

Small knots of large Heegaard genus

WILLIAM WORDEN

Building off ideas developed by Agol, we construct a family of hyperbolic knots K_n whose complements contain no closed incompressible surfaces and have Heegaard genus exactly n . These are the first known examples of such knots. Using work of Futer and Purcell, we are able to bound the crossing number for each K_n in terms of n .

1. Introduction

Closed incompressible surfaces in irreducible 3-manifolds have been a subject of great interest since the notion was introduced by Haken in [12]. Most notably, Thurston in [28] proved his celebrated geometrization conjecture for the case of closed irreducible 3-manifolds containing an incompressible surface (i.e., Haken manifolds) 20 years before the theorem was proved in full generality by Perelman [22–24]. This piece of history suggests a need to better understand irreducible 3-manifolds containing *no* closed incompressible surfaces (i.e., small manifolds). Answering a question of Reid, Agol in [1] constructed the first examples of small link complements having arbitrarily many components. By Dehn filling such a link, Agol was able to give the first examples of small closed manifolds having large Heegaard genus. Previous to Agol’s result, many examples of small manifolds had been constructed (c.f. [8], [6], [21], [7], [13], [18], [19]), but the largest known Heegaard genus of a small manifold was 3 (in particular, by [5], this is true of many of the punctured torus bundles of [6, 8]).

The small links Agol constructs are n -braids having n components. Such braids are easily seen to have Heegaard genus at least $\frac{n}{2}$, and work of Rieck [25] (which generalized work of Moriah and Rubinstein [20]) implies that one can find filling slopes such that the Heegaard genus does not decrease upon Dehn filling. By work of Hatcher [14], the filling slopes can also be chosen so that the filled manifold remains small. While Agol’s methods can easily be extended to construct small knots which are n -braids, one immediately loses the ability to control Heegaard genus. Nonetheless, Agol in [1] suggests

that it should be possible to generalize his construction to exhibit examples of small knots having large Heegaard genus.

Although demonstrating the existence of small knots with large Heegaard genus is our main goal, we are actually able to prove the following stronger result:

Theorem 1.1. For every $l \geq 1$, there exists a family $\{K_N\}_{N=4}^\infty$ of small l -component links such that $\mathbb{S}^3 \setminus K_N$ has Heegaard genus $n := Nl$. Furthermore, K_N is an n -bridge link, and can be constructed to have fewer than $4\pi n^5$ crossings.

Thus for any l , we can build small l -component links, with Heegaard genus as large as desired. We immediately get the following corollary:

Corollary 1.2. There exists a family $\{K_n\}_{n=4}^\infty$ of small knots such that $\mathbb{S}^3 \setminus K_n$ has Heegaard genus n . Furthermore, K_n is an n -bridge knot, and can be constructed to have fewer than $4\pi n^5$ crossings.

Proving Theorem 1.1 relies fundamentally on Agol's work in [1]. The main novelty of our work lies in our approach to controlling Heegaard genus, for which our main tool is a theorem of Rieck and Sedgwick [26]. One advantage of our approach is that it allows us to determine the Heegaard genus on the nose, in contrast to Agol's bound of $\frac{n}{2}$. By Dehn filling each $\mathbb{S}^3 \setminus K_N$ along an appropriate slope system, one can obtain small closed 3-manifolds having Heegaard genus exactly n .

The bound on crossing number given in Theorem 1.1 is calculated using recent work of Futer and Purcell, which effectivizes the main results of [20] and [26].

Remark. There is an alternate approach to proving Theorem 1.1, suggested to us by Dave Futer. After using Agol's construction to obtain small l -component links, one could then appeal to work of either Biringer-Souto [3] or Bachman-Schleimer [2]. In particular, the link K_N is built by Dehn filling a braid augmented by loops, each of which encircles some of its strands (see Figure 4). Dehn filling these loops gives a fibered manifold with fiber a punctured disk, which is an l -component braid in a solid torus (an $(l+1)$ -component link). Long Dehn fillings can be shown to produce a monodromy for this fibered manifold which has large translation distance in the curve complex, and either of [3] or [2] then implies that the Heegaard genus is $n+1$ (the genus is subsequently reduced to n when the extra component is filled). This approach comes with a caveat, though. Both results [3] and [2] apply,

as written, only to closed hyperbolic 3-manifolds. Thus one would first have to generalize one of these results to cusped hyperbolic 3-manifolds. As our proof is more direct, and more elementary, we feel it is the right approach.

The paper is organized as follows. In Section 2, we cover some needed background on pants decompositions, essential surfaces, and Heegaard splittings. In Section 3 we describe Agol's construction and state a version of the main lemma in Agol's paper. In Section 4 we describe our particular application of Agol's construction to obtain a family of links which we will Dehn fill to get the desired links K_N . Finally, in Section 5 we prove Theorem 1.1. For the sake of clarity, we push most of the proof of Theorem 1.1 into a series of lemmata, before finishing up the proof at the end of the section.

1.1. Acknowledgements

We thank Alan Reid for bringing to our attention Agol's work, and for many helpful conversations throughout the course of this project. We also thank Dave Futer for showing us the alternate approach outlined in the preceding remark, and for other helpful conversations. Finally, we thank the referees for many helpful suggestions that improved the exposition.

2. Preliminaries

Throughout the paper, all surfaces and 3-manifolds will be assumed to be orientable, and all isotopies will assumed to be smooth.

2.1. Pants Decompositions

Given a connected compact surface Σ , we say Σ is of type (g, n) if it is a genus g surface with n boundary components (and no punctures). Let Σ be such a surface, with $\chi(\Sigma) < 0$. We define a **pants decomposition** P of Σ to be a choice of disjoint, simple closed curves $\alpha_1, \dots, \alpha_m$ on Σ that cut Σ into a disjoint union of surfaces of type $(0, 3)$, called **pants**. Our definitions in this section follow those of [15], though we note that pants decompositions appear in the literature much earlier [16]. Any pants decomposition of a surface of type (g, n) consists of $3g - 3 + n$ curves, and has $|\chi(\Sigma)|$ pants. Let α be a curve in a pants decomposition P for Σ . If we cut Σ along all curves *except* α , then the connected component Σ_0 containing α will either be a surface of type $(1, 1)$ (that is, a one-holed torus), or a surface of type $(0, 4)$ (a four-holed sphere). In the first case, let $\alpha' \subset \Sigma_0$ be a simple closed

curve that transversely intersects α once. If we replace α by α' , we get a new pants decomposition P' . In this case we say that P' is obtained from P via a **simple move**, or *S*-move (see Figure 1a). If, on the other hand, Σ_0 is a four-holed sphere, and α' is a curve that transversely intersects α twice (and cannot be made disjoint from α by isotopy), then replacing α with α' again gives a new pants decomposition P' , which we say is obtained by an **associative move**, or *A*-move, on P (see Figure 1b).




Figure 1. Left: An *S*-move. Right: An *A*-move.

2.2. Essential surfaces

Let M be a compact irreducible manifold, and Σ a surface properly embedded in M . A **compression disk** for Σ is an embedded disk D in M with boundary $\partial D = D \cap \Sigma$, such that ∂D does not bound a disk in Σ . If Σ has no compression disks, then it is **incompressible**. Otherwise it is **compressible**. A disk D with $\text{int}(D) \cap \Sigma = \emptyset$ is a **boundary compression disk** if ∂D is a union of two simple arcs α and β with $\alpha \subset \Sigma$ and $\beta \subset \partial M$, such that α is not isotopic into $\partial \Sigma \subset \Sigma$. If Σ has no boundary compression disks, then it is **boundary incompressible**. If Σ is incompressible and boundary incompressible, and is not isotopic into ∂M , then it is **essential**.

Theorem 2.1. [14] Let M be an orientable, compact, irreducible 3-manifold with ∂M a union of n tori. Then the projective classes of curve systems in ∂M which bound incompressible, ∂ -incompressible surfaces in M form a dense subset of a finite (projective) polyhedron in $\mathcal{PL}(\partial M) \cong \mathbb{S}^{2n-1}$ of dimension $\leq n - 1$.

