Knotted Spheres and Graphs in Balls

Hugh Nelson Howards
Wake Forest University
Visiting Researcher
University of Texas, Austin
Department of Mathematics C1200
Austin, TX 78712
e-mail: howards@fireant.ma.utexas.edu

Abstract

Given a properly embedded graph Γ in a ball B and a punctured sphere Σ properly embedded in $B - \Gamma$ we examine the conditions on Γ that are necessary to assure that Σ is boundary parallel.

Keywords: Knotted graphs, rational tangles, Dehn Surgery, Berge Knots.

AMS classification: 57M99

1 Knotted Spheres

This paper explores embedding punctured spheres in balls. I would like to thank Mike Freedman for suggesting the question and for insightful comments. I would also like to thank Cameron Gordon, John Luecke, Martin Scharlemann, and Ying-Qing Wu for helpful comments along the way.

We first must set up some definitions that greatly simplify the statements of the theorems. Throughout the paper let B be the unit ball in $\mathbb{R}^3$. Let S be the boundary of B.

Definition 1.1. *A complete graph on n vertices whose vertices $\{v_1, \ldots v_n\}$ are disjoint disks on S, and whose edges $\{e_{\{i,j\}}, 1 \leq i < j \leq n\}$ are properly embedded geodesics in B with the property that if $s \neq t$, $v_s \cap v_t = \emptyset$, is a* <u>*standard unlinked n-graph*</u>.

Definition 1.2. *A complete graph on n vertices whose vertices* $\{v_1, \ldots v_n\}$ *are disjoint disks on* S *and whose edges* $\{e_{\{i,j\}}, 1 \leq i < j \leq n\}$ *are properly-embedded, disjoint unknotted arcs in* B *with the property that no two edges are linked in the ball, is called a pairwise unlinked n-graph.*

Definition 1.3. *Given a graph* Γ *in* B*, and* Σ *a properly embedded n-holed sphere in* B*, such that* $\partial\Sigma = \cup_{i=1}^{n} \partial v_i$ *and* $\Sigma \cap e_{\{i,j\}} = \emptyset$ *for all* i*,* j*.* Σ *is called an enveloping n-holed sphere for* Γ*.* $B - \Sigma$ *consists of two components. The part containing* $\{e_{\{i,j\}}, 1 \leq i < j \leq n\}$ *is called the inside of* Σ*, the other component is called the outside.*

Definition 1.4. *An enveloping n-holed sphere* Σ *is standard, if it is boundary parallel, meaning there is a product structure on the outside of* Σ *taking the interior of* Σ *to* $S - \cup_{i=1}^{n} v_i$*, but leaving* $\partial\Sigma$ *fixed.*

Definition 1.5. *Let a core of an enveloping n-holed sphere* Σ *be a 1-complex, such that the boundary of a regular neighborhood of the complex co-bounds a product region with* Σ*. Let the star core be the core with exactly* n *edges and one vertex of valence of* n*. We shall call the vertices of the star core* $\{\nu_1, \ldots \nu_n, \nu_{n+1}\}$ *where* $\nu_i \subset v_i$ *and* ν_{n+1} *is the vertex of valence* n*. We call the edges of the star core* $\{\epsilon_1, \ldots \epsilon_n\}$ *where* ϵ_i *is the edge containing* ν_i*.*

2 Pairwise unlinked n-graphs

2.1 The Core Lemmas

In this section we introduce a couple of basic lemmas that will increase our insight to the theorem and will be useful in some of the cases.

Lemma 2.1. *Given* ϵ_i *and* $\epsilon_j, i \neq j$*, two edges of the star core of an enveloping n-holed sphere for a pairwise unlinked n-graph,* $\epsilon_i \cup \epsilon_j$ *may never be a knotted arc in* B*.*

Proof. $\epsilon_i \cup \epsilon_j$ makes up the core of the cylinder that is left over if the n-holed sphere is compressed along disks parallel to all of the uninvolved vertices. If the core of the cylinder were knotted (see Figure 1), then by standard satellite knot theory so is any arc running through the cylinder, but the edge between these two vertices in our graph is unknotted and can be assumed to be inside the cylinder, so this cannot be the case. □

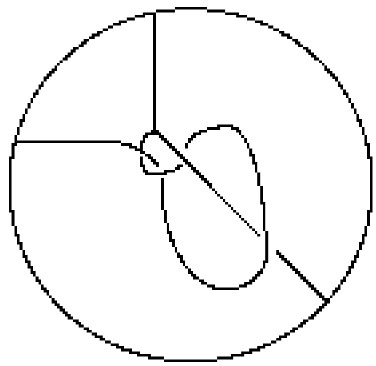

Figure 1: A knotted core

Lemma 2.2. *Let ϵ_i, ϵ_j, and ϵ_k be three edges of a star core of an enveloping n-holed sphere for a pairwise unlinked n-graph. Let ϵ_i^k and ϵ_j^k be edges obtained from ϵ_i and ϵ_j by contracting ϵ_k and perturbing the arcs slightly so their end points are disjoint. Then ϵ_i^k and ϵ_j^k are not linked.*

ϵ_j',

Proof. If ϵ_i^k and ϵ_j^k were linked, then $e_{\{i,k\}}$ and $e_{\{j,k\}}$ (edges of the the original unlinked n-graph running from vertex i to vertex k and j to k respectively) would have to be linked. □

[C] and [GF] are good places to look for an introduction to rational tangles. As explained in [GF] a rational tangle T is assigned a rational number $F(T)$ corresponding to a simple continued fraction. Two tangles T_1, T_2 are isotopic if and only if $F(T_1) = F(T_2)$. Let $N(T)$ be the knot obtained by from T by connecting the ends of the tangle in the manner called the numerator of T and $D(T)$ the denominator (See Figure 2).

