ON THE TREE-WIDTH OF KNOT DIAGRAMS[*]

Arnaud de Mesmay,[†] *Jessica Purcell,*[‡] *Saul Schleimer,*[§] *and Eric Sedgwick*[¶]

ABSTRACT. We show that a small tree-decomposition of a knot diagram induces a small sphere-decomposition of the corresponding knot. This, in turn, implies that the knot admits a small essential planar meridional surface or a small bridge sphere. We use this to give the first examples of knots where any diagram has high tree-width. This answers a question of Burton and of Makowsky and Mariño.

1 Introduction

The *tree-width* of a graph is a parameter quantifying how "close" the graph is to being a tree. While tree-width has its roots in structural graph theory, and in particular in the Robertson-Seymour theory of graph minors, it has become a key tool in algorithm design in the past decades. This is because many algorithmic problems, on small tree-width graphs, can be solved efficiently using dynamic programming techniques.

The algorithmic success of tree-width has also been demonstrated in the realm of topology. Makowsky and Mariño showed that, despite being $\#P$-hard to compute in general, the Jones and Kauffman polynomials can be computed efficiently on knot diagrams where the underlying graphs have small tree-width [14]. Recently Burton obtained a similar result for the HOMFLY-PT polynomial [6]. In a different vein, Bar-Natan used divide-and-conquer techniques to design an algorithm to compute the Khovanov homology of a knot; he conjectures that his algorithm is very efficient on graphs of low cut-width [4], a graph parameter closely related to tree-width. Similarly, many topological invariants can be computed efficiently on manifold triangulations for which the face-pairing graphs have small tree-width; this is nicely captured by a generalization of Courcelle's theorem due to Burton and Rodney [7].

A *diagram* D of a knot K is a four-valent graph embedded in S^2, the two-sphere, with some "crossing" information at the vertices. Thus, we can apply the definition of tree-width directly to D. The *diagrammatic tree-width* of K is then the minimum of the tree-widths of its diagrams. Burton [1, page 2694] and Makowsky and Mariño [14, page 755] pose a

---

natural question: is there a family of knots where the tree-width increases without bound? For example, while the usual diagrams of torus knots $T(p, q)$ (see Figure 1.1) can be easily seen to have tree-width going to infinity with p and q, it could be that other diagrams of these torus knots have a uniformly bounded tree-width.


Figure 1.1: The usual diagram of a torus knot $T(9, 7)$.

In this article, we give several positive answers. Our main tool comes from the analysis of surfaces in three-manifolds. We defer the precise definitions to Section 2.

Theorem 3.7. *Suppose that k is an integer and K is a knot having a diagram with tree-width at most k. Then either*

1. *there exists an essential planar meridional surface for K with at most $8k + 8$ boundary components or*

2. *K has bridge number at most $4k + 4$.*

The main idea in the proof of Theorem 3.7 is as follows. Suppose that D is the given small tree-width diagram of a knot K. In Lemma 3.4 we transform the given decomposition of D into a family of disjoint spheres, each meeting K in a small number of points. This partitions the ambient space into simple pieces. We call this a *sphere-decomposition* of K. In Lemma 3.5 we turn this into a *multiple Heegaard splitting*, as introduced by Hayashi and Shimokawa [10, page 303]. Their thin position arguments, roughly following [9], establish the existence of the surfaces claimed by Theorem 3.7.

To find families of knots with diagrammatic tree-width going to infinity, we use the contrapositive of Theorem 3.7. The following families of knots have neither small essential planar meridional surfaces nor small bridge spheres.

- Torus knots $T(p, q)$ where p and q are large, coprime, and roughly equal. These have no essential planar meridional surface by work of Tsau [26, page 199]. Also, the bridge number is $\min(p, q)$, by work of Schubert [22, Satz 10].

- Knots with a high distance Heegaard splitting, for example following Minsky, Moriah, and Schleimer [17, Theorem 3.1]. A theorem of Scharlemann [19, Theorem 3.1] shows that such knots cannot have essential planar meridional surfaces in their complements with a small number of boundary components. A theorem of Tomova [25, Theorem 1.3] shows that such knots cannot have small bridge spheres.

- Highly twisted plats, for example following Johnson and Moriah [13, Theorem 1.2]. They show that the given bridge sphere has high distance. A theorem of Bachman and Schleimer [3, Theorem 5.1] excludes small essential planar meridional surfaces. A theorem of Tomova [24, Theorem 10.3] excludes bridge spheres with smaller bridge number than the given one.

Corollary 1.2. *For every integer $k > 0$, there is a knot K with diagrammatic tree-width at least k.*

Proof. By Theorem 3.7, the three families of knots above have diagrammatic tree-width going to infinity. $\square$

1.1   Relations with the tree-width of the complement

Two recent works [11, 15] have investigated the interplay between the tree-width of a three-manifold and its geometric or topological properties. Here the tree-width of a three-manifold M is the minimum of the tree-widths of the face-pairing graphs of its triangulations. Interestingly, both articles find it more convenient to work with *carving-width* (also called *congestion*) than with tree-width directly. This is also the case in our paper; see Section 2 for definitions.

