De plus, de la même manière que les courbes non contractiles sont essentielles pour définir les trous dans les surfaces, et sont donc constitutives du comportement topologique des surfaces, les surfaces au sein des 3-variétés sont constitutives de leurs propriétés topologiques. Ainsi, les calculs impliquant des 3-variétés exploitent souvent des surfaces. Il va sans dire que les calculs impliquant des 3-variétés sont importants puisque l'espace dans lequel nous vivons est, à notre échelle, une 3-variété : il s'ensuit que les simulations du monde réel ont besoin de surfaces, et d'une compréhension algorithmique de celles-ci.

Graphes. Comme nous l'avons expliqué, les structures discrètes sont très pratiques pour des considérations algorithmiques. Il est donc naturel de s'intéresser aux graphes, qui occupent une place majeure au sein des mathématiques et de l'informatique ; leur étude est

l'objet de la théorie des graphes. Un graphe se compose d'une paire d'informations : d'une part, un ensemble de points appelés **sommets**, et d'autre part, des liens entre eux, appelés **arêtes**. Les arêtes et sommets peuvent être complétés par d'autres informations, telles que des poids, des couleurs, des orientations ou des étiquettes. Par conséquent, en tant que structures mathématiques, ils sont très flexibles pour englober des informations sur les objets et les relations entre eux, ils ont donc un très large éventail d'applications en tant que modèles. Tout organigramme d'entreprise, carte de métro, réseau ou diagramme de dépendance est une sorte de graphe. Les graphes sont souvent décrits visuellement, en représentant les sommets par des points et les arêtes par des segments entre eux (voir Figure 7.8 par exemple). En fait, nous avons déjà représenté certains graphes, puisque la Figure 7.6 peut être vue comme un graphe dont les sommets sont les segments composant le bord de chaque triangle tandis que les arêtes sont les liens gris représentant leurs identifications.

Le papier fondateur de la théorie des graphes est un papier publié en 1736, écrit par Leonhard Euler, sur les sept ponts de Königsberg [39]. Son objectif était de prouver l'insolubilité d'un vieux défi mathématique lancé par les habitants de Königsberg. Ce défi, qui demande de trouver un parcours de la ville empruntant chacun de ses ponts exactement une fois, n'a pu être prouvé insoluble avant d'être formalisé dans le cadre de la théorie des graphes. Comme nous voulons mettre l'accent sur les liens entre la théorie des graphes et la topologie, il est significatif que cet article soit aussi souvent cité comme étant fondamental pour la naissance de la topologie puisque les questions dont il traite sont de nature topologique.

Les graphes étant des objets plutôt compréhensibles, certains résultats célèbres de théorie des graphes sont également accessibles. Le théorème des 4 couleurs est l'un d'entre eux : il stipule que toute carte géographique peut être colorée avec quatre couleurs de telle sorte que deux régions limitrophes aient des couleurs différentes. Ce théorème, qui s'exprime naturellement dans le cadre de la théorie des graphes, a résisté à des tentatives de preuves pendant plus de 100 ans avant qu'une preuve assistée par ordinateur ne réussisse en 1969 [59]. La théorie des graphes est actuellement un domaine de recherche très actif qui présente des interactions importantes avec d'autres domaines des mathématiques et de l'informatique tels que l'algèbre, les probabilités, la science des données et, plus important encore pour nous, la topologie.

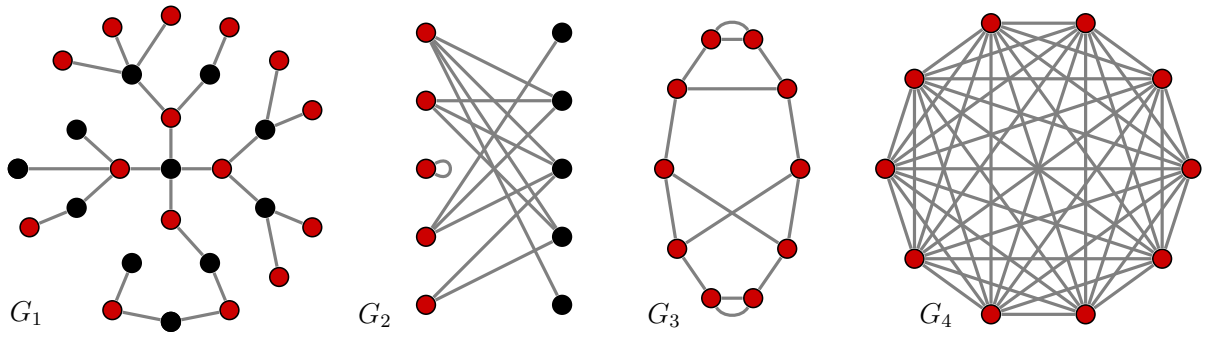

FIGURE 7.8 : Quelques exemples de graphes. Parmi eux, G_1 est un arbre et G_4 est le graphe complet K_{10}.

Plongeons maintenant dans quelques définitions et propriétés des graphes en utilisant les 4 graphes de la Figure 7.8. Un **chemin** sur un graphe est une séquence de sommets distincts

tels que chacun d'entre eux est relié au suivant par une arête. Comme en topologie, on dit qu'un graphe où toute paire de points peut être reliée par un chemin est **connexe**. Tous les graphes de la Figure 7.8 sont connexes, à l'exception de G_2 dont un sommet est isolé. Un **cycle** est un chemin, sauf au niveau de ses extrémités qui doivent être identiques. Vu comme un espace continu, c'est-à-dire si les arêtes sont vues comme des segments qui sont collés sur leurs extrémités quand ils partagent un sommet, un cycle est une courbe non contractile plongée dans le graphe. Certains problèmes peuvent être réduits à la recherche d'un cycle spécifique dans un graphe donné : si l'on cherche une visite exhaustive et passionnante d'un pays (visiter deux fois la même ville est ennuyeux), cela revient à chercher un cycle visitant chaque sommet exactement une fois.

Permettez-nous d'ouvrir une parenthèse sur la complexité algorithmique. Nous pouvons doter les arêtes précédentes de la longueur correspondant à la route qu'elles représentent. Dans ce graphe, le problème de la recherche d'un cycle de longueur minimale passant par chaque sommet est connu comme un problème informatique particulièrement difficile à résoudre, appelé problème du voyageur de commerce. Ici, difficile ne signifie pas trouver un moyen de calculer la solution, puisqu'il existe un algorithme plutôt naïf qui résout le problème : cet algorithme opère en énumérant toutes les séquences de visite des sommets possibles et garde en mémoire le cycle le plus court. Le temps d'exécution d'un algorithme est le nombre d'opérations élémentaires effectuées par l'algorithme, qui est fonction de la taille de l'entrée. Plus ce nombre est élevé, plus l'exécution de l'algorithme prend du temps. Trouver un algorithme ayant le meilleur temps d'exécution possible, et classer les problèmes en fonction de ce meilleur temps d'exécution possible, est l'un des principaux objectifs de l'informatique. Pour notre algorithme naïf, étant donné qu'il y a $n!$ ordres de visite où n est le nombre de sommets, il faudra énormément de temps pour l'exécuter car chaque ordre doit être considéré. Le problème est par conséquent appelé difficile, car nous ne connaissons pas encore d'algorithme s'exécutant en un temps meilleur qu'exponentiel qui résoudrait ce problème dans le cas général.

Pour en revenir aux cycles dans les graphes, le graphe G_4, appelé **graphe complet** sur 10 sommets, contient une arête entre chaque paire de sommets, et donc beaucoup de cycles. On désignera par K_n le graphe complet sur n sommets. À l'opposé, les graphes peuvent aussi ne pas avoir de cycles, comme G_1 par exemple. Un graphe qui est à la fois connexe et sans cycle est appelé un **arbre**. À cet égard, les arbres ont une topologie très simple. Ils sont donc très intéressants et utiles : comme les sphères, ils n'ont pas de courbes non contractiles. Une autre de leurs propriétés est qu'il existe un chemin unique entre chaque paire de leurs sommets. Comme exemple de leurs utilisations, les arbres sont souvent employés en informatique comme structure de données efficace. On peut aussi vouloir organiser des tâches à accomplir pour un processus dans un graphe de dépendance : les sommets sont les tâches à accomplir et les arêtes sont les dépendances entre elles. Ici, les arêtes sont orientées : si une première tâche est liée à une autre par une arête orientée, cela signifie que la première tâche doit être achevée avant que la suivante ne commence. Il est donc crucial qu'aucun cycle n'apparaisse structurellement lors de la création d'un tel graphe de dépendance.

Les graphes peuvent avoir plusieurs arêtes entre deux sommets, comme G_3 ou même une arête allant d'un sommet à lui-même, comme le sommet isolé de G_2, une telle arête est appelée une **boucle**. Les graphes peuvent être **bipartis** : leurs sommets peuvent être séparés entre

deux ensembles, de sorte qu'il n'y ait pas d'arête entre deux sommets du même ensemble. Les graphes G_1 et G_2 sont tous deux bipartis, les couleurs des sommets décrivant la division susmentionnée. Tout arbre est biparti : si l'on fixe d'abord un sommet v dans l'arbre, et que l'on colore les autres en fonction de la parité du nombre d'arêtes sur l'unique chemin qui les relie à v, on obtient la partition souhaitée. Remarquez que les cycles des graphes bipartis doivent être de longueur paire. Comme il y a autant d'arêtes que de sommets dans les cycles, et que ces sommets alternent entre les deux ensembles de la partition, il y a un nombre pair de sommets, et donc un nombre pair d'arêtes aussi. Enfin, nous pouvons définir le graphe biparti complet $K_{n,m}$, c'est le graphe biparti où A a n sommets, B a m sommets, et il y a une arête entre chaque sommet de A et chaque sommet de B (voir Figure 7.9).

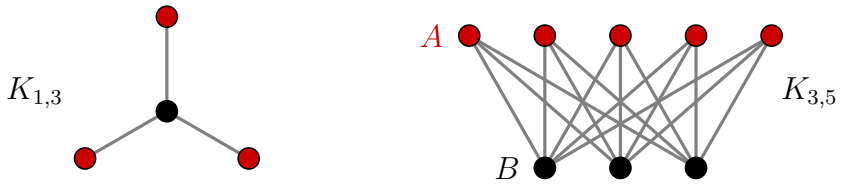

FIGURE 7.9 : Deux exemples de graphes bipartis complets.

D'un point de vue algorithmique, la propriété structurelle d'être biparti présente certains avantages algorithmiques. Par exemple, il est plus facile de calculer un couplage maximal dans un graphe biparti, c'est-à-dire un ensemble maximal d'arêtes du graphe qui ne partagent aucun sommet. De même, de nombreux problèmes algorithmiques sont beaucoup plus faciles à résoudre dans les arbres, puisqu'ils contiennent peu d'arêtes et pas de cycles. Une idée fructueuse, d'un point de vue informatique, a été d'identifier les graphes qui ressemblent à des arbres et d'y généraliser des algorithmes sur les arbres. C'est pourquoi de nombreuses mesures ont été mises au point pour déterminer à quel point un graphe est proche d'un arbre. La *largeur arborescente* (*treewidth*) est l'une de ces mesures. Intuitivement, les éléments d'un graphe de largeur arborescente k peuvent être rangés dans des *sacs* de taille $k+1$ et présentés sous la forme d'un arbre pour former une *décomposition arborescente*. La figure 7.10 montre une décomposition arborescente de largeur 3.

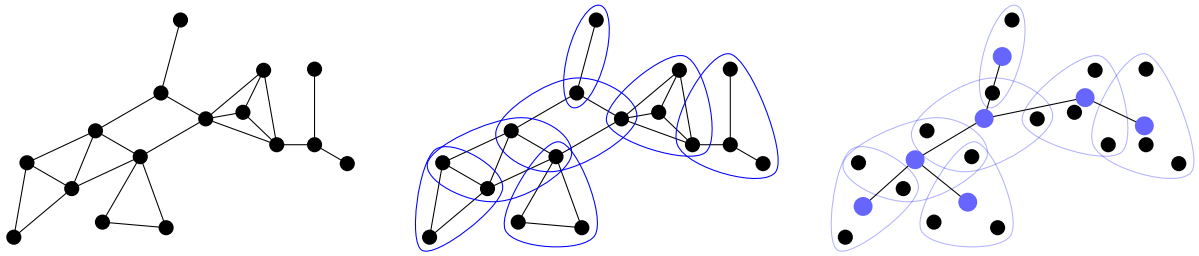

FIGURE 7.10 : Gauche : un graphe G. Milieu : une décomposition des sommets de G en sacs. Droite : une décomposition arborescente de G.