Lemma 2.3. *If T is a rational tangle with $D(T)$ an unknot, then $F(T) = p, p \in \mathbb{Z}$. Therefore T may be undone by merely twisting two of the vertices around each other p times leaving the other two vertices fixed.*

Proof. It is well known that if $F(T) = p/q$, then $D(T)$ produces a p/q 2-bridge link, which is trivial if and only if q=1. Similarly, of course, $N(T)$ is an unknot if and only if $p = 1$ so in that case $F(T) = 1/q, q \in \mathbb{Z}$ (the picture is just rotated ninety degrees). [M] section 9.3 and [BZ] sections 12.A and 12.B have expositions on the these facts. The classification was first done in [S].

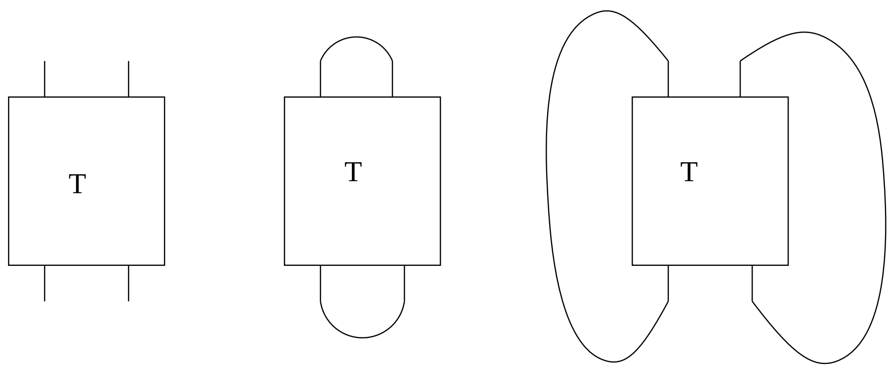

Figure 2: A representation of tangle T is given on the left. The numerator of T is demonstrated in the center, the denominator on the right.

$\square$

2.2 The case, $n = 3$.

Theorem 2.4. *Every enveloping 3-holed sphere for a pairwise unlinked 3-graph is standard.*

Proof. By Lemma 2.1 the star core is not knotted, so we may assume it consists of one straight edge $\epsilon_1 \cup \epsilon_2$ running from ν_1 the north pole to ν_2 at the south pole and another edge ϵ_3 meeting this edge at ν_4 at the origin and winding around in some manner before ending up at a ν_3 somewhere on the Southern hemi-sphere (as in Figure 3).

Act upon ϵ_3 by an ambient isotopy leaving $\epsilon_1 \cup \epsilon_2$ fixed until ϵ_3 lies entirely below the equatorial disk except for at ν_4 (Figure 4). Now if we examine the two edges of the core in B' the Southern hemi-ball (the ball we get by cutting the original ball in half along the equatorial disk) and pull ϵ_2 and ϵ_3 apart slightly so that they no longer intersect to get ϵ_2' and ϵ_3', the core's edges cannot be linked by Lemma 2.2 (this is the same as contracting ϵ_1 and looking at ϵ_2^1 and ϵ_3^1. This together with the fact that neither ϵ_2 nor ϵ_3 can be knotted in B' by Lemma 2.1, means that $\epsilon_2' \cup \epsilon_3'$ is just a rational tangle in B'.

Now by Lemma 2.1 we know that the rational tangle to which they correspond, must be an unknot if vertices on the equatorial disk are connected as are the vertices on the sphere. By Lemma 2.3 we may assume that the rational tangle is obtained merely by twisting the two vertices on the southern

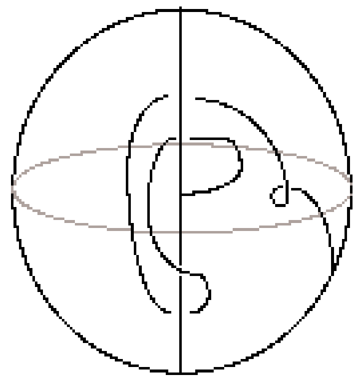

Figure 3: Straightening the core

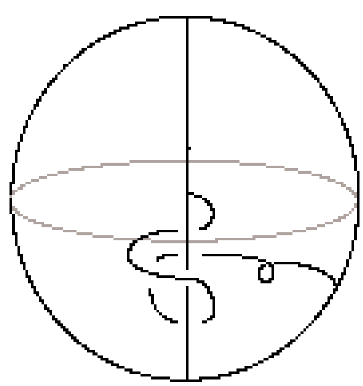

Figure 4: An isotopy acting upon the straightened core

hemisphere around each other, and therefore that it may be untangled without affecting the rest of the core (Figure 5). Therefore the core is standard, as must be the enveloping n-holed sphere. $\square$

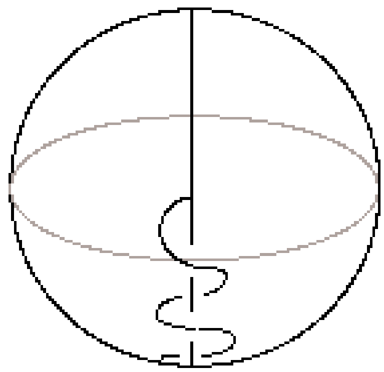

Figure 5: The only possible complication for $n = 3$.

2.3 The Case $n = 4$

The argument for the case $n = 4$ is particularly interesting, because the methods for the previous case fall slightly shy of working, but the counterexample from the case $n \geq 5$ also just fails to disprove it. I would like to thank Ying-Qing Wu for ideas that were particularly helpful in this section.

Theorem 2.5. *Given $n \leq 4$ every enveloping n-holed sphere for a pairwise unlinked n-graph Γ is standard.*

Proof. Assume there is a counterexample for $n = 4$ and examine what it must look like. B minus the interior of the non-standard ball, bounded by the non-standard enveloping 4-holed sphere is homeomorphic to $S^2 \times I$ minus an open neighborhood of four arcs $\{\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4\}$ which run from the inner sphere to the outer sphere. Since the enveloping 4-holed sphere is non-standard, the arcs cannot each simultaneously be isotoped in $S^2 \times I$ to be $pt \times I$.

