Topology and its Applications 158 (2011) 325–343

Contents lists available at ScienceDirect

Topology and its Applications

www.elsevier.com/locate/topol

Generating disjoint incompressible surfaces

Hugh Nelson Howards

Wake Forest University, Department of Mathematics, Winston-Salem, NC 27109, United States

ARTICLE INFO

Article history:
Received 1 February 2008
Accepted 6 November 2010

Keywords:
Incompressible surfaces
Boundary compressible
Tunneling
Parallel surfaces

ABSTRACT

We show that one can embed an arbitrarily large collection of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces in any compact, orientable 3-manifold with at least one boundary component of genus greater than or equal to two.
We also provide an answer to the open question, Question III.16, from Jaco's book, Lecture Notes on 3-Manifold Topology.

1. Incompressible surfaces in 3-manifolds with a boundary component of genus $n \geqslant 2$

The main result of this paper, Theorem 4.1, is to show that one can embed an arbitrarily large collection of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces in any compact, orientable 3-manifold with at least one boundary component of genus greater than or equal to two. This contradicts Theorem III.24 of Jaco's book [5]. We also provide an answer to the open question, Question III.16, from the same book.

If M does not contain a boundary component of genus greater than or equal to two then with the exception of annuli, any incompressible surface is also boundary-incompressible. The work of Kneser (1929) shows that for any triangulation of a compact 3-manifold the number of pairwise disjoint, distinct, normal surfaces is bounded. This in turn leads to the conclusion that for any compact 3-manifold the number of pairwise disjoint, incompressible, boundary-incompressible, non-parallel surfaces is bounded. The number of disjoint, non-boundary parallel, incompressible annuli in a manifold is also bounded by Jaco [5] and later Freedman and Freedman [1]. This shows that if all the boundary components of M are spheres or tori then the number of disjoint non-boundary parallel, non-parallel incompressible, surfaces can be bounded. Thus compact, orientable 3-manifolds contain arbitrarily many such surfaces if and only if they have a genus two or greater boundary component.

Freedman and Freedman's result contrasts nicely with Theorem 4.1. They show that in a given compact, orientable 3-manifold M and given a number n, one can bound the number of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces if all the surfaces are required to have Euler Characteristic at least n. Thus the surfaces in our paper must have larger and larger negative Euler Characteristic.

Note that our result contradicts Theorem III.24 of Jaco's book [5], which claims that you cannot embed an arbitrarily large collection of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces in any compact, orientable 3-manifold. He says it "is a new theorem due to P. Shalen and [himself]." He states it in general, but only proves a special case. The special case he proves requires that all the surfaces are annuli. Theorem III.24 is true in this special case and the proof in [5] is correct. The annulus result also follows from Freedman and Freedman [1] (although Shalen and Jaco's theorem predates Freedman and Freedman). The first counter example to Theorem III.24 was found in William Sherman's unpublished PhD thesis at UCLA [6]. He showed that the theorem fails in the manifold that is obtained by taking a genus two surface crossed with an interval. His work shows that the special case of the two holed torus crossed with the unit interval admits

E-mail address: howards@wfu.edu.

0166-8641/$ – see front matter
doi:10.1016/j.topol.2010.11.008

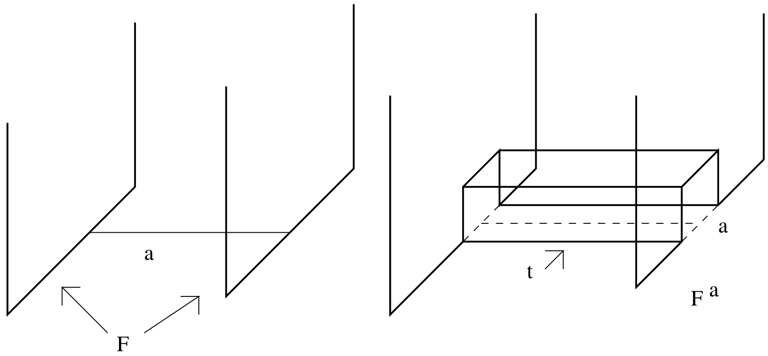

Fig. 1. Tunneling is the inverse of boundary compressing. Here the surface on the right is obtained from the surface on the left by tunneling along arc *a*.

the embedding of an arbitrarily large family of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces. His argument is long and requires many detailed steps quite specific to his chosen manifold. Our argument is more direct and shares little in common with Sherman's argument.

Interestingly, Theorem 4.1 also contrasts with, but does not contradict the result from [4] that it is impossible to embed an infinite number of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces in any compact, orientable 3-manifold M.

This paper also answers an open question from Jaco's book. Question III.16 asks if a genus 2 handlebody contains a separating incompressible surface of arbitrarily high genus. We show in Theorem 7.1 that the answer is yes, and that the answer is also yes for any manifold with a boundary component of genus 2 or higher.

A new algebraic result follows from this theorem. Corollary 7.2 states that the free group on two generators may be split into a free product with amalgamation over two arbitrarily large free groups.

Section 2 gives definitions. Section 3 is the heart of the paper. In it we establish numerous situations in which one can create a new incompressible surface by tunneling an incompressible surface to itself. Section 4 introduces two related graphs from [1] and [6] that are used in the construction. Section 5 contains the algorithm used to construct the surfaces. Sections 6 provides a technical proof necessary to start the algorithm. Section 7 answers Question III.16 from Jaco. Finally Appendix A provide technical details from known results necessary for proofs in this paper.

The author would like to thank the referee for many excellent comments which strengthened the paper's presentation and readability.

2. Definitions

In this section we review a few definitions which can be found in most introductory texts on 3-manifolds. We rely heavily on Hempel's versions in [3] and Jaco's definitions in [5]. From this point on when we refer to a surface in a three-manifold, it will be a properly embedded compact, orientable surface unless otherwise noted. Taken from Jaco, two surfaces F_1 and F_2 in a 3-manifold M are parallel if there exists an embedding $\Gamma(F_1 \times I) \to M$, such that $\Gamma|F_1 \times \{0\} : F_1 \times \{0\} :\to F_1$ and $\Gamma|F_1 \times \{1\} : F_1 \times \{1\} :\to F_2$ are homeomorphisms and $\Gamma|\partial F_1 \times I : \partial F_1 \times I \to \partial M$ is an embedding. A surface F_1 in a 3-manifold M is boundary parallel if there exists an embedding $\Gamma(F_1 \times I) \to M$, such that $\Gamma|F_1 \times \{0\} : F_1 \times \{0\} :\to F_1$ is a homeomorphism and $\Gamma|F_1 \times \{1\} \cup \partial F_1 \times I : F_1 \times \{1\} \cup \partial F_1 \times I \to \partial M$ is an embedding.

We recall the definition of boundary compressible taken directly from Jaco [5]. A surface F is boundary compressible in a three-manifold M if either

1. F is a disk and is parallel to a disk in the boundary of M or
2. F is not a disk and there exists a disk $D \subset M$ such that $D \cap F = \kappa$, an arc in ∂D, and $D \cap \partial M = \mu$ is an arc in ∂D with $\mu \cap \kappa = \partial\mu = \partial\kappa$ and $\mu \cup \kappa = \partial D$, and either κ (μ) does not separate F ($\partial M - \partial F$) or κ (μ) separates F ($\partial M - \partial F$) into two components and the closure of neither is a disk. (See Fig. 1.)

Otherwise, F is boundary incompressible.

Because there are no essential arcs on a disk, a disk cannot technically be boundary compressed. There is, however, an obvious generalization of boundary compressing where one eliminates the requirement above that the closure of neither component is a disk. We shall refer to the general version as disk splitting. This concept will be useful at times for replacing a disk we are given with a simpler one.

Tunneling is just the inverse of boundary compressing. Let F be a (not necessarily connected) surface. Let a be an embedded arc contained in ∂M with $a \cap F = \partial a$. Now choose an embedded band $B = I \times I \subset \partial M$ such that $a = \frac{1}{2} \times I$ and $B \cap F = I \times 0 \cup I \times 1$. Let $F^{a'}$ be the band connect sum of F with itself along B, that is, $F^{a'} = F \cup B$.

Let $A = I \times I \times I = B \times I$. Then $A \cap F = (B \cap F) \times I = (I \times \{0\} \times I) \cup (I \times \{1\} \times I)$. Let $t = (\{0\} \times I \times I) \cup (B \times \{1\}) \cup (\{1\} \times I \times I)$. Then we say the properly embedded surface $F^a = F \setminus (B \cap F) \cup t$ is F tunneled to itself along a (see Fig. 1). The image of

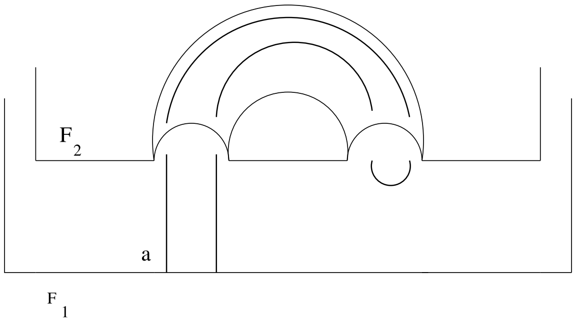

Fig. 2. This tunnel is not essential, and F_1^a is compressible even if F_1 is not. Note that in all of these pictures we identify the right side of the figure with the left, so the horizontal line for F_1, for example, depicts a circle and the region of the boundary between F_1 and F_2 here is a three times punctured sphere.

$B \subset F^{a'}$ in F^a after the isotopy is referred to as t, the tunnel added to F to form F^a. Since up to isotopy an arc a uniquely determines a tunnel t and vice versa we will also sometimes refer to F^a by F^t when it is more natural to use the tunnel in the superscript instead of the arc.

Let a be called a trivial extension arc for F if there is an isotopy of a in ∂M that takes a into ∂F leaving the end points of a fixed throughout the isotopy (see Fig. 2). Otherwise a is said to be non-trivial.

If a has both end points on the same boundary component of F then the new tunnel is said to result in a splitting of ∂F since F ends up with one more boundary component than it previously had. If a has end points on different boundary components of F then the new tunnel is said to result in a gluing of ∂F because F ends up with one less boundary component than it originally had. (Note we will deal almost exclusively with connected surfaces, but these properties hold even if F is not connected.)

Given a surface F in a manifold M, we define M', the manifold obtained from M by splitting M along F to be $M' = M - N(F)$ where $N(F)$ is an open regular neighborhood of F with $N(F) = F \times (0,1) \subset F \times [0,1]$. $F' = F_1' \cup F_2'$ is a two component surface in $\partial M'$ corresponding to the two surfaces $F \times 0$ and $F \times 1$.