Les décompositions arborescentes doivent vérifier plusieurs propriétés, mais nous ne voulons pas trop nous étendre sur les définitions techniques ici. Pour les besoins de cette introduction, il est suffisant de dire que la largeur arborescente est la largeur minimale d'une

décomposition d'arbre et de se concentrer sur les exemples suivants. Certains graphes ont une largeur arborescente élevée, comme les graphes complets ou les grilles (voir la partie droite de la Figure 7.11) : ils sont loin d'être des arbres. D'autres graphes, comme les arbres ou les graphes comme celui de la partie gauche de la Figure 7.11 ont une faible largeur arborescente.

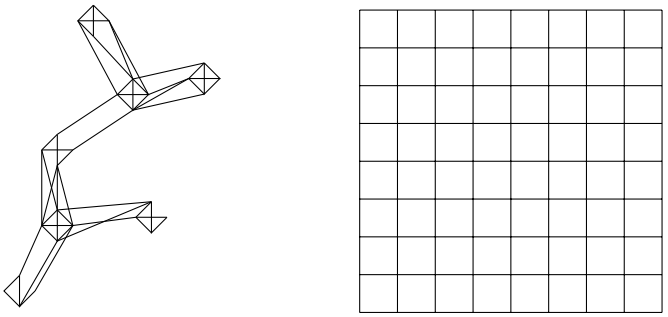

FIGURE 7.11 : Gauche : un graphe de haute largeur arborescente. Droite : un graphe de faible largeur arborescente.

Quand nous sommes confrontés à un problème algorithmique sur un graphe de largeur arborescente k, il est possible de d'abord trouver une décomposition arborescente de largeur k. Puis, il est possible de résoudre le problème sur chaque sac, et ensuite d'essayer de reconstituer une solution pour le graphe entier en exploitant l'arbre de la décomposition. Cette approche s'est révélée être très fructueuse pour de nombreux problèmes algorithmiques.

Graphes planaires. Une question très naturelle se pose à quiconque essaie de dessiner un graphe de manière à ce qu'il soit facile à lire. Les croisements entre les arêtes constituent un obstacle à la lisibilité, de sorte que l'on essaie généralement de dessiner le graphe sans croisements. Par exemple, les deux dessins de la Figure 7.12 représentent le même graphe, mais celui de droite n'a pas de croisements et est donc plus facile à lire.

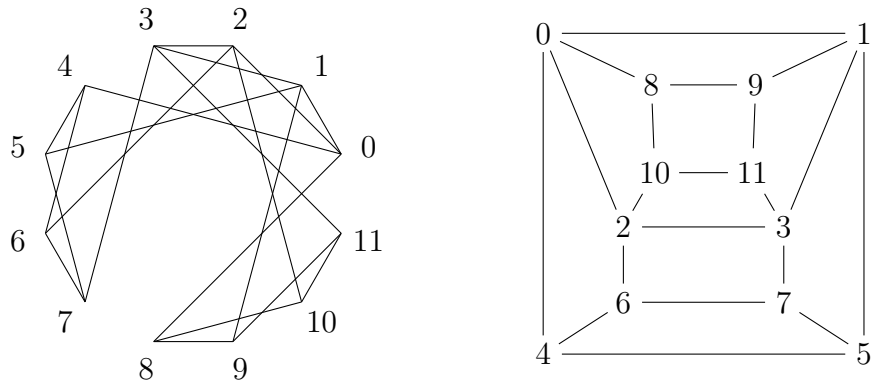

FIGURE 7.12 : Deux représentations du même graphe, dont une est planaire. Deux sommets sont adjacents si leurs écritures en binaire diffèrent d'exactement un bit.

D'où la définition suivante : un graphe est dit **planaire** s'il existe un dessin de ce graphe dans le plan tel que deux arêtes ne se croisent jamais. Ici, une classe de graphes est définie par

une propriété topologique, à savoir être plongeable dans le plan ou, puisque c'est équivalent, dans la 2-sphère.

Cette propriété est particulièrement utile pour certaines applications où les arêtes représentent des connexions qui doivent être construites. Les circuits imprimés, où les sommets sont des composants et les arêtes des circuits gravés sur une plaque, ou plus simplement la conception de routes, où les sommets sont des villes et les arêtes les routes qui les relient, sont des exemples de ce type d'application. Dans les deux cas, on veut éviter les croisements, car ils sont coûteux à gérer ; dans nos exemples, ils sont gérés respectivement par un composant électrique ou un pont.

Voici une célèbre énigme mathématique appelée « énigme des trois maisons », dont la première mention remonte à 1913 par H. Dudeney [35] qui l'appelait déjà un vieux problème. L'objectif est de relier, sans croisement, trois maisons à trois services publics, par exemple le gaz, l'électricité et l'eau.

Nous encourageons le lecteur qui n'a jamais essayé de résoudre ce problème à tenter de le faire en griffonnant sur une feuille de papier, au moins jusqu'à ce qu'il doute de sa faisabilité. En fait, toute tentative finira par ressembler à la partie gauche de la Figure 7.13, où toutes les arêtes sauf une sont dessinées. À ce stade, la dernière arête ne peut pas être ajoutée entre la maison la plus à gauche et l'installation de gaz. En effet, la première maison est enfermée dans un cercle composé de 4 arêtes (celles entre les maisons du milieu et de droite et les installations d'eau et d'électricité). En effet, le célèbre théorème de Jordan stipule que tout chemin entre un point à l'extérieur du cercle (l'usine à gaz) et un point à l'intérieur (la première maison) croisera nécessairement une arête.

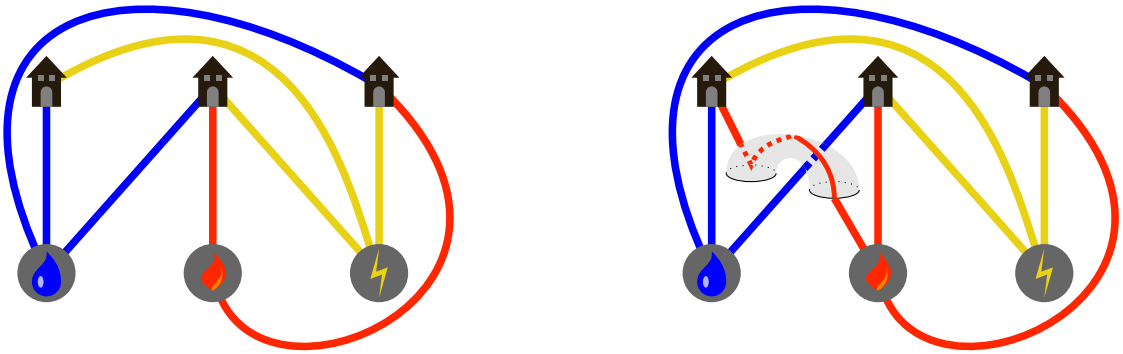

FIGURE 7.13 : À gauche : une tentative de résolution de l'énigme des trois maisons. À droite : la seule façon de résoudre ce problème, i.e., en attachant un tube au plan.

Il existe cependant un moyen de résoudre ce problème, en travaillant non pas sur une feuille de papier, mais sur une tasse. Une fois de plus, nous encourageons le lecteur qui n'a jamais essayé de résoudre ce problème sur une tasse et qui n'a pas peur de la laver ensuite à dessiner dessus avec un marqueur. L'essentiel est que les extrémités de la poignée soient placées sur des faces distinctes. Ainsi, à l'instar de la partie droite de la Figure 7.13, où un tube est attaché sur ses bords à deux trous faits sur le plan, on pourra dessiner la dernière arête sans qu'elle ne croise les autres. Ici, en modifiant la surface sur laquelle on cherche à tracer le graphe, le problème devient résoluble.

En topologie, la formulation de ce problème est : « est-ce que $K_{3,3}$ est plongeable dans une 2-sphère ? ». Il se trouve que la topologie fournit des outils pour répondre à ce problème. Les arguments présentés ci-dessus ne sont pas vraiment une preuve, puisqu'ils reposent sur le fait que « toute tentative aboutira à ceci ». Essayons une approche plus formelle, et prouvons la Proposition 7.1.

Proposition 7.1. *Le graphe $K_{3,3}$ est non planaire : il ne peut pas être plongé sur une 2-sphère.*

Démonstration. Supposons qu'il existe un plongement de $K_{3,3}$ dans le plan. Nous appelons faces les composantes connexes du plan lorsque le plongement en est retiré. Notons respectivement par V, E et F le nombre de sommets, d'arêtes et de faces du plongement. La formule d'Euler [40] (une liste de preuves est disponible dans [38]), stipule que, pour tout plongement d'un graphe connexe dans le plan, l'égalité $V - E + F = 2$ est vérifiée. Dans notre cas, puisque $V = 6$ et $E = 9$, il découle de cette formule que $F = 5$. Si, pour chaque face f, nous comptons le nombre δ_f d'arêtes qui la délimitent, alors chaque arête est comptée deux fois. Comme tout cycle du graphe a au moins 4 arêtes, puisque le graphe est biparti, on en déduit alors que $18 = 2E = \sum_f \delta_f \geq 4F = 20$, ce qui est absurde. Par conséquent, $K_{3,3}$ n'est pas planaire, et l'énigme des trois maisons n'a pas de solution dans le plan. □

Nous avons ainsi établi qu'il existe des graphes planaires et des graphes non planaires. Une question naturelle qui suit est de savoir s'il existe une caractérisation de la planarité pour les graphes. Kuratowski a prouvé qu'une telle caractérisation existe en 1930 [77], et Wagner a prouvé en 1937 une caractérisation similaire [139] : un graphe est planaire si et seulement si on ne peut pas « trouver » K_5 ou $K_{3,3}$ en lui (voir Figure 7.14). Le terme « trouver » ici est volontairement flou, en effet il y a une subtilité dans sa définition dans les deux théorèmes. Les graphes dont l'absence caractérise la propriété de planarité des graphes dans le théorème de Wagner, sont appelés **mineurs interdits**. Ces théorèmes sont très intéressants pour nous, car une propriété structurelle, c'est-à-dire une information sur les sous-structures présentes dans les graphes, caractérise l'une de leurs propriétés topologiques.

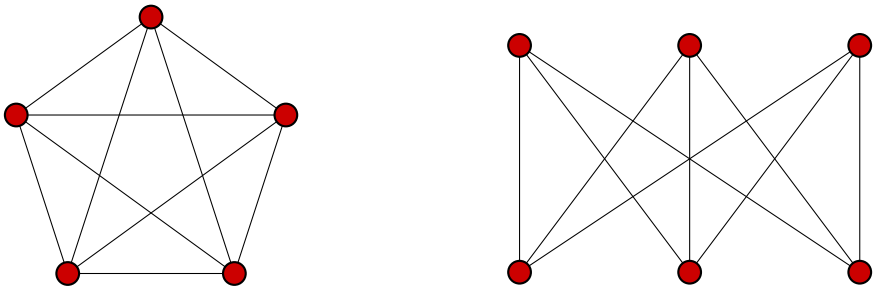

FIGURE 7.14 : Les deux mineurs interdits caractérisant les graphes planaires : K_5 et $K_{3,3}$.

Lorsqu'un graphe planaire est plongé dans $\mathbb{R}^2$, on peut définir le **graphe dual** de ce plongement de la manière suivante : ses sommets sont les faces du plongement, et pour chaque arête partagée par deux faces, il existe une arête duale entre les faces la traversant (voir Figure 7.15). Une telle définition repose en grande partie sur le plongement en question.

Par exemple, le graphe représenté dans la Figure 7.15 a un sommet incident à une seule arête, ce sommet implique l'existence d'une boucle dans le graphe dual dont l'extrémité est fonction de la face sur laquelle ce sommet est plongé (et pour vérifier que le graphe obtenu est différent, il est facile de vérifier que le nombre maximal d'arêtes incidentes à une face-sommet est bien fonction du plongement). Le graphe dual d'un graphe, comme on peut le voir, est également planaire.