Huszár, Spreer and Wagner [11, Theorem 2] give a result similar to ours, but for three-manifolds. They prove that a certain family of manifolds $\{M_n\}$, constructed by Agol [2], has the property that any triangulation of M_n has tree-width at least n. Their result differs from ours in two significant ways. For one, Agol's examples are non-Haken, while we work with knots and their complements, which are necessarily Haken. Furthermore, there is an inequality between the diagrammatic tree-width of a knot and the tree-width of its complement which cannot be reversed. That is, given a knot diagram of tree-width k, there are standard techniques to build a triangulation of its complement of tree-width $O(k)$. For example, one may use the method embedded in SnapPy [8, `link_complement.c`]; see the discussion of [15, Remark 6.2].

To see that this inequality cannot be reversed recall that Theorem 3.7 gives a sequence of torus knots with tree-width going to infinity. On the other hand we have the following.

Lemma 2.10. *The complement of the torus knot $T(p,q)$ admits a triangulation of constant tree-width.*

This is well-known to the experts; we include a proof in Section 2. There is much more to say about the tree-widths of triangulations of knot complements; however in this article we will restrict ourselves to knot diagrams.

Acknowledgements

2 Background

2.1 Graph Theory

Unless otherwise indicated, all graphs will be simple and connected. The tree-width of a graph measures quantitatively how close the graph is to a tree. Although we provide a definition below for completeness, we will not use it in the proof of Theorem 3.7, as we rely instead on a roughly equivalent variant.

Definition 2.1. Suppose $G = (V, E)$ is a graph. A *tree-decomposition* of G is a pair $(\mathcal{T}, \mathcal{B})$ where $\mathcal{T}$ is a tree, $\mathcal{B}$ is a collection of subsets of V called *bags*, and the vertices of $\mathcal{T}$, called *nodes*, are the members of $\mathcal{B}$. Additionally, the following properties hold:

- $\bigcup_{B \in \mathcal{B}} B = V$

- For each $uv \in E$, there is a bag $B \in \mathcal{B}$ with $u, v \in B$.

- For each $u \in V$, the nodes containing u induce a connected subtree of $\mathcal{T}$.

The *width* of a tree-decomposition is the size of its largest bag, minus one. The *tree-width* of G, denoted by $\mathrm{tw}(G)$ is the minimum width taken over all possible tree-decompositions of G.

Tree-width has many variants, which all turn out to be equivalent up to constant factors. For our purposes, we will rely on the concept of carving-width; it has geometric properties that are well-suited to our setting.

Definition 2.2. Suppose that G is a graph. A *carving-decomposition* of G is a pair $(\mathcal{T}, \phi)$ where $\mathcal{T}$ is a binary tree and ϕ is a bijection from the vertices of G to the leaves of $\mathcal{T}$.

For an edge e of $\mathcal{T}$, the *middle set* of e, denoted $\mathrm{mid}(e)$ is defined as follows. Removing e from $\mathcal{T}$ breaks $\mathcal{T}$ into two subtrees S and T. The leaves of S and T are mapped via ϕ to vertex sets U and V. Then $\mathrm{mid}(e)$ is the set of edges connecting a vertex of U to a vertex of V.

The *width* of an edge e is the size of $\mathrm{mid}(e)$. The width of the decomposition $(\mathcal{T}, \phi)$ is the maximum of the widths of the edges of $\mathcal{T}$. Finally the *carving-width* $\mathrm{cw}(G)$ is the minimum possible width of a carving-decomposition of G.

The carving-width of a graph of constant degree is always within a constant factor of its tree-width, as follows.

Theorem 2.3. *[5, page 111 and Theorem 1] Suppose that G is a graph with degree at most d. Then we have:*

$$\frac{2}{3} \cdot (\mathrm{tw}(G) + 1) \leq \mathrm{cw}(G) \leq d \cdot (\mathrm{tw}(G) + 1).$$ $\square$

Recall that a *bridge* in a connected graph G is an edge separating G into more than one connected component. A *bond carving-decomposition* is one where, for any edge e in the tree $\mathcal{T}$, the associated vertex sets U and V induce connected subgraphs. One of the strengths of the notion of a carving-decomposition is the following theorem of Seymour and Thomas (see also the discussion in Marx and Piliczuk [16, Section 4.6]).

Theorem 2.4. *[23, Theorem 5.1] Suppose that G is a simple connected bridgeless graph with at least two vertices and with carving-width at most k. Then there exists a bond carving-decomposition of G having width at most k.* $\square$

Suppose that G has a bond carving-decomposition. Suppose also that G is *planar*: it comes with an embedding into S^2. Then each middle set $\mathrm{mid}(e)$ for G gives a Jordan curve $\gamma_e \subset S^2$ separating the two vertex sets U and V. One can take, for example, the simple cycle in the dual graph G^* corresponding to the cut. After a small homotopy of each, we can assume that these Jordan curves are pairwise disjoint. We say that a family of pairwise disjoint Jordan curves, crossing G transversely, *realizes* a carving-decomposition $(\mathcal{T}, \phi)$ if the partitions of the vertex sets that it induces correspond to those of $(\mathcal{T}, \phi)$. Theorem 2.4 and the above discussion yield the following.

