AGT

Algebraic & Geometric Topology

Volume 26 (2026)

A structural approach to tree decompositions of knots and spatial graphs

CORENTIN LUNEL AND ARNAUD DE MESMAY

msp

msp

Algebraic & Geometric Topology **26**:6 (2026) 2015–2047
DOI: 10.2140/agt.2026.26.2015
Published: June 22, 2026

A structural approach to tree decompositions of knots and spatial graphs

CORENTIN LUNEL AND ARNAUD DE MESMAY

Knots are commonly represented and manipulated via diagrams, which are decorated planar graphs. The treewidth of a graph is a parameter quantifying how close a graph is to a tree. When a knot diagram has low treewidth, parameterized graph algorithms can be leveraged to ensure the fast computation of many invariants and properties of the knot. 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. In this work, 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. We then show the existence of this obstruction for knots of high representativity, which include, for example, torus knots, providing a new and self-contained proof that those do not admit diagrams of low treewidth. This last step is inspired by a result of Pardon on knot distortion.

1 Introduction

A (tame) **knot** is a polygonal embedding of the circle S^1 into $\mathbb{R}^3$, and two knots are considered equivalent if they are **isotopic**, i.e., if they can be continuously deformed one into the other without introducing self-intersections. The **trivial knot**, or **unknot**, is, up to equivalence, the embedding of S^1 as a triangle. The investigation of knots and their mathematical properties dates back to at least the nineteenth century [1] and has developed over the years into a very rich and mature mathematical theory. From a computational perspective, a fundamental question is to figure out the best algorithm testing whether a given knot is the unknot. It is not obvious that this problem is decidable, and in fact, this was famously posed as an open question by Turing [50]. The current state of the art on this problem is that it lies in NP [20] and co-NP [28], i.e., there exist polynomial-size certificates, that can be verified in polynomial times, certifying that a knot is trivial (NP) or not trivial (co-NP). Furthermore, a quasipolynomial-time algorithm has been announced [29] 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 [30].

In recent years, many attempts have been made to attack such seemingly hard problems via the route of parameterized algorithms. In particular, the **treewidth** of a graph is a parameter quantifying how close a graph is to a tree (we refer to Section 2 for the precise technical definitions of all the concepts discussed

MSC2020: primary 05C10; secondary 57K10, 57M15.
© 2026 The Authors, under license to MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via Subscribe to Open.

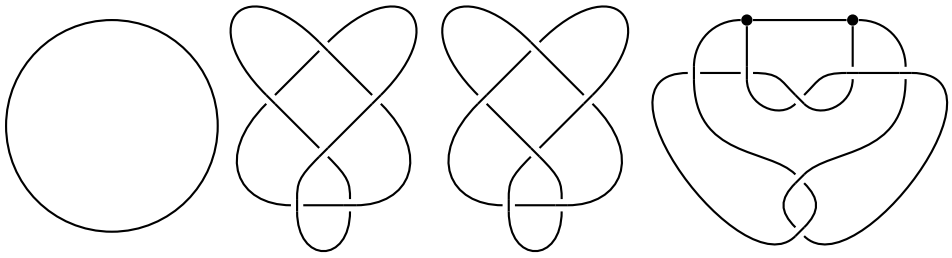

Figure 1: Diagrams of two trivial knots on the left, a bowline knot and a knotted spatial graph.

in this introduction), 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 instance. The concept of **branchwidth**, which we also use below, is somewhat equivalent and always within a constant factor of treewidth [39]. One approach is to study a knot via one of its **diagrams** (see Figure 1), that is, a decorated graph obtained by a planar projection where it is indicated on each vertex which strand goes over or under. Then, if such a diagram has low treewidth, one can apply these standard dynamic programming techniques to solve seemingly hard problems very efficiently. While this approach has not yet been successful for unknot recognition beyond treewidth 2 [5], it has proved effective for the computation of many knot invariants, including Jones and Kauffman polynomials [32] (which are known to be $\#P$-hard to compute in general [27]), HOMFLY-PT polynomials [6], and quantum invariants [8; 33]. Since any knot admits infinitely many diagrams, these algorithms naturally lead to the following question raised by Burton [7, page 2694], and Makowsky and Mariño [32, page 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.

The treewidth of knot diagrams echoes similar notions on the knot theory side, aiming at studying the "best" way to sweep a knot while trying to minimize the number of intersections in this sweep-out. One of the oldest knot invariants, the bridge number, can be seen through this lens (see, for example, [47]). A key concept in knot theory, introduced by Gabai in his proof of the property R conjecture [15], is the notion of **thin position** which more precisely quantifies the best way to place a knot to minimize its width. It is at the core of many advances in knot theory (see, for example, Scharlemann [43]). Taking inspiration from this idea, the previous question was answered recently by de Mesmay, Purcell, Schleimer and Sedgwick [35] 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 [22] on the *thin position* of *multiple Heegaard splittings*.

Our results The main purpose of this work is to provide new techniques to characterize which knots, or more generally which **spatial graphs** (polygonal embeddings of graphs into $\mathbb{R}^3$, considered up to isotopy; see, for example, Figure 1), do not admit diagrams of low treewidth. Our starting point is similar to the one in [35]: 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 2.2). 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 call this a **sphere decomposition**,[1] and the resulting measure (maximal number of intersections) the **spherewidth** of the knot.

Thus, in order to lower bound the treewidth of all the diagrams of a knot, it suffices 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 [41], handling families of graphs with bounded treewidth turns out not to be too hard [38], and in contrast, a large part is devoted to analyzing the structure shared by graphs of large treewidth. There, a fundamental contribution is the concept of a **tangle**.[2] We refer to Diestel [12] or Grohe [17] 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, as is proved in [39], for any graph G, the maximal possible order of a tangle is exactly equal to its branchwidth. 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 1.1 *For any knot or spatial graph K, the maximum order of a bubble tangle for K is equal to the spherewidth of K.*

This 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 disk** is a disk properly embedded in $\mathbb{R}^3 \smallsetminus \Sigma$ whose boundary is a noncontractible 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 disk. The **compression-representativity** of a knot or spatial graph is the supremum of that quantity over all embeddings on surfaces (this was originally defined by Ozawa [36]). Our second theorem is the following.

Theorem 1.2 *For any knot or spatial graph K embedded on a surface Σ in $\mathbb{R}^3$, there exists a bubble tangle of order $\frac{2}{3}$ of the compression-representativity of the embedding. Therefore, for any knot or spatial graph K, there exists a bubble tangle of order $\frac{2}{3}$ of the compression-representativity of K.*

[1] Our sphere decompositions are different from the ones in [35] but functionally equivalent for knots.

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

Combining together Theorems 1.1 and 1.2 with Proposition 2.2 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 [35], without relying on deep knot-theoretical tools.

Corollary 1.3 *A torus knot $T(p,q)$ has spherewidth at least $\frac{2}{3}\min(p,q)$, and thus any diagram of $T(p,q)$ has treewidth at least $\frac{1}{3}\min(p,q)$.*

Related work and proof techniques The results in this article and many of their proof techniques stem from two very distinct lineages in quite distant communities, the first one being knot theory or more generally low-dimensional topology, and the second one being structural graph theory. While there have been some recent works aiming at building bridges between combinatorial width parameters and topological quantities (for example, the aforementioned [35], but also [24; 25; 34] for related problems in 3-manifold theory), the main contribution in this article is that we dive deeper in 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 [9, preface]).

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 generalized Heegaard splittings [44; 45]; see also [22; 23; 49]). The concept of compression-representativity of a knot or a spatial graph finds its roots in the works of Ozawa [36], and Blair and Ozawa [4] 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 1.2 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 1.1. 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.

The duality theorem of Robertson and Seymour between branchwidth and tangles in [39] has been generalized many times since its inception, for example, in order to encompass other notions of decompositions and their obstructions [2; 31], to apply more generally to matroids [16] and to the wide-ranging setting of abstract separations systems [10; 11]. The key difference in our work, and why it does not fit into these generalizations, 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 [48] 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 [48] as well as to sphere-cut decompositions and their algorithmic applications (see, for example, [13]). Our sphere decompositions are the generalization one dimension higher of these sphere-cut decompositions, and Theorem 1.1 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 [42].

The **representativity** (also called **facewidth**) of a graph embedded on a surface S is the smallest number of intersections of a noncontractible curve with that graph. Theorem 1.2 will not come as a surprise for readers accustomed to graph minor theory, as Robertson and Seymour proved a very similar-looking theorem in [40, Theorem 4.1], showing that the branchwidth of a graph embedded on a surface is lower bounded by its representativity, which they prove by exhibiting a tangle. The key difference is that our notion of compression-representativity only takes into account the length of cycles bounding compression disks, instead of all the noncontractible 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 1.2; instead we have to rely on more topological arguments.

Finally, an important inspiration for our proof of Theorem 1.2 comes from a seemingly unrelated breakthrough of Pardon [37], who solved a famous open problem of Gromov [18] 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 proofs 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 [19, Lemma 4.2], where the simplicial map is similar to our sphere decompositions, up to a constant factor. Then our proof of Theorem 1.2 is inspired by the second part of Pardon's argument, with a quantitative strengthening to obtain the $\frac{2}{3}$ factor, whereas his argument would only yield $\frac{1}{2}$.

Organization of this paper After providing background and defining our key concepts in Section 2, we prove Theorem 1.1 in Section 3, and Theorem 1.2 in Section 4. We provide examples in Section 5.