The above result of Hatcher relies fundamentally on a result of Floyd and Oertel [9], which shows that any such surface is carried by finitely many branched surfaces in M . The slopes that must be avoided in Theorem 2.1 are precisely those that are carried by the boundaries of the branched surfaces given by Floyd–Oertel.

We are interested in the case in which M has $m + 1$ boundary components, m of which are filled along slopes of the form $\frac{1}{s}$, $0 \neq s \in \mathbb{Z}$. Since we will only be concerned with surfaces that do not meet the unfilled boundary component of M , the subset that we need to avoid has dimension $\leq m - 1$. The space of such fillings, on the other hand, consists of integer lattice points in an affine subspace of $H_1(\partial M; \mathbb{R}) \cong \mathbb{R}^{2m+2}$ of dimension m , which in $\mathcal{PL}(\partial M)$ accumulates on a subspace of dimension $m - 1$. Thus infinitely many of these projected lattice points are outside *any* subspace of dimension $m - 1$, and we can choose our filling slope system to avoid the polyhedron given by Hatcher's theorem. In fact, more is true: if $\gamma = (\frac{1}{s_1}, \dots, \frac{1}{s_m}, \infty)$ is a curve system in Hatcher polyhedron, then for some i , increasing s_i by 1 will give a curve system outside the polyhedron. This is because there are m such linearly independent s_i to choose from, and the Hatcher polyhedron has dimension $m - 1$.

2.3. Heegaard splittings

The main tool that will allow us to extend Agol's result to knots in \mathbb{S}^3 is a result of Rieck and Sedgwick, which describes how Heegaard surfaces that appear after Dehn filling lead to essential surfaces in the unfilled manifold. To state this theorem we will need some terminology.

Let M be a compact connected manifold with boundary a union of tori. A **Heegaard surface** $H \subset M$ for M is an embedded closed surface that cuts M into two **compression bodies**, thus giving a **Heegaard splitting** for M . For our purposes, a compression body is defined to be a connected, compact 3-manifold obtained by attaching 1-handles to the $T^2 \times \{1\}$ boundaries of a disjoint union of thickened tori $T^2 \times I$, or to a ball (in which case we get a **handlebody**). The **Heegaard genus** of M , which we will denote by $g(M)$, is the minimal genus over all Heegaard surfaces for M .

Now let T be a boundary torus of M , and let α be a Dehn filling slope on T . Denote by M_α the manifold resulting from Dehn filling along α . Let H be a Heegaard surface for M_α , and let ρ be the core of the filling solid torus. Then one of the following holds:

- (1) ρ is isotopic into H and H is a Heegaard surface for M .
- (2) ρ is isotopic into H and H is *not* a Heegaard surface for M , or
- (3) ρ is *not* isotopic into H and H is *not* a Heegaard surface for M .

Our goal is to construct knots having large Heegaard genus, by Dehn filling m components of an $m + 1$ -component link complement, which has

Heegaard genus $\geq \frac{m}{2}$. Therefore we would like to avoid (3) and, when possible, (2). By [25, Corollary 4.2], only finitely many filling slopes need to be excluded to ensure that (3) is avoided, so if we are filling multiple boundary components, we only need to avoid finitely many slopes on each. The main obstacle then is (2). When (2) happens, there is a unique (isotopy class of) curve β on $T \subset \partial M$ such that β is isotopic onto H , regarded as a surface in M . If we cut H along β we get a surface H^* with two boundary components on T of slope β . In this case we will have $g(M_\alpha) = g(M) - 1$, and this will be true for *every* filling slope on T that intersect β once. Such a line of Dehn surgery slopes is called a **destabilization line** for T . Since we will always be filling along a slope of the form $\frac{1}{s}$, our main concern is when the longitude of a boundary torus defines a destabilization line, since in this case every slope of this form results in a reduction in Heegaard genus.

The surface H^* described above is called an **almost Heegaard** surface. It requires only a single stabilization, which increases genus by 1, to become a Heegaard surface for M . When a filling is of type (3), it is called a **bad** filling. For bad fillings, one has no control over how much the genus may drop. Fillings of type (1) and (2) are called **good** fillings, and are further characterized by the following theorem of Rieck and Sedgwick:

Theorem 2.2. [26, Lemma 3.4, Theorem 5.1] Let M be a compact, orientable, acylindrical manifold with boundary a union of tori, and let (M_α, H) be a good filling. Then one of the following holds:

- (1) H is a Heegaard surface for M (perhaps after an isotopy in M_α), and H is boundary compressible.
- (2) the slope of the almost Heegaard surface H^* is the boundary slope of a separating essential surface of genus less than or equal to that of H^* .

The essential surface given by (2) will have exactly two boundary components, though it may be disconnected. In fact, such a surface is obtained by compressing H^* , then throwing away components that do not have boundary, as is immediately clear from the proof given in [26].

3. Agol's construction

Let F be a surface of type (g, n) , and let $\varphi : F \rightarrow F$ be a homeomorphism. Denote by T_φ the mapping torus $(F \times I) / \{(x, 0) \sim (\varphi(x), 1)\}$. Let P be a pants decomposition of F , and let $\varphi(P)$ be the image of P under the homeomorphism φ . Let $C = \{P = P_0, P_1, \dots, P_m = \varphi(P)\}$ be a path in the pants graph from

P to $\varphi(P)$, such that no (isotopy class of) curve α on F is contained in every P_i . Let β_i be the curve of P_i that is replaced when passing from P_i to P_{i+1} , $0 \leq i \leq m - 1$. The requirement that no curve α is contained in every P_i is equivalent to requiring that every curve in P appears as some β_i . For each curve β_i , let $L_i = \beta_i \times \{i/m\}$, and remove a neighborhood $\mathcal{N}(L_i)$ of L_i from T_φ to obtain a compact manifold $M_C = T_\varphi \setminus \bigcup_i \mathcal{N}(L_i)$. Define the **horizontal** boundary of M_C to be those boundary components coming from drilling out neighborhoods $\mathcal{N}(L_i)$, and define the **vertical** boundary to be the boundary components coming from T_φ .

Here is another way to build M_C . First, let T_1 be a torus with one disk removed, and let B_S be obtained by thickening T_1 to get $T_1 \times I$, then retracting the annulus $\partial T_1 \times I$ to its core curve. Similarly, for S_4 a sphere with four disks removed, let B_A be the result of thickening S_4 then retracting the annuli $\partial S_4 \times I$ to their core curves. Begin with the surface F , marked by the curves of the pants decomposition $P = P_0$. Suppose $P_0 \rightarrow P_1$ is an S -move, so that the curve β_0 lies on a subsurface of F which is a one-holed torus with boundary some other curve β' of P_0 . Now glue a copy of B_S to this subsurface along $T_1 \times \{0\}$, so that the pinched $\partial T_1 \times I$ glues to β' , and mark $T_1 \times \{1\}$ with the curve that replaces β_0 . If $P_0 \rightarrow P_1$ is an A -move, the operation is similar, except the subsurface is a four-punctured sphere. Continue in this way for each move $P_i \rightarrow P_{i+1}$. Since every curve in P_0 is eventually replaced, the result is a copy of $F \times I$ with embedded curves $\{\beta_i\}_{i=0}^m$, which can be isotoped so that β_i is on the fiber $F \times \{\frac{i}{m}\}$. The curves on the bottom glue to the curves on the top via φ , and we obtained M_C by drilling out a neighborhood of each curve. This shows that M_C can be decomposed into **A-blocks** and **S-blocks** by cutting along pants. An A -block is defined here to be the complex obtained by removing from B_A neighborhoods of two simple closed curves, one on $S_4 \times \{0\}$ and one on $S_4 \times \{1\}$, which intersect twice as curves on S^4 , and a neighborhood of the pinched $\partial S_4 \times I$ (see Figure 2, bottom left). An S -block is defined similarly. In general, we will call both A -blocks and S -blocks **pants blocks**.