Since the original edges of Γ were unknotted and pairwise unlinked, they could be thought of as rational tangles. Examine the three pairs of edges of Γ corresponding to the three possible pairings of the vertices, $(e_{\{1,2\}}, e_{\{3,4\}})$, $(e_{\{1,3\}}, e_{\{2,4\}})$, $(e_{\{1,4\}}, e_{\{2,3\}})$. Pick one pair, say $(e_{\{1,2\}}, e_{\{3,4\}})$. Because the pair is a rational tangle, there is a disk D_1 in B that separates the edges.

Assume D_1 intersects Σ minimally. An innermost curve of intersection on D_1 yields a subdisk D_1' which compresses Σ into two annuli $A_{\{1,2\}}$ and $A_{\{3,4\}}$ which are just the boundary of a regular neighborhood of $e_{\{1,2\}}$ and $e_{\{3,4\}}$ respectively. The same arguments can be made for the other two pairs. Thus, there are three disks that can be added to $S^2 \times 0$ that separate the end points of the ϵ_i into pairs, each yielding a rational tangle.

If we take the branched double cover of $S^2 \times I$ over the four arcs $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4$, we get a manifold M with boundary two tori T_0, T_1 because the branched double cover of a sphere over four points is a torus. Since adding D_1' to B minus the "inside" of Σ, thought of as $S^2 \times I$ minus a neighborhood of the four branching arcs yields the exterior of a rational tangle, and the branched double cover of a ball over a rational tangle is, of course, just a solid torus. Thus, D_1' lifts to a disk that gives a filling of T_0 that turns M into a solid torus, so M is just a solid torus minus a knot. Note that the three pairings of the vertices gives three different fillings of T_0 each of which yields a solid torus.

Theorem 2.6. [**B Corollary 2.9**] *If k is a nontrivial knot in $D^2 \times S^1$ such that k is not parallel into $\partial D^2 \times S^1$ and there exists more than one nontrivial surgery on k yielding $D^2 \times S^1$, then k is equivalent either to $W_3^1 W_7^{-3}$ or its mirror image $W_3^{-1} W_7^3$.*

$W_3^{-1} W_7^3$ is the $(-2, 3, 7)$ pretzel knot embedded as shown in Figure 8. We refer to the this as the Berge knot. Since the arcs were not standard, we must lift to the Berge knot or a knot k parallel into T_1, the boundary of $D^2 \times S^1$.

Proposition 2.7. *If k is parallel into the boundary of $D^2 \times S^1$ then k fails to produce a counter example to Theorem 2.5.*

The proof of the proposition requires two steps. First we prove that k can be assumed to have the standard embedding in $D^2 \times S^1$ by showing there is a unique strong inversion on k. Second we prove that the punctured sphere produced by quotienting out by the $\mathbb{Z}_2$ symmetry is either boundary parallel or else violates Lemma 2.1 and therefore is not a counterexample.

Proof. To prove that k must be embedded in the standard manner we examine an annulus A that runs from T_0, the torus boundary component corresponding to k in the exterior of k, to T_1, the boundary of the solid torus.

A can be assumed to be embedded (see [CF]). Let F be the involution of $D^2 \times S^1$. Let $F(A) = A'$.

Lemma 2.8. *There is a unique strong inversion of k, a torus knot parallel to the boundary of a solid torus.*

Proof. Our first goal is to show that A can be chosen such that $A = A'$. We choose A to have a minimal number of intersections with A'. It is clear that ∂A can be assumed to be fixed by F, so perturbing A slightly we can assume that $\partial A \cap \partial A' = \emptyset$. Now $A \cap A'$ must consist of simple closed curves. An innermost disk argument suffices to show that each of the simple closed curves is essential on A and A', so the intersection consists of disjoint simple closed curves $c_1, c_2, \ldots c_n$ parallel to ∂A on A. Let c_1 be the curve of intersection closest to the boundary component of A on T_1, c_2 be the second, and so on increasing the index as the curves move towards the boundary component on T_0. Likewise on A' the intersection consists of parallel essential circles $c'_1, c'_2, \ldots c'_n$ labeled in the same manner. Let A_i be the sub-annulus of A running from the boundary component on T_1 to c_i, and A'_i be the sub-annulus of A' running from the boundary component on T_1 to c'_i. Let c_j be the curve of intersection on A that corresponds to the intersection with A' at c'_1. If we cut and paste A replacing A_j by a push off of A'_1 we reduce the number of intersections of A and A' by at least one. In order to preserve the property that $F(A) = A'$ we must also replace A'_j by a push off of A_1. This, however, cannot increase the number of intersections, so we have a new annulus running from T_1 to T_0 that has fewer intersections with its image under F, contradicting minimality. Thus, we can assume that $A \cap A' = \emptyset$.

This, however, implies that restricting to the solid torus between A and A', F takes $A \times I$ to itself, exchanging $A \times 0$ with $A \times 1$. This in turn implies that there must be an annulus in $A \times I$ that is fixed by F.

Now that we know that A is fixed, we use it to show that we have a standard inversion of k. Let D be a meridional disk for T_1 that is fixed by F. Examine $D \cap A$. In general the intersection pattern on D will look something like Figure 6 with a collection of arcs running from T_0 (which punctures D several times) to T_1 and a collection that run from one of the punctures from T_0 to another. We can remove the arcs running from T_0 to T_0 by picking an outermost arc on A that runs from one component of ∂A to itself. This small disk gives an isotopy of A together with k that reduces the number of intersections of A with D. because $F(A) = A$ we can simultaneously do a second isotopy of A that also reduces the number of intersections of A with

D and preserves the symmetry of k. Thus, we can assume that $D \cap A$ consists solely of arcs running from k to T_1. This, however, shows that we have the standard symmetry for k because cutting the solid torus along D turns it into a cylinder and A becomes bands running from the top of the cylinder to the bottom in the unique way possible. $\square$