3. Incompressibility

This section contains the heart of the paper. We prove a sequence of lemmas showing that certain tunnels can be added to incompressible surfaces without making them compressible. To prove the theorems in this paper it suffices to prove them for irreducible 3-manifolds, so from this point on all 3-manifolds in proofs will be assumed to be irreducible unless otherwise specified.

Let F_1 and F_2 be distinct, properly embedded incompressible surfaces in a compact orientable three manifold M. Assume F_1 and F_2 each have a boundary component contained in T, a boundary component of M of genus at least 2. Let F_2 have a boundary compressing disk B_2 ($\partial B_2 \subset F_2 \cup T$). Let a be an arc in T which is disjoint from F_2 and intersects F_1 in exactly the end points of a (i.e. an extension arc for F_1 that is disjoint from F_2). We say a is an essential arc with respect to B_2 if there is no isotopy of a in T that makes a disjoint from B_2, but leaves its end points fixed throughout the isotopy. See Fig. 2 for an example of an arc that is not essential. All other extension arcs pictured in the paper are essential.

Conjecture 3.1. *Given F_1 and B_2 as above, if a is essential with respect to B_2 then F_1^a is incompressible.*

This conjecture was made by Mike Freedman and we believe it is true, but combinatorial arguments such as those employed in this paper become more and more complicated as the number of intersections of a with B_2 increases. We instead will prove the specific cases listed below.

Throughout this section, let F_1 and F_2 be incompressible and let B_2 be the boundary compressing disk as pictured in Fig. 3 for F_2 (and B_1 the boundary compressing disk for F_1 if such a disk exists where the boundary of B_i consists of one arc κ in F_i, (or F_i^a) and one arc μ on ∂M). Assume that in Fig. 4 the parallel, horizontal lines represent annular regions, so the boundary region (the portion of $T \subset \partial M$) between ∂F_1 and ∂F_2, containing a is always either a 3 or 4 times punctured sphere depending on no ∂-compression on F_1 or a ∂-compression on F_1 respectively.

In this section we will generally assume F_1^a is compressible and that D is a compressing disk for F_1^a that intersects one of B_1 or B_2 minimally. In each case we will show that this implies that D is disjoint from B_1 (or B_2) and therefore that F_1 is also compressible, yielding a contradiction.

We are intersecting two disks and a standard innermost loop argument shows that there are no simple closed curves of intersect in any of the arguments below, so we will assume that $D \cap B_1$ (or $D \cap B_2$) consists only of arcs.

Lemma 3.2. *F_1^a in Fig. 3, and $F_1^{a_1}$, $F_1^{a_2}$, and $F_1^{a_3}$ in Fig. 4 are incompressible.*

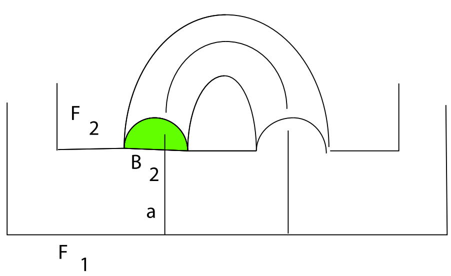

Fig. 3. Adding a tunnel to F_1.

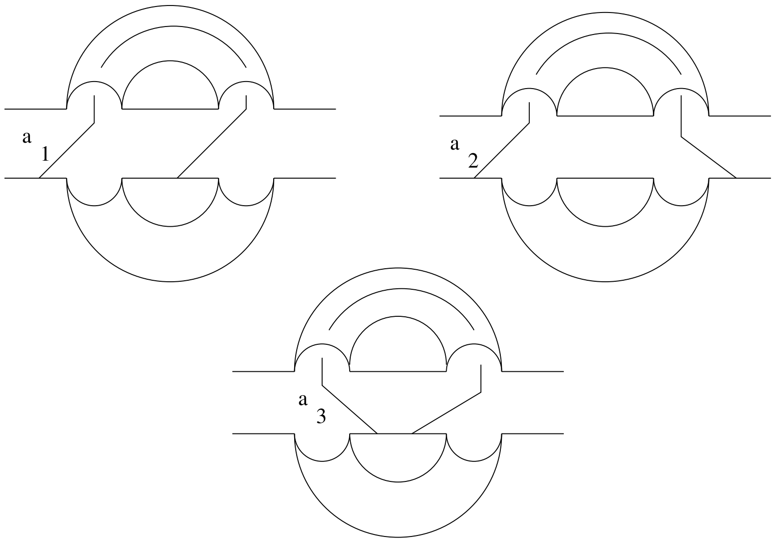

Fig. 4. Three similar options for adding tunnels.

Proof. In Fig. 3 the boundary region is a three times punctured sphere and in each of the examples in Fig. 4 the boundary region is a four times punctured sphere, but the same proof works in all cases. For simplicity's sake we refer to F_1^a in Fig. 3 and $F_1^{a_i}$ in Fig. 4 simply as F_1^a in this proof. $B_2 \cap F_1^a$ consists of a single, properly embedded arc, call it γ_1.

Suppose D were a compressing disk for F_1^a and among all possible such disks, the disk D is chosen so that the number of components of $D \cap B_2$ is minimal. Consider $D \cap B_2$. We shall show that if D were to exist, then F_1 would necessarily not be incompressible. By the choice of D there are no simple closed curve intersections; hence, all intersections are arcs properly embedded in D. For a contradiction, assume there is an arc of intersection. There must be an arc on exactly one of the components of $B_2 - \gamma_1$ (since F_1^a is orientable, arcs of intersection can only be on one side of F_1^a and thus on only one side of γ_1). Call the closure of the component of $B_2 - \gamma_1$ containing the arc of intersection B_2'. Choose an outermost arc of $B_2' \cap D$ on B_2'. Disk splitting D using this subdisk produces two disks D' and D'', both properly embedded in F_1^a. At least one of these disks must have a boundary that is essential in F_1^a since ∂D was essential in F_1^a. This, however, produces a compressing disk for F_1^a that intersects B_2 fewer times than D does. This contradiction shows that D must be disjoint from B_2 and in turn that D may be chosen to have its boundary disjoint from the tunnel added along a. Thus, $\partial D \subset F_1$ and therefore D is also a compressing disk for F_1. Since F_1 is incompressible this is a contradiction showing that no compressing disk can be found for F_1^a and therefore that F_1^a must be incompressible. $\square$

As stated above, we let F_1 and F_2 be incompressible.

Lemma 3.3. *The gluing, F_1^a, is incompressible in Fig.* 5.

Proof. Suppose F_1^a is not incompressible and D is a compressing disk for F_1^a such that D has a minimal number of intersections with B_1, the boundary compressing disk for F_1 (note we used a boundary compressing disk for F_2, not F_1 in the previous lemma). By definition the boundary of B_1 consists of one arc κ in F_1, (or F_1^a) and one arc μ on ∂M. Because t, the tunnel added along a passes once through B_1, we know $B_1 \cap F_1^a$ consists of two arcs; one is κ and the other is a single properly embedded arc γ_1 with end points on μ as in Fig. 6 and there are no simple closed curves in the intersection. γ_1 splits t into two parts. If an arc of $D \cap B_1$ runs from γ_1 to itself (or κ to itself), then the argument from the previous lemma again leads to a contradiction.

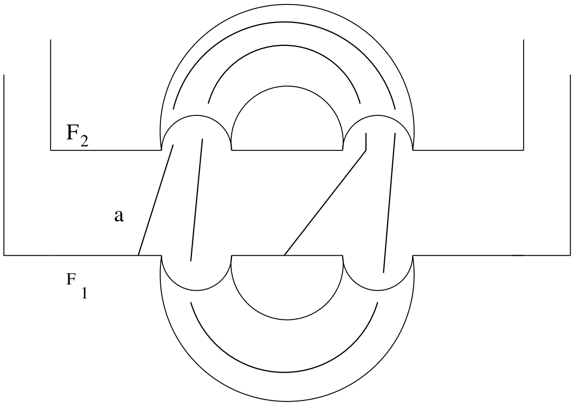

Fig. 5. A gluing of F_1.

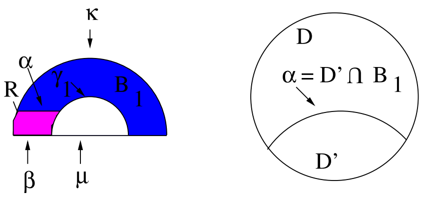

Fig. 6. The arcs α and β are depicted on B_1 on the left and α is shown on D on the right.

Thus, all arcs in $D \cap B_1$ run from κ to γ_1. Choose an outermost arc α on D. α cuts off a subdisk D' of D. $\partial D'$ consists of the arc α and an arc in F_1^a. The arc in F_1^a runs from one end point of α in γ_1 along F_1^a leaving γ_1 along one of the two components of $t \setminus \gamma_1 = t \setminus (t \cap B_1)$ and eventually returning along F_1 to the other end point of α in κ. $\mu \subset \partial B_1$ is split into three pieces by $\partial \gamma_1$. Let β be the sub-arc of μ that runs from F_1 to γ_1 and that is parallel to α on the closure of the component of $B_1 - \gamma_1$ containing α as in Fig. 6 (there are actually two arcs parallel to α and it does not matter which one we use. Also since α is outermost on D, it does not matter if α was outermost on this subset of B_1).

Now $\partial D'$ is essential on $(F_1^a)^\beta$ (F_1^a tunneled to itself along β). It remains so on the surface that results if we boundary compress $(F_1^a)^\beta$ on the component of $t - \gamma_1 = t - (t \cap B_1)$ that D' does not cross. The combination of adding the two tunnels and then boundary compressing is, of course, equivalent to adding just one tunnel t' to F_1 that runs along β and then along the appropriate component of $a - (a \cap B_1)$. D' shows that $F_1^{t'}$ also must be compressible. There are four options (really only three since two of the options are isotopic) for what the $F_1^{t'}$ initially looks like, depending on which subarc of μ was chosen for β and which component of $t - (t \cap B_1)$ was compressed. t' must be the result of tunneling along one of the arcs a_1, a_2 or a_3 in Fig. 4. In each case $F_1^{a_i} = F_1^{t'}$, however, is incompressible by Lemma 3.2 above. Thus D' and in turn D cannot exist and F_1^a is incompressible completing the proof. □

Lemma 3.4. *The splitting of ∂F_1^a, is incompressible in Fig. 7.*

Proof. Examine how t intersects B_1. We label the intersections γ_1, γ_2, and γ_3 in the order they occur as we traverse t as in Fig. 8 (note this is not the same order as would be obtained by just reading them left to right along the boundary of B_1).

Assume F_1^a is compressible and choose a compressing disk D for F_1^a with a minimal number of intersections with B_1. Again there are, of course, only arcs of intersection in $D \cap B_1$ and no simple closed curves. Label the ends of the arcs, 1, 2, 3, or κ, depending on whether they lie on the tunnel at component γ_1, γ_2, or γ_3 or on the surface F_1 respectively.