En outre, la propriété d'être planaire a de fortes implications algorithmiques. L'idée est que de tels graphes peuvent être découpés en morceaux plus petits, de taille à peu près équivalente, de telle sorte que les solutions à des problèmes peuvent être déduites de celles sur les morceaux plus petits. Ils ont en fait une largeur arborescente d'environ $\mathcal{O}(\sqrt{n})$, où n est le nombre de sommets, de sorte que la méthode récursive présentée plus haut fournit souvent des algorithmes efficaces. Il s'ensuit que certains problèmes, qui peuvent être difficiles à résoudre dans le cas général, admettent une solution efficace lorsque l'entrée est planaire. L'existence de décompositions arborescentes de faible largeur est déjà un résultat fort en soi.

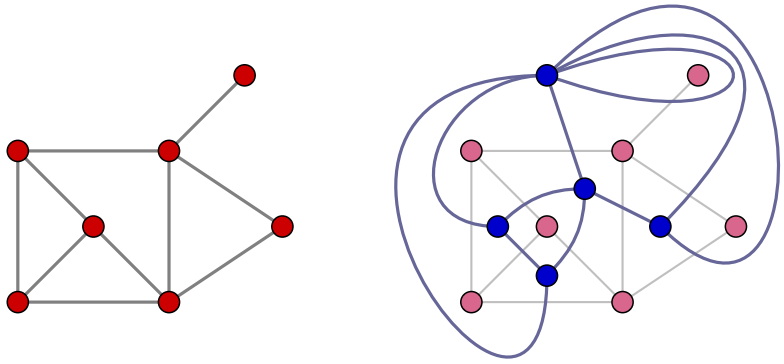

FIGURE 7.15 : Un plongement planaire d'un graphe dans $\mathbb{R}^2$ et le plongement dual associé en bleu.

Remarquez que les cycles d'un graphe dual sont des courbes de $\mathbb{R}^2$ qui intersectent le graphe initial au niveau de ses arêtes. Chercher un ensemble d'arêtes dont le retrait rend le graphe non connexe revient à trouver un cycle du dual. Ce comportement est un premier exemple des allers-retours qui peuvent être faits entre un graphe et son dual. En fait, l'existence d'un dual pour les graphes planaires, ainsi que le fait que ce dual est plongé dans $\mathbb{R}^2$ sont ce qui permet le calcul rapide des décompositions arborescentes équilibrées susmentionnées. Par conséquent, la planarité et la dualité ont un impact fort sur les caractéristiques algorithmiques de ces graphes.

Graphes linkless. En poursuivant sur les propriétés topologiques qui définissent des classes de graphes, nous présentons ici la classe des graphes linkless. Deux objets de l'espace sont dits **séparés** s'il existe une sphère telle que chacun de ces objets se trouve d'un côté différent de la sphère, ils sont dits **entrelacés** sinon. Par exemple, deux cercles dans l'espace peuvent être séparés ou non (voir les côtés gauche et droit de la figure 7.16 respectivement).

Similairement aux graphes planaires, qui sont des graphes qui peuvent être plongés sur une sphère, les **graphes linkless** sont des graphes qui peuvent être plongés dans l'espace, c'est-à-

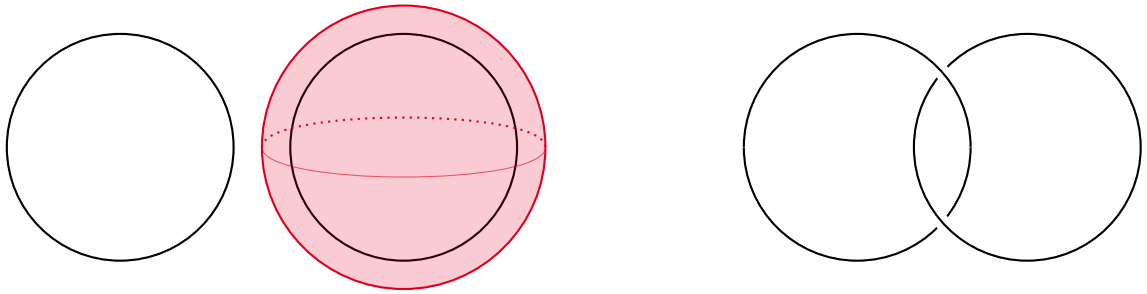

FIGURE 7.16 : Gauche : deux cercles de l'espace séparés, ils sont séparés par une sphère. Droite : deux cercles entrelacés.

dire $\mathbb{R}^3$ ou $\mathbb{S}^3$, de telle façon qu'il n'y ait aucune paire de cycles disjoints qui soient entrelacés. Par exemple, tous les graphes planaires sont des graphes linkless, car un plongement sur une sphère ne peut avoir deux cycles disjoints entrelacés (une petite perturbation de la sphère peut séparer les cycles). Il s'agit d'une autre classe de graphes définie par une propriété topologique et, comme pour les graphes planaires, nous verrons que ces graphes sont caractérisés par des propriétés structurelles.

Dans ce qui suit, nous allons montrer dans les grandes lignes que le graphe complet sur 6 sommets, K_6, ainsi que certains graphes qui lui sont associés, sont tous non linkless. En d'autres termes, tout plongement de K_6 dans l'espace aura au moins une paire de cycles disjoints entrelacés. Dans la Figure 7.17, nous présentons un plongement de K_6, noté $\mathcal{K}$, pour lequel il y a exactement une paire de cycles disjoints qui sont entrelacés. La présentation utilisée ici est ce que nous appellerons grossièrement une *projection* : un dessin à partir d'un point de vue fixe de l'espace du plongement du graphe dans l'espace, où les ambiguïtés concernant quelle partie est au-dessus de l'autre sont levées par des informations supplémentaires. Nous présentons sur cette figure les seules paires de cycles disjoints qui pourraient être liées, puisqu'une telle paire doit avoir au moins un croisement dans la projection.

Pour notre preuve, nous nous inspirons de [25, 125], et définissons pour chaque paire de cycles disjoints $\{C, C'\}$ d'une projection fixée la quantité $\delta(C, C')$, qui est le nombre de croisements entre C et C' où C est au dessus.

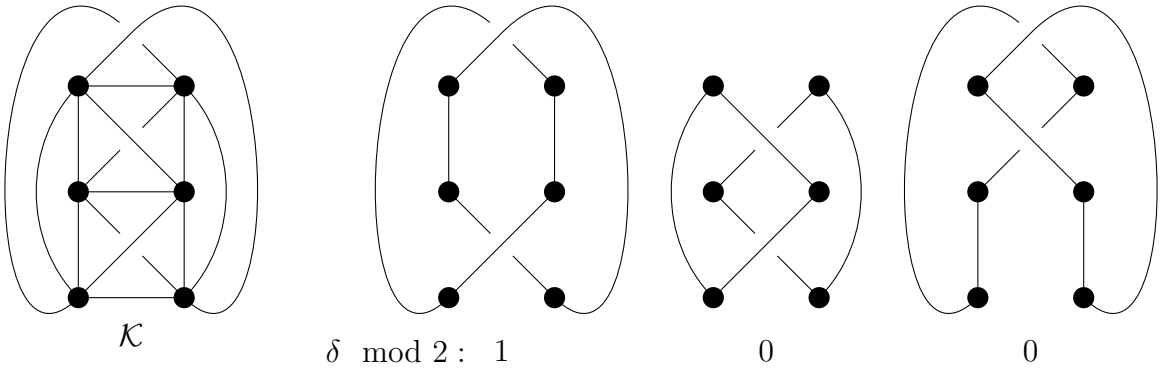

FIGURE 7.17 : Un dessin d'un plongement $\mathcal{K}$ de K_6 où exactement une paire de cycles disjoints est entrelacée.

Proposition 7.2. *Le graphe K_6 n'est pas linkless.*

Sketch de preuve : Dans la suite, nous utiliserons $\mathcal{K}$ pour désigner à la fois le plongement décrit par la Figure 7.17 et sa projection. Similairement, considérons maintenant un autre plongement de K_6, et nous désignerons à la fois ce plongement et sa projection par $\mathcal{K}'$. Il y a $\frac{1}{2}\binom{6}{3} = 10$ paires de cycles disjoints dans K_6. Pour une projection K, nous allons utiliser une quantité que nous noterons $D(K) \equiv \sum_{C_1,C_2} \delta(C_1, C_2) \mod 2$ où $\{C_1, C_2\}$ sont les paires de cycles disjoints de K. Seules les paires de cycles $\{C_1, C_2\}$ illustrées par la Figure 7.17 peuvent vérifier $\delta(C_1, C_2) \mod 2 \neq 1$. Ainsi, par sommation, nous déduisons que $D(\mathcal{K}) \equiv 1$.

Considérons maintenant une transformation continue, une *homotopie*, qui transforme $\mathcal{K}'$ en $\mathcal{K}$. Cette transformation est autorisée à faire se croiser des paires d'arêtes et des arêtes elles-mêmes. À petites perturbations près, nous supposons également que ces croisements ne se produisent pas simultanément, que la projection est *régulière* tout au long de la transformation (les croisements sur la projection peuvent se produire entre 3 arêtes au plus, et de tels croisements impliquant 3 arêtes sont immédiatement résolus) et que les arêtes ne croisent pas les sommets. Si ce dernier cas se produit, nous poussons l'arête traversante de manière à ce qu'elle croise les arêtes incidentes au sommet à la place.

Étudions comment $D(\mathcal{K}')$ est modifié lorsque deux arêtes non incidentes e, e' se croisent au cours de cette transformation ; appelons cet évènement une traversée d'arête. Ces deux arêtes sont présentes dans deux paires de cycles $\{C_1, C_2\}, \{C_1', C_2'\}$, en fonction de l'association de chacun des deux sommets restants avec les deux arêtes considérées.

La traversée d'arête va modifier $\delta(C_1, C_2)$ par ± 1 par définition. Chaque cycle contenant e ou e' n'est pas modifié par la transformation. Par conséquent, $D(\mathcal{K}')$ n'est pas affecté par cet évènement (-2 ou 2 est ajouté à la somme modulo 2). La quantité reste aussi la même lorsque la traversée d'arête prend place entre deux arêtes incidentes. Comme $D(\mathcal{K}') \equiv 1 \equiv D(\mathcal{K})$ à la fin de la transformation, et que $D(\mathcal{K}')$ n'a pas été modifié par celle-ci, nous en déduisons qu'au moins une paire de cycles disjoints était entrelacée au début de la transformation. Ainsi, tout plongement de K_6 dans l'espace a au moins une paire de cycles disjoints qui sont entrelacés : K_6 n'est pas linkless. □

Dans [125], H. Sachs mentionne une transformation issue de l'ingénierie électrique, qu'il appelle « star-triangle-transformation » qui préserve la propriété d'un graphe d'être linkless. Cette opération, illustrée par la Figure 7.18, et plus communément appelée transformation ΔY en théorie des graphes, consiste à remplacer une 3-clique (un triangle), par une griffe $K_{1,3}$ (un « Y ») et vice-versa.

FIGURE 7.18 : La transformation ΔY.

La transformation ΔY préserve le nombre d'arêtes du graphe initial. Ainsi, il y a un nombre fini de graphes pouvant être obtenus en appliquant des transformations ΔY à un

graphe. Dans le cas de K_6, les graphes obtenus de cette manière sont appelés la famille de Petersen (voir Figure 7.19), dont le nom vient du célèbre graphe de Petersen qui en fait partie (graphe en bas de la Figure 7.19). Tous ces graphes sont non linkless, il est possible de vérifier que les propriétés pertinentes de $\mathcal{K}$ pour la preuve de la Proposition 7.2 sont aussi satisfaites par tous les dessins de la Figure 7.19 de telle sorte que la preuve peut être adaptée. En effet, il y a une bijection entre les paires de cycles disjoints avant et après une transformation ΔY. Le principal argument étant qu'une telle paire est soit disjointe du sous-graphe affecté par la transformation, soit qu'un seul cycle de la paire utilise une arête du sous-graphe susmentionné. De plus, si le sous-graphe de départ est $K_{1,3}$, exactement 2 arêtes sont utilisées.