2 Preliminaries

We include the most relevant definitions, but some familiarity with low-dimensional topology will help; see, for example, in the textbook of Schultens [46]. We refer to Diestel [9] for a nice introduction to graph theory and in particular its structural aspects.

Low-dimensional topology Following standard practice, instead of considering knots and spatial graphs within $\mathbb{R}^3$, we compactify it and work within $\mathbb{S}^3$. As is standard in low-dimensional topology, we work

in the piecewise-linear (PL) category, which means that *all the objects that we use in this article are assumed to be piecewise-linear*, i.e., made of a finite number of linear pieces with respect to a fixed triangulation of $\mathbb{S}^3$. This allows us to avoid pathologies such as wild knots or the Alexander horned sphere. An **embedding** of a compact topological space X into another one Y is a continuous injective map, and it is **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, [3, Theorem XIV.1]). A **knot** is 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 $\mathbb{S}^3$. All these objects are considered equivalent when they are ambient isotopic, i.e., when there exists a continuous deformation preserving the embeddedness. 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 emphasize that we do not thicken edges, which might be considered nonstandard. 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 disks. *From now on, all the graph embeddings will be thickened, and thus to ease notation we will omit the word thickened.*

As mentioned in the introduction, for Σ a surface embedded in $\mathbb{S}^3$, a **compression disk** is a properly embedded disk D in $\mathbb{S}^3 \setminus \Sigma$ such that the boundary ∂D is a noncontractible curve on Σ. 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. A **compressible curve** γ of Σ is the boundary of a compressing disk of Σ. 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 $G \hookrightarrow \Sigma \hookrightarrow \mathbb{S}^3$.

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 disks (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 2. Such a double bubble defines three spheres, which, by the PL Schoenflies theorem, bound three balls.

Two surfaces (respectively 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 (respectively 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

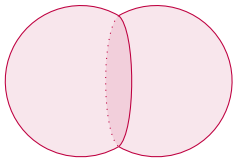

Figure 2: A double bubble: two spheres that intersect in a single disk.

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.

Spherewidth In this paragraph, we introduce sphere decompositions, which are the main way that we use to sweep knots and spatial graphs using spheres. Following standard notation in graph theory, we denote by $V(G)$, $E(G)$, and $L(G)$ the vertices, edges and leaves (degree-one vertices) of a graph G.

Definition 2.1 (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 with at least one edge:

- For all $x \in L(T)$, $f^{-1}(x)$ is a point disjoint from G.
- For all $x \in V(T) \smallsetminus L(T)$, $f^{-1}(x)$ is a PL 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:

$$\mathrm{sw}(G) = \inf_{f:\mathbb{S}^3 \mapsto T} \sup_{x \in \mathring{e} \in E(T)} |C(f^{-1}(x) \cap G)|.$$

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, and the spherewidth is a measure of how well we can sweep a graph G using sphere decompositions; see Figure 3. 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 1.1.

The point of using thickened embeddings instead of usual ones is that this allows disjoint spheres of a sphere decomposition to intersect a same vertex of a graph embedding. This is motivated by the following proposition, which provides a bridge between sphere decompositions and tree decompositions of diagrams of knots and spatial graphs.

For the sake of completeness, we include a definition of treewidth and of the neighborly concept of sphere-cut decompositions, but the reader unfamiliar to graph theory can safely skip these technicalities and treat Proposition 2.2 as a black-box for the rest of the paper. The **treewidth** of a graph G is defined as a min-max taken over a set of decompositions called tree decompositions. A tree decomposition of G

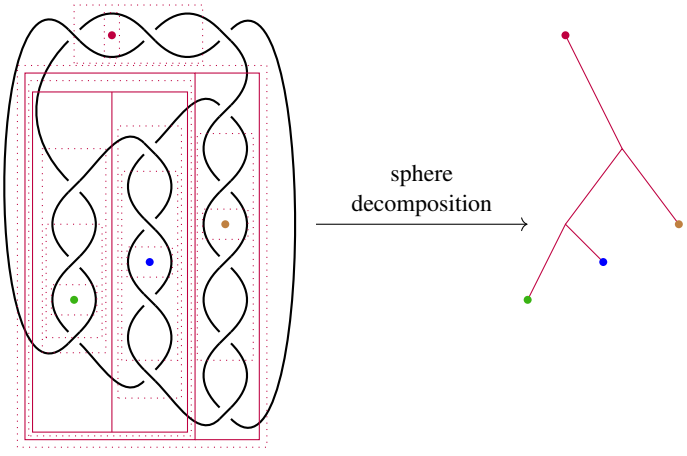

Figure 3: A width-4 sphere decomposition of a knot.

is a tree T whose vertices $X_1, \dots, X_t$, called bags, are subsets of $V(G)$. Furthermore, it must satisfy the following conditions:

- Each vertex is contained in at least one bag: $\bigcup_i X_i = V(G)$.
- If $v \in V(G)$ is contained in both X_i and X_j, every bag X_k in the path between X_i and X_j in T contains v.
- Each edge e of G is contained in a bag: there exists a bag X of T such that $e \subset X$.

The width $w(T)$ of a tree decomposition T is $w(T) = \max\{|X| - 1 \mid X \in V(T)\}$. The treewidth of G is then the minimal width of a tree decomposition of G: $\mathrm{tw}(G) = \min\{w(T) \mid T \text{ tree decomposition of G}\}$.

When G is a planar graph, the concept of sphere-cut decompositions [13] provides more structure than mere tree decompositions. 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 parameterized 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 [39] and the results of Seymour and Thomas on bond carving decompositions [48] (see, for example, [26, Lemma 2] for the exact statement that we use connecting branchwidth and sphere-cut width).

Proposition 2.2 *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 tree-width of D.*

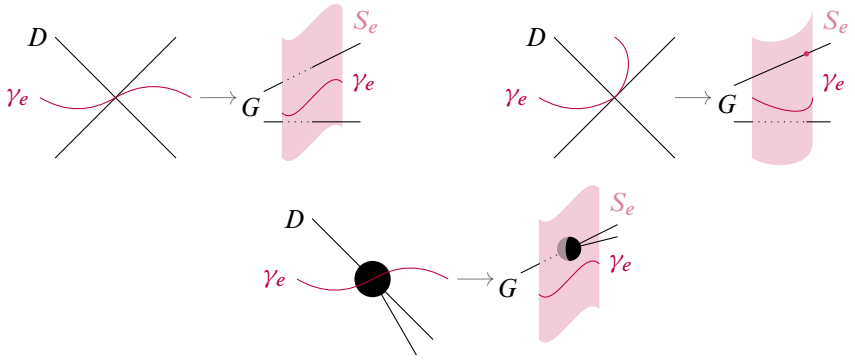

Figure 4: Three cases for the intersection between γ_e and D and their lift in $\mathbb{S}^3$.

The proof is very similar to that of Lemma 3.4 in [35].

Proof of Proposition 2.2 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 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 disks, and so that the Jordan curves γ_e intersect the interior of these disks. 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 disks. 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 disk above and one disk below the projection. The resulting spheres, which we denote by 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 two or one of the edges of G (depending on how γ_e splits the crossing; see Figure 4).

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

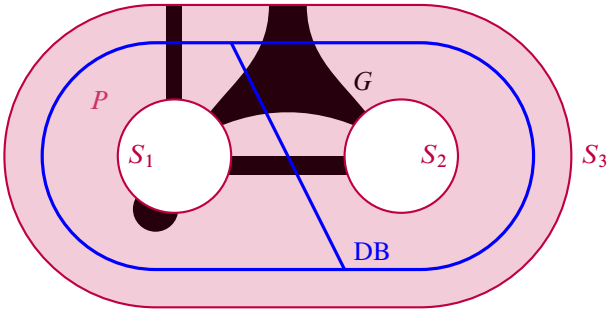

Figure 5: Definition of a double bubble DB within a solid pant P.

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 5).

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. □

Bubble tangle Bubble tangles are our second main concept in this article. 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 6). 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 $B(0,1) \setminus [-1,0] \times \{(0,0)\} \subset \mathbb{R}^3$.

We can now introduce bubble tangles.

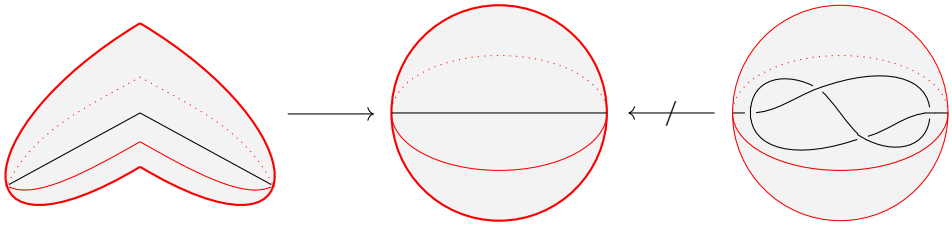

Figure 6: Representation of a G-trivial ball and a non-G-trivial ball.

Definition 2.3 Let G be an embedding of a graph in $\mathbb{S}^3$ and n be an integer such that $n \geq 1$. A **bubble tangle $\mathcal{T}$ of order** n 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 $\bar{B}_1$ or $\bar{B}_2$ is in $\mathcal{T}$, where $\mathbb{S}^3 \smallsetminus S = \{B_1, B_2\}$.