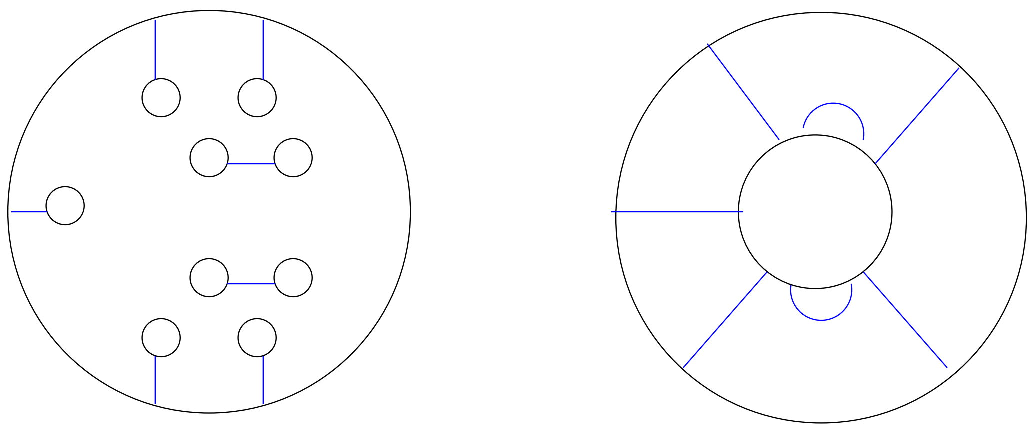

Figure 6: $D \cap A$. D is pictured punctured nine times by T_0 on the left and A is pictured on the right.

Now we need only argue that torus knots with standard embeddings fail to give a counterexample. Let k be a torus knot embedded in $D^2 \times S^1$, fixed by an involution F of $D^2 \times S^1$. Let A be the fixed annulus above. Let M be $S^2 \times I$ with four branching arcs $\epsilon_1, \ldots \epsilon_4$, the quotient of the exterior of k in $D^2 \times S^1$ by F. In the quotient, T_0 maps down to $S^2 \times 0$ which we will designate S_0 and T_1 maps to $S^2 \times 1$ designated S_1.

Lemma 2.9. *Either $S_0 \cup \epsilon_1 \cup \epsilon_2 \cup \epsilon_3 \cup \epsilon_4$ violates Lemma 2.1 or it is standard and therefore in either case is not a counterexample to Theorem 2.5.*

Proof. Because $\epsilon_1, \ldots \epsilon_4$, each run from S_0 to S_1 the two arcs a_1, a_2 that are fixed in $D^2 \times S^1$ by F must each intersect k in exactly one point. In the knot exterior a_1 and a_2 are then broken each into two arcs with one end point of each arc on T_1 and one end point on T_0. In M the a_i then become the four branching arcs $\epsilon_1 \ldots \epsilon_4$.

Under the quotient, the annulus A becomes a rectangular disk D which without loss of generality runs down ϵ_3 along S_0 up ϵ_4 and back along S_1 and is disjoint from ϵ_1 and ϵ_2. This is clear because $A \cap (a_1 \cup a_2)$ consisted of the preimage of ϵ_3 and ϵ_4 but was disjoint from the preimage of ϵ_1 and ϵ_2 (recall that A, like the a_i is fixed by F and runs from T_0 to T_1). This means that M looks exactly like Figure 7, where ϵ_1 and ϵ_2 form a rational tangle inside the ball designated T.

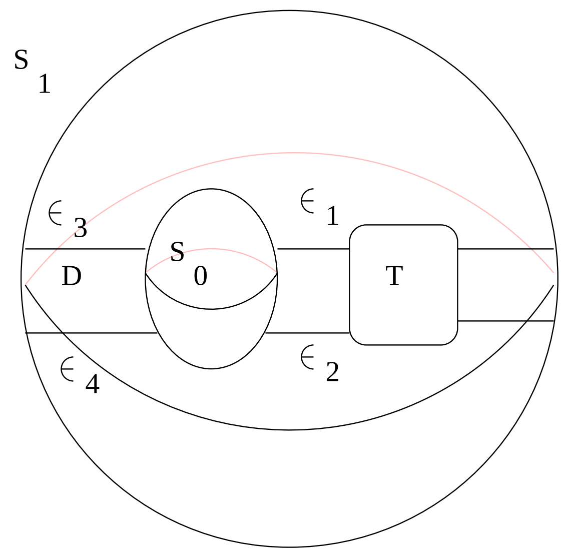

Figure 7: The branching arcs of M consist of two standard arcs ϵ_3 and ϵ_4 and a rational tangle running from S_0 to S_1 designated T consisting of arcs ϵ_1 and ϵ_2.

By Lemma 2.3 we see that $\epsilon_1 \cup \epsilon_2$ must be the tangle T that results from two horizontal arcs whose eastern vertices are twisted n times around each other or $\epsilon_1 \cup S_0 \cup \epsilon_2$ will be knotted violating Lemma 2.1. On the other hand, if $\epsilon_1 \cup \epsilon_2$ is the tangle T above, then twisting the eastern portion of M, n times shows M is homeomorphic to the standard picture and therefore again fails to be a counterexample.

$\square$

Note: One can in fact prove that (if k is a torus knot other than the unknot) k never produces a standard enveloping sphere, but instead always creates one that violates Lemma 2.1.

Lemma 2.9 completes the proof that k cannot be parallel to T_1 and therefore must be a Berge knot if it is to produce a counterexample. $\square$

We may transform the traditional picture of the Berge Knot to a symmetric one as in Figure 8. Snappea [W] tells us that this knot (entered as a link) has exactly one $\mathbb{Z}_2$ symmetry, which we can now see.

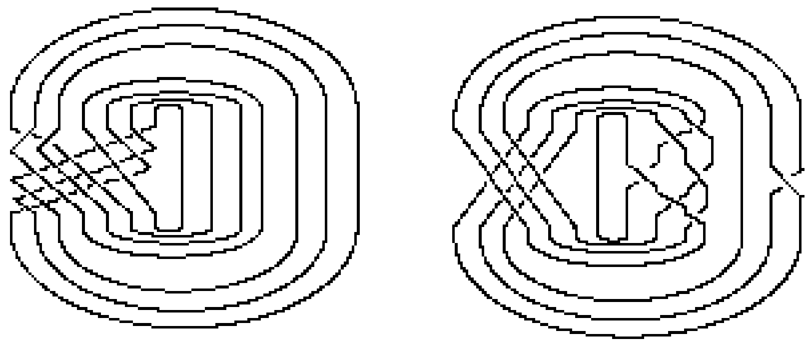

Figure 8: The Berge knot in $D^2 \times I$ and an embedding after an isotopy in $D^2 \times I$ that reflects the knots $\mathbb{Z}_2$ symmetry ($D^2 \times I$ is not drawn, but is the obvious choice for the initial braid).