Corollary 2.5. *Suppose that G is a bridgeless planar graph with at least two vertices and with carving-width at most k. Then there exists a family of pairwise disjoint Jordan curves realizing a bond carving-decomposition of G of width at most k.* $\square$

2.2 Knots

We refer to Rolfsen's book [18] for background on knot theory. Here we will use the equatorial two-sphere S^2, and its containing three-sphere S^3, as the canvas for our knot diagrams. In this way we avoid choosing a particular point at infinity. If one so desires, they may place the point at infinity on the equatorial two-sphere and so obtain the xy–plane in $\mathbb{R}^3$.

A *knot K* is a regularly embedded circle S^1 inside of S^3. One very standard way to represent K is by a *knot diagram*: we move K to be generic with respect to geodesics from the north to south poles and we then project K onto the equatorial two-sphere. We add a label at each double point of the image recording which arc was closer to the north pole.

Any knot has infinitely many knot diagrams, even when considering diagrams up to isotopy. The four-valent graph underlying a knot diagram need not be simple, but it is always connected. To avoid technical issues we may subdivide edges, introducing vertices of valence two, to ensure our graphs are simple. Since knot diagrams are four-valent, this makes at most a constant difference in any relevant quantity.

Figure 2.8: Three tangles represented by two-dimensional projections. The first one is trivial in a three-ball, the second one is non-trivial in a three-ball and the third one is flat in the solid pants.

Definition 2.6. Suppose that K is a knot with diagram D. The *tree-width* of D is defined to be the tree-width of the underlying graph (after any subdivision needed to ensure simplicity). We define the *carving-width* of D similarly.

The *diagrammatic tree-width* (respectively *diagrammatic carving-width*) of K is defined to be the minimum tree-width (respectively carving-width) taken over all diagrams of K.

Remark 2.7. It is relatively easy to find knots with small diagrammatic tree-width. For example, any connect sum of trefoil knots has uniformly bounded diagrammatic tree-width [18, Section 2.G]. For example, any two-bridge knot has uniformly bounded diagrammatic tree-width [18, Section 4.D and Exercise 10.C.6]. On the other hand, it is an exercise to show that any knot (even the unknot) has *some* diagram of high tree-width. This then motivates the question of Burton [1, page 2694] and of Makowsky and Mariño [14, page 755]: is there a family of knots $\{K_n\}$ so that *every* diagram of K_n has tree-width at least n?

2.3 Tangles

Suppose that M is a compact, connected three-manifold. A *tangle* $T \subset M$ is a properly embedded finite collection of pairwise disjoint arcs and loops. We call (M, T) a *three-manifold/tangle pair*, see Figure 2.8 for examples. We call the components of T the *strands* of the tangle. We take $N(T) \subset M$ to be a small, closed, tubular neighborhood of T. Let $n(T)$ be the relative interior of $N(T)$. So if $\alpha \subset T$ is a strand then $n(\alpha)$ is homeomorphic to an open disk crossed with a closed interval (or with a circle). We define the *tangle complement* to be $M_T = M - n(T)$. Now suppose that $F \subset M$ is a properly embedded surface (two-manifold) which is transverse to T. In particular, $T \cap \partial F = \emptyset$. Shrinking $n(T)$ if needed, we define $F_T = F - n(T)$. Note that F_T is properly embedded in M_T.

For any surface F, a simple closed properly embedded curve $\alpha \subset F$ is *essential* on F

Figure 2.9: Left: a compressing disk for a planar meridional surface. Right: a meridional disk for the boundary of a solid torus.

if it does not cut a disk off of F. A properly embedded arc $\alpha \subset F$ is *essential* on F if it does not cut a *bigon* off of F: a disk $D \subset F$ with $\partial D = \alpha \cup \beta$ and with $\beta \subset \partial F$.

Suppose that (M, T) is a three-manifold/tangle pair. A *meridian* for a tangle T is an essential simple closed curve in $\partial N(T) - \partial M$ that bounds a properly embedded disk in $N(T)$. The given disk is called a *meridional disk*. A surface F_T is *meridional* if ∂F is empty and thus every component of ∂F_T is a meridian of T. A surface is *planar* if it is a subsurface of the two-sphere. For example, if (B^3, T) is a three-ball/tangle pair, the boundary of the three-ball is a planar meridional surface. Similarly, a meridian for a solid torus $D \times S^1$ is an essential simple closed curve on the boundary $S^1 \times S^1$ bounding a properly embedded disk. For example, any disk of the form $D \times \{pt\}$, with product structure as above, is a meridional disk, see Figure 2.9, right.

With the notion of meridians in hand, we can give the promised proof of Lemma 2.10.

Lemma 2.10. *The complement of the torus knot $T(p, q) \subset S^3$ admits a triangulation of constant tree-width.*

We prove this by building an explicit triangulation of the complement of torus knots which is very path-like. We use Jaco and Rubinstein's layered triangulations [12].