If R is an A -block or S -block of M_C , then we will call $R \cap \text{int}(M_C)$ an **ideal** A - or S -block (or, in general, an ideal pants block). Each ideal pants block can be obtained by gluing either one or two ideal octahedra, as is shown in Figure 2. If these octahedra are given the structure of regular ideal hyperbolic octahedra, then the boundary of each ideal pants block is totally geodesic, and gluing them along their geodesic faces gives $\text{int}(M_C)$ a complete, finite volume, hyperbolic structure, as was observed by Agol [1].




Figure 2. An A -block (lower left) is obtained by doubling a vertex-truncated octahedron (lower right) along some of its faces. Alternatively, an ideal A -block is obtained by doubling an ideal octahedron along some of its faces. Since a truncated octahedron can be identified with a tetrahedron minus a neighborhood of its one-skeleton, an A -block can also be viewed as a doubling of such a tetrahedron.

3.1. Tubing pants

In this section, M_f will be a manifold obtained by Dehn filling any number of the horizontal boundary components of M_C (possibly none). Boundary components of M_C have a natural longitude and meridian, which induce a framing on the boundary components of M_f . We can also refer to horizontal and vertical boundary components of M_f , according to whether the corresponding boundary component of M_C is horizontal or vertical. We will say a surface $\Sigma \subset M_f$ has **horizontal** boundary if every component of $\partial\Sigma$ is isotopic to either the meridian of a vertical boundary component, or the longitude of a horizontal boundary component.

An annulus $A \subset M_f$ with interior disjoint from Σ is called a **compression annulus** for Σ if it has one boundary component ∂A^+ in $\text{int}(\Sigma)$ and the other boundary component ∂A^- on ∂M_f is horizontal, and cannot be isotoped into Σ via an isotopy through which $\partial A^- \subset \partial M_f$. If A is a compression annulus for Σ then we can **compress along** A by attaching $(A \times I)$ to Σ , then removing $A \times \text{int}(I)$. When a (possibly disconnected) surface $(\Sigma, \partial\Sigma) \subset (M_f, \partial M_f)$ has two horizontal boundary components α_1 and α_2 on a boundary torus $T \subset \partial M_f$, we can **tube** Σ along T by attaching an annulus along the α_i and pulling it away from ∂M_f . This introduces a compression annulus, along which we can **untube** Σ by compressing the annulus (see Figure 3a). Note that tubing involves a choice of annulus, but we will

see later that for our applications both choices will result in the same surface up to isotopy.




Figure 3. Left: compressing Σ along A (i.e., un tubing) by cutting along the dotted intersection produces two new boundary components. Conversely, tubing along a boundary torus produces a compression annulus. Right: two normal quads Q_1 and Q_2 glue together to give a twice punctured annulus A_2 in an A -block. In this case A_2 has a compression annulus, shown here as a red dome.

The following lemma shows that compressing along annuli preserves incompressibility for M_f with hyperbolic interior.

Lemma 3.1. Let M_f be a filling of M_C with hyperbolic interior, and let Σ be an incompressible surface. Then the surface Σ' resulting from compressing Σ along a compression annulus A is also incompressible.

Proof. Let ∂A^+ be the boundary component of A lying on Σ . Suppose there is a (non-boundary) compressing disc D for Σ' . We may assume that ∂D is bounded away from $\partial \Sigma'$, so that $\partial D \cap \partial A^+ = \emptyset$. Now suppose that ∂D bounds a disk D_Σ in Σ . Since ∂A^+ is disjoint from ∂D , we must have that ∂A^+ is contained in $\text{int}(D_\Sigma)$ or $\partial A^+ \cap D_\Sigma = \emptyset$. In the first case we must have ∂D isotopic to ∂A^+ and hence ∂D is isotopic through A to an essential curve in ∂M_f . Since M_f has hyperbolic interior this is impossible [28]. Therefore we are in the second case, so when we compress the annulus A to get Σ' , D_Σ is preserved as a disk in Σ' , contradicting our assumption that D was a compression disk. Thus if D is a compression disk in Σ' , it is also a compression disk in Σ , but this is impossible since Σ was assumed to be incompressible. \square

The following lemma is due to Agol [1]. Since our statement of this lemma is somewhat different than in [1], and since it is of central importance to the proof of Theorem 1.1, we reproduce Agol’s proof below.

Lemma 3.2 ([1]). Let M_f be a filling of M_C with hyperbolic interior. Then every incompressible surface in M_f that is isotopic into M_C and has horizontal boundary is obtained by tubing pants in M_C .

Proof. Let $(\Sigma, \partial\Sigma) \subset (M_f, \partial M_f)$ be an incompressible surface in M_f that is isotopic into M_C and has horizontal boundary. First isotope Σ into $M_C \subset M_f$. With this done, we will view Σ from now on as a subset of M_C . We will untube Σ along compression annuli in M_C , and show that the resulting surface is a disjoint union of pants.

Since components of $\partial\Sigma$ must be parallel on ∂M_C to pants boundaries, we can isotope Σ away from the pants in a neighborhood of $\partial\Sigma$. It follows that $\Sigma \cap \{\text{pants}\}$ is a disjoint union of essential simple closed curves. If α is such a curve, then α is parallel to a boundary component ∂Y of a pair of pants Y , since all closed curves on pants are boundary parallel. Thus α forms one boundary curve of an annulus $A \subset Y$ whose other boundary curve is on ∂Y . Therefore we can untube Σ along A , and the resulting surface will have one fewer curves of intersection with the pants (Figure 3a may be helpful). Doing this surgery at every such α results in a surface Σ' that is disjoint from pants. Since M_C is finite volume and hyperbolic, we may assume that no component of Σ' is an annulus. By Lemma 3.1, Σ' is incompressible in M_C .

Since the pants partition M_C into pants blocks, each component of Σ' is contained in such a block. An ideal pants block is made up of ideal octahedra, so a pants block is made up of vertex-truncated octahedra. A truncated octahedron is the same as a tetrahedron with a neighborhood of its 1-skeleton removed, as shown in Figure 2. With this point of view, we consider the shaded faces (blue in Figure 2) to be the faces of the tetrahedra. Since pants blocks are obtained by gluing the shaded faces of the truncated octahedra, we can see them as manifolds triangulated by either one or two tetrahedra, with a neighborhood of the 1-skeleton of the triangulation removed. That is, a pants block is the exterior of the 1-skeleton of a triangulation of a closed manifold. Since Σ' is incompressible in the pants block, it is normal (after isotopy) in this closed manifold by [27, Claim 1.1], and hence in the truncated tetrahedra it decomposes into (truncated) normal triangles and quads having normal arcs on shaded faces (see [11] or [17] for background on normal surface theory). Since an A -block is obtained by doubling a truncated tetrahedron along its faces, normal triangles glue to other triangles to give pants, which are isotopic to boundary pants. For the same reason, quads glue to other quads. The surfaces coming from pairs of quads are either, (1) isotopic to $S_4 \times \{1/2\}$, and so are a tubing of two 3-punctured spheres

isotopic into the boundary of the A -block, or (2) a twice punctured annulus A_2 with boundary the two curves on the S_4 boundaries of the A -block. In the latter case one can find a compression annulus, as shown in Figure 3b. For S -blocks there is one quad that is isotopic to $T_1 \times \{1/2\}$, and this is the only quad that can glue up coherently as part of a properly embedded surface. Triangles must glue up in pairs, and result in surfaces that have compression annuli, and untube to copies of $T_1 \times \{1/2\}$. We leave further details for the S -block case to the reader, as these will not play a part going forward. This shows that Σ' is a disjoint union of pants, and that Σ was obtained by tubing together pants. \square

4. Augmented braids

In this section we use Agol’s construction to obtain explicit manifolds M_C , which will be used in Section 5 to build the links of Theorem 1.1. These manifolds are very similar to those used in Agol’s proof, with the main difference being that we choose the monodromy φ so that the link resulting from filling the horizontal boundary has l components. Although we describe the path in the pants graph differently than Agol, ours is essentially the same as his, but traced out (almost) twice.