We quotient out by the symmetry to either get a counterexample, or proof that there is none. We get $S^2 \times I$ minus four arcs $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4$. We will show that three of the arcs violate Lemma 2.2 (See Figure 9).

Label the edge omitted from the picture ϵ_4. Label the horizontal edge that has one of its vertices on the eastern side of the sphere ϵ_3. Contraction of ϵ_3 takes edges ϵ_1 and ϵ_2 to ϵ_1^3 and ϵ_2^3 as pictured in Figure 10.

To see that ϵ_1^3 and ϵ_2^3 are in fact linked we call on the following well known facts about rational tangles. (See, for example, [M] **Theorem 9.3.1**.)

Theorem 2.10. *1. A 2-bridge knot (or link) is the denominator of some rational tangle*

2. Conversely, the denominator of a rational tangle is a 2-bridge knot (or link).

Corollary 2.11. *If k is the denominator of a rational tangle, then k is prime.*

The denominator of the tangle in Figure 10 is the connect sum of a trefoil and a figure eight knot and therefore it is not a rational tangle by Corollary 2.11. Since neither of the arcs is knotted, they must be linked violating

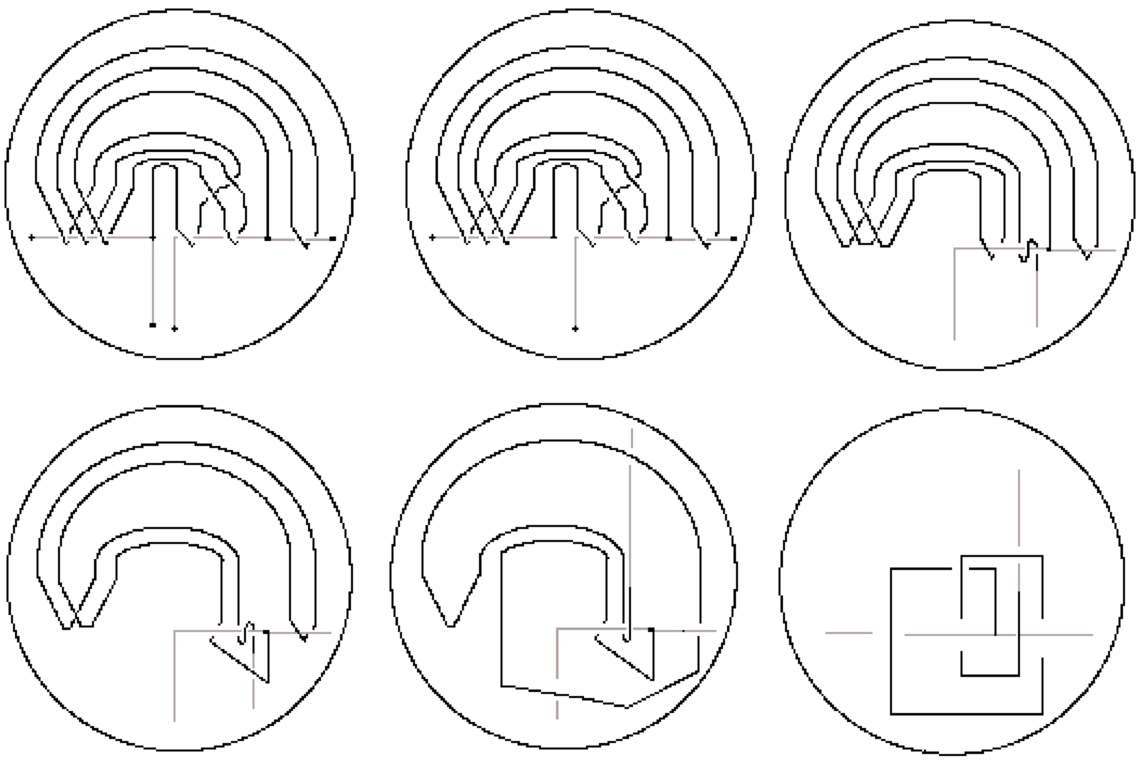

Figure 9: The dark edge is the interior sphere in $S^2 \times I$ obtained from quotienting out by the $\mathbb{Z}_2$ symmetry of the Berge knot in $D^2 \times I$. The light edges are the four singular arcs $\{a_1, a_2, a_3, a_4\}$. After dropping one of the a_i and an isotopy, the three other edges appear to be linked.

Lemma 2.2. Since this was the only possible counterexample, Theorem 2.5 must be true.

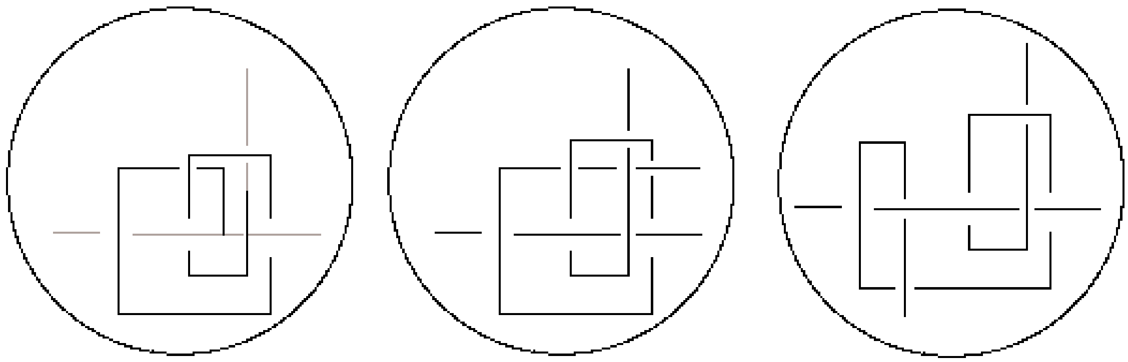

Figure 10: The edges violate Lemma 2.2.