Because D intersects B_1 minimally, it may be assumed that $D \cap t$ is a set of parallel disjoint arcs that run across the tunnel t. Since a tunnel is simply connected, these arcs are, of course, unique up to isotopy relative to their end points. We have labeled $\partial D \cap B_1$ with 1, 2, 3, or κ based on whether the point is on γ_1, γ_2, γ_3, or F_1. We may thus conclude that as we read the labels off clockwise around the boundary of D as it intersects B_1, the result is a word made up exclusively of a random collection of the strings “123” (from running one way across the tunnel), “321” (from running the other way) and “κ” (from the intersections that are not on the tunnel). One example is “$123\kappa\kappa321321\kappa$.” Because there is a unique way to traverse the tunnel (up to direction), we never see strings such as 132 or $12\kappa3$ on ∂D.

Now let us examine the arcs of intersection of $D \cap B_1$.

Claim 3.5. *No arc can have end points with the same label.*

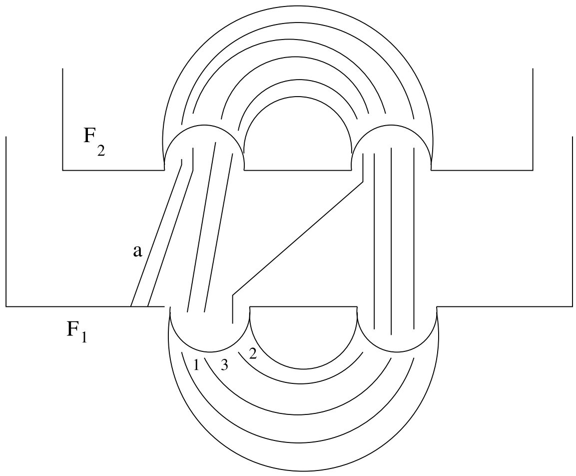

Fig. 7. A splitting of ∂F_1.

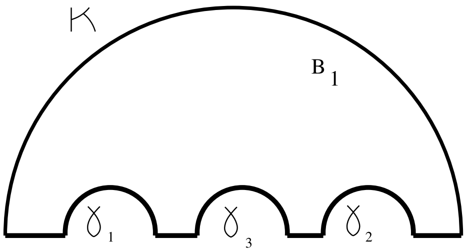

Fig. 8. Labeling γ_1, γ_2, γ_2, and κ on the boundary-compressing disk, B_1.

Proof. Assume there is an arc on B_1 with both endpoints with the same label. As before we could find an outermost arc on B_1 along which to boundary compress, yielding a compressing disk with fewer intersections with B_1 than D has, a contradiction. □

It is clear that an outermost arc of $D \cap B_1$ on D must connect adjacent labels on D (such possibilities are limited to $(\kappa, 1)$, $(1, 2)$, $(1, 3)$, $(3, 2)$, or $(3, \kappa)$), so an outermost arc α (again cutting off a subdisk D' of D) could never connect a 2 and a κ. Outermost arcs also cannot have label $(\kappa, 1)$, as we could then use the same end game we employed in the previous lemma, using the arc β, that runs along the boundary of B_1 from F_1 to γ_1, forming $(F_1^a)^\beta$, then boundary compressing $(F_1^a)^\beta$ along the portion of the tunnel that resulted from a that is missed by $\partial D'$, leading to the same contradiction of Lemma 3.2.

The argument is similar, but slightly more subtle to show that an arc cannot be labeled 3 and κ. If there were such an outermost arc we could call the corresponding outermost portion of $\partial D' \subset \partial D$, s. Now, s by definition misses γ_1 and γ_2. We replace F_1^a by a new surface by first boundary compressing F_1^a along a boundary compressing disk that runs once over γ_3 and once over the boundary of the manifold, which means there now exists an isotopy between the newly boundary compressed surface and F_1 (but we do not do the isotopy now, we will only do it to one portion of the compressed tunnel). Take an isotopy that removes the portion of the negated tunnel that was disjoint from s to create a surface that we will call G_1^a and that looks like Fig. 9 (G_1^a is, of course, as we just observed isotopic to F_1, not F_1^a, since we have added a tunnel to F_1 to form F_1^a and then done the inverse by boundary compressing returning to F_1, up to isotopy). G_1^a is therefore, incompressible.

Create $(G_1^a)^\beta$ by tunneling along an arc β that runs from κ to γ_3 as in Fig. 10. Since the portion of the tunnel containing γ_2 and γ_1 was removed via boundary compressing and isotopy there are two choices for β on B_1 with the interior of β disjoint from G_1^a and either will work fine. In Fig. 10 we arbitrarily make a choice for β, but the other choice would work, too, just as in Lemma 3.3. Now D' yields a compressing disk for $(G_1^a)^\beta$ pictured in Fig. 10 but $(G_1^a)^\beta$ is incompressible by Lemma 3.2.

In previous examples we tunneled first and boundary compressed second to get our contradiction. This time the added complication that we had to do the boundary compression first creates no problems since $s \subset (G_1^a)$ and $s \subset \partial D' \subset (G_1^a)^\beta$ and $\partial D'$ is essential in $(G_1^a)^\beta$ giving a contradiction to the existence of D.

We will now show that ζ cannot have ends labeled 2 and 3. Let ζ be the subset of ∂D cut off by the outermost arc α on D and contained in $\partial D'$ (the subdisk whose interior is disjoint from B_1). Now $\zeta \cup \alpha = \partial D'$. If it had such labels, up to isotopy there is only one arc connecting points labeled 2 and 3 on F_1^a and disjoint from γ_1 (recall the interior of ζ is

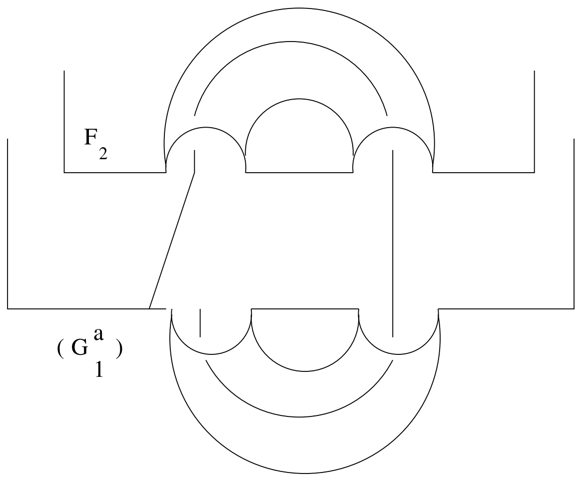

Fig. 9. After boundary compressing at γ_3 and an isotopy we have a surface G_1^a that is isotopic to F_1.

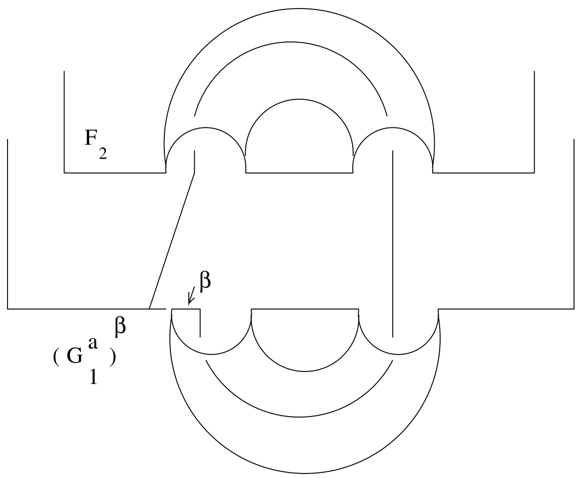

Fig. 10. Tunneling G_1^a to itself along arc β results in a surface equivalent to one in Fig. 4.

disjoint from B_1 and thus misses all points labeled 1). In this case, let $\beta \subset \partial B_1$ be the arc that runs from γ_2 to γ_3. Now $\partial D' = \zeta \cup \alpha$ not only can be embedded on $(F_1^a)^\beta$, but may be embedded on the annulus A in Fig. 11 because β and ζ are both on this annulus. (Note that A is isotopic to the annulus that results if we boundary compress $(F_1^a)^\beta$ at both ends where the tunnel corresponding to a attaches to F_1.)

An easy minimal intersection argument with B_2 shows the annulus is incompressible proving that the outermost arc cannot have ends labeled 2 and 3. ζ also cannot have labels 1 and 2 because similar to the argument above, this implies the annulus in Fig. 12 is compressible, which again easily leads to a contradiction (clearly no essential curve on the annulus bounds a disk disjoint from F_2).

Thus, all outermost arcs must be labeled 1 and 3. By assumption, D must run across t, so there must be arcs on D with one end point labeled 2. Let $\{\rho_i\}$ be the collection of such arcs. Without loss of generality let ρ_1 be an outermost arc on D within that collection (ρ_1 is, of course, NOT outermost on D if all arcs of intersection are included). ρ_1 breaks D into two pieces. Because it is outermost among $\{\rho_i\}$ the outer subdisk cut off by ρ_1, D' must have no 2's inside of it. Therefore the string of labels on D' running from one end point of ρ_1 to the other must be a subset consisting of the first several labels (in order) of one of the following strings $21\kappa\kappa\ldots\kappa12$, $23\kappa\kappa\ldots\kappa32$, $21\kappa\kappa\ldots\kappa32$, $23\kappa\kappa\ldots\kappa12$, 2112, 2332, 2132, or 2312. Note that an outermost arc must exist on D' that connects adjacent labels. Of all the possible strings and subsets, only the final two strings 2132 and 2312 have labels 1 and 3 adjacent to each other and thus only these two options are possible. This, however, implies that ρ_1 has both endpoints labeled 2 violating Claim 3.5. $\square$

Lemma 3.6. *F_1^a is incompressible in Fig.* 13.

Proof. $B_2 \cap F_1^a$ consists of two properly embedded arcs, call them γ_1 and γ_2. Assume D intersects B_2 minimally and look at $D \cap B_2$. As before we can have no simple closed curves of intersection and by Claim 3.5 no arcs may have both endpoints on

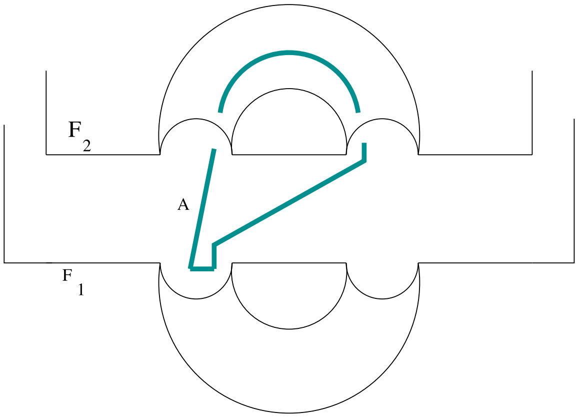

Fig. 11. An annulus that would have to be compressible if F_1 were compressible.