Proof of Lemma 2.10. Recall that the two-torus $T \cong S^1 \times S^1$ has a triangulation with exactly one vertex, three edges, and two triangles. We call this the *standard triangulation* of T; we label the edges by a, b, c. We form $P = T \times [-1, 1]$ and triangulate it with two triangular prisms, each made out of three tetrahedra. Without changing the upper or lower boundaries $(T \times \{\pm 1\})$ we subdivide the triangulation of P to make $n(a \times \{0\})$ an open subcomplex. Finally, we form $Q = P - n(a \times \{0\})$, together with its triangulation, by removing those open cells.

Since the torus knot $T(p, q)$ is connected, the integers p and q are coprime. Using Bézout's identity, pick u and v such that $pv - qu = 1$. Following Jaco and Rubinstein [12, Section 4], build two layered solid tori U and V of slopes p/u and q/v: these are one-vertex triangulations of $D \times S^1$ so that

Figure 2.11: The daisy chain graph.

- the boundary of each is a standard triangulation of the two-torus,

- the face-pairing graph for the triangulation of U and V is a daisy chain graph (see Figure 2.11), and

- the meridian $\mu \subset \partial U$ crosses the edges (a, b, c) in ∂U exactly $(p, u, p + u)$ times while

- the meridian $\nu \subset \partial V$ crosses the edges (a, b, c) in ∂V exactly $(q, v, q + v)$ times.

We now glue U to P by identifying ∂U with the upper boundary of P, respecting the a, b, c labels. Similarly we glue V to P along the lower boundary, again respecting labels. Let $M = U \cup P \cup V$ be the result. Since the upper and lower halves of P are products, we may isotope μ down, and ν up, to lie in the middle torus $T \times \{0\}$. These curves now cross $|pv - qu| = 1$ times: that is, once. We deduce that M is homeomorphic to S^3; see for example Rolfsen [18, Section 9.B]. Now, the edge $a \times \{0\} \subset P \subset M$ crosses these copies of μ and ν exactly p and q times, respectively. So this edge gives a copy of $T(p, q)$ in S^3.

We now form $X(p, q) = S^3 - n(T(p, q))$ by drilling $a \times \{0\}$ out of M. We do this by replacing the triangulation of P by that of Q. The resulting triangulation of $X(p, q)$ has a standard triangulation in Q and a daisy chain triangulation in each of U and V. So the overall triangulation has constant tree-width, with the model tree being a path. $\qquad\square$

2.4 Compression bodies

Before defining sphere-decompositions of knots, we recall a few more notions from three-manifold topology.

Definition 2.12. A *one-handle* is a copy of the disk cross an interval, $D \times [-1, 1]$, with *attaching regions* $D \times \{-1\}$ and $D \times \{1\}$.

Suppose that F is a disjoint (and possibly empty) union of closed oriented surfaces. We form a *compression body* C by starting with $F \times [0, 1]$, taking the disjoint union with a three-ball B, and attaching one-handles to $(F \times \{1\}) \cup \partial B$. We attach enough one-handles to ensure that C is connected. Then $\partial_- C = F \times \{0\}$ is the *lower boundary* of C. Also, $\partial_+ C = \partial C - \partial_- C$ is the *upper boundary* of C. If $\partial_- C = \emptyset$ then C is a *handlebody*.

Here is a concrete example which we will use repeatedly.

Definition 2.13. Fix $n > 0$. Suppose that F is a disjoint union of $n - 1$ copies of S^2. We thicken, take the disjoint union with a three-ball B, and attach $n - 1$ one-handles in such a way to obtain a compression body $C(n)$. Note that $\partial_+ C(n) \cong S^2$. We call $C(n)$ an *n–holed three-sphere*. Equally well, $C(n)$ has exactly n boundary components, all spheres, and $C(n)$

embeds in S^3. Equally well, $C(n)$ is obtained from S^3 by deleting n small open balls with disjoint closures.

The four simplest n-holed three-spheres are the three-sphere (by convention), the three-ball, the shell $S^2 \times [0,1]$, and the *solid pants*: the thrice-holed three-sphere $C(3)$.

Suppose that B is the closed unit three-ball, centered at the origin of $\mathbb{R}^3$. Remove the open balls of radius $1/4$ centered at $(\pm 1/2, 0, 0)$, respectively to obtain the *standard model* for $C(3)$. The *equatorial* pair of pants $P \subset C(3)$ is the intersection of $C(3)$ and the xy–plane.

Definition 2.14. Suppose that $C(3)$ is the standard model of the solid pants. Suppose that $T \subset C(3)$ is a tangle. We say T is *flat* if

- T is properly embedded in $P \subset C(3)$, the equatorial pair of pants,

- T has no loop component, and

- all strands of T are essential in P.

Definition 2.15. Suppose that C is a compression body. The construction of C gives us a homeomorphism $C \cong N \cup H$ where N is either a three-ball or a copy of $\partial_- C \times [0,1]$ and where H is a disjoint union of one-handles. Suppose that $\alpha \subset C$ is a properly embedded arc.

- We call α a *vertical* arc if α lies in N and there has the form $\{\text{pt}\} \times [0,1]$.

