

Generating disjoint incompressible surfaces

Hugh Nelson Howards
Wake Forest University
Department of Mathematics
Winston Salem, NC 27109
e-mail: howards@wfu.edu

February 4, 1999

Abstract

We show that one can embed an arbitrarily large collection of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces in any 3-manifold with at least one boundary component of genus greater than or equal to two. We also answer in the affirmative Jaco's question "is it possible to find an incompressible separating surface of arbitrarily high genus in handlebodies of genus $n \geq 2$?"

Keywords: Incompressible Surface, Haken Finiteness, Parallel Surfaces, Tunnel, Boundary Compressible.

AMS classification: 57M99

1 Embedded, Separating Incompressible Surfaces in Handlebodies

In his book Lectures on 3-Manifolds Topology [2] Jaco poses the open question, "is it possible to find an incompressible separating surface of arbitrarily high genus in handlebodies of genus $n \geq 2$?" We answer the question in the affirmative for any compact orientable 3-manifold with a boundary component of genus $n \geq 2$. As a warm up we first give an explicit construction for the case $n \geq 3$ in a handlebody. The simple construction for this case

was discovered jointly with Ian Agol. It is quite easy to picture. The general case is more complicated and uses a recursive argument so it quickly becomes impossible to see the actual surface.

Jaco constructs a non-separating, incompressible surface of arbitrarily high genus for a genus two handlebody in example III.14(a) of his book. Since his description of the original non-separating surface is quite complete, and this is merely a constructive example that is not as strong as the result we will develop later, we will assume familiarity with his example.

Since Jaco's surface is non-separating we can make a parallel copy and have a set of surfaces that separates as a pair. This, of course, does not count as an example, because it is not a connected surface. A natural way to connect two properly embedded surfaces, leaving them properly embedded is to tunnel them together. Tunneling is the inverse of boundary compressing, it takes the boundary connect sum of two surfaces (or perhaps a surface with itself) along an arc. (See Figure 1). We are tempted to make the two parallel surfaces one by connecting the two components with a tunnel, since tunneling preserves whether a set of surfaces separates or not. Unfortunately tunneling does not necessarily preserve whether a set of surfaces is incompressible or not. The immediately obvious ways of tunneling the two surfaces together makes them compressible in this example.

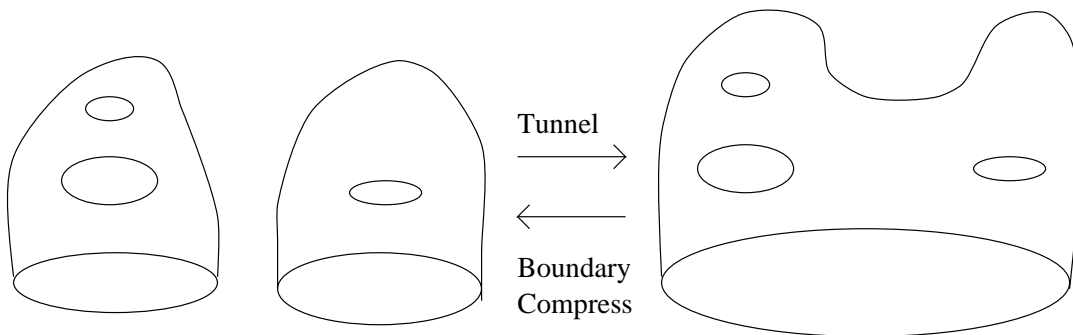


Figure 1: Tunneling is the inverse of boundary compressing.

Instead we alter the ambient manifold by adding a handle in the product region between the two parallel surfaces getting a genus 3 handlebody. Now tunnel the two surfaces together going once over the new handle. (See Figure 2) We now have a connected, separating surface of arbitrarily high genus (equal to twice the genus of the original surfaces). We need only to show that it is incompressible.

Assume there exists a compressing disk D for the new surface that intersects a compressing disk B for the new handle minimally. A traditional innermost loop argument eliminates simple closed curves, but then an outermost arc on B will boundary compress D into two compressing disks that intersect B fewer times (if either of the pieces is not a compressing disk for the surface then clearly $D \cap B$ was not minimal). This means that $D \cap B = \emptyset$ but then one of the original surfaces was compressible, which is our contradiction.