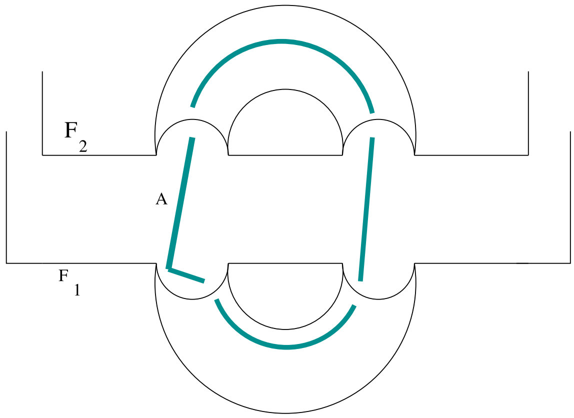

Fig. 12. An annulus that would have to be compressible if an outermost arc had labels 1 and 2.

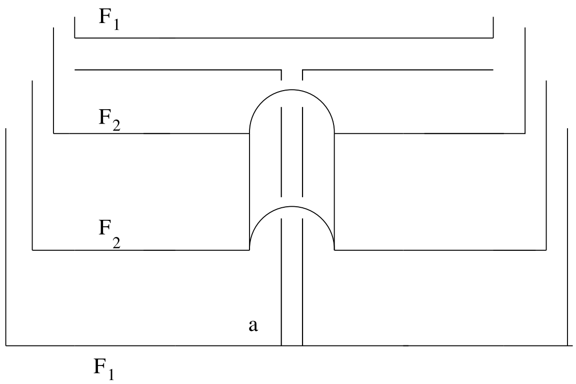

Fig. 13. A gluing of ∂F_2 becomes a splitting of ∂F_1. (Recall that the left side of the picture is identified with the right and thus a is a single connected arc.)

γ_1 or both on γ_2. Since F_2 is incompressible we may assume D is disjoint from it and therefore all arcs of intersection must have one end on γ_1 and the other on γ_2 ∂D is divided into sub-arcs by $\partial D \cap B_2$. Half of these are completely contained on t, the tunnel added to form F_1^a from F_1. Label these $\{\alpha_1, \alpha_2, \ldots, \alpha_s\}$ (see Fig. 14). These arcs must be parallel on t and have endpoints on γ_1 and γ_2. $D \cap B_2$ is a sequence of s parallel arcs on B_2 that also run from γ_1 to γ_2. Thus an outermost arc of $D \cap B_2$ on B_2 is also outermost on D (in fact every arc of $D \cap B_2$ is outermost on D).

Let β be the arc of ∂B_2 in T that runs between γ_1 and γ_2. Take an outermost arc of $D \cap B_2$ on B_2. Disk splitting D along B_2 using the outermost arc yields a compressing disk D' for $(F_1^a)^\beta$ that runs once along β and once along $\alpha_i \subset t$.

We know that $t - (\gamma_1 \cup \gamma_2) = t - (t \cap B_2)$ consists of three components, two of which are connected to F_1. If we boundary compress each of the two components of $t \subset (F_1^a)^\beta$ that connect to F_1, the result is a two component surface consisting of

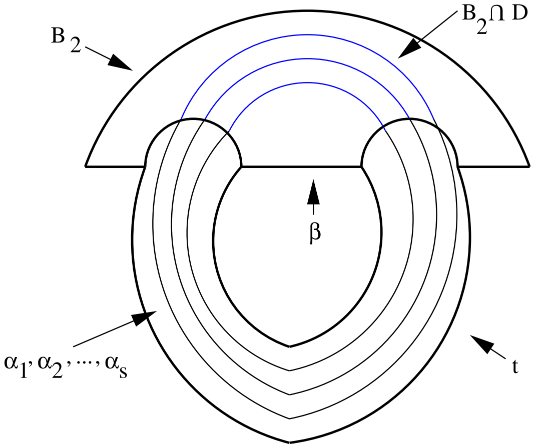

Fig. 14. Each arc of $B_2 \cap D$ cuts off an outermost arc on D that runs along t.

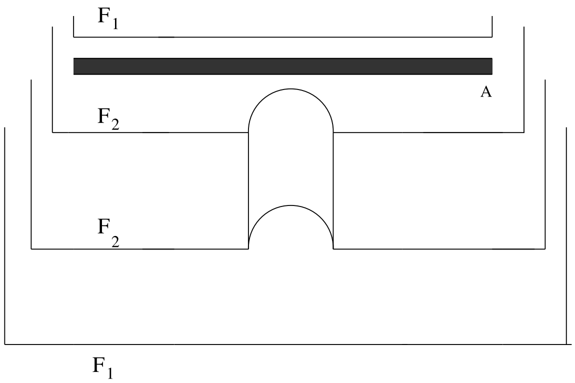

Fig. 15. An incompressible annulus.

components F_1 as well as A a boundary parallel annulus as in Fig. 15 (if we had done this to F_1^a instead of $t \subset (F_1^a)^\beta$ the result would have been F_1 plus a disk that resulted from chopping off t). D' shows that A is compressible. However this is clearly not true, as that would imply that ∂A bounds a disk, but ∂A is parallel to a boundary component of F_1. Since F_1 is not a disk this means F_1 is compressible, which is a contradiction, and the lemma is proven. □

In the final lemma of this section we will let F be a not connected surface consisting of 2 connected components F_1 and F_2. To form F^a we will tunnel F_1 to F_2 along an arc a to yield a single, connected surface F^a as in Fig. 16. F_1 and F_2 are assumed to be incompressible. The portion of $\partial M - (\partial M \cap (\partial F_1 \cup \partial F_2))$ containing a is a four times punctured sphere. This lemma is not necessary for the main theorem, but is used to resolve an open question in [5].

Lemma 3.7. *F^a is incompressible in Fig.* 16.

Proof. Assume F^a is compressible. Choose a compressing disk D for F^a that intersects B_1 minimally. Since neither F_1 nor F_2 is compressible, the disk must run across the tunnel t that we add along a to connect the two surfaces. t intersects B_1 in a single arc, which we shall call γ_1. As before, $D \cap B_1$ contains no simple closed curves and no arcs connecting γ_1 to itself or F_1 to itself by Claim 3.5. Therefore all arcs of intersection have one endpoint on γ_1 and the other on F_1. Let α be an outermost arc on D, cutting off an outermost disk D'. Let β be an arc in the portion of ∂B_1 that is contained in ∂M that is parallel to α on $B_1 - \gamma_1$ as in Fig. 6. Let R be the rectangle between α and β. Disk splitting D along R turns D' into a compressing disk for $(F^a)^\beta$.

The portion of $\partial D' \subset \partial D$ must have one end point on $\gamma_1 \subset t$ and one on F_1 (because it intersects β at those two spots) and thus must run over the component of $t - \gamma_1$ that connects to F_1. Thus boundary compressing $(F^a)^\beta$ along the component of $t - \gamma_1 = t - (t \cap B_1)$ that attaches to F_2 results in a disconnected surface consisting of F_2 and a second component that contains $\partial D'$ and thus is compressible. This component, however is equivalent to one of the surfaces in Fig. 4. Those surfaces, however, are incompressible by Lemma 3.2 yielding a contradiction, completing the proof of the lemma. □

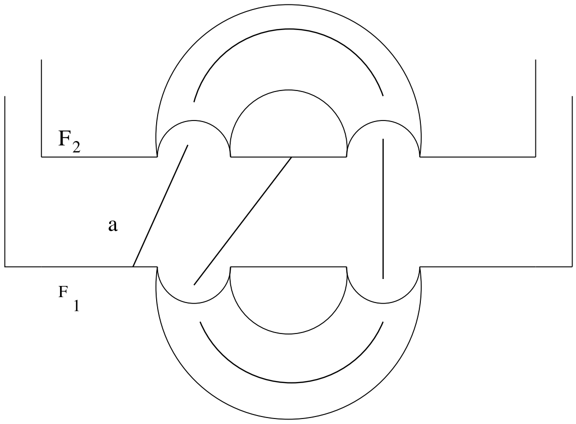

Fig. 16. This tunnel connects the incompressible surfaces F_1 and F_2 yielding the incompressible surface F^a.

4. The set up

We set out to prove:

Theorem 4.1. *One can embed arbitrarily many disjoint, non-parallel, non-boundary parallel, incompressible surfaces in a compact orientable* 3*-manifold M if and only if M has at least one boundary component, T, of genus greater than or equal to two.*

In order to generate an arbitrarily large collection of surfaces we must first start with a single non-separating surface. The existence of such a surface is well known and is often proven in a graduate 3-manifold topology class. The existence is informally known as "the half lives half dies theorem." Although it is fairly well known we could not find it in the literature so we have attached the proof in Appendix A. The result we will use follows.

Corollary A.4 *(Half Lives, Half Dies Surfaces). Let M be a compact orientable* 3*-manifold with a boundary component of genus greater than one. Then we can find a properly embedded incompressible surface $F \subset M$ that intersects a boundary component T of genus greater than one in a collection of parallel curves* (*with orientation induced from F so that they are equivalent homologically in T*).

An elementary example of such a surface is a properly embedded non-separating disk in a genus 2 handlebody. A more curious example is a non-trivial, non-separating vertical annulus in $F \times I$. In a manifold such as $F \times I$, each such surface must have boundary in a component distinct from T. Even if this happens, though it will have no effect on our proof.

Let F be a minimal genus incompressible half lives half dies surface for T, one of the boundary components of M of genus greater than or equal to two (Corollary A.4), such as the disk mentioned above. Let $F_1, \dots, F_n$, be n pairwise disjoint parallel copies of F, where n is at least the number of surfaces we want to create. In the case of the handlebody above we would just have n parallel copies of the non-separating disk. We will add a series of tunnels to the surfaces in such a way that we leave them disjoint and incompressible, but no longer parallel.

We construct two graphs that help to keep track of the details of the proof including how many non-parallel copies of the surface we have constructed. The interior graph $\mathcal{G}$ is used to examine the three-dimensional regions into which the surfaces cut the three manifold. It is the graph defined in Section 2 of Freedman and Freedman [1] with the labeling scheme slightly simplified.

Let $\{F_1, \dots, F_n\}$ be the set of surfaces. The closure of a component of $M \setminus \{\bigcup F_i\}$ corresponds to a vertex of $\mathcal{G}$; the set of vertices is designated $\{v_i\}$. Each F_j yields an edge e_j of $\mathcal{G}$ joining the two vertices (or possibly one vertex to itself) that correspond to the regions of $M \setminus \{\bigcup F_i\}$ that contain F_j on their boundary.