- We call α a *bridge* if there is an embedded disk $D \subset C$ so that $\partial D = \alpha \cup \beta$ with $\beta \subset \partial_+ C$. We call D a *bigon* for α.

A tangle $T \subset C$ is *trivial* if

- every strand of T is either vertical or is a bridge and

- every bridge $\alpha \subset T$ has a bigon D so that $D \cap T = \alpha$.

Note that when C is a handlebody, a trivial tangle consists solely of bridges. We refer to Figure 2.8 for examples. Now we can begin to define the two types of surfaces that appear in Theorem 3.7.

Definition 2.16. Suppose that K is a knot in S^3. A two-sphere $S \subset S^3$ is a *bridge sphere* for K if

- S is transverse to K and

- the induced tangles in the two components of $S^3 - n(S)$ are trivial.

The *bridge number* of K with respect to S is the number of bridges in either trivial tangle.

Definition 2.17. Suppose that M is a compact connected oriented three-manifold. Suppose that $F \subset M$ is a properly embedded two-sided surface. An embedded disk $(D, \partial D) \subset (M, F)$ is a *compressing disk* for F if $D \cap \partial M = \emptyset$, if $D \cap F = \partial D$, and if ∂D is an essential loop in F, see Figure 2.9 left for an illustration. If F does not admit any compressing disk, then F is *incompressible*.

A *boundary compressing bigon* is an embedded disk $D \subset M$ with boundary $\partial D = \alpha \cup \beta$ being the union of two arcs where $\alpha = D \cap F$ is an essential arc in F and where $\beta = D \cap \partial M$. If F does not admit any boundary compressing disk, then F is *boundary-incompressible*.

A closed surface F, embedded in the interior of M is *boundary parallel* if F cuts a copy of $F \times [0, 1]$ off of M.

A surface $F \subset M$ that is incompressible, is boundary-incompressible, and is not boundary-parallel is called *essential*.

2.5   Multiple Heegaard splittings

Here we recall the concept of a multiple Heegaard splitting, which is central to the work of Hayashi and Shimokawa [10]. We also state one of their theorems that we will rely on.

Suppose that M is a compact connected oriented three-manifold. A *Heegaard splitting* $\mathcal{C}$ of M is a decomposition of M as a union of two compression bodies C and C', disjoint on their interiors, with $\partial_+ C = \partial_+ C'$ and with $\partial M = \partial_- C \cup \partial_- C'$. Following [20, 21], a *generalized Heegaard splitting* $\mathcal{C}$ of M is a path-like version of a Heegaard splitting; it is a decomposition of M into a sequence of compression bodies C_i and C_i', disjoint on their interiors, where $\partial_+ C_i = \partial_+ C_i'$, where $\partial_- C_i' = \partial_- C_{i+1}$, and where $\partial M = \partial_- C_1 \cup \partial_- C_n'$.

Following Hayashi and Shimokawa [10, page 303] we now generalize generalized Heegaard splittings. We model the decomposition of M on a graph instead of just a path. We also must respect a given tangle T in M. Here is the definition.

Definition 2.18. Suppose that (M, T) is a three-manifold/tangle pair. A *multiple Heegaard splitting* $\mathcal{C}$ of (M, T) is a decomposition of M into a union of compression bodies $\{C_i\}$, with the following properties.

1. For each index i there is some $j \neq i$ so that $\partial_+ C_i$ is identified with $\partial_+ C_j$.

2. For each index i and for each component $F \subset \partial_- C_i$ either F is a component of ∂M or there is an index j so that F is attached to some component $G \subset \partial_- C_j$. Here we allow $j = i$ but require $G \neq F$.

3. The surfaces $\cup \partial_\pm C_i$ are transverse to T. The tangle $T \cap C_i$ is trivial in C_i.

When discussing the union of surfaces we will use the notation $\partial_\pm \mathcal{C} = \cup \partial_\pm C_i$. Again, following [10, page 303] we give a complexity of multiple Heegaard splittings.

Definition 2.19. Suppose that (M, T) is a three-manifold/tangle pair. Suppose that $F \subset M$ is a connected closed surface embedded in the interior of M. Suppose also that F is transverse to T. The *complexity* of F is the ordered pair $c(F) = (\text{genus(F)}, |F \cap T|)$.

Now, suppose that $\mathcal{C}$ is a multiple Heegaard splitting of (M, T).

- The *cost* of $\mathcal{C}$ is the maximal number of intersections between a component of $\partial_+\mathcal{C}$ and T. That is, the cost is

$$\max\{|F \cap T| : F \text{ is a component of } \partial_+\mathcal{C}\}.$$

- The *width* of $\mathcal{C}$ is the ordered multiset of complexities

$$w(\mathcal{C}) = \{c(F) : F \text{ is a component of } \partial_+\mathcal{C}\}.$$

Here the complexities are listed in lexicographically non-increasing order.

The *width* of a pair (M, T) is the minimum possible width of a multiple Heegaard splitting of (M, T), ordered lexicographically.

A multiple Heegaard splitting of (M, T) is *thin* if it achieves the minimal width over all possible multiple Heegaard splittings of (M, T).