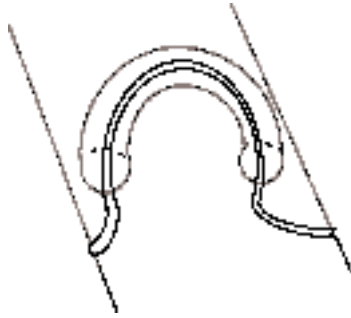


Figure 2: Connecting parallel surfaces across an added handle

Corollary 1.1 *The free group on three generators may be split into a free product with amalgamation over two arbitrarily large rank free groups.*

Note: We will refine this result to the free group on two generators, but that will be more complicated and much harder to picture.

2 Incompressible Surfaces in 3-manifolds with a Boundary Component of Genus $n \geq 2$

The main result of this paper is to show that one can embed an arbitrarily large collection of disjoint, incompressible, non-parallel, non-boundary-parallel surfaces in any 3-manifold with at least one boundary component of genus greater than or equal to two. Thus all compact, orientable 3-manifolds may be classified as supporting arbitrarily many surfaces if and only if they have a genus two or greater boundary component because, (with the exception of annuli) any incompressible surface in a manifold with no boundary

component of genus $n \geq 2$ is also boundary-incompressible so Haken's theory applies.

Note that this result contrasts with, but does not contradict the result from [5] that it is impossible to ever embed an infinite number of such surfaces in any compact, orientable 3-manifold M .

We also note that William Sherman's Ph.D. thesis [6] at UCLA showed that the two holed torus crossed with the unit interval supports arbitrarily many 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.

3 The Set Up

We set out to prove:

Theorem 3.1 *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.*

We start by taking a minimal genus incompressible “half lives half dies” surface F . (In a handlebody F would just be a nonseparating disk). We take a large collection of parallel copies of F , $F_1 \dots F_n$. It is only important that n is slightly over twice as large as the number of surfaces we want to create, but we might as well assume that it is several times larger. Examine the two outside copies, F_1 and F_n , and tunnel each of these surfaces to itself in opposite directions in such a way that they remain disjoint and incompressible (see Figure 3).

Lemma 3.2 *It is always possible to tunnel F_1 and F_n to themselves in opposite directions in such a way that they remain disjoint and incompressible*

Proof: This essentially follows for Theorem 10 of [4]. Adding such a tunnel is possible in general as we can see if we imagine splitting the manifold along F , so $F' = F'_1 \cup F'_2$ (the surface after splitting) is now a two component incompressible subsurface in the boundary of M' (the manifold after splitting). F has $n \geq 1$ boundary components, so F'_1 will have n

boundary components, $n - 1$ of which bound an annulus in $\partial M'$ with $n - 1$ of the boundary components of F'_2 and the last of which together with the final boundary components of F'_2 cobound a subsurface of $\partial M'$ of positive genus, which we will call \hat{F}' . Let these last two boundary components be called α'_1 and α'_2 . Now M' is clearly not $F' \times I$ since F' is not connected but M' is. If \hat{F}' is parallel into F' we may without loss of generality assume that \hat{F}' is parallel into F'_1 . Either the entire product fixes α'_1 or it does not. If \hat{F}' is parallel into F'_1 leaving α'_1 fixed, then F' was not minimal genus, because the product can be pulled back into M and the image of \hat{F}' in F_1 could be replaced by an annulus, (essentially the annulus that α_2 traces out under the product structure, plus the annulus running from α_1 to α_2 , if $n = 1$ so $\alpha_1 = \alpha_2$ then the second annulus is, of course trivial). This leaves a new surface of lower genus (if the new surface is not embedded it will be easy to cut and paste along the curves of intersection to make it embedded).

Thus, we may assume that under the product α'_1 cobounds a non-boundary parallel annulus A with another curve on F' . If a nontrivial extension arc γ for F' makes F'_γ (F' extended along γ) compressible then we can choose a compressing disk D which intersects A minimally. Since F' is incompressible it is easy to eliminate both essential and inessential circles on A . Since A is not boundary parallel, we see that A also cannot have an essential arc of intersection with D as any arc of intersection that ran from one boundary component on A to another would imply that A is boundary compressible into a nonboundary parallel disk with boundary on F' , showing either F' is compressible or that M is reducible. Since the lemma is clearly true for all M if it is true for all irreducible M , we may assume M is irreducible to force a contradiction. Now we see that the only possible arcs of intersection on A are arcs with both endpoints on the same boundary component. This, however, means that we can compress D along A to get at least one disk that is either a compression disk for F' or for \hat{F}' . Since F' is incompressible, we may assume the latter, but if \hat{F}' has a compressing disk D' , then it is easy to see that F' can be arc extended by choosing an arc that intersects D' essentially in one or two points.