The vertices of $\mathcal{G}$ are labeled N, P, C_g or C_s. In $\mathcal{G}$ a vertex $v_i \in N$ if the closure of $M \setminus \bigcup F_i$ corresponding to v_i is not a product region; $v_i \in P$ if the closure of $M \setminus \bigcup F_i$ corresponding to v_i is homeomorphic to $F_i \times I$ for some F_i; $v_i \in C_g$ or $v \in C_s$, which preempts an N if the closure of $M \setminus \bigcup F_i$ corresponding to v_i is not a product region, but there is a single boundary compression that turns the N into a P. We use C_g, called a self-gluing cusp, if the boundary compression increases the number of boundary components, and C_s, called a splitting cusp, if it decreases the number. Initially since the surfaces are parallel we will have $n-1$ vertices labeled P (for the regions between F_i and F_{i+1}, $1 \leqslant i < n$) and one labeled N for the region between F_n and F_1 (see Fig. 17). One should note that F_1 and F_n are parallel surfaces, but the product between them occurs only in one component of $M - (F_1 \cup F_n)$ and the vertex labeled N corresponds to the other component. We also note that throughout the evolution of the surfaces, the interior graph remains a circle, only the labeling changes. This is true because our tunnels always connect a surface to itself, either splittings or self-gluings, and neither of these affect the topology of the graph.

Our objective is to send a cusp in each direction out of the N vertex by adding tunnels. The first cusp is advanced by adding a tunnel to F_1 through the region labeled N, then F_2 through the new cusp region and so on. The other cusp is

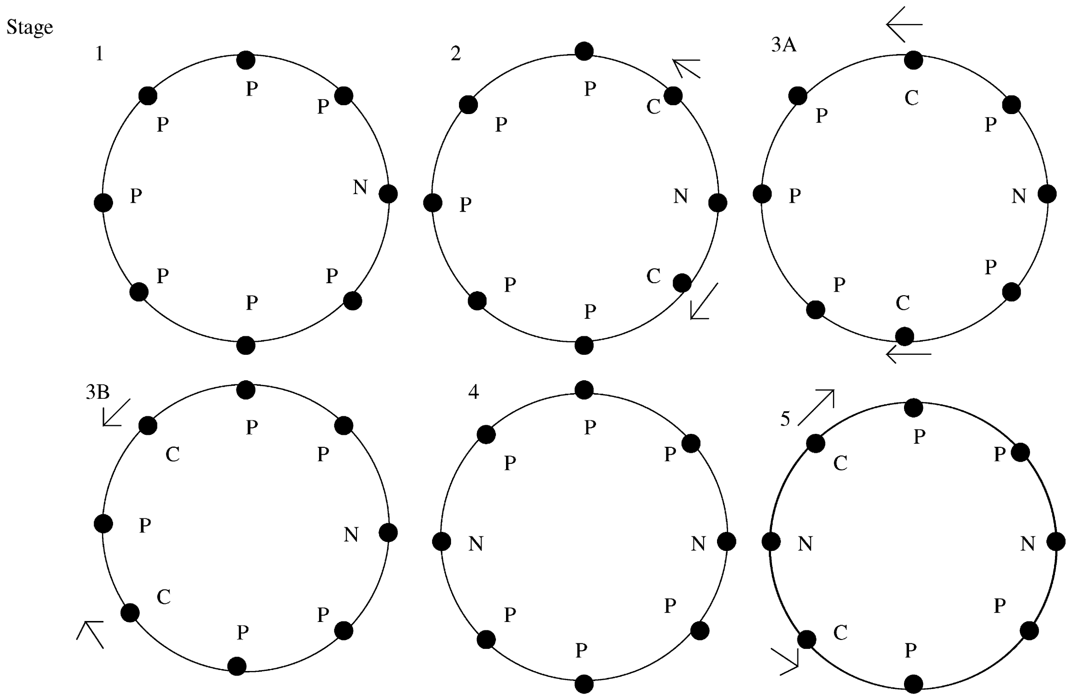

Fig. 17. The interior graph $\mathcal{G}$.

advanced by adding a tunnel to F_n through the region labeled N, then adding a tunnel to F_{n-1} through this new cusp region and so on. The cusps will travel in opposite directions around the graph until they arrive at the same vertex. This occurs when we have added a tunnel in one direction to F_i and another to F_{i+1} in the other direction. At this time we say the cusps collide since they have just arrived at the same vertex and changed that vertex's label into an N. We repeat the process carefully making sure the new N remains an N. We are free to add more tunnels in one direction than the other if we like so the cusps may be forced to meet between F_t and F_{t+1} for any t, $1 \leqslant t < n$. Each time we repeat the process we gain another N on the interior graph. In turn we are gaining another non-parallel surface (parallel surfaces must have a P vertex between them). See Fig. 17.

The boundary graph is the two-dimensional analog of the interior graph. It is used to examine the way the surfaces meet the boundary. It is essentially the same graph used in Sherman [6]. To define the graph, examine the way the surfaces meet T (the selected genus $\geqslant 2$ boundary component of M). Let the curves $\{T \cap F_i\}$ be $\{c_{ij}\}$. The closure of a component of $\{T \setminus \bigcup_i \{c_{ij}\}\}$ corresponds to a vertex. Each element of $\{c_{ij}\}$ corresponds to an edge joining the vertices that are assigned to the regions on either side of the curve. To ensure that tunneling does not affect the topological type of the (thickened) graph we also add r edges connecting a vertex to itself, for any vertex that corresponds to a region of the boundary with genus r. For example in the case of the genus 2 handlebody with n parallel, non-separating disks mentioned above, all but one of the boundary regions are annuli and the other is a twice punctured torus. We add on an extra loop to the vertex that corresponds to the punctured torus. In the case of a genus 3 handlebody with disks, the non-annular region would be a twice punctured genus 2 surface and we would add 2 loops.

In our application, if T has genus $m > 2$, then the boundary graph will consist of m circles, but $m - 2$ of the circles are extraneous to the proof and remain unchanged throughout the proof so we will just draw the graph and refer to it as if T had genus exactly two. Thus we start with a boundary graph that is the wedge of two circles (see Fig. 18).

In this paper the boundary graph serves a minor role. It verifies that whenever the cusps arrive at the same vertex of the interior graph to produce a region labeled N, that the boundary region created by the collision is a single four times punctured sphere and not two separate three times punctured spheres (although the interior region is known to be connected, we want to make sure the boundary region is, too, otherwise we could not add the next tunnel without creating another cusp).

5. The algorithm

With the hard work behind us we are now ready to design the algorithm that will build our arbitrarily large collection of surfaces.

Although the argument works in any manifold with a boundary component, T, of genus greater than or equal to two, everything interesting in the proofs happens in a neighborhood of T, so one may imagine M is a closed surface of genus greater than equal to 2 crossed with an interval and the surfaces are vertical, non-separating annuli without missing any

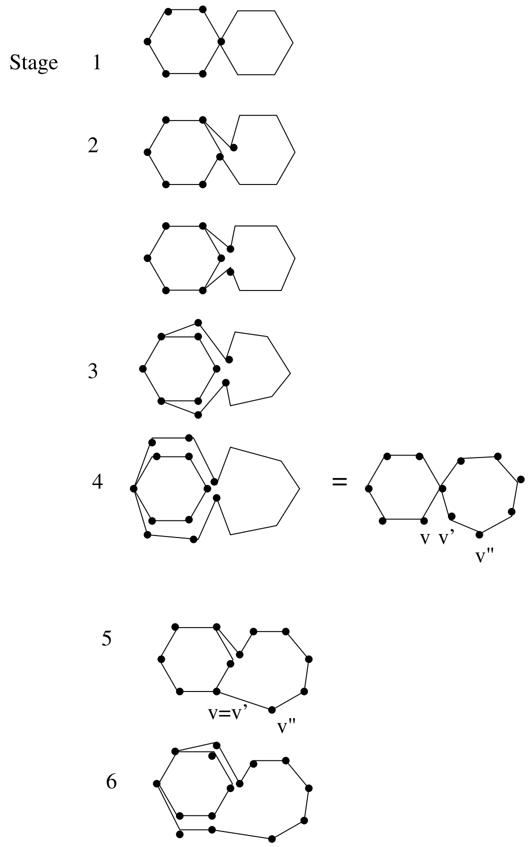

Fig. 18. The evolution of the boundary graph is pictured. As we go from stage 4 to 5 we see a gluing cusp at vertices v and v' as well as a splitting cusp above that.

major subtleties of the arguments in the paper. One could also picture the genus 2 handlebody and then initially the F_i are parallel, non-separating disks.

5.1. Step 1

Take $\{F_1, \ldots, F_n\}$, n parallel, non-separating least genus, properly embedded, incompressible surfaces as described above in Section 4. The interior graph looks like Fig. 17, Stage 1. The boundary graph looks like Fig. 18, Stage 1.

5.2. Step 2

Examine F_1 and F_n, the 2 surfaces that adjoin the region labeled N in the interior graph. We must tunnel F_1 to itself along an arc a_1 in such a way that $F_1^{a_1}$ is incompressible and disjoint from $\{F_2, \ldots, F_n\}$. We must then tunnel F_n to itself along an arc a_n in such a way that $F_n^{a_n}$ is incompressible and disjoint from $\{F_1^{a_1}, F_2, \ldots, F_{n-1}\}$ (see Fig. 20). We prove this is possible in Theorem 6.1, deferring the technical details to the next section. It is essentially an application of Theorem 10 [2]. The graphs move to Stage 2 in this step. We should note that $F_1^{a_1}$ and $F_n^{a_n}$ have the same genus. It is in theory possible that the original N from this region could turn into a P as a result of the two new tunnels. As we shall see, this is the only time in the construction when an N created in the interior graph is not guaranteed to remain an N for the rest of the algorithm.