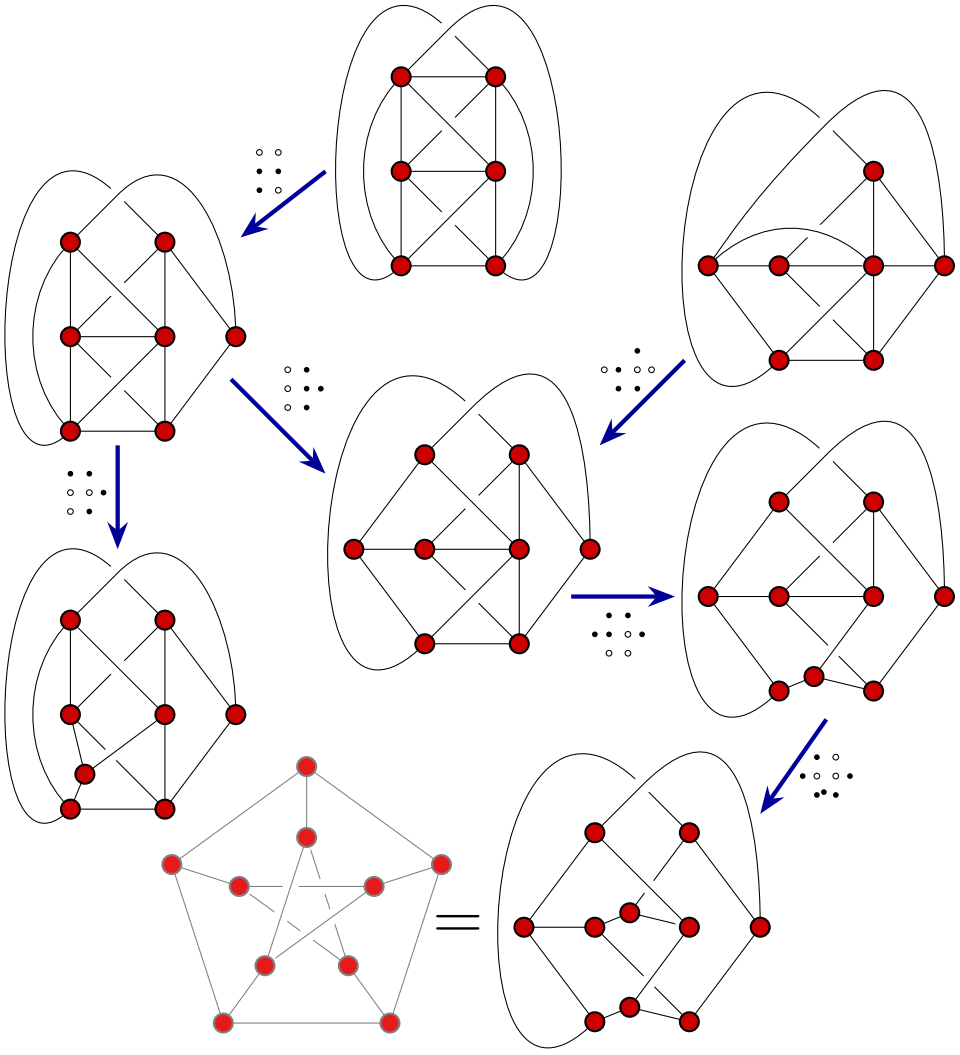

FIGURE 7.19 : La famille de Petersen : les graphes qui peuvent être obtenus de K_6 par transformation ΔY. Les 3 sommets blancs à côté de chaque flèche bleue indiquent quel 3-cycle est transformé en un graphe $K_{1,3}$. Le graphe de Petersen en bas est aussi présenté sous sa forme commune en couleurs plus claires.

Remarque 7.3. *Comme dit plus haut, la Figure 7.19 présente des plongements des graphes de la famille de Petersen faits à la fois pour adapter la preuve de la proposition 7.2 et comprendre quel triangle est transformé en griffe. Cependant, il n'est pas évident en regardant cette image que choisir un autre triangle pour les transformations ΔY donne le même graphe. Ceci est vrai, sauf pour le graphe en haut à gauche où le 3-cycle induit par les 3 sommets de degrés 5 produit un graphe non homéomorphe aux graphes induits par les autres triangles. Il se trouve que les graphes de cette famille présentent de très nombreuses symétries intrinsèques, ce qui explique le rôle symétrique joué par presque tous les triangles. Une présentation de cette famille soulignant ces symétries peut être trouvée ici [2].*

Robertson, Seymour et Thomas ont annoncé [113] puis prouvé [114], que cette famille est exactement l'ensemble des mineurs interdits qui caractérise les graphes linkless. C'est un autre exemple de classe de graphes où l'ensemble des mineurs interdits la caractérisant est connu.

D'un point de vue algorithmique, les classes de graphes caractérisées par des mineurs interdits sont intéressantes. Nous avons vu que les graphes planaires admettent des décompositions très utiles permettant à des problèmes d'être résolus efficacement sur eux. En fait, cette propriété est satisfaite par toute classe de graphes caractérisée par des mineurs interdits [110]. Néanmoins, en comparaison aux graphes planaires, peu de problèmes sont connus avoir une amélioration conséquente de leurs temps d'exécution sur les classes de graphes caractérisées par mineurs interdits. Sur les graphes linkless par exemple, les propriétés topologiques associées aux plongements sont plus dures à exploiter (une exception notable étant [132]).

Théorie des nœuds La dernière classe de graphes que nous avons abordé est définie par des propriétés topologiques de leurs plongements dans $\mathbb{S}^3$ ou $\mathbb{R}^3$. Nous plongeons maintenant plus profondément sur cette idée en étudiant les **nœuds**, qui sont des plongements de $\mathbb{S}^1$ dans $\mathbb{R}^3$ (ou $\mathbb{S}^3$ de manière équivalente) considérés à une déformation continue près, appelée *isotopie ambiante* (voir Figure 7.20 pour des exemples de tels plongements). Une isotopie ambiante est, grossièrement parlant, une transformation continue non dégénérée qui ne s'auto-intersecte pas.

Intuitivement, si nous prenons une corde et la nouons, puis considérons toutes les transformations continues qui peuvent être appliquées dessus, il est toujours possible de la dénouer en inversant les mouvements qui ont été faits pour nouer. En conséquence, nous collons les extrémités de la corde de telle sorte que différents nœuds apparaissent : il n'est plus possible de systématiquement dénouer de tels nœuds. L'idée est qu'un nœud reste le même tant que nous manipulons la corde, nous pouvons l'agrandir, la courber, la tordre tant que nous ne la coupons pas pour la recoller plus tard (cette dernière partie est une opération non continue qui modifie le nœud).

Les nœuds sont tous homéomorphes au cercle $\mathbb{S}^1$, et c'est un bon moment pour souligner que toutes les propriétés topologiques intrinsèques de deux objets homéomorphes sont les mêmes. Cependant, les nœuds sont des plongements sur lesquels les déformations continues que nous considérons sont distinctes des homéomorphismes : ces déformations préservent les propriétés topologiques des plongements, qui ne sont pas des propriétés intrinsèques du cercle.

FIGURE 7.20 : Deux nœuds, i.e., deux plongements de $\mathbb{S}^1$ dans $\mathbb{R}^3$, le nœud à gauche est appelé nœud de huit.

À l'instar des graphes, qui peuvent modéliser un nombre incroyablement élevé d'objets physiques ou abstraits, les nœuds sont le modèle naturel à adopter pour comprendre et étudier tout objet réel qui « ressemble » à une 1-variété. La théorie des nœuds trouve donc des applications dans de nombreux domaines scientifiques. Par exemple, l'ADN peut être étudié sous cet angle puisqu'il s'agit d'une chaîne en double hélice de nucléotides qui se nouent sous l'action d'enzymes [111, 63], dont le pliage peut aussi être étudié par la théorie des nœuds [64]. En physique, les lignes de champ peuvent être vues comme des nœuds, et la théorie quantique des champs s'avère bénéficier d'une compréhension des nœuds [68]. La liste pourrait continuer avec l'impression 3D, la chimie ou la conception de matériaux. Pour en revenir à notre objectif plus théorique, la compréhension des nœuds est un premier pas vers la compréhension des plongements d'objets en dimension supérieure, ce qui est à la base de nombreux autres problèmes.

Cependant, l'étude des nœuds n'est pas facile. Leur définition seule est à l'origine de la question fondamentale de la théorie des nœuds : deux nœuds sont-ils les mêmes ? En d'autres termes, étant donné deux plongements de $\mathbb{S}^1$, existe-t-il une isotopie ambiante qui transforme l'un en l'autre ? Cette question semblerait immédiatement très difficile à quelqu'un qui doit déterminer si deux pelotes de laine données représentent le même nœud (en supposant avoir collé les extrémités ensemble). En effet, il faudrait manipuler les fils suffisamment longtemps pour que les deux nœuds concrets correspondent visuellement. Et ceci ne fonctionne que si les pelotes représentent effectivement le même nœud. En effet, si ce n'est pas le cas, comment certifier que manipuler plus longtemps les nœuds ne finira pas par les faire correspondre ? Trouver une isotopie ambiante entre deux nœuds est mathématiquement difficile. En réalité, c'est aussi difficile pour les ordinateurs. En outre, cette dernière affirmation soulève une question : comment donner un nœud en entrée à une machine de Turing ou un ordinateur ?

Tout comme les surfaces qui peuvent être encodées par un nombre fini de triangles, les nœuds peuvent admettre une représentation finie : une ligne brisée et fermée dans l'espace. Par exemple, le nœud de huit de la Figure 7.20 peut être transformé en la ligne brisée de la Figure 7.21. Un tel plongement est appelé **plongement polygonal**. En théorie des nœuds, il est courant de se restreindre aux nœuds qui admettent un plongement polygonal, ce que nous ferons dans la suite.

Les plongements polygonaux sont une représentation pratique des nœuds, et de plus,

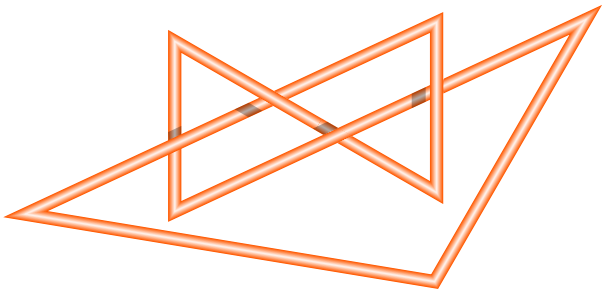

FIGURE 7.21 : Un plongement polygonal du nœud de huit.

les nœuds admettant de tels plongements admettent également des **diagrammes**. Intuitivement, les diagrammes sont des dessins non ambigus de nœuds dans le plan. Le dessin est obtenu grâce à une projection *régulière* (une telle projection peut être vue dans la Figure 7.22), qui est une projection où les points à antécédents multiples sont en nombre fini et viennent de croisements de brins dans l'espace. L'absence d'ambiguïté vient d'une information supplémentaire ajoutée aux croisements désignant quel brin se trouve au-dessus de l'autre, ceci est généralement représenté par un espace vide à la place du brin inférieur autour de chaque croisement (comme illustré en bas de la Figure 7.22, et qui a déjà été utilisé dans la Figure 7.17 pour représenter les plongements dans l'espace). Un diagramme peut être vu comme un graphe planaire plongé tel que chaque sommet a 4 arêtes incidentes (les sommets proviennent de brins croisés dans la projection), un tel graphe est appelé 4-**régulier** (similairement, le graphe G_3 de la Figure 7.8 est 3-régulier).

Reidemeister a prouvé en 1927 [138] le théorème éponyme : deux nœuds, représentés par deux diagrammes, sont équivalents par isotopie ambiante si et seulement si une séquence de **mouvements de Reidemeister** peut transformer l'un des diagrammes (vu comme un graphe planaire plongé) en l'autre. Les mouvements de Reidemeister sont des modifications locales des brins représentées dans la Figure 7.23. Intuitivement, RI crée ou supprime une torsade sur un brin, RII crée ou supprime un chevauchement entre deux brins, et RIII représente le fait qu'un brin peut aller et venir au-dessus d'un autre croisement.