Thus, we may assume that \hat{F}' is not parallel into F' . Now we can apply Theorem 10 from [4], which states

Theorem 10 [4]: Let F', \hat{F}', M be as in Lemma 3.2. Suppose M is not a product $F' \times I$, and suppose \hat{F}' is not parallel into F' . Then \hat{F}' contains an extension arc γ of F' with endpoints on any prescribed components of $\partial F'$.

to see that F' may be extended and therefore F_1 and F_n may each be tunneled to itself (the proof of the theorem is strong enough that we may assume that the two copies of F may both have tunnels added in opposite directions remaining disjoint since both components of F' can be extended in the boundary of M' in this manner).

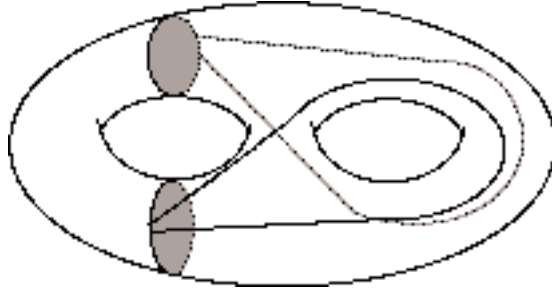


Figure 3: The first two tunnels added.

Now we are ready to begin.

4 The Air Graph

The air graph \mathcal{G} is used to examine the three dimensional regions into which the surface cuts the three manifold. We essentially use the graph defined in Section 2 of Freedman [3].

Let $\{F_1, \dots, F_N\}$ be the set of surfaces. The closure of a component of $M \setminus F_i$ is a vertex; the set of vertices is designated $\{v_i\}$. Each component of F_i yields an edge e joining the vertices (or vertex).

The vertices of \mathcal{G} are labeled N , P , C_g or C_s . In \mathcal{G} a vertex v is assigned an N if it does not correspond to a product region. If the region for v is homeomorphic to a product with $[-1, 1]$ (between the two surfaces), it is assigned a P . The label C is another option for a non-product region which preempts an N if 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. (Note: this is slightly simplified from the graph in [3], but it suffices for our purposes.)

In our case the graph starts as a circle with one or two vertices labeled N , and the rest P (See Figure 4). Throughout the evolution of the surfaces, the

air graph remains a circle, only the labeling changes. This is true because our tunnels are always either splittings or self-gluing, and neither of these affect the topology of the graph.

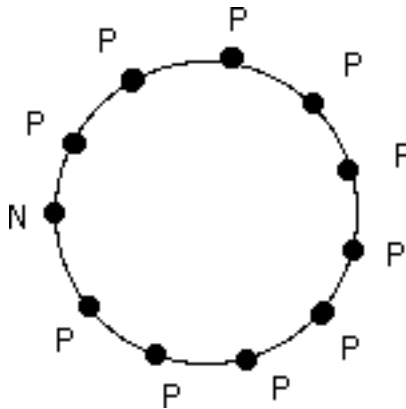


Figure 4: The air graph \mathcal{G}

Essentially what we will do is shoot a cusp in each direction out of the N vertex making sure to leave it as an N . The cusps will then travel in opposite directions around the graph until they collide. If they collide in the same boundary region yielding another N we repeat the process. If they are in different boundary regions then they continue around without a true collision.

Each time we repeat the process we gain another N on the air graph. In turn we are gaining another non-parallel surface (parallel surfaces must have a P vertex between them).

5 The Floor Graph

The floor graph is used to examine the way the surfaces meet the boundary. In this paper it serves a minor role. It verifies that although in theory whenever the cusps collide in the air graph the tunnels could be in different floor regions, in reality the collisions regularly occur in the same floor region just as one would expect. We need to eliminate the improbable possibility that the new air region marked N always cuts the boundary of M in two disks punctured two times each, instead of one disk punctured three times.