(T3) For every triplet of balls B_1, B_2, and B_3 of $\mathbb{S}^3$ 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 $|C(\partial B \cap G)| < n$,[3] then $B \in \mathcal{T}$.

In the following sections, we sometimes use bubble tangles of fractional order n, by which we mean a bubble tangle of order $\lceil n \rceil$ (we omit the $\lceil \cdot \rceil$ notation for clarity).

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 $\mathbb{S}^3$.

Remark 2.4 Tangles in graph theory are often endowed with an additional axiom, specifying that small sides should be stable under inclusion; see, e.g., [16, 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 1.1) and the connection to compression-representativity (Theorem 1.2). Whether such an axiom can be additionally enforced in our definition of bubble tangle while preserving these properties is left as an open problem. Note, however, that the compression bubble tangles that we will introduce in Section 4 do satisfy this stronger axiom.

3 Obstruction and duality

In this section, we prove Theorem 1.1: 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 other cases will be treated separately as special cases at the very end of this section.

The proof is split into two inequalities: Propositions 3.1 and 3.5 which together imply most cases of Theorem 1.1.

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

[3] The condition $|C(\partial B \cap G)| < n$ implies that when $n = 1$, the only G-trivial balls in $\mathcal{T}$ are the empty balls, and when $n = 2$, only the empty balls and balls such that $B \smallsetminus G$ is homeomorphic to $B(0,1) \smallsetminus [-1,0] \times \{(0,0)\} \subset \mathbb{R}^3$ are in $\mathcal{T}$.

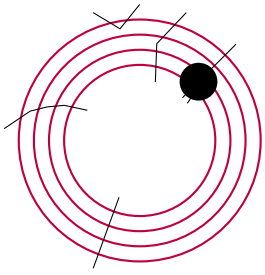

Figure 7: The three innermost spheres are braid-equivalent, but not the fourth one.

The proof of this proposition is similar to its graph-theoretical counterparts showing that tangles are an obstruction to branchwidth (see, e.g., [39]). The main difference with these proofs lies in the continuous aspects of our sphere decomposition, which we control using Lemmas 3.2 and 3.4.

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. The intuition behind this definition is that it means that G forms a braid between S and S'; see Figure 7. The following lemma explains how braid-equivalent spheres interact with a bubble tangle.

Lemma 3.2 *Let $\mathcal{T}$ be a bubble tangle and S, S' be two braid-equivalent spheres. 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 3.2 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 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 disks on $\mathbb{S}^2$, each one covering a hole of S_k: for all $i \in [\![1,k]\!]$, the space $\mathring{D}_i \setminus t_i$ is an open annulus. For $i \in [\![1,k]\!]$ notice that $h((D_i \setminus t_i) \times [0,1])$ is homeomorphic to $B(0,1) \setminus [-1,1] \times \{(0,0)\}$ (see Figure 8).

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

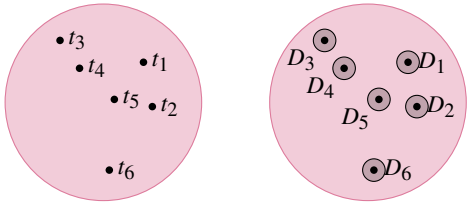

Figure 8: Disks $D_1, \dots, D_k$ over holes of $\mathbb{S}^2$.

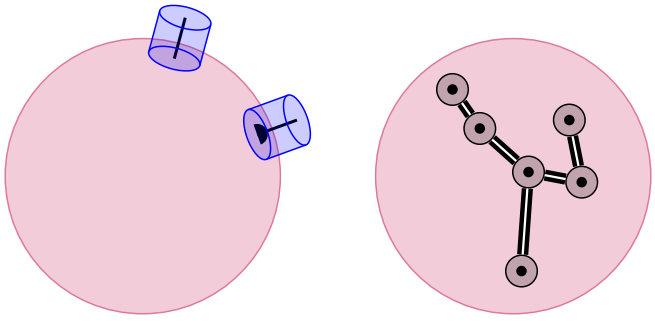

Figure 9: A G-trivial ball covering a strand and an example of tree connecting the disks D_i.

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 the edges of G by continuity: we define φ to send the boundary components of $h(S_k \times \{x\})$ to $\{t_i\} \times \{x\}$. There only remains 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 disks such that each disk is bounded by a circle induced by one of the boundary components of $h(S_k \times \{x\})$. Then we send each disk to $\{t_i\} \times \{x\}$ with φ.

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 $\varphi^{-1}(D_i \times [0,1]) \smallsetminus G$ equals $h((D_i \smallsetminus t_i) \times [0,1])$; see left of Figure 9. Hence β_i belongs to $\mathcal{T}$. Notice that B_1 intersects β_i on a disk 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, $\big(B_1 \cup \bigcup_{j\in[\![1,i]\!]} \beta_j\big) \cap \beta_{i+1} = D_{i+1}$. Hence, these balls induce a double bubble. By induction, we obtain that $B_1 \cup \bigcup_{j\in[\![1,k]\!]} \beta_j$ is in $\mathcal{T}$.

At that point, we have that $\big(B_1 \cup \bigcup_{j\in[\![1,k]\!]} \beta_j\big) \cap S' = h\big(\bigcup_{j\in[\![1,k]\!]} D_i \times \{1\}\big)$. Now, we consider a tree T connecting the disks D_i as in the right of Figure 9. Then, we thicken the edges of that tree to replace each edge e joining two disks 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 $B_1 \cup \bigcup_{j\in[\![1,k]\!]} \beta_j$ on a U shaped disk 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 $B_1 \cup \bigcup_{j\in[\![1,k]\!]} \beta_j$ to get $B \in \mathcal{T}$ such that B intersects S' on a single connected component: $h\big(\big(\bigcup_{e\in T} r_e \cup \bigcup_{i\in[\![1,k]\!]} D_i\big) \times \{1\}\big)$. Finally, we add to B the ball $\overline{B'_1 \smallsetminus B}$ disjoint from G which intersects S' on a disk : the complementary of $h\big(\big(\bigcup_{e\in T} r_e \cup \bigcup_{i\in[\![1,k]\!]} D_i\big) \times \{1\}\big)$. Since its intersection with B is a disk, they induce a double bubble and $B'_1 \in \mathcal{T}$ by (T3). □

Remark 3.3 Lemma 3.2 essentially states that small sides of braid-equivalent spheres are stable by containment. On the other hand, they might not be stable under containment in the reverse direction, i.e., with the same hypotheses as in the lemma, we might have $B'_1 \in \mathcal{T}$ but $B_1 \notin \mathcal{T}$. Here is an illustrative example. Consider a graph G consisting of two vertices u and v, a trefoil-knotted loop at u, a trefoil-knotted loop at v and a single straight edge $e = (uv)$ (see Figure 10). Furthermore, let x be a point the middle of u. We define a bubble tangle of order 3 as follows: a closed ball B belongs to the bubble tangle if $|C(\partial B \cap G)| < 3$ and either

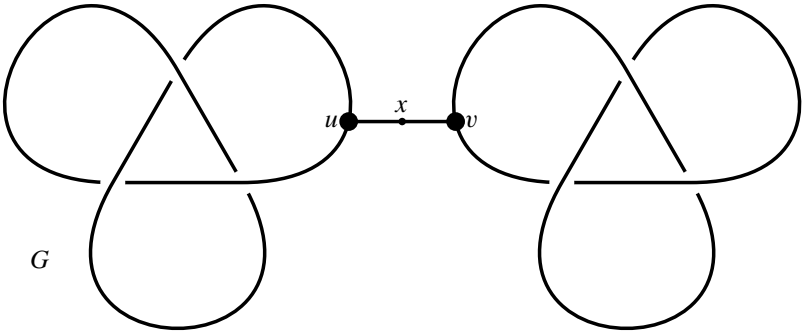

Figure 10: A graph for which small sides of braid equivalent spheres are not stable under isotopy.

(i) B is G-trivial,

(ii) B contains u and x in its interior,

(iii) B contains v and x in its interior, or

(iv) B contains u in its interior and x on its boundary.

One can check that the four axioms of bubble tangles are satisfied. However, braid-equivalent spheres may lead to different small sides depending on their location with respect to x. The point is that the axiom (T3) allows one to extend the small side of a bubble tangle to larger braid-equivalent spheres, but *not* to extend the big side to larger braid-equivalent spheres.

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 neighbor 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 notation: 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 3.4 *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 to the same tree such that o is consistent on T.*

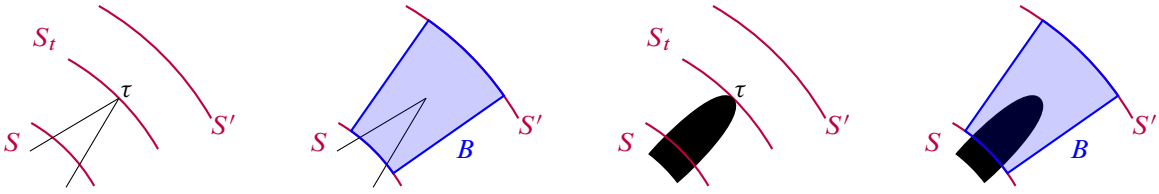

Figure 11: The tangency τ on S_1 and a cover by a G-trivial ball.