□

2.4 The Case $n \geq 5$

Theorem 2.12. *There exist pairwise unlinked n-graphs with non-standard enveloping n-holed spheres for all n, $n \geq 5$.*

Proof. Figure 11 shows a θ_n curve in a ball. Such a graph is well known not to be standard, but every subgraph is standard. Let a θ_n curve be the star core of an enveloping n-holed sphere. Although it is not standard, it supports a pairwise unlinked n-graph Γ. If we think of the enveloping sphere as bounding a central (round) ball with n tentacles running to the boundary, we can picture the pairwise unlinked graph as being a standard unlinked graph in the central ball which is extended by a product down each of the tentacles. Now since $n \geq 5$ any two edges of Γ miss at least one vertex, but there is an isotopy of any $n-1$ arcs of a θ_n curve that makes those arcs appear standard. Likewise since the edges of Γ completely miss one of the tentacles, they may be pictured as being embedded in a ball bounded by a standard enveloping $(n-1)$-holed sphere. The arcs remain standard within the central ball and are extended by a product down the tentacles throughout the entire process, so clearly the edges are not linked pairwise. (See Figure 11).

□

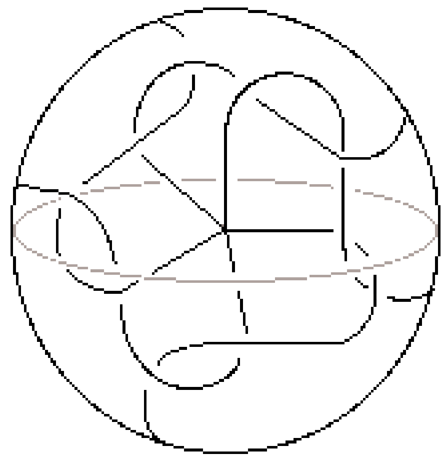

Figure 11: The θ_n ($n = 5$ here) curve is a non-standard core but supports a pairwise unlinked graph for $n \geq 5$

3 Standard unlinked n-graphs

Theorem 3.1. *Given a standard unlinked n-graph* Γ*, every enveloping n-punctured sphere* Σ *is standard.*

Note: Since the edges of our graph are geodesics in this case, Morse theory assures us that we can assume that there is an embedded disk D in B whose interior is disjoint from all of the edges of Γ and whose boundary consists of two arcs α and β, where α is one of the edges of the graph, β is strictly contained in ∂B, $\partial\alpha = \partial\beta$ and the interior of β is disjoint from all of the edges of Γ.

Let's examine how D meets Σ.

Claim 3.2. *We may assume that* $D \cap \Sigma$ *contains no simple closed curves.*

Proof. Assume D is chosen with a minimal number of intersections with Σ. Examine an innermost curve δ on D. If δ is not essential on Σ, then there is an obvious isotopy through which this intersection could have been eliminated, so we may assume it is essential on Σ.

Therefore the innermost loop gives us a compressing disk for Σ. Homology is enough to assure us that the disk must be on the inside of Σ (the component of $B - \Sigma$ containing $\{e_{\{i,j\}}, 1 \leq i < j \leq n\}$). Since δ is assumed to be essential in Σ, it must separate the vertices of Σ into two non-empty sets. With no loss of generality, let v_1 be in one set and v_2 be in the other. Since δ separates the vertices (as in Figure 12), the disk it bounds does too, and e_{12} must intersect it. This, however, is a contradiction since the interior of D is disjoint from the edges of Γ. $\square$

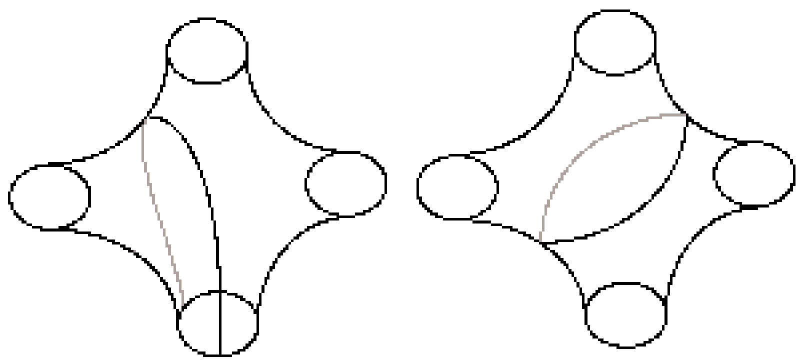

Figure 12: An essential arc and curve separate the remaining vertices into two sets.

Now we examine an outermost arc γ on D. If γ runs from one vertex of Σ to a different one, then it is obvious that the corresponding subdisk of D is on the outside of Σ. This gives us a compression disk that allows us to complete the proof by induction. (See Figure 13).

We may therefore assume that γ connects a vertex, say v_n to itself. Because $D \cap \Sigma$ is assumed to be minimal γ must be essential on Σ. If the disk γ cuts off on D is on the inside of Σ, then once again it must separate $\{v_1, \ldots v_{n-1}\}$ into two sets and the argument proceeds as in the simple closed curve case. (See Figure 12).