It is essentially the same graph used in Sherman [6]. We examine the way the surfaces meet T . Let the curves $\{T \cap F_i\}$ be $\{c_{ij}\}$. The closure of a component of $\{T \setminus \{c_{ij}\}\}$ is a vertex. Each element of $\{c_{ij}\}$ is an edge joining the vertices (or vertex). To ensure that tunneling does not affect the topological type of the (thickened) graph we also add n edges connecting a vertex to itself, if it corresponds to a region with genus n . For example if a region were a punctured torus, an additional edge would be added to the graph starting and ending at the corresponding vertex.

In this case we start with a floor graph that looks like a figure eight (see Figure 5). We will split the non-product region corresponding to the vertex of valence four by one splitting and one self-gluing, changing the graph to the second stage in Figure 5. We propel the cusps around, doing the gluing and splittings.

Just before the collision we stop self-gluing and start splitting so we end with the final stage of Figure 5. At this time we repeat the process.

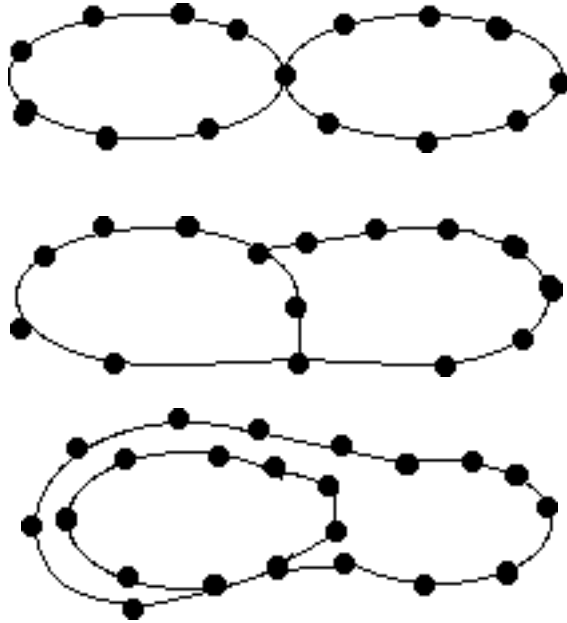


Figure 5: The evolution of the floor graph

6 Verifying Incompressibility

We now check the details of the procedure. We must make sure that we can indeed resolve the collisions with one self-gluing and one splitting leaving behind a non-product region in the air graph and leaving both surfaces incompressible.

Propelling a cusp leaves surfaces parallel (incompressible) if and only if they began parallel (incompressible) by [6] or [3], but it is worth noting that there is also an easy argument using the tools we are about to develop, in the proof of Lemma 7.

7 Resolving the Collision

We can control the relative rate at which the cusps are propelled. So essentially by general position we may assume that the collisions always occur at vertices in the air graph that were labeled P just before the collision.

Let F_1 and F_2 be the surfaces on either side of the new N region. Resolve the region with tunnels t_1 and t_2 as in Figure 6, yielding surfaces F'_1 and F'_2 . Note that we have assumed that we have two splittings meeting at a product region in the air graph, and we have one vertex of valence four in the floor graph, so this is the exact picture we want, where compressing F_1 along the boundary compressing disk B_1 and compressing F_2 along the boundary compressing disk B_2 leaves two parallel surfaces. (B_1 and B_2 are the boundary compressing disks for the large tunnels in Figure 6).

This means that F_1 and F_2 have the same genus (they are obtained from two parallel surfaces by splittings), but F'_1 and F'_2 do not, as they are obtained from two surfaces of the same genus, one 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 the surfaces cannot be parallel and that we have left behind an N . Note that sending a cusp generically through this region in the future will affect the genus of the two surfaces the same way, so the N s we create in the air graph in this manner will remain as long as we allow all future collisions and all future conversions of a gluing cusp to a splitting cusp only in product regions, which, as we mentioned earlier, we are fully at liberty to do.