The following definitions, still from Hayashi and Shimokawa [10, page 304], provide a notion of essential surfaces in the setting of a three-manifold/tangle pair. Recall that if $X \subset M$ is transverse to T then X_T denotes $X - n(T)$.

Definition 2.20. Let X be a compact orientable 3-manifold, T a 1-manifold properly embedded in X, and F a closed orientable 2-manifold embedded transversely to T in X. Let $\tilde{X}$ be the 3-manifold obtained from X by capping off all the spherical boundary components disjoint from T with balls. An embedded disk Q is said to be a *thinning disk* of F if $T \cap Q = T \cap \partial Q = \alpha$ is an arc and $Q \cap F$ contains the arc $cl(\partial Q \setminus \alpha) = \beta$ as a connected component. The surface F is *T–essential* if

1. F_T is incompressible in X_T,

2. F has no thinning disk,

3. no component F_0 of F cobounds with a component F_1 of ∂X in $\tilde{X}$ a 3-manifold homeomorphic to $F_0 \times I$, possibly intersecting T in vertical arcs, where $F_0 = F_0 \times \{0\}$ and $F_0 = F_1 \times \{1\}$, and

4. no sphere component of F bounds a ball disjoint from T in $\tilde{X}$.

Note that when $\partial X = \emptyset$, which will be the case throughout this paper, a T–essential surface F is incompressible and boundary-incompressible in X_T. We can now state the necessary result from [10].

Theorem 2.21. *Suppose that $\mathcal{C}$ is a thin multiple Heegaard splitting of a three-manifold/tangle pair (M, T). Then every component of $\partial_-\mathcal{C}$ is T–essential in (M, T).*

*Proof.* This follows from Theorem 1.1 and Lemma 2.3 of [10]. □

Figure 3.3: The pretzel knot $P(-2, 3, 7)$ with a sphere-decomposition made of three spheres S_1, S_2 and S_3[2].

3   Sphere-decompositions

We connect multiple Heegaard splittings to tree-width by turning the Jordan curves of Corollary 2.5 into spheres, and broadening these into (almost) a multiple Heegaard splitting.

Definition 3.1. Let K be a knot in S^3. A *sphere-decomposition* $\mathcal{S}$ of K is a finite collection of pairwise disjoint, embedded spheres in S^3 meeting K transversely, and so that every component X of $S^3 - n(\mathcal{S})$ is either

- a three-ball, and $(X, X \cap K)$ is a trivial tangle, or

- a solid pants, and $(X, X \cap K)$ is a flat tangle.

We now define various notions of complexity of a sphere-decomposition.

Definition 3.2. The *weight* of a sphere S in a sphere-decomposition $\mathcal{S}$ is the number of intersections $S \cap K$. The *cost* of $\mathcal{S}$ is the weight of its heaviest sphere. The *width* of $\mathcal{S}$ is the list of the weights of its spheres, with multiplicity, given in non-increasing order.

As an example, a bridge sphere S for a knot in S^3 gives a sphere-decomposition with one sphere; the width is $\{2b\}$, where b is number of bridges of K on either side of S.

As another example consider a pretzel knot, as in Figure 3.3, or even more generally a Montesinos knot consisting of three rational tangles. Then the three two-spheres about the three rational tangles give a sphere-decomposition of width $\{4, 4, 4\}$.

A carving-decomposition of a knot diagram yields a sphere-decomposition of the knot, as follows.

Lemma 3.4. *Suppose that K is a knot. Suppose that D is a diagram of K with carving-width at most k. Then there is a sphere-decomposition of K with cost at most k.*

Proof. Let S be the equatorial sphere containing the graph D. Note that D is bridgeless. Thus, by Corollary 2.5, there is a family Γ of pairwise disjoint Jordan curves in S realizing a minimal width carving-decomposition of D. By assumption this has width at most k.

We obtain the knot K by moving the overstrands of D slightly towards the north pole of S^3 while moving the understrands slightly towards the south pole. Every Jordan curve $\gamma \in \Gamma$ can be realized as the intersection of a two-sphere $S(\gamma)$ with the equatorial sphere S: we simply cap γ off with disks above and below S. We do this in such a way that resulting spheres $S(\gamma)$ are pairwise disjoint and only meet K near S.

We claim that the resulting family of spheres $\mathcal{S} = \{S(\gamma)\}$ is a sphere-decomposition of K of cost at most k. First, the bound on the cost follows from the bound on the width of the initial carving-decomposition realized by Γ. The arborescent and tri-valent structure of the carving-decomposition translates into the fact that every sphere in $\mathcal{S}$ is adjacent, on each side, to either a pair of spheres (respectively to no other sphere) as the corresponding half-edge is not (respectively is) a leaf of the carving-tree. Since we are in S^3, this implies that the connected components of $S^3 - n(\mathcal{S})$ are either three-balls (at the leaves) or solid pants. At the leaves, the tangle in the three-ball is a small neighborhood of a single crossing in the original diagram D (or of a valence two vertex). Thus the tangles at the leaves are trivial. Inside each solid pants C, the tangle lies in the equatorial pair of pants P. This is because there are no crossings of D in P. Finally, the tangle is flat in P as otherwise we could find a carving of lower width. $\qquad\square$

A sphere-decomposition is not in general a multiple Heegaard splitting, as the example of Figure 3.3 shows. However, the following lemma shows that a sphere-decomposition of small cost can be upgraded to a multiple Heegaard splitting of small cost.