Ce théorème est très puissant, car il relie la nature tridimensionnelle des nœuds à leurs diagrammes, qui sont des graphes planaires. Il ouvre également des pistes pour résoudre la question initiale consistant à déterminer si deux nœuds sont identiques : on peut chercher une séquence de mouvements de Reidemeister à appliquer qui certifie l'équivalence entre les deux nœuds. Avec cette méthode, il ne reste plus qu'à trouver une telle séquence. Il s'agit en fait d'une question difficile. Il n'existe pas de méthode simple pour y parvenir ; intuitivement, il faut essayer tous les mouvements possibles et espérer que les diagrammes finissent par correspondre. Cependant, cette recherche ne fait sens que si l'on limite les mouvements possibles, sinon il est toujours possible d'avoir de plus en plus de croisements par la répétition des mouvements RI ou RII, ce qui amène à un diagramme très complexe. L'approche naturelle consiste alors à essayer uniquement les mouvements qui vont diminuer la complexité du diagramme, correspondant intuitivement à son nombre de sommets. Cependant, il est prouvé que la complexité de certains diagrammes devra d'abord augmenter avant qu'elle puisse être réduite [21]. Ce comportement illustre à quel point ce problème peut être difficile. Actuelle-

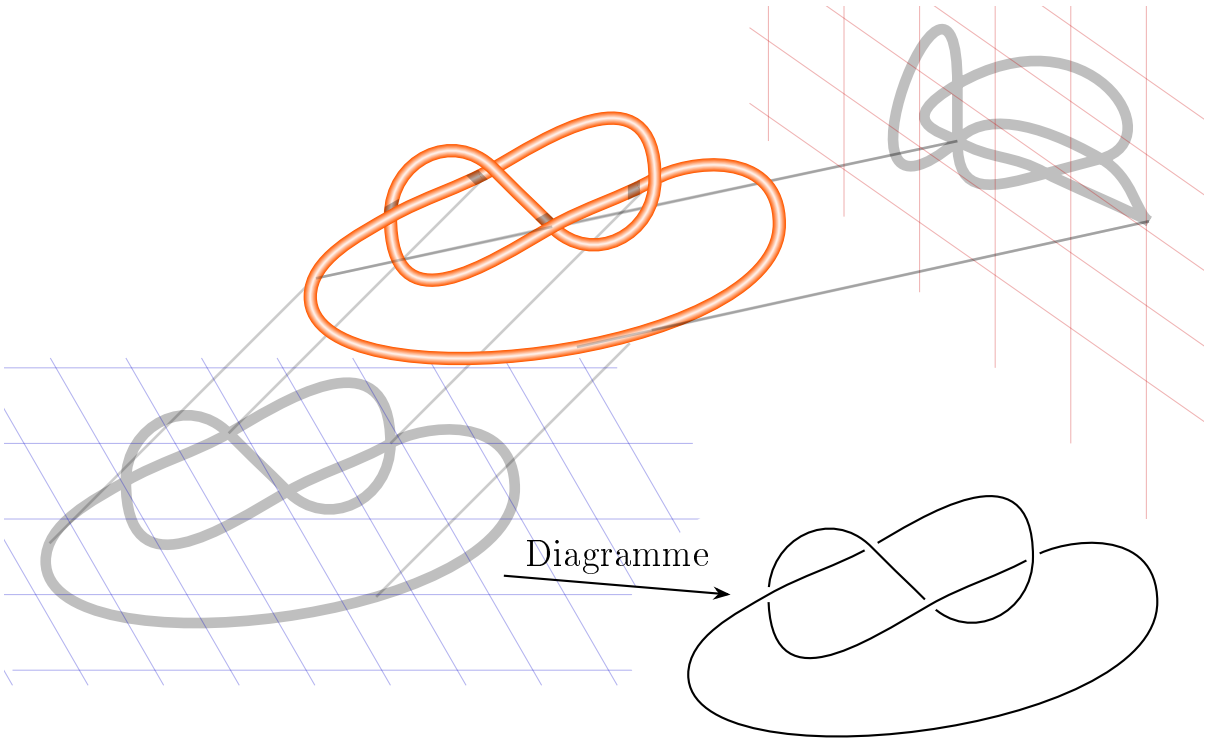

FIGURE 7.22 : Une projection régulière du nœud de huit dans le plan bleu, et une non régulière dans le rouge : deux points de la projection ont plus de deux antécédents. En bas à droite : le diagramme issu de la projection régulière.

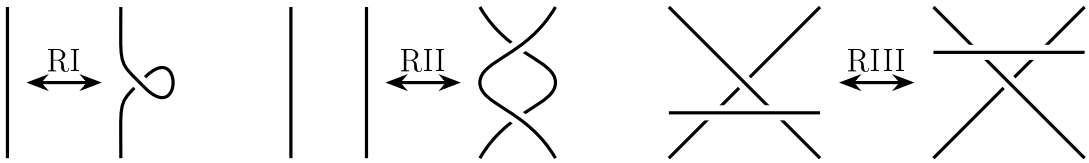

FIGURE 7.23 : Les trois mouvements de Reidemeister.

ment, les meilleurs algorithmes connus pour décider si deux nœuds sont identiques n'utilisent pas les mouvements de Reidemeister et exploitent plutôt la topologie de « l'extérieur » du nœud. Cependant, le temps d'exécution de ces algorithmes est incroyablement long. Pour faire écho à nos parties précédentes, soulignons que ces algorithmes utilisent des surfaces pour mener à bien cette étude.

C'est pourquoi les théoriciens des nœuds ont tendance à chercher des moyens de distinguer les nœuds : des moyens de certifier que deux nœuds sont différents. Pour ce faire, ils étudient des **invariants** des nœuds. Ces invariants sont des quantités, ou des objets, qui sont les mêmes pour toutes les présentations possibles d'un nœud. Ainsi, si deux plongements ne coïncident pas sur un invariant, nous savons qu'ils sont différents. La réciproque n'est pas vraie en général, deux nœuds partageant un invariant peuvent être différents. L'invariance de certaines quantités peut être clairement déduite de leurs définitions. Le **nombre de croisements** est l'une d'entre elles : c'est le nombre minimum de croisements qu'un diagramme de nœuds peut

avoir. Comme ce nombre ne dépend pas d'un plongement fixe mais de l'ensemble de la classe d'équivalence , il est clair que cette quantité est un invariant. Certains autres invariants sont définis par un calcul ou une propriété d'un plongement ou d'un diagramme fixe. Ainsi, pour prouver leur invariance, il faut prouver que le calcul aboutit au même résultat pour chaque plongement ou diagramme.

Développons un exemple simple : la tricolorabilité. Un brin d'un diagramme est un arc continu d'un diagramme lorsqu'il est dessiné en utilisant la convention selon laquelle la partie inférieure du nœud est remplacée par un blanc à chaque croisement comme dans les Figures 7.23 et 7.22 (dans celle-ci, le nœud de huit a 4 brins). Un nœud est dit **3-colorable** si l'un de ses diagrammes est **3-colorable**, c'est-à-dire s'il est possible de colorer, à l'aide de trois couleurs, les brins d'un diagramme de telle sorte que chaque croisement soit incident soit aux 3 couleurs, soit à seulement 1 (les trois couleurs doivent également être utilisées au moins une fois au total). Le théorème de Reidemeister montre que toute paire de diagrammes d'un même nœud sont reliées par une séquence de mouvements de Reidemeister. Il suffit donc de montrer que si un diagramme est 3-colorable, alors tous les diagrammes de la séquence, y compris le dernier, le sont aussi. Par contradiction, il est clair que si un diagramme de nœud n'est pas 3-colorable, aucun de ses diagrammes ne l'est.

Proposition 7.4. *La tricolorabilité est un invariant de nœud.*

Démonstration. Poursuivons la preuve ébauchée ci-dessus et considérons une 3-coloration initiale d'un diagramme de nœud, et un mouvement de Reidemeister. Nous ne modifierons que la couleur des brins créés par le mouvement de Reidemeister de manière à ce que les propriétés d'une 3-coloration soient toujours satisfaites aux croisements non impliqués. En particulier, les couleurs des brins s'étendant à l'extérieur de chaque motif de mouvement de Reidemeister resteront les mêmes. En fait, la preuve peut être résumée par la Figure 7.23 lorsqu'une seule couleur est impliquée et la Figure 7.24 dans le cas contraire.

•Si le mouvement est RI, un brin est séparé en deux, ou deux sont fusionnés. Dans les deux cas, ils doivent être de la même couleur (voir la partie la plus à gauche de la Figure 7.23 où la couleur est le noir, par symétrie).

•Si le mouvement est RII, supposons que plus d'une couleur est impliquée (sinon il n'y a rien à dire). L'image est alors, à changement de couleur près, la partie supérieure gauche de la Figure 7.24. Si des croisements sont impliqués, ils satisfont la propriété d'un 3-coloriage.

•Si le mouvement est RIII, supposons à nouveau que plus d'une couleur est impliquée (sinon il n'y a rien à dire). Nous nous référerons à la partie pertinente de la Figure 7.24. Fixons d'abord la couleur du brin supérieur (disons vert dans notre cas), elle sera inchangée pendant le mouvement. La partie intermédiaire du nœud est constituée de deux brins qui ont soit une, soit deux couleurs (respectivement en haut à droite et en bas). Le premier cas étant simple, concentrons-nous sur le second. Soit le croisement n'impliquant pas le brin supérieur n'a qu'une couleur, auquel cas il s'agirait du cas en bas à gauche ; soit il en a 3 et le cas est traité en bas à droite. Dans tous ces cas, la propriété de 3-coloration est préservée.

□

Comme nous l'avons vu, l'invariance de la tricolorabilité admet une preuve assez élémentaire et accessible. Utilisons la sur la Figure 7.25 où les trois nœuds les plus simples sont représentés : de gauche à droite, le **nœud trivial** (le plongement standard du cercle), le

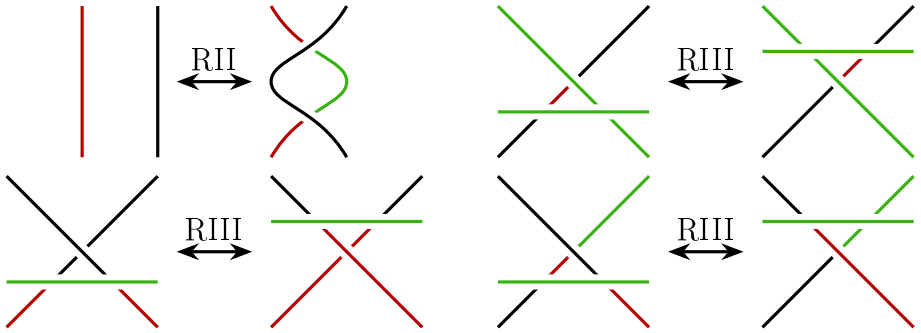

FIGURE 7.24 : Maintien de la propriété de 3-colorabilité pendant le mouvement de Reidemeister.

nœud de trèfle (seul nœud avec nombre de croisements égal à 3), et le nœud de huit (seul nœud avec nombre de croisements égal à 4). Parmi eux, le nœud de trèfle peut être distingué à la fois du nœud trivial et du nœud de huit en utilisant cet invariant, puisque le diagramme du nœud trivial n'a qu'une seule couleur : il n'est pas 3-colorable. De même, chaque paire de brins du nœud de huit du diagramme fourni a deux croisements en commun. Puisqu'il y a 4 brins, au moins deux brins auront la même couleur, et une contradiction est obtenue à l'un de leurs croisements communs puisque le brin restant doit être d'une couleur différente. Comme dit plus haut, la Proposition 7.4 et la Figure 7.25 prouvent que le nœud trèfle est en effet différent du nœud trivial et du nœud de huit, mais cet invariant n'est pas suffisant pour distinguer ces deux derniers nœuds.

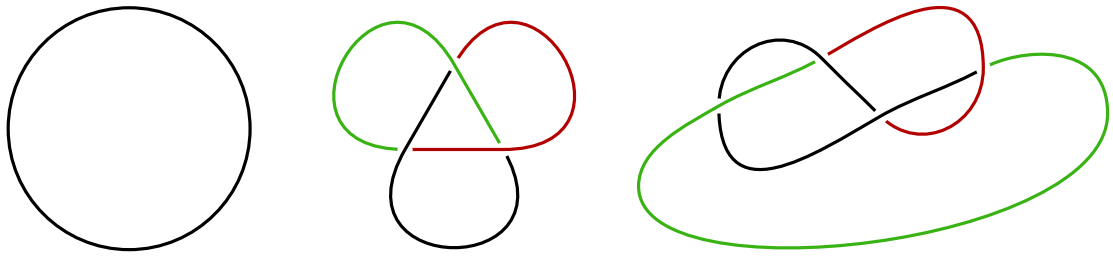

FIGURE 7.25 : Trois diagrammes de nœuds. Parmi eux, seul le nœud de trèfle, au milieu, est 3-colorable.