5.3. Step 3

Add cusps to the surfaces in sequence so that as we examine the interior graph one C moves around clockwise and the other moves around counterclockwise (the figure for the interior graph moves to Stage 3) until the cusps converge on the same vertex forming a new N in the interior graph as in Fig. 17, Stage 4. We do this by adding a tunnel to F_2 that makes it parallel to $F_1^{a_1}$ and do the same to F_{n-1} with respect to $F_n^{a_n}$ as in Fig. 3. We continue for F_3 and F_{n-2} and so on. Without loss of generality let the collision occur between surfaces F_t and F_{t+1} forming surfaces $F_t^{a_t}$ and $F_{t+1}^{a_{t+1}}$. We may choose to have one of the cusps progress through $1 \leqslant t < n$ vertices by adding a tunnel to t surfaces then forcing the cusps to meet at that vertex by sending the other cusp through $n - t$ vertices by adding tunnels to the other $n - t$ surfaces in the other direction. Because we are free to pick any t we like as long as $1 \leqslant t < n$ we have quite a bit of flexibility deciding at which

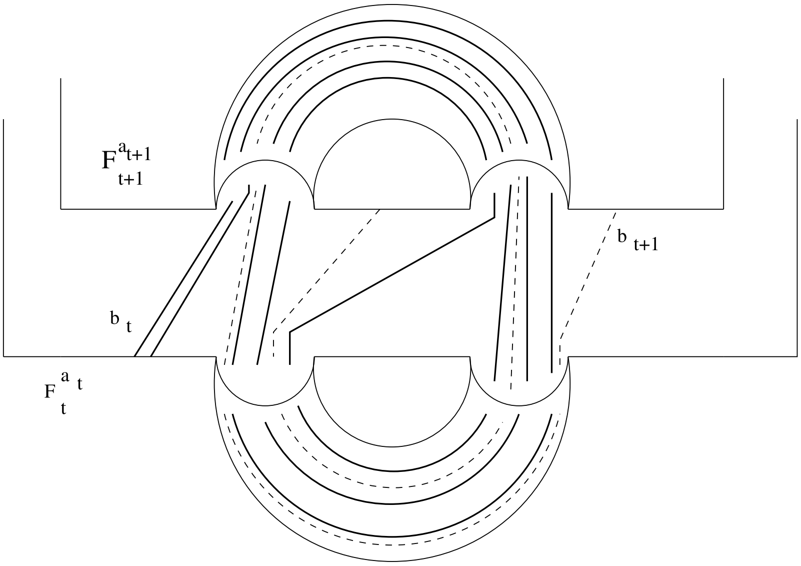

Fig. 19. When two cusps meet at the same vertex of the interior graph, we create two disjoint new surfaces that are not parallel to each other via one gluing and one splitting.

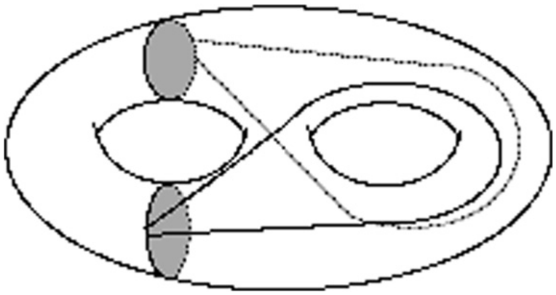

Fig. 20. The first two tunnels added. The darker path is on the top of $\partial M'$ and the lighter one is on the bottom.

vertex the cusps should meet. Note that in Stage 4 the boundary graph has returned to the wedge of two circles. These tunnels create incompressible surfaces by Lemma 3.2.

5.4. Step 4

Split the non-product region corresponding to $F_t^{a_t}$ and $F_{t+1}^{a_{t+1}}$ (the vertex of valence four in the boundary graph) by one splitting of $\partial F_t^{a_t}$ and one self-gluing of $\partial F_{t+1}^{a_{t+1}}$, along arcs we will call b_t and b_{t+1} as in Fig. 19 to form $(F_t^{a_t})^{b_t}$ and $(F_{t+1}^{a_{t+1}})^{b_{t+1}}$. $(F_t^{a_t})^{b_t}$ and $(F_{t+1}^{a_{t+1}})^{b_{t+1}}$ are two new incompressible surfaces by Lemmas 3.4 and 3.3 respectively. The interior graph moves to Stage 5. The boundary graph again looks like Stage 2. Since $F_t^{a_t}$ and $F_{t+1}^{a_{t+1}}$ have the same number of boundary components (they are obtained from two parallel surfaces by splittings), $(F_t^{a_t})^{b_t}$ and $(F_{t+1}^{a_{t+1}})^{b_{t+1}}$ do not have the same number of boundary components, as one is obtained by splitting and one by self-gluing, (the Euler Characteristic of each has gone down by one, but one now has one more boundary component than before and the other now has one less). This assures us that $(F_t^{a_t})^{b_t}$ and $(F_{t+1}^{a_{t+1}})^{b_{t+1}}$ cannot be parallel and that we have left behind a non-product region. The interior graph now has one more vertex labeled N than it did before. Continue doing the gluings and splittings that cause one cusp to move around the interior graph clockwise and the other counterclockwise. Note that if the number of boundary components of two surfaces is different, then sending a cusp through the region between the two surfaces will leave them with an unequal number of boundary components, so the N's created in the interior graph will remain N's throughout the construction.

5.5. Step 5

Some time before the cusps arrive at the same vertex of the interior graph, we may choose to convert the self-gluing cusp into a splitting cusp as in Fig. 13. The surfaces remain incompressible by Lemma 3.6. Again this results in two surfaces with a different number of boundary components and thus another P is replaced by an N in the interior graph. We convert

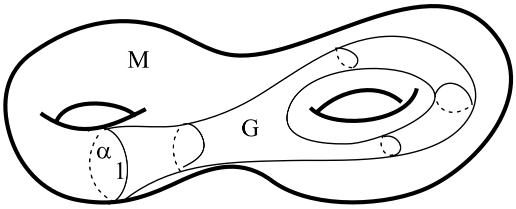

Fig. 21. Let M be the manifold with a genus 2 boundary component pictured above and G be the properly embedded once punctured torus contained in M. Then $\hat{G} = \partial M - \partial G$ is parallel into G. We will show that any time this happens it implies that G is not a least genus half lives, half dies surface.

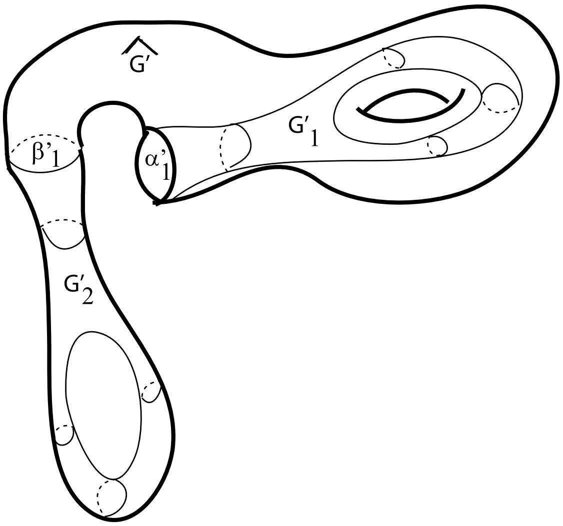

Fig. 22. We split M along G to form M' and G'. G'_1 is the image of G now on the boundary of M' that appears to be on the "inside" of M', G'_2 is the other component derived from G.

the gluing to a splitting in order to ensure that when the two cusps collide they both correspond to splittings. Continue tunneling, advancing the cusps (both of which now correspond to splittings) through the graph until they meet at a new vertex of the interior graph. This creates a new vertex labeled N. Make sure the new collision occurs at a vertex in the interior graph labeled P (this is, of course, easy to do).

5.6. Step 6

Return to Step 4. Repeat Steps 4 and 5 until there are as many regions labeled N as desired. Recall that if we have n N's then we have at least $n-1$ surfaces that are not parallel, since parallel surfaces must bound a region labeled P.

In this manner we can generate as many disjoint, non-parallel, non-boundary parallel, incompressible surfaces as we like.

6. Adding the tunnels in Step 2

Now we need only check that we can indeed do Step 2 of the proof and add the first tunnel to each of the original surfaces while leaving them incompressible and then the proof of the theorem will be complete.

Recall that we took F, a minimal genus, incompressible, half lives half dies surface for one of the boundary components of M of genus greater than or equal to two. We then took n parallel copies of F, $F_1, \ldots, F_n$. Examine the first and last copies, F_1 and F_n, the two surfaces that bound the region initially labeled N in the interior graph.

Theorem 6.1. *It is always possible to tunnel F_1 to itself along an arc a_1 in such a way that $F_1^{a_1}$ is incompressible and disjoint from $\{F_2, \ldots, F_n\}$ and to tunnel F_n to itself along an arc a_n in such a way that $F_n^{a_n}$ is incompressible and disjoint from $\{F_1^{a_1}, F_2, \ldots, F_{n-1}\}$.*

Proof. Let G be a least genus, incompressible, half lives half dies surface for T. Finding an arc a_1 along which to extend G is possible in general as we can see if we imagine splitting the manifold along G, so $G' = G'_1 \cup G'_2$ is now a two component incompressible subsurface in the boundary of M' (the manifold after splitting) as in Figs. 21 and 22. A quotient map q exists identifying G'_1 to G'_2 that takes M' back to M and G' to G. G has $m \geqslant 1$ boundary components $\{\alpha_1, \ldots, \alpha_m\}$, so G'_1 will have m boundary components $\{\alpha'_1, \ldots, \alpha'_m\}$, and G'_2 will have m boundary components $\{\beta'_1, \ldots, \beta'_m\}$ such that $q(\alpha'_i) = q(\beta'_i) = \alpha_i$.

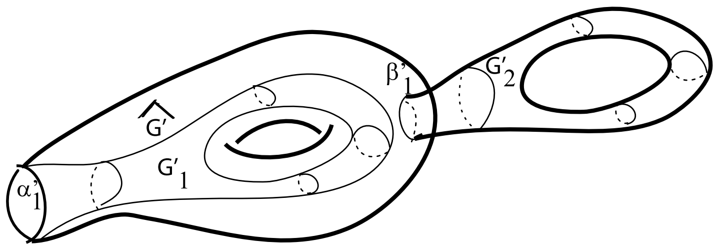

Fig. 23. We take an isotopy of M' that shows more clearly that $\hat{G}' = \partial M' - (G_1' \cup G_2')$ is parallel into G_1'.

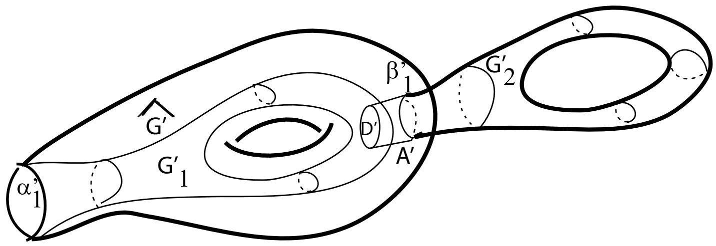

Fig. 24. We see an annulus A' in M' that runs from $\beta_m' = \partial G_2' \subset \partial M'$ to G_1' given by the parallel structure and a disk D' on G_1' whose boundary is one of the boundary components of A'. Let $D'' = A' \cup D'$.

Because ∂G divides T into one region of positive genus between α_m and α_1 and $m-1$ annuli between α_i and α_{i+1} for $1 \leqslant i < m$, $\partial G' = \partial G_1' \cup \partial G_2'$ divindes $\partial M'$ into G_1' and G_2', as well as a region of positive genus called $\hat{G}'$ between β_m' and α_1' and annuli between β_i' and α_{i+1}' for $1 \leqslant i < m$ as in Fig. 25 (that example results from T and G of genus 2 and $m = 3$).