Lemma 3.5. *Suppose that $K \subset S^3$ is a knot, and suppose that K admits a sphere-decomposition $\mathcal{S}$ of cost at most k. Then (S^3, K) admits a multiple Heegaard splitting $\mathcal{C}$ of cost at most $2k$, where all components of $\partial_\pm \mathcal{C}$ are spheres.*

Proof. The construction simply adds a "thick" sphere into every component of $S^3 - n(\mathcal{S})$. Let C be a solid pants of $S^3 - n(\mathcal{S})$; label its boundaries U, V, and W and set $u = |K \cap U|$ and so on. Suppose that α is a strand of $K \cap C$ so that, relabelling as needed, α meets U and V. We form a two-sphere $S(C)$ by tubing U to V using an unknotted tube parallel to the strand α but such that α is outside of the tube. See Figure 3.6 for an example.

The sphere $S(C)$ meets K in $u + v$ points. By the definition of the cost, we have $\max\{u, v, w\} \leq k$, thus $u + v \leq 2k$. We perform this tubing construction in every solid pants C of $S^3 - n(\mathcal{S})$; we place the sphere $S(C)$ into the set $\mathcal{R}$.

For every three-ball component B of $S^3 - n(\mathcal{S})$ we take $S(B)$ to be a push-off of ∂B into B. We place the sphere $S(B)$ into the set $\mathcal{R}$. We claim that $\mathcal{C} = S^3 - n(\mathcal{R} \cup \mathcal{S})$ is a multiple Heegaard splitting, where $\partial_+ \mathcal{C} = \mathcal{R}$ and where $\partial_- \mathcal{C} = \mathcal{S}$. Note that every component $C' \subset \mathcal{C}$ is either a three-ball, a shell, or a solid pants. So there is a unique (up to

---

[2]This picture and the next one are adapted from a public domain figure from the Wikimedia commons by Sakurambo.

Figure 3.6: The pretzel knot $P(-2, 3, 7)$ with its natural sphere-decomposition, and with the addition of a "thick" sphere S' in the solid-pants component, obtained by tubing S_1 to S_2. The spheres of $\mathcal{S}$ are dashed and the thick sphere is dotted (compare with Figure 3.3).

isotopy) compression body structure on C' that has $\partial_+ C' \subset \mathcal{R}$ and $\partial_- C' \subset \mathcal{S}$. Let C be the component of $S^3 - n(\mathcal{S})$ containing C'.

- Suppose that C' is a three-ball. Then $K \cap C'$ is either one or two boundary parallel bridges, by the construction of $\mathcal{S}$.

- Suppose that C' is a shell and C is a three-ball. Then $K \cap C'$ is four vertical arcs, because $\partial_+ C'$ is parallel to ∂C.

- Suppose that C' is a shell and C is a solid pants. We use the notation of the construction: $\partial C = U \cup V \cup W$ and $\partial_+ C'$ is obtained by tubing U to V. The tangle $K \cap C'$ consists of vertical arcs (coming from strands of $K \cap C$ connecting U or V to W) and bridges parallel into the upper boundary (coming from strands of $K \cap C$ connecting U to V, or U to itself, or V to itself).

- Suppose that C' is a solid pants and thus C is a solid pants. Then $K \cap C'$, by construction, consists of vertical arcs only.

This concludes the proof.   □

We are now ready to prove our main result.

Theorem 3.7. *Suppose that k is an integer and K is a knot having a diagram with tree-width at most k. Then either*

1. *there exists an essential planar meridional surface for K with at most $8k + 8$ boundary components or*

2. *K has bridge number at most $4k + 4$.*

Proof. Let D be a diagram of K that has tree-width at most k. Then by Theorem 2.3, it has carving-width at most $4k + 4$. Applying successively Lemmas 3.4 and 3.5 to this carving-decomposition, we obtain a multiple Heegaard splitting $\mathcal{C}$ of (S^3, K) of cost at most $8k + 8$ and in which all the surfaces are spheres.

Now let us consider a thin multiple Heegaard splitting $\mathcal{C}'$ of (S^3, K). Since $\mathcal{C}'$ is thin, its width is less than or equal to that of $\mathcal{C}$. Thus all of the surfaces of $\partial_+ \mathcal{C}'$ must be spheres, as genus is the first item in our measure of complexity. We next deduce that the cost of $\mathcal{C}'$ is at most that of $\mathcal{C}$, so is at most $8k + 8$.