Pour revenir à notre point de vue informatique, les diagrammes de nœuds sont des graphes planaires. Il est donc naturel d'essayer d'appliquer les résultats et les méthodes de la théorie des graphes pour comprendre les nœuds au travers de leurs diagrammes. Par exemple, certains invariants de nœuds qui sont difficiles à calculer dans le cas général peuvent être calculés efficacement lorsque l'un de leurs diagrammes ressemble à un arbre.

Nous avons présenté plus haut une méthode pour traiter des problèmes algorithmiques sur les graphes à l'aide de décompositions arborescentes. Esquissons un algorithme pour décider de la tricolorabilité d'un nœud de petite largeur arborescente. Naïvement, si un diagramme a n croisements, il possède n brins qui peuvent donc être colorés de 3^n façons. La vérification de la validité de chaque coloration conduit à un algorithme exponentiel. Supposons alors que le diagramme ait une largeur arborescente d'au plus k pour faire mieux. Une propriété

des graphes planaires est qu'il est possible de trouver une décomposition en arbre où les sacs sont délimités par des cercles dans le plan (une telle décomposition est montrée dans la Figure 7.10). On peut alors énumérer au plus 3^k colorations pour les brins de chaque sac et vérifier la compatibilité avec la coloration des sacs voisins de l'arbre le long des cercles de la décomposition. Cet algorithme grossièrement présenté est polynomial lorsque k est fixé, contrairement à l'algorithme naïf exponentiel.[2]

Thème de cette thèse. Précédemment, nous avons montré comment les propriétés topologiques des graphes définissent des classes intéressantes de graphes. Nous avons également décrit comment les propriétés topologiques de planarité et d'être linkless pour les graphes sont caractérisées par des sous-structures dans les graphes. Ces propriétés et leurs conséquences font partie d'un sous-domaine de la théorie des graphes appelé **théorie structurelle des graphes**. Nous avons ensuite souligné comment les graphes peuvent éclairer notre compréhension de la théorie des nœuds. Cette thèse porte sur cette dernière partie : s'inspirer des résultats et méthodes de la théorie structurelle des graphes pour pousser plus loin notre compréhension de la théorie des nœuds.

Présentons brièvement quelques-uns de nos principaux résultats. Comme nous l'avons illustré ci-dessus avec la tricolorabilité, lorsqu'un nœud admet un diagramme de faible largeur arborescente, cette structure arborescente peut être exploitée pour fournir des algorithmes efficaces pour calculer des invariants. Cependant, tous les nœuds admettent des diagrammes de grande largeur arborescente puisqu'il est toujours possible de manipuler une partie du nœud pour la faire apparaître comme une grille, comme dans la Figure 7.26.

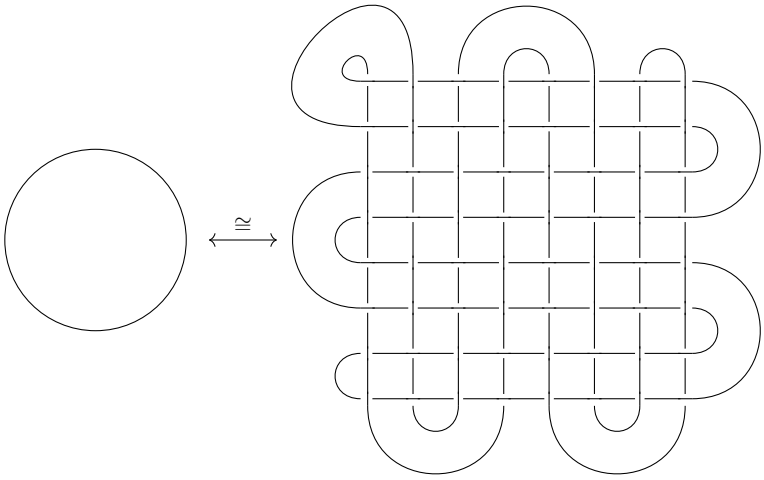

FIGURE 7.26 : Un diagramme du nœud trivial avec haute largeur arborescente.

Par conséquent, une question naturelle, posée par Burton [22], Makowski et Mariño [86], est de savoir s'il existe des nœuds qui n'admettent pas de diagrammes de faible largeur arborescente. En d'autres termes, ils ont demandé s'il existe des nœuds pour lesquels les

[2]En fait, décider de la tricolorabilité peut toujours être fait en temps polynomial en calculant la valeur du polynôme d'Alexander à -1 [100]. Cependant, la tricolorabilité peut être généralisée aux invariants de coloration dont le calcul est prouvé impossible en temps polynomial (modulo ses conjectures standards) [76].

algorithmes exploitant la largeur arborescente ne peuvent pas être utilisés. Cette question a reçu une réponse positive en 2018 [28] en utilisant des résultats complexes [57] sur la forme d'une décomposition de l'espace par des surfaces. Nous développons une théorie plus élémentaire reprouvant ce résultat qui s'inspire de techniques associées à la largeur arborescente, donc issues de la théorie structurelle des graphes. Plus précisément, nous définissons une mesure d'à quel point un nœud ressemble à un arbre (c'est aussi un invariant de nœud), ainsi qu'une *obstruction* à cette mesure. L'invariant, appelé *spherewidth*, est ce qui relie le nœud à la largeur arborescente de son diagramme : si la spherewidth est élevée, la largeur arborescente de chaque diagramme est élevée. La spherewidth quantifie de manière informelle la meilleure façon de balayer l'espace avec des sphères imbriquées de manière arborescente, tout en minimisant leur nombre d'intersections avec le nœud. L'obstruction, quant à elle, fournit une borne inférieure à la spherewidth.

De plus, nous prouvons que notre obstruction existe dès qu'un nœud peut être plongé sur une surface. Pour ce faire, nous exploitons les interactions entre les sphères de nos balayages et cette surface, et en particulier la topologie découlant de ces intersections. Cela nous permet de prouver que tous les diagrammes d'une certaine famille de nœuds ont haute largeur arborescente. Une telle famille est appelée la famille des nœuds toriques (un nœud torique peut être vu sur la Figure 7.27). En outre, nos résultats s'appliquent également aux **entrelacs** et aux **graphes spatiaux**. Il s'agit dans les deux cas de généralisations des nœuds. Les entrelacs sont une union disjointe de nœuds, chacun de ces nœuds étant appelé **composante d'entrelacs**. Par exemple, la partie droite de la Figure 7.16 représente un entrelacs appelé entrelacs de Hopf constitué de deux composantes d'entrelacs, chacune étant un nœud. Les graphes spatiaux sont des plongements de graphes dans $\mathbb{S}^3$. Ces plongements peuvent être vus comme une généralisation des nœuds et des entrelacs : ils ressemblent localement à des nœuds, sauf aux sommets où ils peuvent se ramifier. Il s'ensuit que leur étude est au moins aussi complexe que celle des nœuds et des entrelacs.

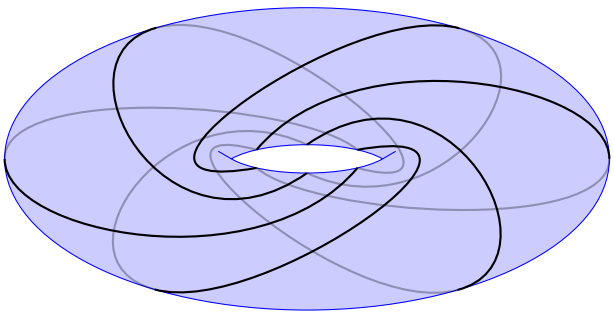

FIGURE 7.27 : Un exemple de nœud torique.

Notre obstruction peut également être utilisée pour résoudre un autre problème qui, en apparence, pourrait sembler sans rapport. Nous avons expliqué précédemment qu'il faut parfois d'abord augmenter le nombre de croisements avant d'espérer progresser dans le processus de démêlage de certains diagrammes de nœuds par des mouvements de Reidemeister. En fait, ce comportement peut apparaître sur n'importe quel problème traité par l'application itérée de mouvements de Reidemeister. Cependant, aucun résultat n'est connu sur le nombre de croisements qui doivent être ajoutés lors de la résolution d'un problème de cette manière. Une

autre de nos contributions est de tirer parti de notre obstruction pour fournir des exemples de diagrammes nécessitant un nombre arbitrairement élevé de croisements ajoutés pour être simplifiés.

7.2 Contributions de cette thèse

Tandis que la section précédente a été écrite pour être abordée avec relativement peu de connaissances en topologie et théorie des graphes, il n'en va pas de même pour ce qui suit. Nous nous référons aux manuels de Schultens sur les 3-variétés [130] (il couvre également les bases des surfaces), de Diestel sur la théorie des graphes [30], de Rolfsen [120] pour la théorie des nœuds, et de Cormen, Leiserson, Rivest, et Stein pour des bases sur les algorithmes [26].

Nous passons maintenant à une description plus précise des contributions de cette thèse, sans pour autant plonger trop profondément dans leurs spécificités. Chaque objet défini ici sera redéfini correctement dans le chapitre correspondant. Notre travail apporte des contributions à la théorie des nœuds en utilisant diverses notions de décompositions arborescentes reposant sur des surfaces et sur la façon dont ces surfaces interagissent avec les nœuds. En outre, l'inspiration et les méthodes utilisées dans nos résultats proviennent de la théorie structurelle des graphes.

Décidabilité du défaut de genre sur les entrelacs de Hopf arborescents. Comme nous l'avons expliqué dans la section précédente, le problème de savoir si deux nœuds sont équivalents est difficile à résoudre, tant du point de vue théorique que du point de vue informatique. C'est pourquoi la théorie des nœuds a recours à des invariants. Parmi eux, un invariant classique est leurs *genre*. C'est le genre minimal possible parmi ses *surfaces de Seifert*, c'est-à-dire les surfaces orientées plongées dans $\mathbb{S}^3$ ayant le nœud comme bord. Par exemple, le nœud trivial est le seul nœud de genre 0, c'est-à-dire que le nœud trivial est le bord d'un disque plongé (cette propriété est une définition courante du nœud trivial). Plusieurs algorithmes pour calculer le genre d'un nœud sont connus, et la complexité de son calcul est assez bien comprise : nous savons que le problème est à la fois dans **NP** et dans **co-NP** [4, 82].

En considérant $\mathbb{S}^3$ comme le bord de $\mathbb{B}^4$, nous considérons une généralisation du genre des nœuds que nous appellerons 4-*genre* : grossièrement, c'est le genre minimal d'une surface orientée intégrée dans $\mathbb{B}^4$ qui a pour bord le nœud plongé dans $\partial\mathbb{B}^4$. Dans le contexte de la topologie en dimension 4, les plongements lisses et topologiquement plats sont différents. Cela donne naturellement lieu à deux notions différentes de 4-genre : nous nous référons au Chapitre 3 pour des définitions précises. Comme nos méthodes et nos résultats s'appliquent aussi bien aux deux, nous continuerons à parler de 4-genre dans la suite de cette introduction.

D'un point de vue algorithmique, la topologie en dimension 4 est difficile et mal comprise. D'une part, de nombreux problèmes topologiques fondamentaux, tels que décider si deux variétés sont homéomorphes, sont connus pour être indécidables en dimension 4 [89]. D'autre part, la décidabilité de nombreux problèmes fondamentaux, tels que la reconnaissance de la sphère quadridimensionnelle, est largement ouverte, et aucun cadre général n'est connu pour aborder ces questions. Le 4-genre fait partie de ces problèmes dont la décidabilité n'est pas connue dans le cas général. Cet invariant est essentiel à l'étude des nœuds *bordant* (*slice*),

c'est-à-dire des nœuds de 4-genre 0, et à des conjectures majeures de la théorie des nœuds comme la conjecture bordant-ruban (slice-ribbon conjecture) [41].

Nous étudions une classe de nœuds et *d'entrelacs* (qui sont des unions de nœuds) que nous appelons *entrelacs arborescents de Hopf*. Il s'agit de nœuds et d'entrelacs qui peuvent être décomposés en un *plombage* arborescent de bandes de Hopf, où une bande de Hopf est un anneau plongé dont le bord est un entrelacs de Hopf (voir la gauche de la Figure 7.28, et la droite pour une telle décomposition). Le *défaut de genre* est la différence entre le genre et le 4-genre. Nous prouvons le Théorème A qui stipule que, pour tout k, il existe un algorithme pour décider si un entrelacs arborescent de Hopf a un défaut d'au plus k. Ainsi, nous résolvons la question de la décidabilité du calcul du défaut sur cette classe de nœuds et d'entrelacs.