Fix an integer $l \geq 1$, and let $n \in \{lN \mid N \in \mathbb{Z}_{\geq 1}, lN \geq 4\}$. Let $D_{z,r}$ be the disk of radius r centered at $z \in \mathbb{C}$, and let $D_n = D_{0,2} \setminus \left(\bigcup_{k=0}^{n-1} D_{e^{2\pi i k/n}, \frac{1}{n}} \right)$, where $i = \sqrt{-1}$. That is, D_n is a disk with n small holes, equally spaced around a circle. Let p_k be the boundary component of D_n coming from removing the disk $D_{e^{2\pi i k/n}, \frac{1}{n}}$, and let q be the boundary component of D_n coming from the boundary of $D_{0,2}$. Define a homeomorphism $\varphi : D_n \rightarrow D_n$ by $z \mapsto e^{-2\pi i l/n} z$, and let $T_\varphi = (D_n \times I) / (x,0) \sim (\varphi(x),1)$ be the associated mapping torus.

By our choice of φ , it is clear that T_φ has exactly $l + 1$ boundary components, since we chose n to be a multiple of l . Let L_q be the boundary component coming from q , which is fixed by φ , and let $L_p = L_1 \cup \dots \cup L_l$ be the union of boundary components coming from the p_k . We may think of T_φ as the exterior $\mathbb{S}^3 \setminus (L_q \cup L_p)$ of a closed braid in \mathbb{S}^3 , augmented with an additional unknotted component L_q . Alternatively, we can view T_φ as a closed braid in a solid torus, where L_q is the boundary of the solid torus. We will generally take the latter view, and this is reflected in the figures. When we glue $D_n \times \{0\}$ to $\varphi(D_n) \times \{1\}$, we can add a twist by any multiple of 2π without affecting the homeomorphism type of T_φ . For the purpose of exposition, it will be convenient to include an extra twist by 4π in the

counter-clockwise direction. With this convention, for $n = 6$ and $l = 1$, T_φ is the gray braid shown in Figure 4 (ignore for now the blue loops in the figure).

For $1 \leq i \leq 4n$, and $2 \leq j \leq n - 2$, $j \neq i \pmod 2$, define $\beta_{i,j}$ to be the simple closed curves on D_n that bounds a disk punctured by p_k , for $\frac{i-(j-1)}{2} \leq k \leq \frac{i+(j-1)}{2}$. Here we take p_k indices modulo n . For i even (so that j is odd), $\beta_{i,j}$ is the curve that bounds j punctures, centered at the puncture $p_{i/2}$. For i odd (j even), $\beta_{i,j}$ bounds j punctures, centered between the two punctures $p_{\frac{i \pm 1}{2}}$.

For $1 \leq i \leq 2n$, Let P_i be the pants decomposition of D_n defined by the collection of curves

$$\{\beta_{2i-1,2}, \beta_{2i-1,4}, \dots, \beta_{2i-1,n-2}, \beta_{2i,3}, \beta_{2i,5}, \dots, \beta_{2i,n-1}\}.$$

For the pants decomposition P_i and its curves, indices are *not* taken modulo n . This means that $\beta_{i,j}$ and $\beta_{i+n,j}$ describe the same curve on D_n , though we considered these to be distinct copies of this curve. Similarly, P_i and P_{i+n} define the same pants decomposition of D_n , but will be regarded as distinct copies of it.

Figure 5a shows P_1 for the case $n = 6$; we get any other P_i by rotating the curves of P_1 by an angle of $\frac{2\pi i}{n}$. Note that P_i contains $n - 2$ simple closed curves. For $0 \leq k \leq n - 2$, let P_i^k be the pants decompositions consisting of the first k curves of P_{i+1} union the last $(n - 2) - k$ curves of P_i . For $n = 6$, the sequence of pants decompositions

$$P_1 = P_1^0, P_1^1, \dots, P_1^4 = P_2$$

is shown in Figure 5. In Figure 5b we show a simplified picture to save space—only the subdisk of D_n cut out by the convex hull of the points p_k is shown, and the $\beta_{i,j}$ curves are represented by transversely oriented arcs.




Figure 4. T_φ , with loops $L_{i,j}$.




Figure 5. The sequence of pants decompositions from P_1 to P_2 , for $n = 6$. The drawing of P_1 on the left shows how to interpret the simplified drawings of P_1^k , $1 \leq k \leq 4$, on the right.

The transverse orientation shows how to complete the arc to a curve—it should be an inward normal for the curve, and the rest of the curve should lie outside the hexagon.

The sequence of pants decompositions

$$P_i = P_i^0, P_i^1, \dots, P_i^{n-2} = P_{i+1}$$

is a path C_i in the pants graph from P_i to P_{i+1} , of length $n - 2$. The path C_i is obtained by shifting each of the curves of P_i , one at a time and in order, by a rotation of $2\pi/n$. Let C be the composition of the paths C_i for $1 \leq i \leq 2n + 1 - l$. Note that the composition of paths C_1, \dots, C_n is a loop of length $n(n - 2)$ that begins at P_1 and ends at P_{n+1} , and the composition C_{n+1}, \dots, C_{2n} is a path of length $(n - l)(n - 2)$ beginning at P_{n+1} and ending at P_{2n+1-l} . Thus C is a path from P_1 to P_{2n+1-l} of length $m := (2n - l)(n - 2)$. Since P_1 and P_{n+1} are two copies of the same pants decomposition, we see that the path C results from traversing a loop in the pants graph, then partially traversing the same loop again, by an amount depending on l .

Since $\varphi(P_1) = P_{2n+1-l}$, the construction of Section 3 can be applied to the path C , giving a compact manifold M_C . In particular, at each step of the path C , a curve $\beta_{i,j}$ is replaced by another curve $\beta_{i',j'}$. If $\beta_{i,j}$ is replaced at step k , then let $L_{i,j} = \beta_{i,j} \times \{\frac{k}{m}\}$ for $1 \leq k \leq m$. Then $M_C = T_\varphi \setminus \bigcup_{i,j} \mathcal{N}(L_{i,j})$. We will denote by $B_{i,j}$ the boundary component of M_C coming from removing $\mathcal{N}(L_{i,j})$ from T_φ . Actually, it will be convenient for the exposition to come to modify the $L_{i,j}$ by an isotopy in T_φ . In particular, we arrange so that $L_{i,j}$ lies in the fiber $D_n \times \{\frac{i}{2(2n+1-l)}\}$, and do the same for

the $B_{i,j}$ boundary components of M_C . This isotopy only moves curves within a pants decomposition P_k past each other, and such an isotopy exists since curves in P_k are disjoint in D_n . The remaining boundary components of M_C , corresponding to the augmentation circle L_q and the braid components L_p of L , will be denoted by B_q and B_p , respectively. Note that B_p is potentially a union of multiple boundary components, but we will rarely need to refer to them individually. Figure 4 shows T_φ with the $L_{i,j}$ loops colored blue (alternatively, we can view these loops as representing boundary components $B_{i,j}$ in M_C). Note that a loop $L_{i,j}$ encircles j strands of the braid L_p , and is at height i relative to the other loops. We say that $B_{i,j}$ has **width** j . We will take the width of B_q to be n , and the width of each component of B_p to be 1.

Figure 6a shows the pairs of pants along which M_C decomposes into A -blocks, for part of M_C (note that M_C has no S -blocks). In this drawing, as in Figure 4, we think of B_q as the boundary of the solid torus containing M_C . The pants meeting B_q are the outermost ones (in blue). Each of the pants is either a twice punctured disk or a punctured annulus, and the same pattern of pants continues in the rest of M_C . Here we are using the term *puncture* somewhat non-standardly, to refer to a boundary component that meets a meridian of a component in B_p .