Lemma 7.1 *Adding the tunnel t_1 to F_1 in Figure 6 yields an incompressible*

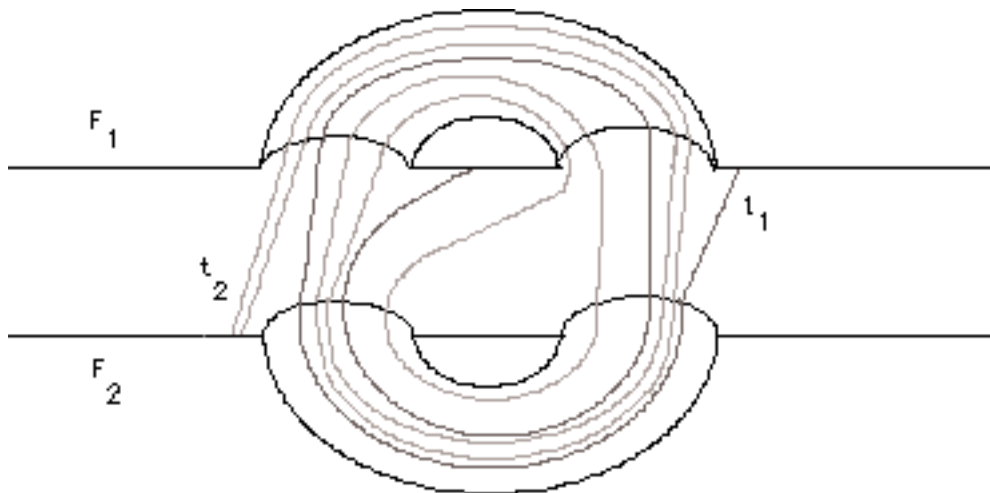


Figure 6: Resolution of a collision of cusps

surface, F'_1

Proof: Assume not. Choose a compressing disk D for F'_1 that has a minimal number of intersections with B_1 . We know that $D \cap B_1$ is nontrivial or else F_1 was already compressible, which by assumption it is not. Since we are intersecting two disks, we may use a simple innermost loop argument to assure that there can be no simple closed curves in the intersection.

If an arc of $D \cap B_1$ ran from t_1 to itself (or F_1 to itself), then we could choose an outermost arc on B_1 , boundary compress D along the corresponding subdisk, and we would get two disks. Both are compressing disks with fewer intersections with B_1 , and this contradicts our minimal assumption.

Thus, all arcs in $D \cap B_1$ run from t_1 to F_1 . Choose an outermost arc on B . This almost gives a compressing disk D' for F'_1 , but it takes a shortcut along B_1 , so it does not remain properly embedded. It runs from F_1 across half of t_1 and then along the shortcut back to F_1 . We can add another tunnel from F_1 to t_1 so that D' is now properly embedded. It remains so if we boundary compress the half of t_1 that D' does not cross. This gives us a new tunnel t'_1 and a new surface F''_1 that is also compressible as in Figure 7. There are four options (really only three since two of the options are isotopic) for what the new surface initially looks like, depending on which side the new tunnel was added and which half of t_1 was compressed. For all four, however, we can use the same argument to prove that they are incompressible.

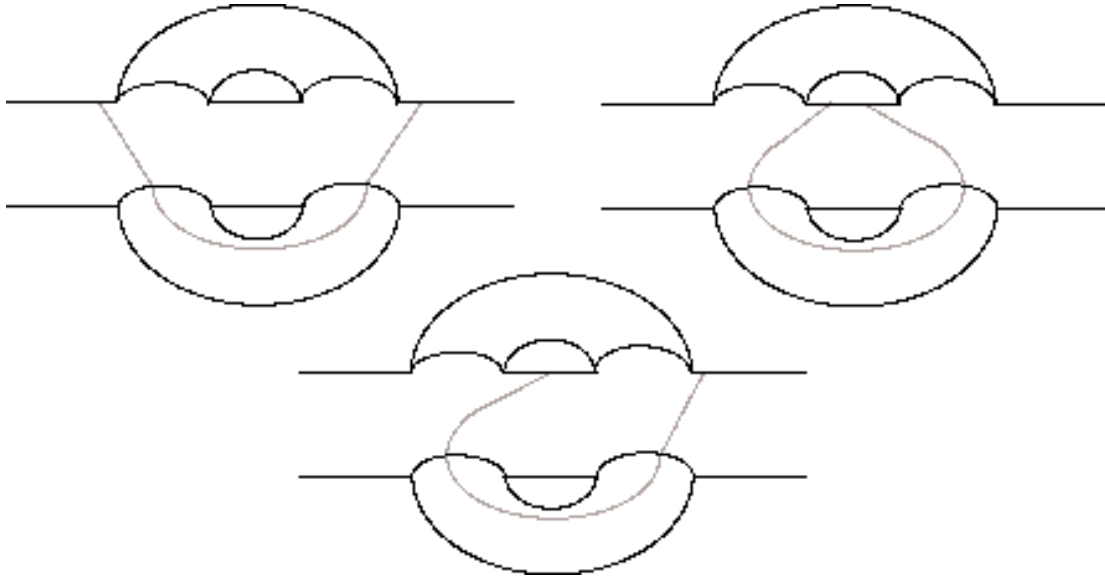


Figure 7: If F'_1 were compressible, then one of these would be, too.