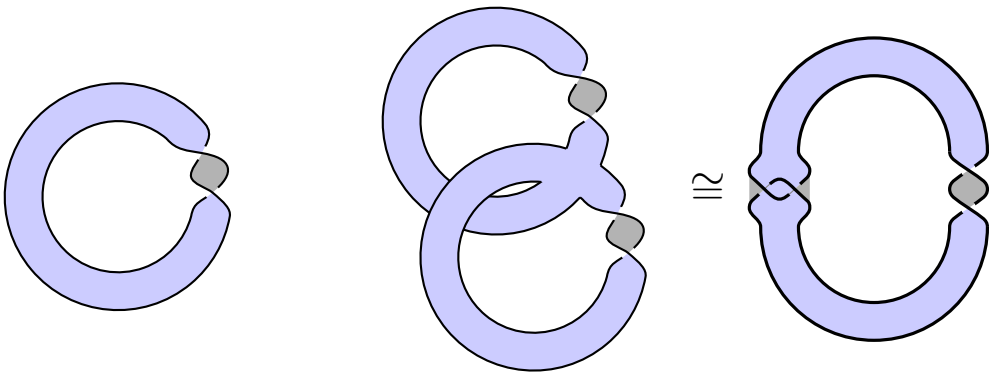

FIGURE 7.28 : À gauche : une bande de Hopf. À droite : un plombage de deux bandes de Hopf réalisant un nœud de huit.

Pour prouver ce théorème, nous nous inspirons de la théorie structurelle des graphes en définissant une relation de mineur, appelée mineur d'entrelacs (link-minor), sur cette classe d'entrelacs et en prouvant qu'il s'agit d'un **bel ordre**. Si un ensemble est ordonné par un bel ordre, toute suite infinie de cet ensemble contiendra au moins deux éléments comparables. Cela implique qu'il n'existe aucune **antichaîne** infinie dans l'ensemble, i.e. aucune suite infinie d'éléments non comparables deux à deux. Ainsi, tout ensemble défini par une propriété stable pour le bel ordre peut être caractérisé par l'antichaîne **finie** des éléments minimaux de son complément. Cela donne un algorithme pour vérifier si un élément x satisfait cette propriété lorsque la relation est décidable : il suffit de décider qu'aucun élément minimal n'est en relation avec x.

Notre technique de preuve suit cette stratégie et consiste à montrer que le défaut de genre est stable pour la relation de mineur d'entrelacs et à fournir un algorithme pour décider cette relation. La relation de mineur d'entrelacs repose sur une définition précise des plombages de Hopf vus comme opérations de construction effectuées le long d'arbres. Grâce à cette définition, nous pouvons associer un arbre et une surface de Seifert particulière à chaque entrelacs arborescent de Hopf. Tout d'abord, nous utilisons le théorème de Kruskal [73] sur ces arbres pour prouver que la relation de mineur d'entrelacs est un bel ordre. Ensuite, nous fournissons un algorithme pour décider de la relation de mineur d'entrelacs. Puis, nous prouvons la stabilité du défaut par mineur d'entrelacs en étudiant une seconde relation,

appelée mineur de surface, qui est plus faible que notre relation de mineur d'entrelacs mais qui se comporte bien par rapport au défaut de genre.

Le mineur de surface est une relation d'inclusion entre les surfaces de Seifert particulières associées à nos entrelacs. De plus, comme la relation de mineurs d'entrelacs est un bel ordre plus fin que la relation de mineur de surfaces, nos techniques de preuve produisent également le théorème B qui stipule que la relation de mineur de surface est un bel ordre sur les surfaces de Seifert associées à nos entrelacs.

- Nous définissons une classe de nœuds et d'entrelacs obtenus par plombages itérés de bandes de Hopf.
- Nous prouvons la décidabilité du défaut sur la classe des entrelacs arborescents de Hopf. Voir Théorème A.
- Nous prouvons que la relation de mineur de surface est un bel ordre sur la classe des entrelacs arborescents de Hopf. Voir Théorème B.

Invariant de largeur d'entrelacs et de graphes spatiaux inspiré de la théorie structurelle des graphes. Nous développons davantage le contexte et les concepts abordés dans la dernière partie de la section précédente. La recherche d'invariants de nœuds pouvant être calculés efficacement est un moyen de contourner la difficulté rencontrée par la reconnaissance des nœuds. Une autre méthode efficace pour s'attaquer aux problèmes difficiles consiste à développer des algorithmes, appelés *soluble à paramètre fixé* (*Fixed-Parameter-Tractable* abrégé en FPT [27]), dont la conception dépend d'une information supplémentaire de l'entrée : le paramètre. Leur principal intérêt est que la complexité de ces algorithmes est faible lorsque le paramètre est fixé. La largeur arborescente est un paramètre essentiel à cet égard : la conception d'algorithmes sur les graphes bénéficie grandement de la structure arborescente sous-jacente des graphes de petite largeur arborescente (voir pour référence cet exposé de Bodlaender [14]). L'application de cette méthode aux diagrammes de nœuds de faible largeur arborescente permet de calculer efficacement de nombreux invariants de nœuds (voir par exemple [86, 19, 87]), qui sont connus pour être difficiles à calculer dans le cas général. Comme expliqué précédemment, cet état de fait a conduit à la question [22, 86] de savoir s'il existe une famille de nœuds pour laquelle tous les diagrammes ont une largeur arborescente élevée ; laquelle a été répondue par la positive [28]. Notre travail se concentre sur reprouver et généraliser cette réponse en s'inspirant de la théorie structurelle des graphes.

La largeur arborescente est également un concept au cœur de la preuve du théorème des mineurs de Robertson et Seymour [118]. Ce paramètre a conduit à la naissance de nombreux autres invariants de largeur connexes qui peuvent présenter des caractéristiques à la fois théoriques et pratiques pour la résolution de problèmes. L'un d'entre eux est la *largeur en branches* (*branchwidth*), qui est équivalente à la largeur arborescente à un facteur constant près. La largeur arborescente et la largeur en branches sont toutes deux définies comme le minimum du maximum d'une mesure prise sur un ensemble de décompositions. Par essence, ces invariants sont difficiles à minorer puisque, pour ce faire, il faut prouver que chaque décomposition a une largeur élevée. La largeur en branches est particulièrement intéressante

pour nous car elle peut être interprétée géométriquement par des balayages arborescents utilisant des cercles d'une sphère sur laquelle un graphe planaire est plongé (rappelons que les diagrammes de nœuds sont des graphes planaires) [124]. En outre, il admet une *obstruction* optimale, appelée *tangle*, qui présente des aspects topologiques et dont l'existence permet de minorer la largeur en branches. Pour être plus précis, un tangle a un *ordre* qui représente sa taille, et un graphe admet un tangle d'ordre k si et seulement si la largeur en branches du graphe est plus grande que k. Par conséquent, le fait de fournir une obstruction de taille k garantit également que k est une borne inférieure sur la largeur en branches. Ces deux faits nous ont amenés à concevoir notre invariant de largeur sur les nœuds, appelé **spherewidth**, et inspiré de la largeur en branches. Il repose sur des **décomposition en sphères** de $\mathbb{S}^3$, et peut être considéré comme une généralisation dans $\mathbb{S}^3$, utilisant des sphères, des balayages de $\mathbb{S}^2$ susmentionnés. Formellement, une décomposition en sphères est une application continue $\mathbb{S}^3 \to T$ où T est un arbre binaire tel que : l'antécédent de chaque feuille est un point, les antécédents de chaque point dans chaque arête forment une 2-sphère, et les antécédents de chaque sommet intérieur forment une double bulle. Une **double bulle** est constituée de deux sphères qui s'intersectent sur un disque et représente intuitivement le moment où deux sphères fusionnent (voir Figure 7.29 pour une double bulle et une représentation bidimensionnelle d'une décomposition en sphères). La largeur de la décomposition est alors le nombre maximal d'intersections entre une sphère de la décomposition et le nœud (on peut compléter la décomposition en sphères illustrée en Figure 7.29 pour qu'elle ait une largeur de 4). La spherewidth est alors l'infimum de la largeur de toutes les décompositions en sphères.

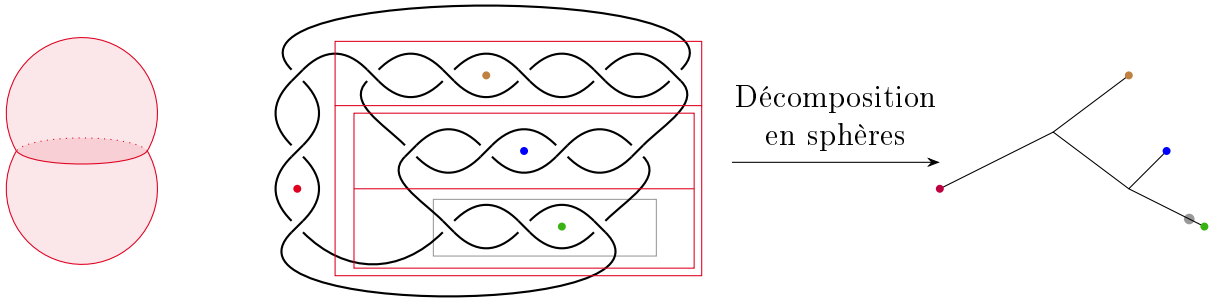

FIGURE 7.29 : À gauche : une double bulle. À droite : un exemple de décomposition en sphère où les antécédents de chaque sommet interne sont représentés en rouge, et les antécédents d'un point interne à une arête sont en gris.

Puisque les balayages des diagrammes de nœuds par des cercles peuvent être « relevés » en des décompositions en sphères, il s'ensuit que la spherewidth est une borne inférieure sur la largeur en branches des diagrammes, qui est elle-même une borne inférieure de la largeur arborescente des diagrammes. De manière similaire à la théorie des graphes, nous concevons une obstruction, appelée bubble tangle, imitant celle définie sur les graphes et minorant la spherewidth. Nous prouvons par le Théorème C que cette obstruction est également optimale : pour tout k, il existe soit une décomposition en sphères de largeur k, soit un bubble tangle d'ordre k, où l'ordre représente également la taille de notre obstruction.

Par ailleurs, nous fournissons des outils pour obtenir une telle obstruction, ce qui constitue un atout majeur de notre approche. Le Théorème D stipule qu'un bubble tangle d'ordre $\Omega(r)$

existe dès que le nœud est plongé sur une surface ayant une *représentativité de compression* (*compression representativity*) de r. La représentativité de compression quantifie à quel point le nœud représente la surface sur laquelle il est plongé. Par exemple, le nœud torique $T_{p,q}$, lorsqu'il est plongé sur le tore standard associé à sa définition, a une représentativité de compression $\min(p, q)$. Intuitivement, notre obstruction désigne un petit côté pour chaque sphère ayant un petit nombre d'intersections avec le nœud (ceci est à nouveau inspiré de l'obstruction sur les graphes). Lorsque le nœud est plongé sur une surface de genre non nul, les sphères avec un petit nombre d'intersections avec le nœud couperont des disques du tore d'un côté, tandis que l'autre contiendra la topologie du tore. Le petit côté est alors celui qui contient uniquement les disques (voir par exemple Figure 7.30).

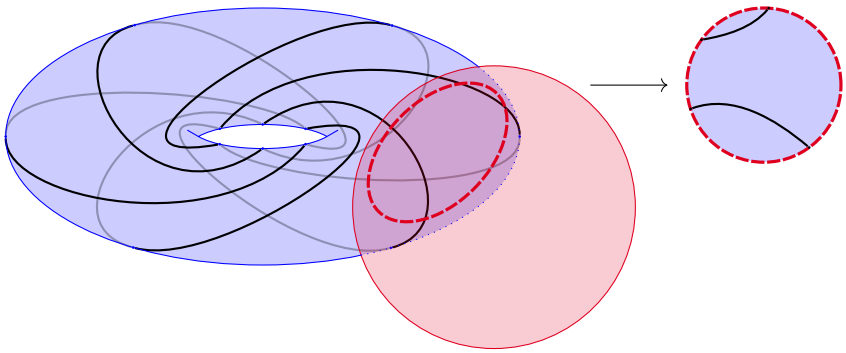

FIGURE 7.30 : Une intersection de petite taille entre un nœud torique $T_{5,6}$ plongé sur un tore et une sphere. Il s'ensuit que la topologie du tore est contenue dans un seul côté, alors que l'autre ne contient que des disques. Ce dernier côté est le petit côté.