Proof of Lemma 3.4 Lemma 3.2 shows that the orientations of spheres are consistent between braid-equivalent spheres. Therefore, in order to prove Lemma 3.4, 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 $t_1, \dots, t_n \in [x, v]^n$ be all the points in $[x, v]$ where $f^{-1}(t_j)$ has a tangency with G. By Lemma 3.2, 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 11 (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 3.2. As we will use an argument similar to the proof of Lemma 3.2, let us introduce similar notation: $f^{-1}(x) = S$, $B_1 = f^{-1}(T_u(x))$, $f^{-1}(t_1) = S_t$, and I the space between S and S': $I = \mathbb{S}^3 \smallsetminus \big(\mathring{B}_1 \cup \mathring{f}^{-1}(T_v(y))\big)$.

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 disk 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 disk 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 disk, and containing g in its interior (see Figure 11). Then B is G-trivial. Indeed, either τ is the point of a "V" shape and $B \smallsetminus G$ is homeomorphic to $B(0,1) \smallsetminus [-1,1] \times \{(0,0)\}$, or $B \smallsetminus G$ is homeomorphic to $B(0,1) \smallsetminus [-1,0] \times \{(0,0)\}$ in the other cases (up to homeomorphism of $B \smallsetminus G$ 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 $\mathcal{T}$.

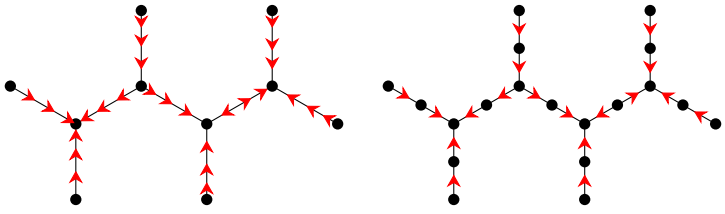

Figure 12: An example of T' from T leading to at least one sink.

At this stage, the space $I \smallsetminus (B \cup G)$ is homeomorphic to $S_k \times [0,1]$ so that we can use a similar technique as in the proof of Lemma 3.2 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 one of 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 3.4 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 3.1.

Proof of Proposition 3.1 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 3.4, 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 12.

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. □

Tightness of the obstruction In this paragraph, we 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 3.5 *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 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 $x \in V(T) \smallsetminus L(T)$, $f^{-1}(x)$ 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 which are not points are called **nontrivial 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 3.5 hinges on the following key lemma. Its proof is similar to branchwidth-tangle duality proofs [39] 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 3.6 *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 3.6 relies on the following preliminary lemma which allows us to slightly move the nontrivial leaves of a partial sphere decomposition.

Lemma 3.7 *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 nontrivial 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 3.7 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 nontrivial leaves. Now, $(I \smallsetminus B) \smallsetminus G$ is homeomorphic to $I \smallsetminus G$, since the subset of $\mathbb{S}^3 \smallsetminus G$ in $B \smallsetminus f^{-1}(\ell)$ 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 $I \smallsetminus B$ and I as is done in the proof of Lemma 3.2, and then we compose f with this homeomorphism to obtain a new partial sphere decomposition f' so that $w(f) = w(f')$ and $f'^{-1}(\ell) = B$. □

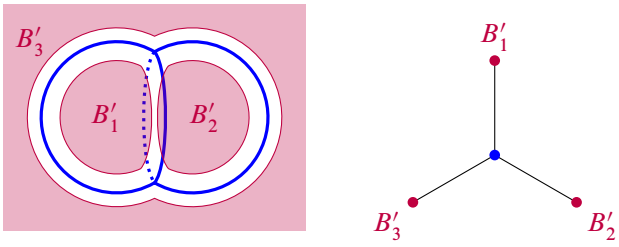

Figure 13: A partial sphere decomposition that conforms to $\mathcal{B}$.

Proof of Lemma 3.6 The proof relies on Zorn's lemma. Assume 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 collection 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 nontrivial 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 exists a triplet of closed balls B_1, B_2, B_3 of the bubble tangle $\mathcal{T}$ 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 bubble are transverse to G. Each of them admits a braid-equivalent sphere in its neighborhood, 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).

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 $h_i : S_k \times [0,1] \to I_i \smallsetminus G$; where $I_i = B_i \smallsetminus \mathring{B}'_i$. As seen in the proof of Lemma 3.2, 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

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

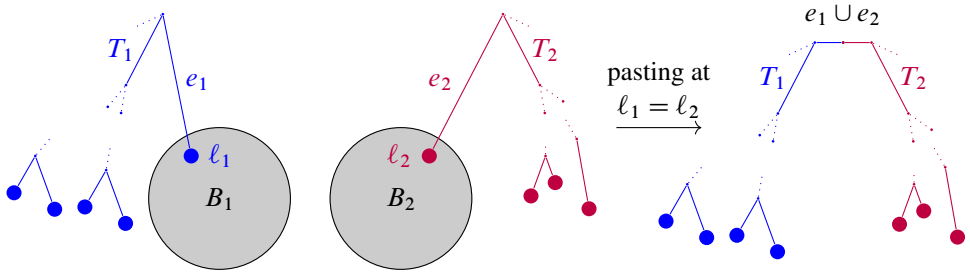

Figure 14: Pasting two partial sphere decompositions.

that conforms to B_1, B_2, B_3 (see Figure 13) on the trivalent tree made of 3 edges

$$(\{v_0, v_1, v_2, v_3\}, \{[v_0, v_1], [v_0, v_2], [v_0, v_3]\}).$$

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 S of weight less than k so that either both balls that it bounds are in $\mathcal{T}$, or none of them are. We get a contradiction with (NPD) by defining a partial sphere decomposition of width less than k that conforms to $\mathcal{T}$. In the first case, let S_1 and S_2 be two disjoint spheres which are braid-equivalent to S and on opposite sides of S. It is straightforward to define a partial sphere decomposition of width less than k for which the tree T consists of a single edge and the preimages of the two leaves are balls B_1 and B_2 so that $\partial B_1 = S_1$ and $\partial B_2 = S_2$. This partial sphere decomposition conforms to $\mathcal{T}$ and thus we have a contradiction. In the second case, if we denote by B_1 and B_2 the two 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 nontrivial 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 nontrivial 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 3.7, 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: we can define a map $f : \mathbb{S}^3 \to T$, where T is obtained by attaching the leaf ℓ_1 in T_1 to the leaf ℓ_2 in T_2, and the image of f within $f_1^{-1}(\ell_1)$ is given by f_2 while the image of f within $f_2^{-1}(\ell_2)$ is given by f_1 (see Figure 14). This map f is now a single partial sphere decomposition of width k which conforms to $\mathcal{T}$. We reach a final contradiction which concludes the proof. □

Remark 3.8 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 3.5.

Proof of Proposition 3.5 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 two, it also satisfies (T1) for k at least three. Therefore, by Lemma 3.6, 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. □

We can finally prove Theorem 1.1.

Proof of Theorem 1.1 Recall that Propositions 3.1 and 3.5 were established under the assumption that there exists a bubble tangle of order 3. There, in order to prove Theorem 1.1 we now establish the remaining special cases:

(1) If G admits a bubble tangle of order 1, then $\mathrm{sw}(G) \geq 1$.

(2) If G admits a bubble tangle of order 2, then $\mathrm{sw}(G) \geq 2$.

(3) If G is such that $\mathrm{sw}(G) \geq 1$, then G admits a bubble tangle of order 1.

(4) If G is such that $\mathrm{sw}(G) \geq 2$, then G admits a bubble tangle of order 2.

For item (1), note that for any nonempty knot or spatial graph G, $\mathrm{sw}(G) \geq 1$ so the statement is trivial.

For item (2), let assume that there exists a sphere decomposition $f : \mathbb{S}^3 \to T$ of width 1. Then, $f|_G$ is an injection and thus G is a forest. We will mimic the proof of Proposition 3.1, i.e., that the orientation o appearing on T as defined before Lemma 3.4 is consistent on T. In the following, we use the notation $S_x = f^{-1}(x)$ where x is a point interior to an edge of T.

The leaves are oriented inward because (T4) applies around them. Then, notice that if B is a closed ball of $\mathbb{S}^3$ such that $|C(B \cap G)| \leq 1$, then B is G-trivial because G is a forest. Hence, walking on T from the leaves, the orientation remains inward: it is consistent. We now focus on what happens when we leave a connected component of $f(G)$. Let y be a point interior to an edge e of T such that $S_y \cap G = \varnothing$ and x be another point interior to E such that there exists a unique tangency between G and a sphere S_t for t between x and y. Then, if we denote by I the space between S_x and S_y, $G \cap I$ is a unique segment g with one endpoint on S_x, since G is a forest. We cover I it by two G-trivial balls, one homeomorphic to $B(0,1) \setminus [-1,0] \times \{0,0\}$ containing g and one disjoint from G; they both intersect both S_x and S_y on a disc. By (T3), applied two times, $o(x) = o(y)$. This process is easily transposable when S_x is disjoint from G and S_y intersects a new connected component of $f(G)$. Thus, we prove that o is constant on T. We get a contradiction by finding a sink in T, in a similar way to Proposition 3.1.

For item (3), note that a bubble tangle of order 1 just contains empty balls, which satisfy the axiom tangles if G is nonempty, so such a bubble tangle always exists.