The desired result essentially follows from Theorem 10 of [2]. Let $G', \hat{G}', M'$ be as above.

Theorem 10. *([2]) Suppose M' is not a product $G' \times I$, and suppose $\hat{G}'$ is not parallel into G'. Then $\hat{G}'$ contains an extension arc γ of G' with endpoints on any prescribed components of $\partial G'$.*

Recall that $\hat{G}'$ is parallel into G' if there is an embedding of a product $\Gamma'(\hat{G}' \times I) \to M'$ such that $\hat{G}' = \Gamma'(\hat{G}' \times 0)$, and $\Gamma'(\hat{G}' \times 1) \subset G'$. Theorem 10 [2] applies immediately if G is a disk, so we may assume throughout the proof that G cannot be chosen to be a disk. M' is clearly not $G' \times I$ since G' is not connected but M' is connected, so now we must only show that $\hat{G}'$ is not parallel into G'. If $\hat{G}'$ is parallel into G' we may without loss of generality assume that $\hat{G}'$ is parallel into G_1'.

Case 1: If the product Γ' that takes $\hat{G}'$ into G' can be chosen so that α_1' is fixed throughout the homotopy then it can be shown that G was not minimal genus (a contradiction). This is the more difficult case and is depicted with $m = 1$ in Figs. 21 through 24 as well as $m = 3$ in Fig. 25.

Let $\hat{G_1} \subset \hat{G}$ be the surface that results by deleting a small regular neighborhood of α_m from $\hat{G}$ and $\hat{G}_1' \subset \hat{G}'$ by deleting a regular neighborhood of β_m'. Let $\beta_{(m,1)}$ be the new boundary component that is parallel to $\alpha_m = q(\alpha_m') = q(\beta_m')$ in $\hat{G}$ and to all of the α_i in T, and let $\beta_{(m,1)}'$ be the analogous component in M' as in Fig. 25. If $m = 1$ we may simply let $\hat{G}_1' = \hat{G}'$, $\beta_{(m,1)} = \alpha_1 = q(\beta_m')$, and $\beta_{(m,1)}' = \beta_1' = \beta_m'$.

Since $\Gamma = q(\Gamma')$ takes $\hat{G}$ into G, fixing α_1, it certainly takes $\hat{G_1}$ (a subset of $\hat{G}$) into G, fixing α_1. Delete the image of $\hat{G_1}$ (i.e. $\Gamma(\hat{G_1} \times 1)$) from G and replace it by the annulus $A = \Gamma(\beta_{(m,1)} \times I)$ and call the resulting surface F (A' such that $q(A') = A$ is the annulus depicted in Fig. 24). We notice that since $\hat{G_1}$ has just two boundary components α_1 and $\beta_{(m,1)} = \Gamma(\beta_{(m,1)} \times 0)$ its image in G, $\Gamma(\hat{G_1} \times 1)$, will have just the two boundary components α_1 and $\Gamma(\beta_{(m,1)} \times 1)$. The boundary of A consists $\beta_{(m,1)} = \Gamma(\beta_{(m,1)} \times 0)$ and $\Gamma(\beta_{(m,1)} \times 1)$. Thus the new surface $F = (G - \Gamma(\hat{G_1} \times 1)) \cup A$ is well defined. Recall that $\partial G = \{\alpha_1, \alpha_2, \ldots, \alpha_{m-1}, \alpha_m\}$. Now $\partial F = \{\alpha_2, \alpha_3, \ldots, \alpha_{m-1}, \alpha_m, \beta_{(m,1)}\}$. These curves are parallel in T, so ∂F is homologous to ∂G in $H_1(T)$, $H_1(\partial M)$, and $H_1(M)$. Note that we needed $\beta_{(m,1)} \neq \beta_m$ for $m > 1$ so that α_m did not appear twice in ∂F preventing F from being a properly embedded surface. This issue does not occur for $m = 1$ as in that case $\partial G = \partial F = \alpha_1$ if $\beta_{(m,1)} = \beta_m = \beta_1$.

F is obtained from G by replacing a portion of G that is of positive genus by an annulus and is therefore of lower genus. Now $\text{genus}(F) = \text{genus}(G) - \text{genus}(\hat{G_1})$. Thus, F is a half lives half dies surface for M of lower genus than G, contradicting our assumption that G was minimal genus.

A specific example is given in Figs. 21 through 24. In this example we have G a punctured torus contained in M, a manifold with a genus two boundary component. The punctured torus is a half lives half dies surface, because its boundary α_1 is essential in $T = \partial M$. We apply the above general argument to show that G is not least genus in order to better understand the general proof.

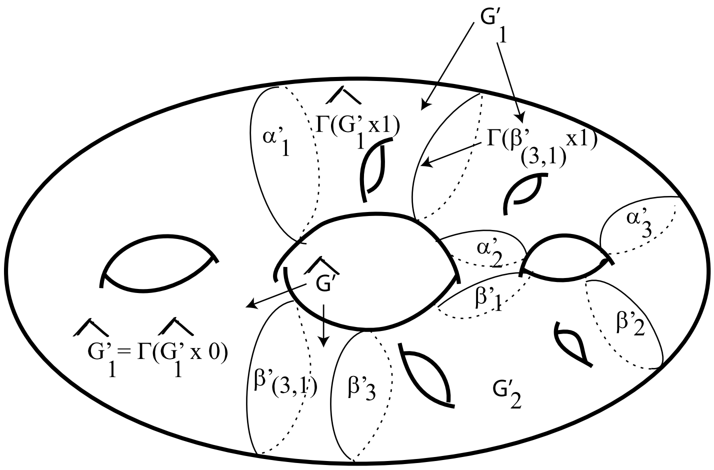

Fig. 25. We depict the boundary component of M' equal to $\hat{G}' \cup G'_1 \cup G'_2$ together with two annuli that results from splitting a manifold M with boundary component T along G, where T is genus 2 and G is a three times punctured genus 2 surface. Here we let $\hat{G}'$ be parallel into G'_1 fixing α'_1 as in the proof of case 1 of Theorem 6.1. We then find a lower genus half lives half dies surface than G in M. The simplification procedure replaces G by $q((G'_1 - \Gamma(\hat{G}'_1 \times 1)) \cup A')$, where A although not pictured is the annulus $\Gamma(\beta'_{3,1} \times I)$. The new surface is a three times punctured genus 1 surface.

We first split M along G resulting in M'. G yields the two component surface $G' = G'_1 \cup G'_2$ in M'. Since G had only one boundary component α_1, the boundary of G' consists of two curves α'_1 and $\beta'_m = \beta'_{(m,1)} = \beta'_1$. In Figs. 22 through 23 we simply take an isotopy of M' to make the fact that $\hat{G}'$ is parallel into G'_1 more clear. In Fig. 23 we see that this is indeed true (M' could be thought of as $G'_1 \times I$ with a 1-handle attached). Note that $\hat{G}'$ is parallel into a subset of G'_1, but since the two surfaces are not homeomorphic it is clear that the map is not onto. The twice punctured torus $\hat{G}'$ is taken into the once punctured torus G'_1.

We are especially interested in the image of $\beta'_m = \beta'_{(m,1)} = \beta'_1$. under the product structure. $\Gamma(\beta'_1 \times I)$ must be an annulus A' as pictured in Fig. 24 with one boundary component equal to β'_1 and the other boundary component on the interior of G'_1. In this case, this second boundary component bounds a disk $D' \subset G'_1$, and in general it will cut off a region of genus equal to genus$(G'_1) -$ genus$(\hat{G}')$. Here $A' \cup D'$ is a disk D''. Our new surface $q(D'')$ is a disk in M with boundary $q(\beta'_1) = \alpha_1$. It is, thus, of lower genus than G finishing the contradiction. In general the proof above will not always result in a disk, but it will always result in a surface of lower genus than G.

Thus, if we do choose G to be minimal genus, $\hat{G}$ will not be parallel into G leaving α_1 fixed (and $\hat{G}'$ will not be parallel into G' leaving α'_1 fixed).

Case 2: We may assume that α'_1 cannot be chosen to be fixed and thus α'_1 traces out an annulus A that is not boundary parallel. If a non-trivial extension arc γ for G' makes G'^{γ} (G' extended along γ) compressible then we can choose a compressing disk D for G' which intersects A minimally. Since G' is incompressible it is easy to see that A, which shares a boundary component with G', must be incompressible, too. $D \cap A$ must be non-trivial or else D would have to be a compressing disk for the incompressible surface G'. Since A is incompressible it is easy to eliminate both essential and inessential circles of intersection of $D \cap A$ on A.

Let γ be an outermost arc on D with respect to $D \cap A$. Since A is not boundary parallel, we see that γ cannot run from one boundary component on A to the other or it would imply that A is boundary compressible into a non-boundary parallel disk with boundary on G', showing either G' is compressible or that M is reducible. In either case we have a contradiction (it is only necessary to prove the theorem true for irreducible manifolds because if we can construct the surfaces for all irreducible manifolds, it is clear that it can easily be extended to reducible manifolds). Now we see that γ must have both endpoints on the same boundary component of A. This, however, means that we can boundary split D along γ to get at least one disk that is either a compression disk for G' or for $\hat{G}'$. Since G' is incompressible, we may assume the latter, but if $\hat{G}'$ has a compressing disk D', then G can be chosen to be a disk for a final contradiction. Thus, Theorem 10 from [2] applies showing that G' may be extended and therefore the surfaces that we are interested in, F_1 and F_n, have arcs a_1 and a_n respectively such that $F_1^{a_1}$ and $F_n^{a_n}$ are incompressible.

The proof of Theorem 10 from [2] is strong enough to show that G' can be extended along all but a finite number of extension arcs (up to isotopy) without becoming compressible, so a_1 and a_n can be chosen so that $F_1^{a_1} \cap F_n^{a_n} = \emptyset$. □

We have now checked the last detail of the proof of Theorem 4.1.

7. Separating surfaces of arbitrarily high genus

To conclude, we turn to Jaco's question: "Is there an incompressible separating surface of arbitrarily-high genus for handlebodies of genus $n \geqslant 2$?" We show that the answer is yes for $n \geqslant 2$. (Agol and Howards constructed a simple example that answered the question in the affirmative for $n \geqslant 3$ by taking parallel copies of Jaco's non-separating surface and tunneling them together, but the proof did not work for $n = 2$.) The current result applies not just to handlebodies, but also to any manifold to which Theorem 4.1 applies.