Our final case therefore is that although γ connects v_n to itself, the disk is on the outside of Σ. Compressing along this disk splits the n-punctured sphere Σ into two pieces Σ_1 and Σ_2. Σ_1 is an r-punctured sphere and Σ_2 is a $n+1-r$-punctured sphere, where $2 \leq r \leq n-1$. By induction we may therefore assume that Σ_1 and Σ_2 are standard. Σ_1 and Σ_2 and their product structures are in different "halves" of B. They are separated by the annulus formed by the boundary compressing v_n. The inverse of the

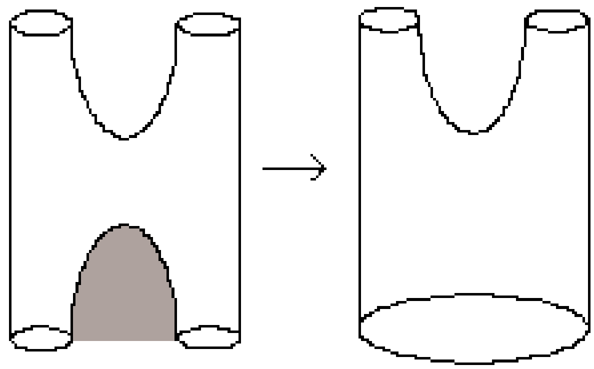

Figure 13: A boundary-compressing disk on the outside between two vertices reduces n to $n-1$

boundary compression is a tunnel connecting the two boundary components of the annulus. Up to isotopy there is a unique arc running across the annulus, so there is a unique tunnel we can add to attain Σ from Σ_1 and Σ_2. Since the product structures on Σ_1 and Σ_2 are to the outside of the annulus, and the tunnel can be added extremely close to the boundary, it is clear that the product structure can be extended across the tunnel to give a product structure on Σ, proving that it is standard.

Note: we needed the full strength of the complete graph here. Figure 14 shows a counter-example for $n = 3$, if one edge is missing from Γ. This counter-example may be generalized for any n to give a counterexample for the complete graph on n vertices, minus one edge.

4 The Infinite Case

Definition 4.1. *$f : F \to \Sigma \subset \mathbb{R}^3$ is a proper embedding of F, a sphere with a Cantor Set worth of punctures, in $\mathbb{R}^3$ if the pre-image of every compact set in $\mathbb{R}^3$ is a compact set in F (H^3 could replace $\mathbb{R}^3$ throughout this section).*

Definition 4.2. *Let $\cup l_\alpha$ be a collection of geodesics in $\mathbb{R}^3$. Let $\cup l_\alpha$ be contained in one connected component of $\mathbb{R}^3 - \Sigma$. Again let the inside of Σ be the component of $\mathbb{R}^3 - \Sigma$ containing Γ and the outside be the remaining component.*

Definition 4.3. *Σ is said to be standard if there exists a product structure*

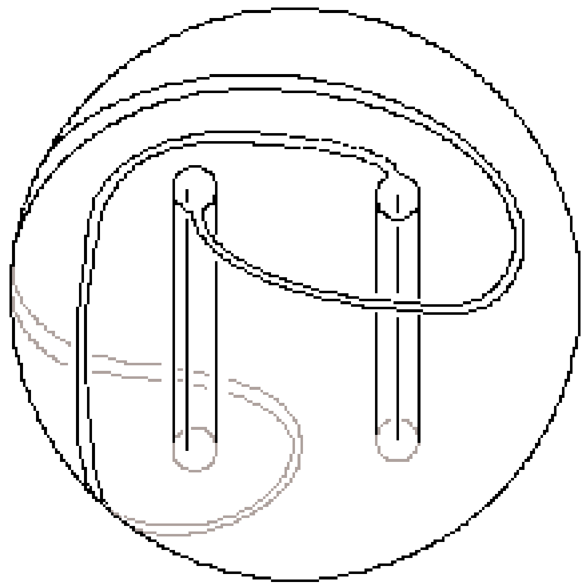

Figure 14: A non-standard sphere for the complete graph on three vertices minus one edge.

on the outside of Σ such that it is a product of the punctured sphere and a half open interval.

We may now ask, how many lines contained inside of Σ it takes to ensure that the embedding of Σ is standard.

Theorem 4.4. *Given Σ, a proper embedding of F, a sphere with a Cantor Set worth of punctures in $\mathbb{R}^3$ and a set of geodesics which are dense in the Cantor Set, we may conclude that the punctured sphere is standard.*

In this context saying that the geodesics are dense in the Cantor Set means that given any two punctures p_1 and p_2 of F and any two neighborhoods of those punctures μ_1 and μ_2 on F, there exists a geodesic that runs from the image of some puncture in μ_1 to the image of some puncture in μ_2.

Note that though there are an uncountable number of points in the cantor set, we are only requiring a countable set of geodesics. Even if we wanted to satisfy the property of being dense for every point on the sphere, we would still only need a countable set of geodesics, since geodesics connecting all the points on the sphere with rational coordinates would suffice and the set of possible pairings of a countable set is itself a countable set.

Perhaps the easiest way to picture the scenario is to imagine the universal cover of a genus-two handlebody in hyperbolic three-space. The punctured sphere would be the boundary of this cover and the lines would be geodesics connecting points at infinity.

It is worthy of note that in the finite case we needed the complete graph, so we needed all possible geodesics, but in the infinite case with an uncountable number of punctures, we only need a countable number of edges.

We now begin the proof of the theorem.

Proof. Choose a point on Σ to be the origin of $\mathbb{R}^3$.

Let S_n be the sphere of radius n centered at the origin in $\mathbb{R}^3$, and let B_n be the ball that it bounds. We may alter Σ slightly if necessary so that we may assume that it intersects each S_i transversally.

Fix i and examine B_i. The pieces of Σ in B_i may be broken into two sets. The first set consists of the connected piece containing the origin Σ_{i1}, and the second set, all the other pieces, $\{\Sigma_{i2}, \ldots \Sigma_{in}\}$. We shall isotope Σ until B_i contains one piece of the first type and none of the second on the "outside" of Σ in B_i. The latter does not prevent us from claiming that the former is boundary parallel in $B_i - \Sigma$, so we do not worry about them. See Figure 15

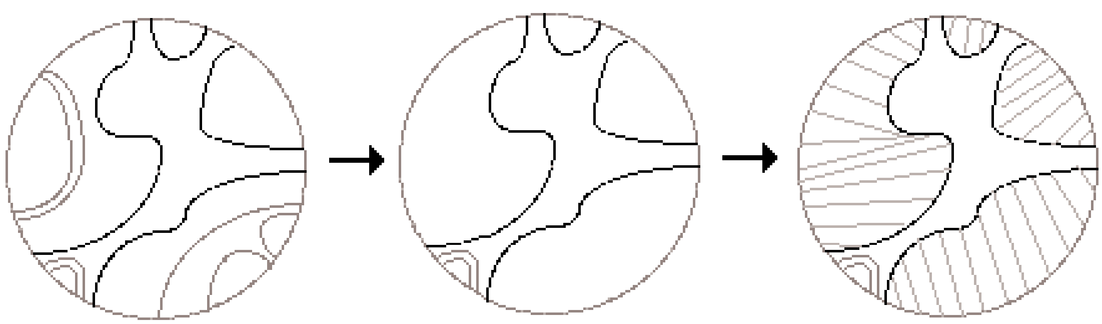

Figure 15: A small ball can be cleaned up leaving a product structure on the outside of the punctured sphere.