Ainsi, notre travail reprouve, en utilisant de nouvelles techniques inspirées par la théorie structurelle des graphes, le fait que tout diagramme d'un nœud torique $T_{p,q}$ a une largeur arborescente de taille $\Omega(\min(p, q))$. Plus généralement, nous établissons une méthode systématique pour aborder de telles questions : par l'étude des surfaces sur lesquelles un nœud, entrelacs ou graphe spatial peut être plongé.

- Nous définissons la **spherewidth** : un invariant de largeur sur les nœuds et plus généralement sur les graphes spatiaux inspiré par la théorie structurelle des graphes. Nous montrons qu'il s'agit d'une borne inférieure sur la largeur arborescente de tous les diagrammes.
- Nous définissons des **bubble tangles** : une obstruction à la spherewidth inspirée par la théorie structurelle des graphes.
- Nous prouvons le Théorème C stipulant que nos bubble tangles sont une obstruction optimale de la sherewidth.
- Nous prouvons le Théorème D affirmant qu'un bubble tangle de taille $\Omega(k)$ existe dès qu'il existe une surface sur laquelle un nœud, entrelacs, ou graphe spatial peut

être plongé avec représentativité de compression k.

- Nous reprouvons, en utilisant des techniques élémentaires issues de la théorie structurelle des graphes, les résultats de [28] stipulant l'existence de nœuds dont tous les diagrammes ont haute largeur arborescente.

Une borne inférieure super constante sur la complexité de ltion des diagrammes d'entrelacs. Comme nous l'avons expliqué dans la section précédente, les mouvements de Reidemeister sont un outil puissant pour étudier les nœuds et entrelacs au travers de leurs diagrammes, grâce au théorème de Reidemeister. Ils apparaissent comme un outil très pratique et naturel pour étudier des considérations élémentaires sur les nœuds, comme leur nombre de croisements ou leur enlacements. Un exemple de choix de tel problème au sein de la théorie des nœuds, est la reconnaissance du nœud trivial, qui est une première instance du problème majeur de la théorie des nœuds : décider si deux nœuds sont équivalents ou non. Pour ce faire, une stratégie naturelle consiste à essayer de démêler un nœud donné en appliquant des mouvements de Reidemeister sur son diagramme de manière exhaustive ou aléatoire jusqu'à ce que le diagramme corresponde à une courbe simple. Cependant, certains diagrammes de nœuds triviaux [21], appelés **nœuds triviaux durs** (hard unknots), présentent un comportement gênant pour cet algorithme : le nombre maximum de croisements d'un diagramme au cours de l'algorithme est plus grand que le nombre initial. Nous *devons* d'abord ajouter des croisements avant de pouvoir atteindre le diagramme démêlé.

Formellement, désignons par $\mathbf{cr}(D)$ le nombre de croisements dans le diagramme D. Ensuite, pour deux diagrammes équivalents D_1, D_2 et une séquence de mouvements de Reidemeister R transformant D_1 en D_2, nous définissons $\mathbf{Top}(D_1, R)$ qui est le maximum de $\mathrm{cr}(D^i) - \mathrm{cr}(D_1)$ tout au long de la séquence de mouvements de Reidemeister où D^i est le diagramme D_1 après avoir effectué les i premiers mouvements de la séquence. La quantité qui nous intéresse est $\mathbf{Add}(D_1, D_2)$ qui est le minimum de $\mathrm{Top}(D_1, R)$ pris parmi toutes les séquences de mouvements de Reidemeister qui transforment D_1 en D_2. Si l'on considère D_2 comme un diagramme objectif, $\mathrm{Add}(D_1, D_2)$ est alors une borne inférieure sur le nombre de croisements à ajouter pendant l'exécution de l'algorithme susmentionné qui applique les mouvements de Reidemeister sur D_1 pour atteindre D_2.

L'étude de $\mathrm{Add}(D_1, D_2)$ s'avère plus délicate qu'il n'y paraît. A part une recherche exhaustive des mouvements de Reidemeister possibles qui devient rapidement irréalisable, aucune méthode n'est connue pour minorer cette quantité. Dans le contexte où D_1 est un diagramme de nœud et D_2 un diagramme de courbe simple, D_1 est un diagramme dur si $\mathrm{Add}(D_1, D_2)$ est positif. En fait, seuls des diagrammes pour lesquels $\mathrm{Add}(D_1, D_2) \leq 2$ sont connus [21] bien qu'il est conjecturé qu'il existe des diagrammes de nœuds triviaux D pour lesquels $\mathrm{Add}(D, D_2)$ est arbitrairement grand.

Nous nous concentrerons sur cette quantité dans le problème de la séparation : décider si un entrelacs L est **séparé**, c'est-à-dire s'il existe une sphère disjointe de L séparant au moins 2 composantes d'entrelacs de L. Si une telle sphère existe, il existe un diagramme d'entrelacs dans lequel deux sous-entrelacs sont séparés et disjoints dans le diagramme : ils sont séparés par un cercle dans le plan. Par conséquent, en termes de mouvements de Reidemeister, nous étudierons $\mathrm{Add}(D_1, D_2)$ où D_2 est un diagramme d'entrelacs d'un entrelacs L où un cercle

séparant le diagramme d'entrelacs peut être dessiné dans le plan, et D_1 est n'importe quel diagramme de L. Nous appellerons un diagramme d'entrelacs D_1 d'un entrelacs L pour lequel $\mathrm{Add}(D_1, D_2) > 0$ un diagramme de séparation difficile. Trouver une sphère dans l'espace séparant deux entrelacs est plus facile que de trouver un disque dont le bord est un nœud. Par conséquent, ce problème, qui est intéressant en soi, a été étudié à plusieurs reprises comme un problème utile et plus facile pour comprendre le problème de reconnaissance du nœud trivial [36, 78].

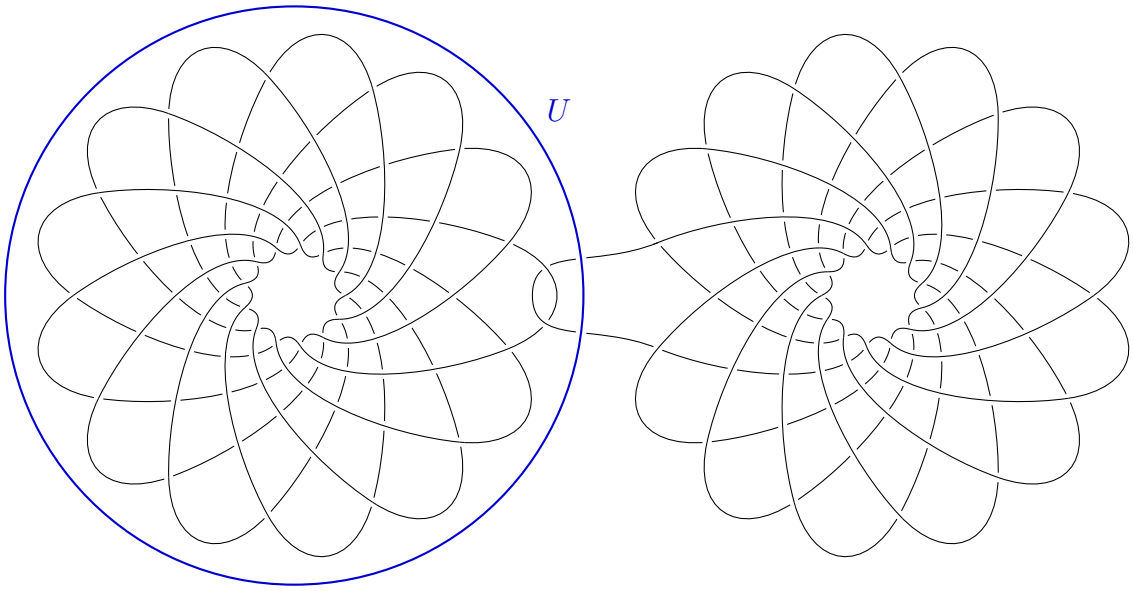

FIGURE 7.31 : Le diagramme d'entrelacs $\mathcal{D}(7, 13)$: deux nœuds toriques $T_{7,13}$ entrelacés et un nœud trivial U.

Nous présentons une famille de diagrammes d'entrelacs $\mathcal{D}(p, q)$ où deux sous-entrelacs sont séparés : le premier sous-entrelacs est constitué de deux nœuds toriques entrelacés et le second est un nœud trivial entourant l'un des nœuds toriques (voir $\mathcal{D}(7, 13)$ de la Figure 7.31 par exemple). En notant $\mathcal{D}'(p, q)$, n'importe quel diagramme d'entrelacs sur lequel U est disjoint des autres composantes d'entrelacs, nous prouvons le Théorème E impliquant que $\mathrm{Add}(\mathcal{D}(p, q), \mathcal{D}'(p, q)) = \Omega(\min(p, q))$. Si nous appelons **complexité de croisement** de D_1 le minimum de $\mathrm{Add}(D_1, D_2) > 0$ parmi tous les diagrammes d'entrelacs séparés D_2 de L, nous obtenons des diagrammes de séparation difficiles dont la complexité de croisement est arbitrairement grande. Plus précisément, le Théorème E montre que pour tout n, il existe un diagramme D_n d'un entrelacs séparé L_n dans $\mathbb{S}^3$ à 3 composantes tel que toute séquence de mouvements de Reidemeister le transformant en un diagramme séparé de L_n passe par un diagramme avec au moins $2n^2 + \frac{2}{3}n$ croisements.

La méthode utilisée ici consiste à exploiter le nœud U présent dans chacun de nos diagrammes et qui est séparé dans $\mathbb{S}^3$ des autres composantes d'entrelacs (voir la composante bleue de la Figure 7.31). Notre approche consiste à montrer que s'il existe une séquence de mouvements de Reidemeister où $\mathrm{Add}(\mathcal{D}(p, q), \mathcal{D}'(p, q))$ reste petit, nous pouvons utiliser l'évolution de U tout au long de ces mouvements pour définir un balayage des deux nœuds toriques entrelacés avec des sphères, où chaque sphère a un petit nombre d'intersections

avec cet entrelacs. Cependant notre obstruction du chapitre 4 a été exactement conçue pour montrer qu'un tel balayage est impossible. Par conséquent, $\mathrm{Add}(\mathcal{D}(p,q), \mathcal{D}'(p,q))$ doit être suffisamment grand.

- Nous définissons plusieurs familles de diagrammes d'entrelacs pour lesquelles le nombre de croisements à ajouter pour séparer l'entrelacs est arbitrairement grand, tel que stipulé par le Théorème E.
- Nous mettons au point une méthode pour fournir des bornes inférieures sur le nombre minimal de croisements à ajouter pendant l'exécution d'algorithmes reposant sur des essais successifs de mouvements de Reidemeister.

7.3 Organisation

Dans le Chapitre 2, nous passons en revue les préliminaires globaux de notre travail. En particulier, nous définissons correctement la plupart des concepts qui n'ont été que vaguement définis dans cette introduction.

Le Chapitre 3 se concentre sur les entrelacs arborescents de Hopf et la décidabilité du défaut de genre. Nous y présentons les théorèmes A et B de notre article [B], écrit avec Pierre Dehornoy et Arnaud de Mesmay, paru dans les Proceedings of the 40th International Symposium on Computational Geometry et invité à un numéro spécial de Discrete & Computational Geometry du Symposium on Computational Geometry 2024.

Le Chapitre 4 se concentre sur les balayages arborescents de $\mathbb{S}^3$ et leurs obstructions. Notre théorème de dualité, Théorème C, et notre théorème d'existence, Théorème D, y sont introduits et prouvés. Ce chapitre est principalement issu de notre article [A], écrit avec Arnaud de Mesmay, paru dans les Proceedings of the 39th International Symposium on Computational Geometry.

Le Chapitre 5 est le fruit d'un projet mené avec Arnaud de Mesmay et Jonathan Spreer. Nous y exploitons l'obstruction développée dans le chapitre 4 pour prouver notre Théorème E.

Enfin, le Chapitre 6 présente les principales conjectures et lignes de recherche restantes découlant de notre travail.

De plus, ce Chapitre 7 est une traduction en français du Chapitre 1.