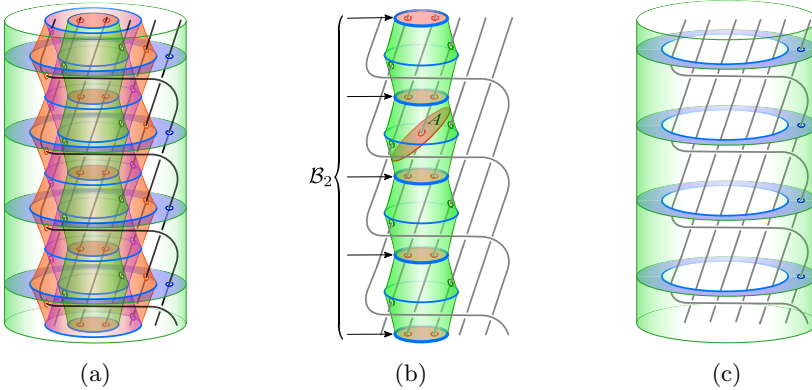


Figure 6. Left: M_C decomposes into A -blocks along pants. With curves in \mathcal{B} omitted, one can also interpret this as the complex \mathcal{C} . Center: the subcomplex $\mathcal{C}_2 \subset \mathcal{C}$, with a compression annulus A . Right: the subcomplex $\mathcal{C}_n \subset \mathcal{C}$. In the center and right picture, we show B_p but omit filled tori that don't meet the subcomplex, to avoid cluttering.

For each pair of pants Y in M_C , let the **size** of Y be the largest width of a horizontal boundary component that Y meets. Thus in Figure 6a the

size 2 pants are the twice-punctured disks, for example. Note that size $k \geq 3$ pants always have another boundary component on a width $k - 1$ loop.

5. Small knots of large genus

In this section we prove Theorem 1.1. For the sake of clarity, we have relegated the bulk of this work to Lemmas 5.1, 5.2, and 5.3. In particular, in Lemma 5.1 we get much of the technical work out of the way, and in Lemmas 5.2 and 5.3 we then address smallness and Heegaard genus. At the end of the section we assemble the pieces to prove the theorem, which we restate below for convenience:

Theorem 1.1. *For every $l \geq 1$, there exists a family $\{K_N\}_{N=4}^\infty$ of small l -component links such that $\mathbb{S}^3 \setminus K_N$ has Heegaard genus $n := Nl$. Furthermore, K_N is an n -bridge link, and can be constructed to have fewer than $4\pi n^5$ crossings.*

The links we are after will result from long Dehn fillings along all $B_{i,j}$ boundary components of M_C , and along the component B_q . Let $\gamma = \{\gamma_q\} \cup \{\gamma_{i,j}\}_{i,j}$ be a system of filling slopes for these boundary components. Predictably, the slope for $B_{i,j}$ is $\gamma_{i,j}$, and the slope for B_q is γ_q . We will require that every slope in the system γ has the form $1/s$ for some $s \in \mathbb{Z} \setminus \{0\}$. We may choose γ so that we avoid boundary slopes given by Theorem 2.1, and so that we avoid bad fillings as described in Subsection 2.3. We may also take all slopes to be sufficiently long so that exceptional (i.e., non-hyperbolic) fillings are avoided. For the sake of convenience, we will refer to such a choice of γ as a **generic** filling system. When we have fixed such a γ , we will denote by M_γ the Dehn filling of M_C along γ .

Let $\mathcal{B}_2 \subset \partial M_C$ consist of all $B_{i,j}$ of width 2, and fix some $\mathcal{B} \subset \mathcal{B}_2$. For such \mathcal{B} , define $L_{\mathcal{B}} = \{L_{i,j} \mid B_{i,j} \in \mathcal{B}\}$, and let $M_{\gamma \setminus \mathcal{B}} = M_\gamma \setminus \bigcup_{L_{i,j} \in L_{\mathcal{B}}} \mathcal{N}(L_{i,j})$. In other words, $M_{\gamma \setminus \mathcal{B}}$ is the manifold obtained by drilling out from M_γ the surgery solid tori for $B_{i,j} \in \mathcal{B}$. Let $\mathcal{B}^c = \partial M_C \setminus (\mathcal{B} \cup \{L_p\})$, viewed as a subset of $M_{\gamma \setminus \mathcal{B}} \supset M_C$. That is, \mathcal{B}^c consists of all the $B_{i,j}$ that are filled in $M_{\gamma \setminus \mathcal{B}}$.

Let $\mathcal{C} \subset M_{\gamma \setminus \mathcal{B}}$ be the 2-complex consisting of the pants in $M_C \subset M_{\gamma \setminus \mathcal{B}}$, and the tori $B_{i,j} \in \mathcal{B}^c \subset M_{\gamma \setminus \mathcal{B}}$ (see Figure 6a). \mathcal{C} carries all surfaces in $M_{\gamma \setminus \mathcal{B}}$ with horizontal boundary that are obtained by tubing pants. Next, define \mathcal{C}_2 to be the subcomplex of \mathcal{C} consisting of all tori $B_{i,j} \in \mathcal{B}^c$ of width 2 and 3, and all pants of size 2 and 3. Define $\mathcal{C}_n \subset \mathcal{C}$ to be the subcomplex consisting

of B_q , all tori $B_{i,j}$ of width $n - 1$, and all pants of size n . The subcomplexes \mathcal{C}_2 and \mathcal{C}_n are shown in Figure ??, respectively.

Lemma 5.1. Fix any generic filling system γ for M_C . Choose $\mathcal{B} \subset \mathcal{B}_2$, and define $M_{\gamma \setminus \mathcal{B}}$ as above. Let $(\Sigma, \partial\Sigma) \subset (M_{\gamma \setminus \mathcal{B}}, \partial M_{\gamma \setminus \mathcal{B}})$ be an essential surface with horizontal boundary, and let Σ_\circ be the surface obtained by compressing Σ along all compression annuli. Then

- (1) Σ is obtained by tubing together pants in M_C ,
- (2) if $\mathcal{B} = \emptyset$ and no pants of Σ is a twice-punctured disk, then Σ_\circ is carried by \mathcal{C}_n ,
- (3) if no pants of Σ has boundary on B_q , then Σ_\circ is carried by \mathcal{C}_2 .

Proof. Let Σ be an essential surface in $M_{\gamma \setminus \mathcal{B}}$. First we show that Σ is isotopic into M_C . Let $B \in \mathcal{B}^c$, and let $\tau(B)$ be the filling solid torus for B . Consider a component α of the intersection of Σ with B , which is necessarily a simple closed curve. First, if α is null-homotopic in B , then since Σ is incompressible and $M_{\gamma \setminus \mathcal{B}}$ is irreducible, the disk α bounds in B must bound a ball in $M_{\gamma \setminus \mathcal{B}}$ with another disk in Σ (we may assume that α is innermost in Σ). In this case we can isotope Σ to remove the intersection α . If α is essential in B , then it cannot bound a disk, since γ was chosen to avoid boundary slopes. In this case there must be another intersection curve α' isotopic to α in B , such that α and α' bound an annulus in $\Sigma \cap \tau(B)$. This annulus cuts off a solid torus with an annulus in B , and can be isotoped out of $\tau(B)$ across this solid torus, thus removing the intersections α and α' (here we assume the annulus in $\Sigma \cap \tau(B)$ is outermost in $\tau(B)$). It follows that Σ is isotopic into M_C , and hence by Lemma 3.2 it is obtained by tubing pants in M_C , thus establishing (1). Since Σ_\circ is obtained by compressing Σ along compression annuli, by Lemma 3.1 it is incompressible. Let Y_1, \dots, Y_k be the pants that tube together to give Σ_\circ (and therefore Σ).

A key observation is that when two pants Y_1 and Y_2 are tubed together along some $B \in \mathcal{B}^c$, the attached annulus can lie to either side of the filling torus, and can be isotoped across it. This is because the filling slope is of the form $\frac{1}{s}$, and therefore intersects each boundary component of the attached annulus once. More concretely, the solid torus is fibered as $D^2 \times \mathbb{S}^1$, and each fiber intersects the annulus in an arc, which can be isotoped across the disk. This shows how to build an isotopy of the annulus.