We examine how D' intersects B_2 (a boundary compressing disk for the large pictured tunnel on F_2). There can be no simple closed curves, and there can be no trivial arcs connecting t'_1 to itself, for the same reasons as in the previous arguments. This time, however, F_2 is a different incompressible surface, so D' does not intersect it, so there are no arcs connecting t'_1 to F_2 , either. This means that $D \cap B_2$ is empty, and this in turn dictates that D does not run over t'_1 and that F_1 was in fact compressible contradicting our assumption.

Lemma 7.2 *Adding the tunnel t_2 to F_2 in Figure 6 yields an incompressible surface, F'_2*

Proof: This argument is slightly more complicated than the previous one. Examine how t_2 intersects B_2 . We label the intersections 1,2, and 3 in the order they occur as we traverse t_2 (see Figure 8).

Once again choose a compressing disk with a minimal number of intersections with B_2 . Again there are of course arcs only and no simple closed curves. Label the ends of the arcs, 1,2,3, or X , depending on whether they lie on the tunnel at component 1,2, or 3 or on the surface F_2 respectively.

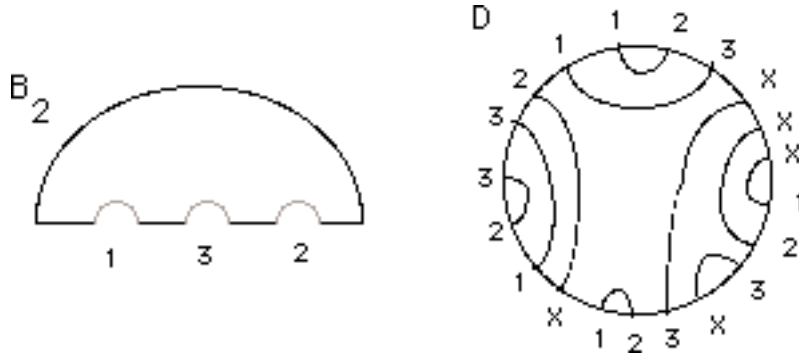


Figure 8: Labeling the boundary-compressing disk, B_2 and the compressing disk D .

Arcs cannot have end points with the same label or once again an outermost one in B_2 yields compressing disks with fewer intersections with B_2 .

As we read the labels off around the boundary of D clockwise, we may assume that it reads something like, “123XX321321X,” always counting directly from 3 down to 1 or 1 up to 3, and if an X appears, it is separating these sequences of numbers from each other, not in the midst of one. It is clear that an outermost arc must connect adjacent labels, so an outermost arc could never connect a 2 and an X . Outermost arcs also cannot have label 1 and X , as we could then use the same end game we employed in the previous Lemma. They also for the same reason cannot be labeled 3 and X .

If an outermost arc ever connected 2 and 3 we could use the shortcut trick to get a compressing disk D' for the annulus in Figure 9. A minimal intersection argument with B_1 shows this is impossible.

Thus, All outermost arcs must be labeled 1 and 3. By assumption, D must run across t_2 , so there must be arcs on D labeled 2 on one side. We now examine what labels they could have on the other side. We define the length of an arc to be the minimum of the number of labels on either of the two subdisks into which the arc splits D . The subdisk outside of an arc with a 2 as one of its labels must at least have one 2 in it, as an arc connecting a 2 to a different 2 is the shortest arc that does not force a contradiction of the rules for outermost arcs. Recall that all outermost arcs must connect a 1 to a 3, but a 2 cannot be connected to a 2, so it must be longer than that (see Figure 10).