For item (4), we first observe that if G is such that $\mathrm{sw}(G) \geq 2$, then it contains at least a cycle C (since it is easy to see that forests have spherewidth 1). No sphere bounding a ball in a bubble tangle of order 2 can intersect the cycle C transversely, since there would be at least two intersections. Therefore we can

define a bubble tangle $\mathcal{T}$ by considering, for any sphere S intersecting G at most once, the ball B_C that does not contain C to be in $\mathcal{T}$. It is immediate to verify that the four axioms are satisfied and thus that $\mathcal{T}$ is a tangle.

These four cases complete the proof of Theorem 1.1. □

4 From compression representativity to bubble tangles

The goal of this section is to show Theorem 1.2: when a graph G is embedded on a compact, orientable, and nonzero 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.*

Under these hypotheses, the idea of the proof is to show that there exists a natural choice of small side for every sphere with fewer intersections with G than the compression representativity. Intuitively, such a sphere will only cut disks 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 4.1 *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 noncontractible curve in the intersection. Then one of the noncontractible 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 disks on S. Let α be a curve of $S \cap \Sigma$ that is innermost in S, i.e., it bounds a disk D in S that does not contain any other curve of $S \cap \Sigma$. If α is noncontractible, then the disk D is a compression disk for α, and thus α is compressible. Otherwise, α bounds a disk D_Σ in Σ (see, for example, Epstein [14, Theorem 1.7]). We deform S continuously by "pushing" D through D_Σ while keeping S embedded (see Figure 15) until α disappears from $\Sigma \cap S$.

Repeating this process on a new innermost curve of S will eventually yield a noncontractible 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 noncontractible. □

Recall that we denote the compression representativity of G on Σ by c-rep(G, Σ). A direct consequence of this lemma is that if G is embedded on a surface Σ, a sphere S intersects Σ, and the intersection has weight less than c-rep(G, Σ), then all the simple closed curves in the intersection are contractible.

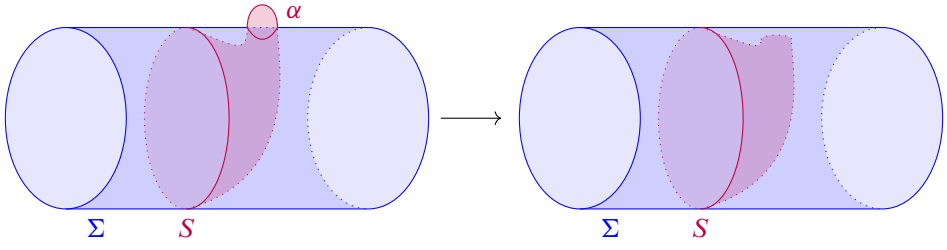

Figure 15: Removing a trivial curve from $S \cap \Sigma$.

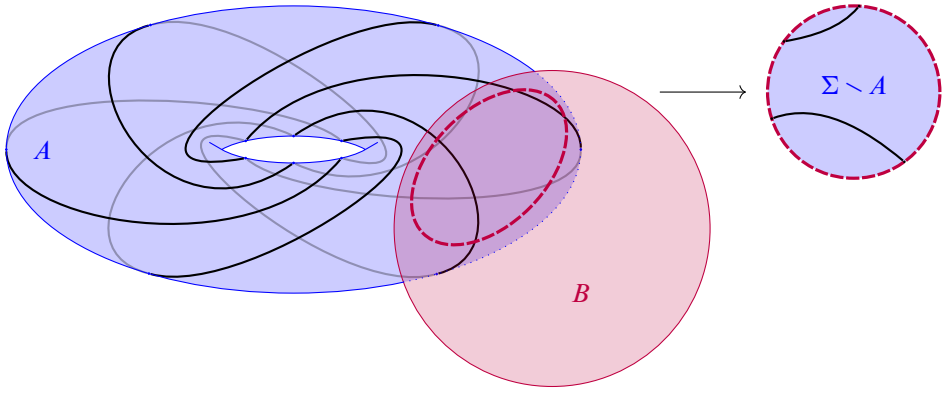

Figure 16: Intersection between a torus knot $T_{6,5}$ embedded on a torus and a sphere. Here the ball B containing the disk on the right is in the compression bubble tangle.

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 16). In order to formalize this, we will rely on fundamental groups (see, for example, Hatcher [21] for an introduction to this concept). Given a choice of basepoint x, the inclusion of a subsurface X on Σ induces a morphism on fundamental groups $i_*: \pi_1(X,x) \to \pi_1(\Sigma,x)$. If this morphism is trivial for every choice of basepoint, we say that X is **π_1-trivial** with respect to Σ.

Definition 4.2 (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 4.1, there is exactly one connected component A of $\Sigma \smallsetminus S$ that is not π_1-trivial. Exactly one of the open balls B of $\mathbb{S}^3 \smallsetminus S$ does not contain A, i.e., contains only π_1-trivial components of Σ; we put this side in c-$\mathcal{T}$: $\overline{B} \in$ c-$\mathcal{T}$.

The main step in the proof of Theorem 1.2 is to prove that a compression bubble tangle on the torus is indeed a bubble tangle. We first assume that c-rep$(G,\sigma) \geq 3$ and deal with the remaining cases at the very end of the proof.

Proposition 4.3 *If c-rep$(G,\Sigma) \geq 3$, a compression bubble tangle is a bubble tangle.*

Note that Proposition 4.3 immediately implies Theorem 1.2 by Proposition 3.1. Therefore, the remainder of this section is devoted to proving Proposition 4.3.

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

Lemma 4.4 *If c-rep$(G,\Sigma) \geq 3$ then for all G-trivial balls B, the intersection $B \cap \Sigma$ is π_1-trivial.*

Proof of Lemma 4.4 Let B be a G-trivial ball. Then its boundary S has weight at most two. By Lemma 4.1 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. If $B = B'$, we are done.

Otherwise, let us denote by K the component of G intersecting B (there is only one because B is G-trivial). We first observe that c-rep(G,Σ) = c-rep(K,Σ): the other components of G are contained in the ball B', and by the π_1-triviality of $B'\cap\Sigma$, they must all be contractible on Σ and thus can be avoided by compressible curves of Σ.

If K only intersects S once, then K is entirely contained in a disk $D\subset\Sigma$, and thus c-rep$(G,\Sigma)=0$, a contradiction. If K intersects S twice, replacing $K\cap B'$ by a single edge either leaves c-rep(K,Σ) unchanged if $K\cap B'$ is connected (since in that case, any intersection of a compressible curve γ with K inside of B' can be pushed on $\partial B'=S$), or possibly increases it if $K\cap B'$ is disconnected. But this transformation turns K into an unknot K', and c-rep$(K,\Sigma)\leq$ c-rep$(K',\Sigma)\leq 1$ (see, for example, [36, Theorem 3.2]), and we obtain here again a contradiction. □

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 disks, and these three disks can even come from a double bubble: we provide an example in the appendix and Figures 21, 23, and 24 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 disks on Σ.

Lemma 4.5 *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 disks.*

Proof of Lemma 4.5 By Lemma 4.1, 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 disk D_c on Σ (the other connected component being a surface with nonzero genus and a puncture), and we can define a **slope**[4]

$$s:\bigcup_{i\in\{1,2,3\}} C(\partial B_i\cap\Sigma)\to\mathcal{P}(\Sigma),\quad c\mapsto\text{the closed disk } D_c \text{ of } \Sigma \text{ such that } \partial D_c=c,$$

that associates the disk D_c to each c.

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 disk D_c in S_i such that all $c'\in C(\mathring{D}_c\cap\Sigma)$ satisfy $s(c')\subset B_i$ (see Figure 17). Denote by $D_{c'}$ the disk associated to each c' on S_i.

The idea is now to push $s(c)$, a disk of Σ, through D_c until $s(c)\subset B_i$ (see Figure 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 neighborhood of c and disjoint from S_i. Then $D_\alpha=s(\alpha)$ satisfies $D_c\subset D_\alpha$.

[4]This terminology mirrors the one of Robertson and Seymour in [40].

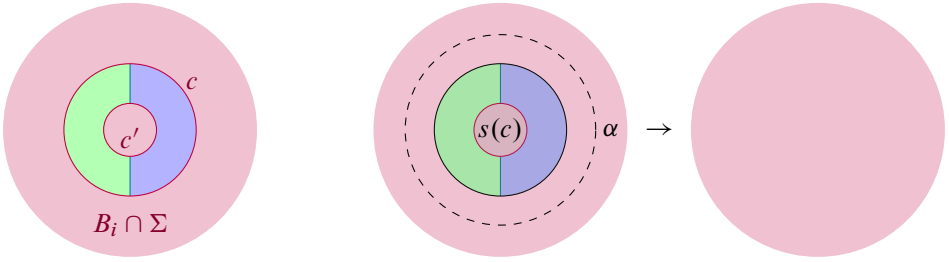

Figure 17: An example of c, innermost with the property that $s(c) \not\subset B_i$. The green and blue parts depict the intersection between Σ and the two other balls induced by the double bubble.

To be more precise, let Δ be the disk 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 disks (any other surface would contradict the fact that c is innermost). These disks 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 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 necessary 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 $A \cup \overset{\circ}{W}$ 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.