Each $B_{i,j}$ of width at least 3 and at most $n - 2$ meets exactly 4 pairs of pants in \mathcal{C} . Of these 4, let Y_- and Y'_- be the two of size j , and let Y_+ and Y'_+ be the two of size $j + 1$. If $B_{i,j} \in \mathcal{B}^c$, let $A_{i,j}^-$ be the twice-punctured

annulus resulting from tubing Y_- and Y'_- along $B_{i,j}$, and let $A_{i,j}^+$ be the result of tubing Y_+ and Y'_+ along $B_{i,j}$ (we will assume that the attaching annuli stay on $B_{i,j}$). Note that by the observation in the above paragraph, the choice of tubing of Y_{\pm} and Y'_{\pm} does not matter, since the attaching annulus can be isotoped across the surgery solid torus. If $B_{i,j} \in \mathcal{B}$, then we define $A_{i,j}^- := Y_- \cup Y'_-$ and $A_{i,j}^+ := Y_+ \cup Y'_+$. For $B_{i,j}$ of width $n - 1$, $B_{i,j}$ only meets 3 pairs of pants. In this case we can define $A_{i,j}^-$ in the same way as above, and identify $A_{i,j}^+$ with the pair of pants meeting $B_{i,j}$ and B_q . For $B_{i,j}$ of width 2, $B_{i,j}$ again meets only 3 pairs of pants, and we can define $A_{i,j}^+$ as above, and identify $A_{i,j}^-$ with the size 2 pants (i.e., the twice punctured disk) meeting $B_{i,j}$.

It follows that if $B_{i,j}, B_{i,j+2} \in \mathcal{B}^c$ then $A_{i,j}^+$ is isotopic to $A_{i,j+2}^-$, by an isotopy across the A -block that they bound (see Figure 7). If, on the other hand, $B_{i,j} \in \mathcal{B}^c$ and $B_{i,j+2} \in \mathcal{B}$, then the same isotopy will take $A_{i,j}^+$ into $A_{i,j+2}^- \cup B_{i,j+2}$, which shows that $A_{i,j}^+$ has a compression annulus. Compressing this annulus results in a pair of once-punctured annuli isotopic to $A_{i,j+2}^-$. Similarly, if $B_{i,j+2} \in \mathcal{B}^c$ and $B_{i,j} \in \mathcal{B}$ then $A_{i,j+2}^-$ has a compression annulus, and after an annular compression is isotopic to $A_{i,j}^+$.




Figure 7. If $B_{i,j} \in \mathcal{B}^c$, then $A_{i,j}^+$ can be isotoped onto $A_{i,j+2}^-$, possibly after an annular compression. Similarly, if $B_{i,j+2} \in \mathcal{B}^c$ then $A_{i,j+2}^-$ can be isotoped onto $A_{i,j}^+$, possibly after a compression.

We are now ready to prove (2) and (3). First, assume that $\mathcal{B} = \emptyset$ and none of the pants Y_1, \dots, Y_k are twice-punctured disks. Our goal will be to reduce \mathcal{C} to the subcomplex \mathcal{C}_n , while ensuring that \mathcal{C}_n also carries Σ_{\circ} . First, since it is assumed that none of the Y_i are twice-punctured disks, we may remove these from \mathcal{C} , and the resulting complex will still carry Σ_{\circ} . With the twice-punctured disks removed, each $A_{i,2}^+$ for $1 \leq i \leq 2n - l$ can be isotoped into $A_{i,4}^-$. The key here is that by removing the twice-punctured disks, we guarantee that the isotopy of $A_{i,2}^+$ extends to an isotopy of \mathcal{C} . In particular, since $B_{i,2} \in \mathcal{B}^c$, $B_{i,2} \setminus A_{i,2}^+$ is an annulus which we can isotope across the

surgery solid torus and into $A_{i,2}^+$, thus allowing us to isotope $A_{i,2}^+$ across its A -block and into $A_{i,4}^-$. Thus we get a new complex that does not contain any $A_{i,2}^+$ or $B_{i,i+2}$. We continue in this way for all $A_{i,k}^+$, $k \leq n - 2$. Note that we must do this in order (with respect to k), since $A_{i,j}^+$ must be removed before $A_{i,j+1}^+$ can be removed. With the above reduction complete, only the subcomplex \mathcal{C}_n remains, and Σ_\circ is still carried since we only changed \mathcal{C} by removing unnecessary pants, and subcomplexes isotopic into \mathcal{C} , or isotopic into \mathcal{C} after an annular compression.

(3) is proved similarly. If no pants of Σ has boundary on B_q , then all pants meeting B_q can be removed from \mathcal{C} , and \mathcal{C} will still carry Σ_\circ . With these pants removed, we can isotope punctured annuli $A_{i,n-1}^-$ down to $A_{i,n-3}^+$, and continue in this way until we are only left with tori of width 2 and 3, and annuli $A_{i,2}^+$, which is exactly the subcomplex \mathcal{C}_2 . Note that since $\mathcal{B} \subset \mathcal{B}_2$, only tori of width 2 can be drilled out, so all of these isotopies are possible (or possibly compressions then isotopies in the case of $A_{i,4}^-$). \square

Lemma 5.2. For a generic filling system γ , M_γ contains no closed incompressible surfaces.

Proof. Suppose Σ is a closed incompressible surface in M_γ . Applying Lemma 5.1 with $\mathcal{B} = \emptyset$, it follows that Σ is isotopic into M_C and obtained by tubing pants. Let Σ_\circ be the surface obtained by compressing Σ along all the compression annuli meeting B_p . This surface is obviously a tubing of pants isotopic into M_C , and by Lemma 3.1 it is also incompressible.

Assume first that one of these pants is a twice-punctured disk. Then Σ_\circ must be a punctured sphere since all pants are either punctured disks or punctured annuli, and they are not tubed along punctures. Since Σ_\circ is isotopic into M_C , it is also isotopic into $M_{\gamma \setminus \{B_q\}}$, which is a fibered manifold with fiber D_n (in particular, it is an augmented braid closure). Although Σ_\circ may intersect every fiber of $M_{\gamma \setminus \{B_q\}}$, there is some finite cyclic cover $\widehat{M}_{\gamma \setminus \{B_q\}}$ in which some fiber F is disjoint from the lifts of Σ_\circ . Since the only incompressible surface in the complement of a fiber is a fiber, it follows that the lifts of Σ_\circ are fibers and hence so was Σ_\circ . But this is impossible since Σ_\circ only has boundary on B_p , and B_q meets every fiber. Thus none of the pants is a twice-punctured disk.

To finish the proof, we will return to viewing Σ_\circ in M_γ . Applying Lemma 5.1 again with $\mathcal{B} = \emptyset$, Σ_\circ must be carried by the complex \mathcal{C}_n since no pants is a twice-punctured disk. Any surface carried by \mathcal{C}_n is obtained by tubing size n pants along B_q and $B_{i,n-1}$. A typical such surface is shown in Figure 8. This surface has an important property, which is shared by any




Figure 8. A surface carried by the subcomplex C_n . Dotted curves show where to compress along annuli meeting B_q and width $n - 1$ loops to decompose into pants in M_C

surface obtained in this way. Namely, it has two annular subsurfaces, one that is isotopic to an annulus on B_q parallel to its meridian, and one that is isotopic to an annulus on B_q parallel to its longitude. It follows that there is an annulus A having one boundary component on Σ_\circ and the other on B_q , such that the boundary on B_q has slope γ_q (the filling slope for B_q). Since γ_q bounds a disk in the surgery solid torus filling B_q , the boundary of A on Σ_\circ bounds a disk in M_γ , and clearly does not bound a disk on Σ_\circ . Thus Σ_\circ is compressible, contradicting our assumption. \square

Lemma 5.3. For a generic filling system γ and $n \geq 4$, M_γ has Heegaard genus $g_\gamma \geq n$.

Proof. We will assume for a contradiction that $g_\gamma \leq n - 1$. Since γ is a generic filling system it contains no bad filling slopes. Consequently, each filling of a boundary component of M_C by a solid torus is good and thus reduces the Heegaard genus by at most one (see Subsection 2.3). Thus, if we drill out some of these solid tori in M_γ , then we increase the Heegaard genus by at most one for each drilling.