We now have two cases, depending on whether or not $\partial_- \mathcal{C}'$ is empty. If it is non-empty, then, taking $T = K$ in Theorem 2.21 tells us that every component of $\partial_- \mathcal{C}'$ is K–essential in S^3. Thus every component of $\partial_- \mathcal{C}'_K$ is incompressible and boundary-incompressible in S^3_K. Furthermore, since all the surfaces of $\partial_+ C'$ are spheres, the components of $\partial_- \mathcal{C}'_K$ are not boundary-parallel, therefore they are essential. Since the cost of $\mathcal{C}'$ is at most $8k + 8$, we deduce that K admits an essential planar meridional surface with at most $8k + 8$ boundary components.

Suppose instead that $\partial_- \mathcal{C}'$ is empty. Thus $\partial_+ \mathcal{C}'$ consists of a single sphere S, which is necessarily a bridge sphere. Again, the cost of $\mathcal{C}'$ is at most $8k + 8$, so S has at most $4k + 4$ bridges. This concludes the proof. $\square$

References

[1] Computational geometric and algebraic topology. Abstracts from the workshop held October 11–17, 2015. *Oberwolfach Rep.*, 12(4):2637–2699, 2015.

[2] Ian Agol. Small 3–manifolds of large genus. *Geom. Dedicata*, 102:53–64, 2003. arXiv:math/0205091.

[3] David Bachman and Saul Schleimer. Distance and bridge position. *Pacific J. Math.*, 219(2):221–235, 2005. arXiv:math/0308297.

[4] Dror Bar-Natan. Fast Khovanov homology computations. *J. Knot Theory Ramifications*, 16(3):243–255, 2007.

[5] Dan Bienstock. On embedding graphs in trees. *J. Combin. Theory Ser. B*, 49(1):103–136, 1990.

[6] Benjamin A. Burton. The HOMFLY-PT polynomial is fixed-parameter tractable. In Bettina Speckmann and Csaba D. Tóth, editors, *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, volume 99 of *LIPIcs*, pages 18:1–18:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.

[7] Benjamin A. Burton and Rodney G. Downey. Courcelle's theorem for triangulations. *J. Combin. Theory Ser. A*, 146:264–294, 2017.

[8] Marc Culler, Nathan M. Dunfield, Matthias Goerner, and Jeffrey R. Weeks. SnapPy, a computer program for studying the geometry and topology of 3-manifolds. Available at http://snappy.computop.org.

[9] David Gabai. Foliations and the topology of 3-manifolds. III. *J. Differential Geom.*, 26(3):479–536, 1987.

[10] Chuichiro Hayashi and Koya Shimokawa. Thin position of a pair (3-manifold, 1-submanifold). *Pacific J. Math.*, 197(2):301–324, 2001.

[11] K. Huszár, J. Spreer, and U. Wagner. On the treewidth of triangulated 3-manifolds. *ArXiv e-prints*, December 2017, arXiv:1712.00434.

[12] W. Jaco and J. Hyam Rubinstein. Layered-triangulations of 3-manifolds. *ArXiv Mathematics e-prints*, March 2006, arXiv:0603601.

[13] Jesse Johnson and Yoav Moriah. Bridge distance and plat projections. *Algebr. Geom. Topol.*, 16(6):3361–3384, 2016.

[14] J. A. Makowsky and J. P. Mariño. The parametrized complexity of knot polynomials. *J. Comput. System Sci.*, 67(4):742–756, 2003. Special issue on parameterized computation and complexity.

[15] C. Maria and J. S. Purcell. Treewidth, crushing, and hyperbolic volume. *ArXiv e-prints*, May 2018, arXiv:1805.02357.

[16] Dániel Marx and Michał Pilipczuk. Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. In *Algorithms—ESA 2015*, volume 9294 of *Lecture Notes in Comput. Sci.*, pages 865–877. Springer, Heidelberg, 2015.

[17] Yair N. Minsky, Yoav Moriah, and Saul Schleimer. High distance knots. *Algebr. Geom. Topol.*, 7:1471–1483, 2007.

[18] Dale Rolfsen. *Knots and links.* Publish or Perish Inc., Houston, TX, 1990. Corrected revision of the 1976 original.

[19] Martin Scharlemann. Proximity in the curve complex: boundary reduction and bicompressible surfaces. *Pacific J. Math.*, 228(2):325–348, 2006.

[20] Martin Scharlemann, Jennifer Schultens, and Toshio Saito. *Lecture notes on generalized Heegaard splittings.* World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2016. Three lectures on low-dimensional topology in Kyoto.

[21] Martin Scharlemann and Abigail Thompson. Thin position for 3-manifolds. In *Geometric topology (Haifa, 1992)*, volume 164 of *Contemp. Math.*, pages 231–238. Amer. Math. Soc., Providence, RI, 1994.

[22] Horst Schubert. Über eine numerische Knoteninvariante. *Math. Z.*, 61:245–288, 1954.

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

[24] Maggy Tomova. Multiple bridge surfaces restrict knot distance. *Algebr. Geom. Topol.*, 7:957–1006, 2007.

[25] Maggy Tomova. Distance of Heegaard splittings of knot complements. *Pacific J. Math.*, 236(1):119–138, 2008.

[26] Chichen M. Tsau. Incompressible surfaces in the knot manifolds of torus knots. *Topology*, 33(1):197–201, 1994.

JoCG 10(1), 164–180, 2019   180