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 $\mathrm{DB} \cap \Sigma' = \mathrm{DB} \cap \Sigma \setminus s(c)$ and the transversality is preserved. Furthermore, $|S_i \cap G|$ did not increase so that the balls B_j still satisfy (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 disks. 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. □

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 disks with Σ and edges

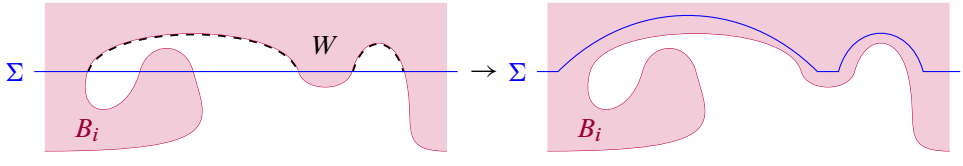

Figure 18: Pushing the surface Σ until c is removed from its intersection with B_i.

are the intersections of Σ with the disks. By Lemma 4.5, we can assume that this graph is trivalent and **cellularly embedded**, i.e., that all its faces are homeomorphic to disks. 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 we will use for the sake of contradiction.

Lemma 4.6 *Assume that we can cover $\mathbb{S}^3$ by three closed balls of c-$\mathcal{T}$ that induce a double bubble transverse to Σ and G. Then the total weight of Γ is less than c-rep(G,Σ):*

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

Proof Since each edge of Γ bounds exactly two faces of Γ, i.e, disks of Σ, and $\Gamma = \mathrm{DB}\cap\Sigma$ we get

$$|C(S_1\cap G)| + |C(S_2\cap G)| + |C(S_3\cap G)| = 2\sum_{e\in E(\Gamma)} |C(e\cap G)|. \tag{1}$$

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

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

Combining (1) and (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 4.7 *Assume that we can cover $\mathbb{S}^3$ by three closed balls of c-$\mathcal{T}$ that induce a double bubble transverse to Σ and G. Then there exists a set of edges X on Γ such that*

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

The proof of Proposition 4.7 is the technical crux of Theorem 1.2. 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 nontrivial way, and thus yield a compressible curve via Lemma 4.1. An additional difficulty is that this curve might be nonsimple in Γ; we circumvent this issue by finding a fractional version of such a curve instead, which will be strong enough to prove Proposition 4.7.

Proof of Proposition 4.7 In order to prove Proposition 4.7, 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 20 with a double bubble intersecting a torus.[5] The point of this merging process is to yield a family of balls, one of which will have a nontrivial intersection with Σ and thus allow us to (almost) find a compressible curve.

[5] 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 the appendix).

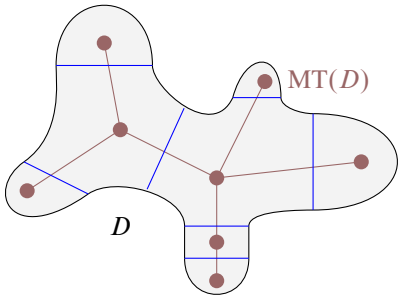

Figure 19: The membrane tree (brown) of a membrane D whose edges of Γ stem from $D \cap \Sigma$ (blue).

The merging process depends on the shape of one specific disk 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 disks 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 $\mathrm{MT}(D)$. Its set of vertices are connected components of $D \setminus \Sigma$, $V(\mathrm{MT}(D)) = C(D \setminus \Sigma)$, and there is an edge between $f, f' \in V(\mathrm{MT}(D))^2$ if $\bar{f} \cap \bar{f}' \neq \varnothing$. In the following we will consider $E(\mathrm{MT}(D)) = C(D \cap \mathrm{MT}(D))$ as there is a natural one-to-one correspondence between the sets (see Figure 19). Note that $\mathrm{MT}(D)$ is indeed a tree since it is the (weak) dual of an outerplanar graph.

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 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 20). All the other faces of Γ are now disks whose boundaries consist of edges with B_3.

We now pick an arbitrary vertex v of $\mathrm{MT}(D)$ 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 20.

For e an edge of $\mathrm{MT}(D)$ 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 20). Indeed, when the edge e of $\mathrm{MT}(D)$ is covered by the growing cylinder, the two surfaces containing the initial disks of B_1 and B_2 incident to e are glued along e. This process is then iterated by adding more edges of $\mathrm{MT}(D)$, thus growing a subtree of $\mathrm{MT}(D)$, until that subtree has fully grown and is $\mathrm{MT}(D)$. This is illustrated in the third and four rows of Figure 20. Eventually, the whole membrane is covered so that $B_{1,2} = B_1 \cup B_2$

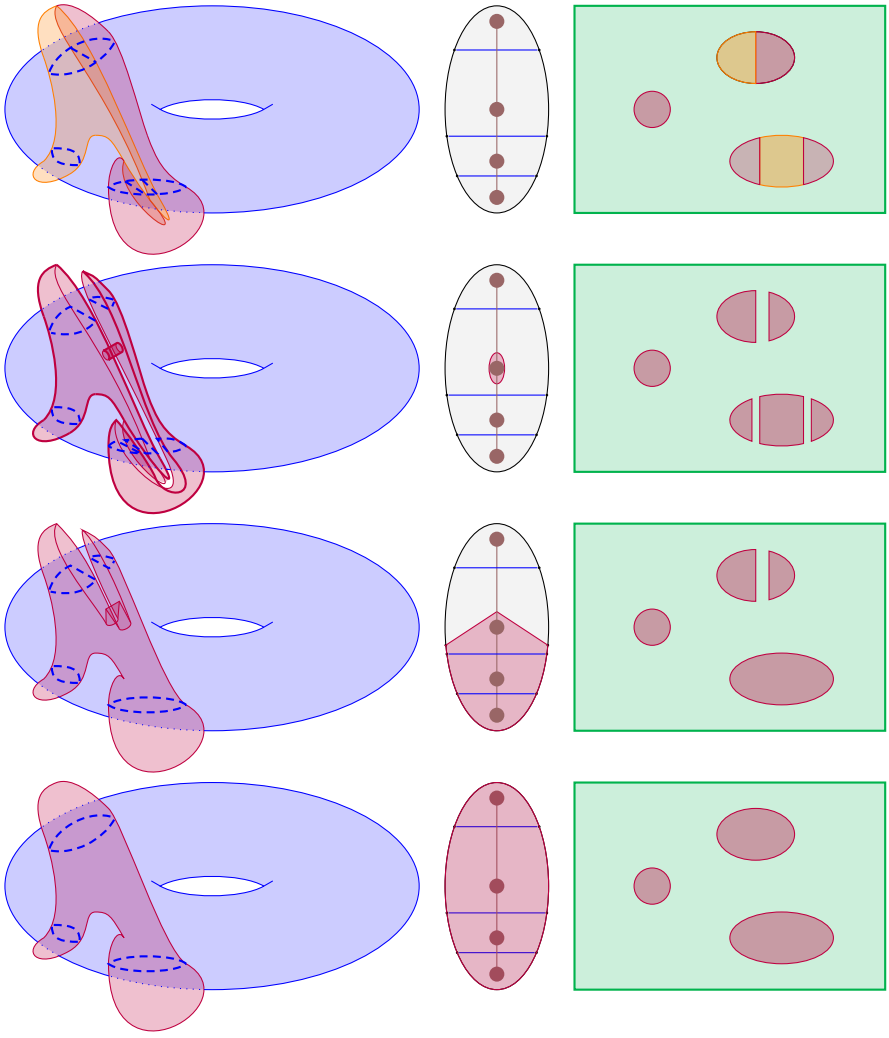

Figure 20: Three views of the merging process. Left: a 3D view of a double bubble (purple and orange) intersecting Σ (blue). Middle: the membrane tree, which is progressively covered by a subtree. Right: The unfolded torus and the graph Γ evolving throughout the merging process. It is first covered by 3 balls (orange, purple, green), then 2 (purple and green) after the cylinder merges two of them. The green ball is not shown on the left view for clarity purpose.

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

$$m : \mathrm{MT}(D)_{\subset} \to \mathcal{P}(\Sigma), \quad T \mapsto B_{1,2}(T) \cap \Sigma,$$

where $\mathrm{MT}(D)_{\subset}$ is the set of subtrees of $\mathrm{MT}(D)$ and $B_{1,2}(T)$ is the closed ball of the merging process parameterized 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 noncontractible 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 disk D_b on Σ that is disjoint from $\mathring{W}$. We define $\bullet(W) = W \cup \bigcup_{b \in C(\partial W), b \text{ contractible}} D_b$, 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 $\mathrm{MT}(D)_{\subset}$ such that $\bullet(m(T))$ is not a union of disks. Note that such a tree exists, since for $T = \mathrm{MT}(D)$, 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 disk. 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 disks. 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 disk of $\bullet(m(T \smallsetminus e))$, so that the merging of e glues together two distinct segments on the boundary of that disk. 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 noncontractible curves of Σ. Lemma 4.1 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'.