In particular, we will drill out the filling solid tori for components $B_{i,2}$, starting with $B_{0,2}$, until the Heegaard genus increases as a result of drilling. This increase in Heegaard genus must happen for some filling, as we now show. Once again, we will let $\mathcal{B} \subset \mathcal{B}_2$ be the set of tori whose surgery solid torus has been drilled out, and define $M_{\gamma \setminus \mathcal{B}}$ as before. Since $M_{\gamma \setminus \mathcal{B}}$ has l boundary component coming from M_γ , if we drill out all $2n - l$ such solid tori then it will have $2n$ boundary components, so its Heegaard genus will be at least $n = 2n/2$. Thus the Heegaard genus of $M_{\gamma \setminus \mathcal{B}}$ must increase for at least one of these drillings. Let $B_{i_0,2}$ be the last torus drilled out, the drilling of which increases the Heegaard genus.

Now, let H be a minimal genus Heegaard surface for the manifold $M'_{\gamma \setminus \mathcal{B}}$ resulting from re-filling $B_{i_0,2}$. By our choice of \mathcal{B} , $\text{genus}(H) = g_\gamma \leq n - 1$. Since the Heegaard genus of $M_{\gamma \setminus \mathcal{B}}$ is strictly larger than that of $M'_{\gamma \setminus \mathcal{B}}$, H is not a Heegaard surface for $M_{\gamma \setminus \mathcal{B}}$. Since γ was chosen to avoid bad fillings, the core of the filling solid torus of $B_{i_0,2}$ is isotopic into H , and by cutting along it we get an almost Heegaard surface $H^* \subset M_{\gamma \setminus \mathcal{B}}$. The surface H^* has two boundary components, and both are isotopic to longitudes of $B_{i_0,2}$. By Theorem 2.2, there is an essential surface Σ in $M_{\gamma \setminus \mathcal{B}}$ such that $\partial \Sigma = \partial H^*$, and $\text{genus}(\Sigma) \leq \text{genus}(H^*)$. As noted following the statement of Theorem 2.2, Σ is obtained by compressing H^* , and may be disconnected. In the case that Σ is disconnected, by $\text{genus}(\Sigma)$ we mean the total genus of both components. Letting $\widehat{\Sigma}$ be the surface obtained by tubing together the boundary components of Σ along $B_{i_0,2}$, it follows that $\text{genus}(\widehat{\Sigma}) \leq \text{genus}(H) \leq n - 1$.

Since Σ has horizontal boundary, Lemma 5.1 implies that it is obtained by tubing pants in M_C . Let Σ_\circ and $\widehat{\Sigma}_\circ$ be the surfaces obtained by compressing Σ and $\widehat{\Sigma}$, respectively, along all compression annuli having a boundary curve on B_p . Note that Σ_\circ and $\widehat{\Sigma}_\circ$ are also tubings of pants in M_C . If a component of $\widehat{\Sigma}_\circ$ contains a pants that is a twice-punctured disk, then by the argument in the proof of Lemma 5.2 that component is a punctured sphere. All other components are punctured tori, since they are made of punctured annuli tubed together only at their boundaries (not their punctures). Let Σ'_\circ be the union of all connected components of Σ_\circ that contain some boundary component of Σ_\circ . Note that by the preceding discussion, Σ'_\circ is either a union of two punctured disks (which tubes to a punctured sphere), or a punctured annulus (which tubes to a punctured torus). Note that Σ_\circ is incompressible by Lemma 3.1, and hence so is Σ'_\circ . Let Y_1, \dots, Y_k be the pants in M_C that tube together to give Σ'_\circ .

Case 1: One of the pants Y_i of Σ'_\circ has boundary on B_q . Since Σ'_\circ has two boundary components on $B_{i_0,2}$, it must contain at least two pants meeting $B_{i_0,2}$. Therefore Σ'_\circ must contain enough pants to tube from $B_{i_0,2}$ up to B_q , then back down to $B_{i_0,2}$. This requires at least 2 pairs of pants of each size j , $3 \leq j \leq n$, for a total of $2(n - 2)$ pairs of pants.

Claim. Σ'_\circ consists of exactly 2 pairs of pants of each size greater than or equal to 3.

Proof of Claim. If Σ'_\circ contains any additional pants, then such pants will introduce at least two additional punctures, for a total of $2n - 2$ punctures (note that there must be an even number of punctures, as they must tube in pairs). If one of the pants in Σ'_\circ is size 2, then it must have two such, so

we get an additional 2 punctures, and total is $2n$, so $\text{genus}(\widehat{\Sigma}) \geq n$. If one of the pants is *not* size 2, then Σ'_\circ is a punctured annulus, whose boundary tubes together to give a punctured torus, so $\text{genus}(\widehat{\Sigma}) \geq \frac{1}{2}(2n - 2) + 1 = n$. Thus we conclude that Σ has genus at least n , contradicting our assertion that $\text{genus}(\widehat{\Sigma}) \leq n - 1$. \square

It immediately follows that each pants of size j for $4 \leq j \leq n - 1$ tubes to pants of size $j - 1$ and $j + 1$. The two size n pants tube to each other, and to the two size $n - 1$ pants, and the two size 3 pants meet the boundary $B_{i_0,2}$ and tube to the two size 4 pants. One possible configurations is shown in Figure 9a. To rule out such a surface, we will need to demonstrate another isotopy of tubed pants. In particular, given a size j pants and a size $j + 1$ pants, with $3 \leq j \leq n - 1$, tubed together along their width j boundary $B_{i,j}$, we have an isotopy which fixes the boundaries of these pants on $B_{i+1,j-1}$ and $B_{i+1,j+1}$, and takes them to two pants which are tubed along $B_{i+2,j}$. This is an isotopy of twice-punctured annuli, as shown in Figure 10. Such an isotopy passes through the twice-punctured annulus that spans $B_{i+1,j-1}$ and $B_{i+1,j+1}$, and that is horizontal in the sense that it is contained in the level $i + 1$ fiber if we view it in T_φ . Note that we are able to isotope the tubed pants across $B_{i,j}$ and $B_{i+2,j}$ since these tori are filled and the slope of the pants boundaries are longitudes.




Figure 9. A subset of \mathcal{C} consisting of exactly two pairs of pants of each size ≥ 3 (left) tubes together to give a surface which can be cut into two pieces Σ_\circ^+ and Σ_\circ^- (center), both isometric to horizontal annuli with $n - 2$ punctures (right). By viewing Σ_\circ in this way we find a boundary compression (right).

Now, let Σ_\circ^+ be the half of Σ'_\circ containing exactly one pair of each size $j \geq 3$ of pants, and containing the size 3 pants having boundary on $B_{i_0+1,3}$. Let Σ_\circ^- be the other half (which also contains one pair of each size of pants). It follows that we can isotope Σ_\circ^+ so that it is a tubing of the aforementioned horizontal twice-punctured annuli spanning $B_{i_0,j}$ and $B_{i_0,j+2}$, for $2 \leq j \leq k$, where k is $n - 2$ if n is even, and k is $n - 3$ if n is odd. In the case where n is odd, we will have in addition to these horizontal twice-punctured annuli the pants of size $n - 1$. Altogether, these tube together to give a horizontal $(n - 2)$ times punctured annulus. We can isotope Σ_\circ^- similarly, so that Σ'_\circ is as shown in Figure 9c.

We pause here to note that there are two possible ways of tubing together the pants meeting B_q , but both of the resulting surfaces are isotopic via an isotopy through the filling solid torus of B_q . So regardless of the choice of tubing, we can isotope Σ'_\circ so that it is as shown in Figure 9c.

Returning to the proof, we find that Σ'_\circ must have a boundary compression disk, as shown in Figure 9c. The demonstrated boundary compression is still a compression if we re-tube Σ_\circ along B_p to obtain Σ , since it may be assumed to avoid a neighborhood of the punctures. Thus Σ is not essential, contradicting our earlier assertion.




Figure 10. When a pants of size j and a pants of size $j + 1$ share a boundary on $B_{i,j}$, and have their other boundaries on tori $B_{i+1,j+1}$, there is an isotopy taking their $B_{i,j}$ boundary to $B_{i+2,j+1}$. Such an isotopy passes through a horizontal twice-punctured disk.