Now take the shortest arc with an end point labeled 2 and notice, that

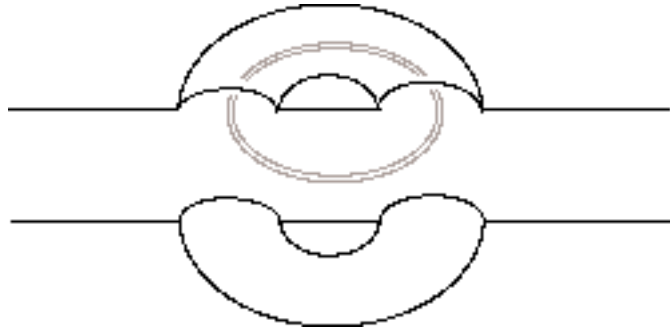


Figure 9: An annulus that would have to be compressible if F_2 were

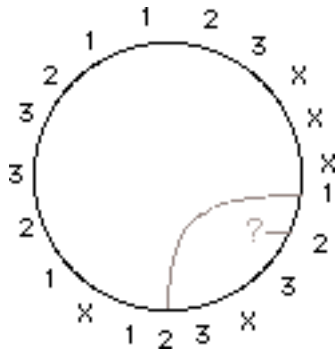


Figure 10: Every arc connected to a 2 must separate the disk so that there is a 2 on each side.

since there is a 2 in the short sector that it cuts off, the other 2 must be connected with a shorter arc, which is of course a contradiction. Therefore, there could not have been a compressing disk in the first place.

8 Changing the Self-Gluing Cusp to a Splitting Cusp

The only thing left to prove is that changing the self-gluing cusp into a splitting cusp does not create a compressible surface. We let the event happen at a set of vertices that were labeled P and that would return to being labeled P if the cusp were allowed to pass through unchanged. Because of the large number of P s relative to the number of N s, this is no problem. Instead of sending the cusp through as a gluing, we turn it into a splitting as in Figure 11. We choose the component of F_1 that causes the splitting to go the direction we want it to in the floor graph.

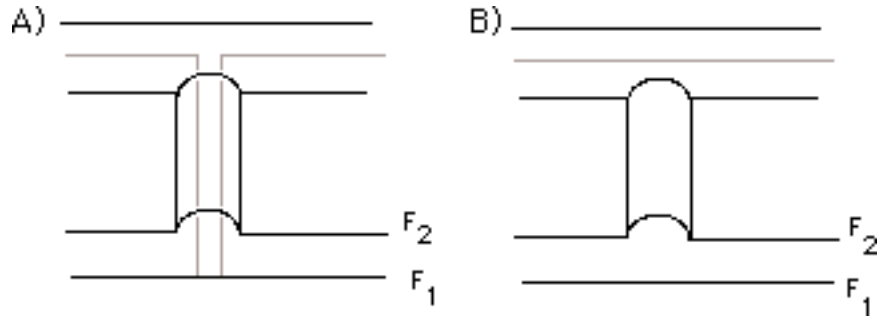


Figure 11: A changes a gluing to a splitting and B shows an incompressible annulus

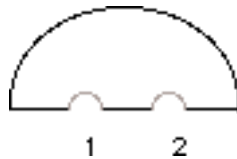


Figure 12: B_2 , a boundary-compressing disk for F_2 .

Lemma 8.1 *Adding the tunnel t_1 to F_1 in Figure 11 yields an incompressible surface, F'_1*

Proof: Again of course the proof requires that we choose a compressing disk D for F'_1 that meets B_2 minimally. Again, $D \cap B_2$ has no simple closed curves for the same reasons. We label the intersections of t_1 with B_2 , 1 and 2 as in Figure 12. As before, there can be no arcs running from 1 to 1 or 2 to 2. There are also no arcs intersecting F_2 , as F_2 is incompressible and is a distinct surface from F_1 .

Therefore, the compressing disk can be chosen disjoint from it. Then all of the arcs of intersection run from 1 to 2 or vice versa. In addition to labeling each of the vertices with a 1 or a 2, we label them with a + or a - according to whether, traversing clockwise around D , one intersects F_1 first or the other intersection region, respectively. We will need the following Lemma.

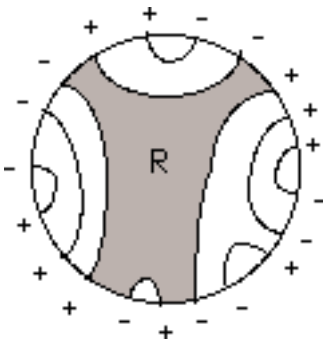


Figure 13: Disks, D Spanned by Edges with One End Labeled + and the other - will always have a region that intersects the boundary of D as region R does.