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

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

This proposition directly implies Proposition 4.3:

Proof of Proposition 4.3 A compression bubble tangle immediately satisfies the bubble tangle axioms (T1) and (T2) by definition, and (T4) by Lemma 4.4. 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 4.5, we can assume the graph Γ induced by the intersection of the double bubble with Σ is cellularly embedded. Then by Proposition 4.7, the total weight of Γ is at least c-rep(G, Σ). This is a contradiction with Lemma 4.6. □

Now, for the final proof:

Proof of Theorem 1.2 If c-rep$(G, \Sigma) \geq 3$, the proof directly follows from Proposition 4.3. Otherwise, if c-rep$(G, \Sigma) = 1$ or 2, we then prove that there exists a bubble tangle of order 2. As in the proof of the special cases of Theorem 1.1, such a bubble tangle exists if the graph G is not a forest. And that is the case since a forest is acyclic, so that for any surface Σ in which G is embedded, the compression-representativity must be zero. Finally, if c-rep$(G, \Sigma) = 0$, then it suffices to prove that there exists a bubble tangle of order 1, which is always the case. □

With regards to Remarks 2.4 and 3.3, we notice small sides of compression bubble tangles are stable by inclusion:

Proposition 4.8 *Let c-$\mathcal{T}$ be a compression bubble tangle of L, induced by an embedding of L into some surface $\Sigma \subset \mathbb{S}^3$, and let k be the order of c-$\mathcal{T}$. Then, for two closed 3-balls A, $B \subset \mathbb{S}^3$, if $A \in$ c-$\mathcal{T}$, $B \subset A$, and $|C(\partial B \cap L)| < k$, then $B \in$ c-$\mathcal{T}$. In other words, c-$\mathcal{T}$ is stable by inclusion up to (T1).*

Proof Let A and B be two closed balls of $\mathbb{S}^3$ such that $|C(\partial A \cap L)|, |C(\partial B \cap L)| < k$, 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 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-$\mathcal{T}$. □

5 Examples

A torus knot $T_{p,q}$ is a knot embedded on a surface of an unknotted torus $\mathbb{T}$ in $\mathbb{S}^3$, for example, a standard torus of revolution. It winds p times around the revolution axis, and q times around the core of the torus. We refer to Figure 16 for an illustration of $T_{6,5}$. We can now combine our results to lower bound the treewidth of any diagram of a torus knot, thus proving Corollary 1.3.

Proof of Corollary 1.3 Combining Proposition 2.2 and Theorems 1.1–1.2, 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 disks, 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 [1, Section 5.2] of any nontrivial knot. We refer to Ozawa [36, Theorem 1.6] for examples of spatial embeddings of any graph with high compression representativity, and thus high spherewidth.

We conclude by observing that the proof of Theorem 1.2 offers more flexibility than what the theorem states and can also be applied in some settings where the compression-representativity is low. For example, a connected sum of two knots $K_1 \# K_2$ has compression-representativity two (see [36, Corollary 1.9]), but if one these two knots, say K_1, has high compression-representativity separately, we can still define a bubble-tangle of high order by considering as big sides the balls containing the surface that K_1 is embedded on.

Appendix Covering a torus with three disks from a double bubble

A torus can be covered by a double bubble which induces three disks on the torus, as pictured in Figure 21. This double bubble is described in Figures 22–25 and one can prove that this is up to homeomorphism the only double bubble realizing such a cover.

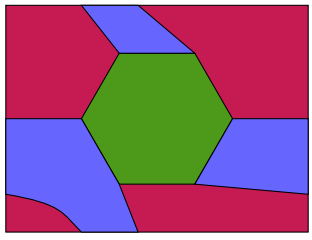

Figure 21: Covering a torus with three disks from a double bubble.

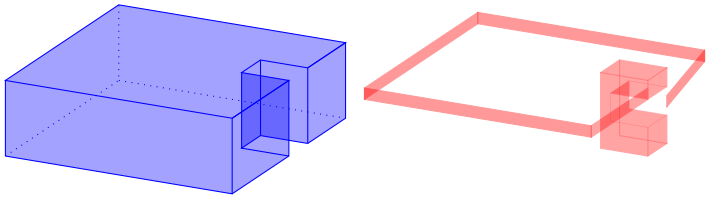

Figure 22: The spheres S_1 (blue) and S_2 (red) up to thickening the ribbon part.

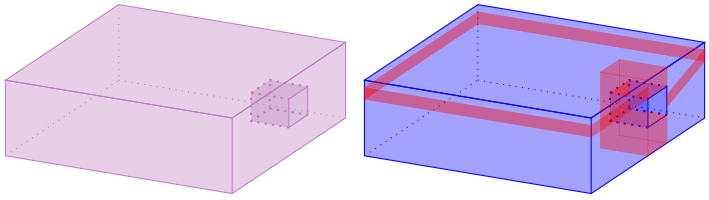

Figure 23: The ball $B_1 \cup B_2$ in purple, B_3 is the complementary.

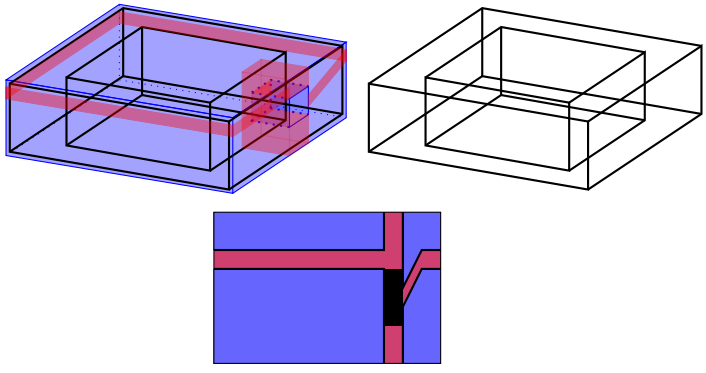

Figure 24: The double bubble and a torus inside, covered by three disks (up to thickening the ribbon part of S_2 by pushing it into S_1).

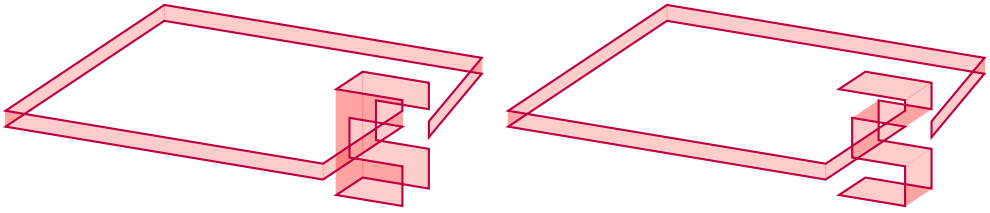

Figure 25: The disks $S_1 \cap S_2$ and $S_3 \cap S_2$.

Acknowledgements

This research was partially supported by the ANR project SoS (ANR-17-CE40-0033). This work was carried out during the PhD thesis of Lunel at Laboratoire d'Informatique Gaspard Monge of Université Gustave Eiffel. We would like to thank Pierre Dehornoy and Saul Schleimer for helpful discussions, and the reviewers for feedback which allowed us to significantly improve the paper.

References

[1] **C C Adams**, *The knot book*: *an elementary introduction to the mathematical theory of knots*, Amer. Math. Soc., Providence, RI (2004) MR

[2] **O Amini**, **F Mazoit**, **N Nisse**, **S Thomassé**, *Submodular partition functions*, Discrete Math. 309:20 (2009) 6000–6008 MR

[3] **R H Bing**, *The geometric topology of* 3*-manifolds*, American Mathematical Society Colloquium Publications 40, Amer. Math. Soc., Providence, RI (1983) MR

[4] **R Blair**, **M Ozawa**, *Height, trunk and representativity of knots*, J. Math. Soc. Japan 71:4 (2019) 1105–1121 MR

[5] **H L Bodlaender**, **B Burton**, **F V Fomin**, **A Grigoriev**, *Knot diagrams of treewidth two*, from "Graph-theoretic concepts in computer science" (I Adler, H Müller, editors), Lecture Notes in Comput. Sci. 12301, Springer (2020) 80–91 MR

[6] **B A Burton**, *The HOMFLY-PT polynomial is fixed-parameter tractable*, from "34th International Symposium on Computational Geometry" (B Speckmann, C D Tóth, editors), LIPIcs. Leibniz Int. Proc. Inform. 99, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2018) art. id. 18 MR

[7] **B A Burton**, **H Edelsbrunner**, **J Erickson**, **S Tillmann**, *Computational geometric and algebraic topology*, Oberwolfach Rep. 12:4 (2015) 2637–2699 MR

[8] **B A Burton**, **C Maria**, **J Spreer**, *Algorithms and complexity for Turaev–Viro invariants*, J. Appl. Comput. Topol. 2:1-2 (2018) 33–53 MR

[9] **R Diestel**, *Graph theory*, 5th edition, Graduate Texts in Mathematics 173, Springer (2017) MR

[10] **R Diestel**, **S-i Oum**, *Tangle-tree duality*: *in graphs, matroids and beyond*, Combinatorica 39:4 (2019) 879–910 MR

[11] **R Diestel**, **S-i Oum**, *Tangle-tree duality in abstract separation systems*, Adv. Math. 377 (2021) art. id. 107470 MR

[12] **R Diestel**, **G Whittle**, *Tangles and the Mona Lisa* (2016) arXiv 1603.06652

[13] **F Dorn**, **E Penninkx**, **H L Bodlaender**, **F V Fomin**, *Efficient exact algorithms on planar graphs*: *exploiting sphere cut branch decompositions*, from "Algorithms—ESA 2005" (G S l Brodal, S Leonardi, editors), Lecture Notes in Comput. Sci. 3669, Springer (2005) 95–106 MR

[14] **D B A Epstein**, *Curves on* 2*-manifolds and isotopies*, Acta Math. 115 (1966) 83–107 MR

[15] **D Gabai**, *Foliations and the topology of* 3*-manifolds, II*, J. Differential Geom. 26:3 (1987) 461–478 MR