Case 2: No pants Y_i has boundary on B_q . By Lemma 5.1 it follows that Σ_\circ is carried by \mathcal{C}_2 . In this case Σ_\circ must contain as a subset a twice-punctured annulus $A_{i,i+3}^-$, for some i . But then Σ_\circ has a compression annulus A which meets B_p , as shown in Figure 6b, contradicting our assumption that Σ_\circ was obtained by compressing it along *all* compression annuli meeting B_p . \square

Proof of Theorem 1.1. Fix $n \geq 4$. Let $\gamma = (\gamma_1, \dots, \gamma_m)$ be a system of Dehn filling slopes for the horizontal boundary components of M_C , with each slope of the form $\frac{1}{s_i}$. For any such system, the resulting manifold M_γ is

an l -component link exterior for the closure of an n -braid K_n . Since n -braid closures have Heegaard genus at most n , we deduce that $g(M_\gamma) \leq n$. If we further assume that $k_i \geq 2\pi(2n - l)$ for each i , then by [10, Theorem 1.1] (M_γ, H) is a good filling for any Heegaard surface of genus at most n (in particular, for any minimal genus Heegaard surface). It follows from the discussion following Hatcher's theorem in Subsection 2.2 that if the system γ is a boundary multi-slope, then for some i increasing s_i by 1 will give a slope that is not a boundary slope. Therefore we can find a filling system that is not a boundary multi-slope, and is a good filling, and such that $2\pi(2n - l) \leq s_i \leq 2\pi(2n - l) + 2$. Such a choice of γ is necessarily an exceptional filling by the 2π -theorem of Gromov and Thurston [4].

Note that while [10, Theorem 1.1] is written in terms of the length $l(\frac{1}{s_i})$ of the filling slope in the induced Euclidean metric on a cusp cross-section, our statement follows since $l(\frac{1}{s_i}) \geq s_i l(\lambda_i) \geq s_i$ for λ_i the longitude of the i^{th} boundary component.

For the choice of γ described above, we have by Lemma 5.2 that M_γ contains no closed incompressible surfaces, and by Lemma 5.3 that $g(M_\gamma) \geq n$. It then follows that the Heegaard genus of M_γ is exactly n . It also follows that K_n must have bridge number n , for if it were smaller the Heegaard genus would have to be smaller as well.

In order to establish the claimed bound on the crossing number of K_n , we first observe that a $\frac{1}{s}$ Dehn filling along a boundary component of width j is equivalent to adding s full twists to the j strands it encircles. One full twist of a band of j strands produces $j(j - 1)$ crossings, so s twists produces $sj(j - 1)$ crossings. Taking each s_i in the range given in the previous paragraph, and taking into consideration that there are $2n - l$ loops of width j for $2 \leq j \leq n - 1$, one loop of width n , and $(n - 1)(2n - l)$ crossings coming from L_p , we get

$$\begin{aligned} c(K_n) &\leq \overbrace{(n - 1)(2n - l)}^{B_p} + \overbrace{(2\pi(2n - 1) + 2)n(n - 1)}^{B_q} \\ &\quad + \sum_{j=2}^{n-1} \overbrace{(2\pi(2n - 1) + 2)j(j - 1)(2n - l)}^{B_{i,j}} \\ &< 4\pi n^2 \left(n + 2 \sum_{j=2}^{n-1} j(j - 1) \right) \\ &= 4\pi n^2 \left(n + \frac{n(n - 1)(2n - 1)}{3} - n(n - 1) \right) \leq 4\pi n^5 \end{aligned}$$

□

References

- [1] I. Agol, *Small 3-Manifolds of Large Genus*, Geometriae Dedicata **102** (2003), no. 1, 53–64.
- [2] D. Bachman and S. Schleimer, *Surface bundles versus Heegaard splittings*, Communications in Analysis and Geometry **13** (2005), no. 5, 903–928.
- [3] I. Biringer and J. Souto, *Ranks of mapping tori via the curve complex*, Journal für die reine und angewandte Mathematik (Crelles Journal) **2019** (2016), no. 748, 153–172.
- [4] S. A. Bleiler and C. D. Hodgson, *Spherical space forms and Dehn filling*, Topology **35** (1996), no. 3, 809–833.
- [5] D. Cooper and M. Scharlemann, *The structure of a solvmanifold’s Heegaard splittings*, in Proceedings of 6th Gökova Geometry–Topology Conference, Tr. J. Mathematics **23** (1999) 1–18.
- [6] M. Culler, W. Jaco, and H. Rubinstein, *Incompressible Surfaces in Once-Punctured Torus Bundles*, Proceedings of the London Mathematical Society **s3-45** (1982), no. 3, 385–419.
- [7] N. M. Dunfield, *Which small volume hyperbolic 3-manifolds are Haken?*, talk given at University of Warwick (1999), slides available at https://faculty.math.illinois.edu/~nmd/slides/haken_slides.pdf.
- [8] W. Floyd and A. Hatcher, *Incompressible surfaces in punctured-torus bundles*, Topology and its Applications **13** (1982), no. 3, 263–282.
- [9] W. Floyd and U. Oertel, *Incompressible surface via branched surfaces*, Topology **23** (1984), no. 1, 117–125.
- [10] D. Futer and J. S. Purcell, *Explicit Dehn filling and Heegaard splittings*, Comm. Anal. Geom. **21** (2013), no. 3, 625–650.
- [11] C. M. Gordon, *The theory of normal surfaces, lecture notes typeset by Autumn Kent* (2001).
- [12] W. Haken, *Some results on surfaces in 3-manifolds*, Studies in Modern Topology (1968) 39–98. (distributed by Prentice-Hall, Englewood Cliffs, NJ).

- [13] J. Hass and W. Menasco, *Topologically rigid non-Haken 3-manifolds*, J. Austral. Math. Soc (Series A) **55** (1993) 60–71.
- [14] A. Hatcher, *On the boundary curves of incompressible surfaces*, Pacific Journal of Mathematics **99** (1982), no. 2, 373–377.
- [15] A. Hatcher, P. Lochak, and L. Schneps, *On the Teichmüller tower of mapping class groups*, Journal für die reine und angewandte Mathematik **2000** (2000), no. 521, 1–21.
- [16] A. Hatcher and W. Thurston, *A presentation for the mapping class group of a closed orientable surface*, Topology **19** (1980) 221–237.
- [17] W. Jaco and J. H. Rubinstein, *PL equivariant surgery and invariant decompositions of 3-manifolds*, Advances in Mathematics **73** (1989) 149–191.
- [18] L. Lopez, *Alternating knots and non-Haken 3-manifolds*, Topology and its Applications **48** (1992), no. 2, 117–146.
- [19] ———, *Small knots in Seifert fibered 3-manifolds*, Mathematische Zeitschrift **212** (1993), no. 1, 123–139.
- [20] Y. Moriah and H. Rubinstein, *Heegaard structures of negatively curved 3-manifolds*, Communications in Analysis and Geometry **5** (1997), no. 3, 375–412.
- [21] U. Oertel, *Closed incompressible surfaces in complements of star links*, Pacific Journal of Mathematics **111** (1984), no. 1, 209–230.
- [22] G. Perelman, *The entropy formula for the Ricci flow and its geometric applications*, arXiv:math.DG/0211159 (2002).
- [23] ———, *Finite extinction time for the solutions to the Ricci flow on certain three-manifolds*, arXiv:math.DG/0307245 (2003).
- [24] ———, *Ricci flow with surgery on three-manifolds* (2003) [UrlarXiv:math.DG/0303109](https://arxiv.org/abs/math/0303109).
- [25] Y. Rieck, *Heegaard structures of manifolds in the Dehn filling space*, Topology **39** (2000) 619–641.
- [26] Y. Rieck and E. Sedgwick, *Persistence of Heegaard structures under Dehn filling*, Topology and its Applications **109** (2001), no. 1, 41–53.
- [27] A. Thompson, *Thin position and the recognition problem for \mathbb{S}^3* , Math. Res. Lett. **1** (1994) 613–630.

- [28] W. P. Thurston, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 3, 357–381.

DEPARTMENT OF MATHEMATICS, RICE UNIVERSITY
HOUSTON, TX 77005, USA
E-mail address: `william.worden@rice.edu`

RECEIVED JANUARY 7, 2020
ACCEPTED OCTOBER 6, 2020