Theorem 7.1. *Any compact orientable* 3*-manifold with a boundary component of genus* $\geqslant 2$ *contains an incompressible separating surface of arbitrarily high genus.*

Proof. It is easy to verify that whether a (not necessarily connected) surface is separating or not remains unchanged under self-tunneling. Since in the previous section all of our surfaces were obtained from parallel non-separating surfaces exclusively by self-tunneling, any one of the connected surfaces is non-separating, but any pair is separating.

Recall that when we resolved a collision of cusps we did so in a manner that drove up the genus of one of the surfaces by one (a self-gluing) and that left the genus of the other unchanged (a splitting). After enough steps, we may assume that the surfaces are of arbitrarily high genus. Call the surfaces bounding the region where the cusps have met F_1 and F_2. Let $F = F_1 \cup F_2$. F is the union of incompressible surfaces and thus incompressible. F^a, the result of connecting the surfaces as in Fig. 16, is incompressible by Lemma 3.7. □

Corollary 7.2. *The free group on two generators may be split into a free product with amalgamation over two arbitrarily large free groups.*

This is a direct result of the Seifert Van Kampen Theorem. It is worthy of note that this requires the amalgamating subgroup to be of large rank.

Appendix A. Existence of surfaces

The "half lives, half dies" theorem is well known, though not universally known in 3-manifold topology. We, however, had a hard time finding it in the literature, so we include both a statement and a proof of the theorem. This particularly clean version of the proof was related by Kenneth Baker based on notes from his 3-manifold class at UT Austin with Cameron Gordon from Spring 2001.

Lemma A.1. *If*

$$0 \to V_1 \to V_2 \to \cdots \to V_n \to 0$$

is an exact sequence of finite dimensional vector spaces then

$$\sum_{i=1}^{n} (-1)^i \dim V_i = 0.$$

Proof. Set $\phi_i : V_i \to V_{i+1}$, $V_0 = 0$, $V_{n+1} = 0$, $\phi_0 = 0$, and $\phi_{n+1} = 0$.

The exactness of

$$0 \to V_1 \to V_2 \to \cdots \to V_n \to 0$$

implies the existence of the short exact sequences

$$0 \to \operatorname{Ker} \phi_i \to V_i \to \operatorname{Im} \phi_i \to 0.$$

Therefore

$$\begin{aligned} \dim V_i &= \dim \operatorname{Ker} \phi_i + \dim \operatorname{Im} \phi_i \\ &= \dim \operatorname{Im} \phi_{i-1} + \dim \operatorname{Im} \phi_i \end{aligned}$$

since $\operatorname{Im} \phi_{i-1} = \operatorname{Ker} \phi_i$. Hence

$$\sum (-1)^i \dim V_i = \dim \operatorname{Im} \phi_0 + \dim \operatorname{Im} \phi_{n+1} = 0. \qquad \square$$

Lemma A.2. *If*

$$0 \to V_1 \to V_2 \to \cdots \to V_n \to W \xrightarrow{\phi} V_n \to \cdots V_2 \to V_1 \to 0$$

is an exact sequence of finite dimensional vector spaces then

$$\dim \operatorname{Ker} \phi = \frac{1}{2} \dim W.$$

Proof. We have the exact sequences

$$0 \to V_1 \to V_2 \to \cdots \to V_n \to\to \operatorname{Ker} \phi \to 0$$

and

$$0 \to \operatorname{Im} \phi \to V_n \to \cdots \to V_2 \to V_1 \to 0.$$

By Lemma A.1 we thus have

$$\dim \operatorname{Ker} \phi = \left| \sum_{i=1}^{n} (-1)^i \dim V_i \right| \quad \text{and} \quad \dim \operatorname{Im} \phi = \left| \sum_{i=1}^{n} (-1)^i \dim V_i \right|.$$

Therefore $\dim \operatorname{Ker} \phi = \dim \operatorname{Im} \phi$. Since $\dim W = \dim \operatorname{Ker} \phi + \dim \operatorname{Im} \phi$,

$$\dim \operatorname{Ker} \phi = \frac{1}{2} \dim W. \qquad \square$$

For the following theorem we will need to recall two results:

- *Poincare–Lefschetz Duality.* Let M be a compact orientable n-manifold. Then for any coefficient group G

$$H_i(M; G) \cong H^{n-i}(M, \partial M; G) \quad \text{and}$$
$$H_i(M, \partial M; G) \cong H^{n-i}(M; G).$$

- *Universal Coefficient Theorem (with field coefficients).* For a topological pair (X, A) and field F,

$$H_i(X, A; F) \cong H^i(X, A; F).$$

Theorem A.3 *(Half Lives, Half Dies Homology). Let M be a compact orientable 3-manifold. Then with field coefficients*

$$\dim\bigl(\operatorname{Ker}\bigl(H_1(\partial M) \to H_1(M)\bigr)\bigr) = \frac{1}{2} \dim H_1(\partial M).$$

Proof. We work using field coefficients. Using the duality

$$H_i(M) \underset{\text{PLD}}{\cong} H^{3-i}(M, \partial M) \underset{\text{UCT}}{\cong} H_{3-i}(M, \partial M)$$

with the homology exact sequence of the pair $(M, \partial M)$

$$\begin{aligned} 0 &\to H_3(M) \to H_3(M, \partial M) \to H_2(\partial M) \to H_2(M) \\ &\to H_2(M, \partial M) \to H_1(\partial M) \xrightarrow{\phi} H_1(M) \to H_1(M, \partial M) \\ &\to H_0(\partial M) \to H_0(M) \to H_0(M, \partial M) \to 0. \end{aligned}$$

Lemma A.2 implies

$$\dim \operatorname{Ker} \phi = \frac{1}{2} \dim H_1(\partial M). \qquad \square$$

Corollary A.4 *(Half Lives, Half Dies Surfaces). Let M be a compact orientable 3-manifold with a boundary component of genus greater than one. Then we can find a properly embedded incompressible surface $F \subset M$ that intersects a boundary component T of genus greater than one in a collection of parallel curves (with orientation induced from F so that they are equivalent homologically in T).*

Proof. Take a collection of curves realizing a non-trivial element of $H_1(\partial M)$ in the $\operatorname{Ker}\phi$ (we get existence from Theorem A.3). Such a collection of curves must bound a surface in M, but not in ∂M. Notice that there is a properly embedded surface (possibly not connected) that meets every boundary component in (homologically) non-trivial curves in that component. Within that collection, let W be the set of surfaces that intersect T in the fewest curves. Call this Minimality Property 1.

Recall that two curves are parallel in T if they bound a product region in T. We will say a set of curves are in the same equivalence class if they are all parallel to each other. Let F' be a surface in W whose intersection with T contains the smallest number of equivalence classes possible. Call this Minimality Property 2. We will show that F' in fact contains only one equivalence class, completing the proof of the corollary.

Assume F' contains more than one equivalence class. Let T' be the closure of a component of $T - (T \cap F')$ that is not a product region. Assume F' is oriented and thus that its boundary inherits an orientation. For any oriented curve c (or collection of curves) let $[c]_T$ denote the homology class of c in $H_1(T)$ and $[c]_M$ denote the homology class of c in $H_1(M)$

Case 1: T' has one boundary component c. In this case c bounds a surface in T and thus $[c]_T = [c]_M = 0$. Thus, there must be a surface whose boundary is not homologically trivial in T but is equal to the boundary of F' minus the curve c, contradicting Minimality Property 1.

Case 2: T' has exactly two boundary components, c_1 and c_2.

Subcase a: If c_1 and c_2 are parallel on T (or are the same curve on T), then we only have one equivalence class of curves and we are done (the two curves cannot have opposite orientation or together they are trivial in $H_1(T)$ and a new surface would exist with a smaller number of boundary curves contradicting Minimality Property 1).

Subcase b: If c_1 and c_2 are not parallel on T, then either $[c_1 + c_2]_T = 0$ or $[c_1 - c_2]_T = 0$ because ignoring orientation, since $c_1 \neq c_2$ on T, $c_1 \cup c_2$ must be separating on T. If the first formula holds, we again contradict Minimality Property 1, as in case 1. If the second formula holds, we see that $[c_1]_T = [c_2]_T$, thus we can find a surface which replaces each curve from F' that is parallel to c_2 with a curve that is parallel instead to c_1 while its boundary remains in the same equivalence classes of $H_1(T)$ and $H_1(M)$ as F'. Thus we now have a new surface with the same number of boundary components as F', but with fewer equivalence classes. This is again a contradiction of Minimality Property 2.

Case 3: T' has at least three boundary components, including c_1, c_2, and c_3. In this case it may not be true that the curves are separating on T, but the banded sum of some pair is homologous to their sum in T.

Subcase a: c_1, c_2, and c_3 all correspond to different curves on T. If so, let c_1 and c_2 be the pair such that their banded sum is homologous to their sum in T. We can now take a properly embedded arc in T' with one end point on c_1 and the other on c_2 and tunnel F' along this arc. This new surface has one less boundary component than F', but since its boundary is homologous to the boundary of F' in $H_1(T)$ and in $H_1(\partial M)$, its boundary must not be null-homologous in ∂M or in T. This contradicts the fact that F' was supposed to satisfy Minimality Property 1.

Subcase b: Two of c_1, c_2, and c_3 correspond to the same curve on T. Without loss of generality assume these two curves are c_1 and c_2. Since the orientations of c_1 and c_2 do not agree with respect to T', one of them must agree with c_3, say c_1. Take a properly embedded arc in T' running from c_1 to c_3. Tunnel along this arc and we again get the same contradiction as we did in Subcase a.

This is the final contradiction, and we must not have more than one equivalence class of curves on T (i.e. $F' \cap T$ is a collection of parallel curves on T). Compress F' as much as possible to yield an incompressible surface. The surface cannot compress to an incompressible annulus parallel into T because of the induced orientation of its boundary curves. Let F be one of the connected components that results that still intersects T. This surface satisfies Corollary A.4, completing the proof. □

References

[1] B. Freedman, M. Freedman, Kneser–Haken finiteness for bounded 3-manifolds locally free groups, and cyclic covers, Topology 37 (1) (1998) 133–147.
[2] Michael Freedman, Hugh Howards, Ying-Qing Wu, Extension of incompressible surfaces on the boundaries of 3-manifolds, Pacific J. Math. 194 (2) (2000) 335–348.
[3] J. Hempel, 3-Manifolds, Ann. of Math. Stud., vol. 86, Princeton Univ. Press, Princeton, NJ, 1976.
[4] Hugh Nelson Howards, Limits of incompressible surfaces, Topology Appl. 99 (1) (1999) 117–122.
[5] W. Jaco, Lecture Notes on 3-Manifold Topology, CBMS Reg. Conf. Ser. Math., vol. 43, Amer. Math. Soc., Providence, RI, 1980.
[6] William Sherman, On the Kneser–Haken finiteness theorem: A sharpness result, PhD thesis, University of California, Los Angeles, 1992.