[16] **J Geelen**, **B Gerards**, **N Robertson**, **G Whittle**, *Obstructions to branch-decomposition of matroids*, J. Combin. Theory Ser. B 96:4 (2006) 560–570 MR

[17] **M Grohe**, *Tangles and connectivity in graphs*, from "Language and automata theory and applications" (A-H Dediu, J Janoušek, C Martín-Vide, B Truthe, editors), Lecture Notes in Comput. Sci. 9618, Springer (2016) 24–41 MR

[18] **M Gromov**, *Filling Riemannian manifolds*, J. Differential Geom. 18:1 (1983) 1–147 MR

[19] **M Gromov**, **L Guth**, *Generalizations of the Kolmogorov–Barzdin embedding estimates*, Duke Math. J. 161:13 (2012) 2549–2603 MR

[20] **J Hass**, **J C Lagarias**, **N Pippenger**, *The computational complexity of knot and link problems*, J. ACM 46:2 (1999) 185–211 MR

[21] **A Hatcher**, *Algebraic topology*, Cambridge Univ. Press (2002) MR

[22] **C Hayashi**, **K Shimokawa**, *Thin position of a pair (*3*-manifold,* 1*-submanifold)*, Pacific J. Math. 197:2 (2001) 301–324 MR

[23] **Q He**, **S A Taylor**, *Links, bridge number, and width trees*, J. Math. Soc. Japan 75:1 (2023) 73–111 MR

[24] **K Huszár**, **J Spreer**, 3*-manifold triangulations with small treewidth*, from "35th International Symposium on Computational Geometry" (G Barequet, Y Wang, editors), LIPIcs. Leibniz Int. Proc. Inform. 129, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2019) art. id. 44 MR

[25] **K Huszár**, **J Spreer**, **U Wagner**, *On the treewidth of triangulated* 3*-manifolds*, J. Comput. Geom. 10:2 (2019) 70–98 MR

[26] **H Jacob**, **M Pilipczuk**, *Bounding twin-width for bounded-treewidth graphs, planar graphs, and bipartite graphs*, from "Graph-theoretic concepts in computer science" (M A Bekos, M Kaufmann, editors), Lecture Notes in Comput. Sci. 13453, Springer (2022) 287–299 MR

[27] **F Jaeger**, **D L Vertigan**, **D J A Welsh**, *On the computational complexity of the Jones and Tutte polynomials*, Math. Proc. Cambridge Philos. Soc. 108:1 (1990) 35–53 MR

[28] **M Lackenby**, *The efficient certification of knottedness and Thurston norm*, Adv. Math. 387 (2021) art. id. 107796 MR

[29] **M Lackenby**, *Unknot recognition in quasi-polynomial time*, lecture slides (2021) Available at http://people.maths.ox.ac.uk/lackenby/quasipolynomial-talk.pdf

[30] **M Lackenby**, *Algorithms in* 3*-manifold theory*, from "Surveys in differential geometry 2020: surveys in 3-manifold topology and geometry" (I Agol, D Gabai, editors), Surv. Differ. Geom. 25, Int. Press, Boston, MA (2022) 163–213 MR

[31] **L Lyaudet**, **F Mazoit**, **S Thomassé**, *Partitions versus sets*: *a case of duality*, European J. Combin. 31:3 (2010) 681–687 MR

[32] **J A Makowsky**, **J P Mariño**, *The parametrized complexity of knot polynomials*, J. Comput. System Sci. 67:4 (2003) 742–756 MR

[33] **C Maria**, *Parameterized complexity of quantum knot invariants*, from "37th International Symposium on Computational Geometry" (K Buchin, E Colin de Verdière, editors), LIPIcs. Leibniz Int. Proc. Inform. 189, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2021) art. id. 53 MR

[34] **C Maria**, **J S Purcell**, *Treewidth, crushing and hyperbolic volume*, Algebr. Geom. Topol. 19:5 (2019) 2625–2652 MR

[35] **A de Mesmay**, **J Purcell**, **S Schleimer**, **E Sedgwick**, *On the tree-width of knot diagrams*, J. Comput. Geom. 10:1 (2019) 164–180 MR

[36] **M Ozawa**, *Bridge position and the representativity of spatial graphs*, Topology Appl. 159:4 (2012) 936–947 MR

[37] **J Pardon**, *On the distortion of knots on embedded surfaces*, Ann. of Math. (2) 174:1 (2011) 637–646 MR

[38] **N Robertson**, **P D Seymour**, *Graph minors, V*: *Excluding a planar graph*, J. Combin. Theory Ser. B 41:1 (1986) 92–114 MR

[39] **N Robertson**, **P D Seymour**, *Graph minors, X*: *Obstructions to tree-decomposition*, J. Combin. Theory Ser. B 52:2 (1991) 153–190 MR

[40] **N Robertson**, **P D Seymour**, *Graph minors, XI*: *Circuits on a surface*, J. Combin. Theory Ser. B 60:1 (1994) 72–106 MR

[41] **N Robertson**, **P D Seymour**, *Graph minors, XX*: *Wagner's conjecture*, J. Combin. Theory Ser. B 92:2 (2004) 325–357 MR

[42] **H Sachs**, *On a spatial analogue of Kuratowski's theorem on planar graphs—an open problem*, from "Graph theory" (Łagów, 1981) (M Borowiecki, J W Kennedy, M M Sysło, editors), Lecture Notes in Math. 1018, Springer (1983) 230–241 MR

[43] **M Scharlemann**, *Thin position in the theory of classical knots*, from "Handbook of knot theory" (W Menasco, M Thistlethwaite, editors), Elsevier B. V., Amsterdam (2005) 429–459 MR

[44] **M Scharlemann**, **J Schultens**, **T Saito**, *Lecture notes on generalized Heegaard splittings*, World Scientific, Hackensack, NJ (2016) MR

[45] **M Scharlemann**, **A Thompson**, *Thin position for* 3*-manifolds*, from "Geometric Topology: Joint US-Israel Workshop on Geometric Topology" (Haifa, Israel), Contemporary Mathematics 164, Amer. Math. Soc., Providence, RI (1994) 231–238

[46] **J Schultens**, *Introduction to* 3*-manifolds*, Graduate Studies in Mathematics 151, Amer. Math. Soc., Providence, RI (2014) MR

[47] **J Schultens**, *The bridge number of a knot*, from "Encyclopedia of Knot Theory" (C Adams, E Flapan, A Henrich, L H Kauffman, L D Ludwig, S Nelson, editors), Chapman and Hall/CRC, Boca Raton (2021) 229–242

[48] **P D Seymour**, **R Thomas**, *Call routing and the ratcatcher*, Combinatorica 14:2 (1994) 217–241 MR

[49] **S Taylor**, **M Tomova**, *Additive invariants for knots, links and graphs in* 3*-manifolds*, Geom. Topol. 22:6 (2018) 3235–3286 MR

[50] **A M Turing**, *Solvable and unsolvable problems*, magazine article (1954)

CORENTIN LUNEL corentin.lunel@inria.fr
Inria Université Côte d'Azur, Sophia Antipolis, France

ARNAUD DE MESMAY arnaud.de-mesmay@univ-eiffel.fr
LIGM, CNRS, Université Gustave Eiffel, Marne la Vallée, France

Received: December 21, 2023 Revised: December 18, 2024

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Vesna Stojanoska
vesna@illinois.edu
University of Illinois at Urbana-Champaign

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Daniel Isaksen	Wayne State University isaksen@math.wayne.edu
Steven Boyer	Université du Québec à Montréal boyer.steven@uqam.ca	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Markus Land	JGU Mainz mland@uni-mainz.de
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Octav Cornea	Université de Montreal cornea@dms.umontreal.ca	Norihiko Minami	OCAMI (Osaka Central Adv. Math. Inst.) norihikominami@gmail.com
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futer	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Kathryn Hess	École Polytechnique Féd. de Lausanne kathryn.hess@epfl.ch		

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2026 is US $795/year for the electronic version, and $1170/year (+$80, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

© 2026 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 26 Issue 6 (pages 1965–2352) 2026

Detecting right-veering diffeomorphisms 1965
MIGUEL ORBEGOZO RODRÍGUEZ

A structural approach to tree decompositions of knots and spatial graphs 2015
CORENTIN LUNEL and ARNAUD DE MESMAY

Punctured JSJ tori and tautological extensions of Azumaya algebras 2049
YI WANG

Combination theorems for geometrically finite convergence groups 2079
ALEC TRAASETH and THEODORE WEISMAN

Modular links: bunch algorithm and upper volume bounds 2123
CONNIE ON YU HUI and JOSÉ ANDRÉS RODRÍGUEZ MIGUELES

Low complexity among principal fully irreducible elements of $\mathrm{Out}(F_3)$ 2157
NAOMI ANDREW, PAIGE HILLEN, ROBERT ALONZO LYMAN and CATHERINE EVA PFAFF

Dirichlet–Voronoi domain and injectivity radius of flag manifolds — equivariant cell structure on $O(3)/O(1)^3$ 2181
ARTHUR GARNIER

Mixed-platonic 3-manifolds 2215
ERIC CHESEBRO, MICHELLE CHU, JASON DEBLOIS, NEIL R HOFFMAN, PRIYADIP MONDAL and GENEVIEVE S WALSH

3d quantum trace map 2255
SAMUEL PANITCH and SUNGHYUK PARK

On the complexity of 2-bridge link complements 2315
JAMES MORGAN and JONATHAN SPREER