We examine the pre-images of the Σ_{ij} in F. We now make an argument to show that we may assume that none of them have a boundary curve which is trivial in the fundamental group of F.

If there were a trivial boundary curve, we could choose an innermost one (on F). This would bound an embedded disk D on Σ that meets S_i in a simple closed curve. The boundary curve splits S_i into two disks, each of which bounds a ball with D. If D is on the interior of B_i, then we choose the ball that does not contain the surface Σ_{i1} (Σ_{i1}, is connected and disjoint from D, so it can only be in one of the two balls). If D is on the exterior of B_i then we choose the smaller of the two balls (it is contained in the other ball). Either way we push D across the ball and through S_i, eliminating at least its intersection with S_i and possibly more extraneous intersections that were contained in the ball through which our isotopy was done.

Note that Σ_{i1} is unchanged away from its boundary and its boundary can only be changed by capping off trivial components. We continue this process until there are no more trivial components in the entire collection of Σ_{ij}.

At the risk of sloppy notation we shall continue throughout to call the new surfaces Σ_{ij} carefully noting at each step that we still have done only a finite number of isotopies to a finite number of pieces.

We now notice that Σ_{i1} fits all of the criterion of the standard finite case. Since the geodesics are dense within the Cantor Set, at any finite stage there will be a complete graph in B_i on the vertices that are given by the intersection of Σ_{i1} and S_i. Thus, by the previously proven finite case, Σ_{i1} is standard in B_i. We can use its inherited product structure to isotope Σ to make sure that it does not intersect B_i on the outside of Σ_{i1}.

Now we repeat the process for some $k > i$. We might worry that this process results in our pushing some piece of Σ an infinite number of times, but this is not the case, as every point in Σ is in some Σ_{k1} for large enough k and our isotopies never affect points of Σ_{k1} that are not near the boundary of B_k. Thus, for each point there is some k such that the point is left alone for good after k steps. Since each step involved only a finite number of isotopies, each point is moved only a finite number of times.

The only thing left for us to check is that the product structure for Σ_{k1} can be chosen to correspond exactly with the product structure we already chose for Σ_{i1}. B_i may be left fixed as we do our operations for Σ_{k1}, so naturally Σ_{i1} remains fixed, too.

Since Σ_{i1} is boundary parallel in B_i, we may substitute part of S_i for it in Σ_{k1} and the resulting surface still contains a complete graph on one side and must be boundary parallel. If we concatenate its product lines in B_k

with the product lines of Σ_{i1} in B_i we see a product structure that suits our desires as in Figure 16.

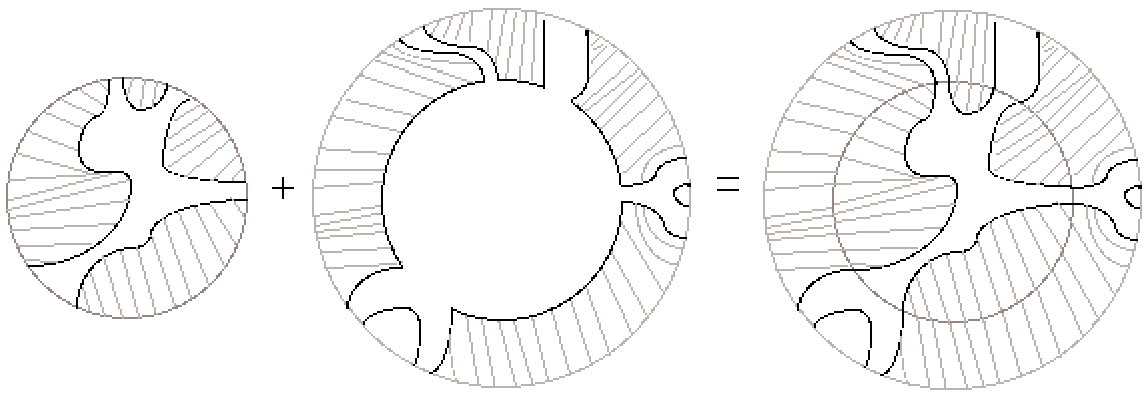

Figure 16: The product extends from one stage to the next

$\square$

5 References

[B] J. Berge, *The knots in $D^2 \times S^1$ which have nontrivial Dehn surgeries that yield $D^2 \times S^1$.* Topology Appl. 38 (1991), no. 1, 1–19.

[BZ] G. Burde and H. Zieschang, *Knots.* Walter de Gruyter, Inc. New York, NY, 1985.

[C] J. H. Conway, *Enumeration of knots and links.* Computational Problems in Abstract Algebra, Pergamon Press, 1970, 329–359.

[CF] J. Cannon and C. Feustel *Essential embeddings of annuli and mobius bands in 3-manifolds.* Transactions of the AMS 215, (1976), 219–239.

[GF] J. Goldman and L. Kauffman, *Rational Tangles.* Advances in Applied Mathematics 18 (1997), 300–332.

[HR] J. Hempel and L. Roeling, *Free factors of handlebody groups.* preprint.

[M] K. Murasugi, *Knot Theory and Its Applications*, Birkhauser, Boston, 1996.

[ST] H. Schubert, *Knoten mit zwei Brucken.* Math. Zeit. 1956, no. 65, 133–170

[ST] M. Scharlemann and A. Thompson, *Detecting unknotted graphs in 3-space.* J. Differential Geom. 1991, no. 2, 539–560

[W] Jeff Weeks, *SnapPea: a computer program for creating and studying hyperbolic 3-manifolds*, available by anonymous ftp from geom.umn.edu.