Lemma 8.2 *Given a disk D and a set of properly embedded arcs $\{a_i\}$ with one end point labeled + and the other -, there will always be a region in $D \setminus \{a_i\}$ that intersects the boundary of D only in segments that when read clockwise run from a - sign to a + sign as in figure 13*

Proof: The proof is by induction. If an outermost arc runs clockwise from - to +, then clearly the condition is satisfied, so the base case of one arc is trivial. Assume the condition holds for $n - 1$. If we add one more edge to a

region R that satisfies our condition, it will split R into two regions, exactly one of which still satisfies the condition. This can be seen because the process is essentially as if we have a disk with alternating $+$ and $-$ marks on it. We split it into two disks with an edge running from the top to the bottom that runs between intervals on the disk that go from $-$ to $+$ clockwise (or possibly from an interval to itself that meets the same description). Putting a $+$ at the top and a $-$ at the bottom makes the new disk on the left satisfy the condition. Doing it the opposite way means the one on the right will.

Lemma 8.2, however, means that by adding a tunnel from 1 to 2 and boundary compressing D turns R into a compressing disk for the new surface that can be made embedded yet does not intersect F_1 ! Thus, the annulus in Figure 11 is compressible. However this is clearly not true, as it is boundary parallel, so its boundary would have to bound a disk, but its boundary is parallel to a boundary curve of the incompressible surface F_1 . Since F_1 is not a disk we have a contradiction, and the Lemma as well as the Theorem are proven.

9 Separating Surfaces of Arbitrarily High Genus

We now return to Jaco's question: "Is there an incompressible separating surface of arbitrarily-high genus for handlebodies of genus $n \geq 2$?" In the first section we saw that the answer was yes for $n \geq 3$. We are about to see that the answer is yes for $n = 2$. The result applies not just to handlebodies, but also to any manifold to which Theorem 3.1 applies.

Theorem 9.1 *Any compact orientable 3-manifold with a boundary component of genus ≥ 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 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). Depending on how fast we let each of the two cusps advance relative to each other and how long we wait before changing the gluing to a splitting, we have virtually complete control over what percentage of the surfaces increases in genus between collisions. After enough collisions, we may assume that the collision is between two surfaces of arbitrarily high genus. We need only prove that connecting the surfaces (as in Figure 14) leaves the surface incompressible.

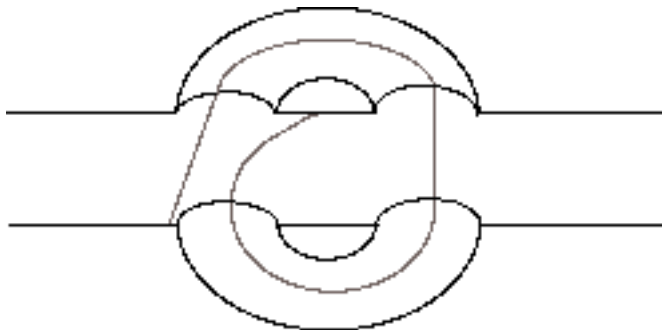


Figure 14: Creating a connected incompressible surface from a collision (F_1 Above, F_2 Below).

Let the new surface be called F' . Choose a compressing disk D for F' that intersects B_1 minimally. Since neither F_1 nor F_2 was compressible, the disk must run across the tunnel t . As before, there are no simple closed curves and no arcs connecting t to itself or F_1 to itself. Therefore all arcs run from t to F_1 . Moving clockwise around the boundary of D , we label the end points of the arcs of intersection corresponding to t , 1) with a - if after hitting them we hit F_1 before F_2 and 2) with a + if the opposite is true. We label the other end point of the arc that is not on t with the opposite sign.

By Lemma 8.2 there must be a region that intersects the boundary of D only in arcs running clockwise from - to +. This region is exactly what we need to see that we can use the shortcut trick to yield either the second or third result pictured in Figure 7. By construction, F'_1 has a compressing disk, but we have already seen many times the argument, using B_2 , that shows that this is impossible.

Note: essentially what we are doing here is coloring the portions of ∂D that are on the half of the tunnel t that is attaches to F_1 green, and the half

that attaches to F_2 red, and the parts that are on neither tunnel white and using Lemma 8.2 to see that some region of $D \cap B_2$ will have no red on its boundary.

Corollary 9.2 *The free group on two generators may be split into a free product with amalgamation over two arbitrarily large free groups.*

10 References

[1] Hempel, J., *3-Manifolds (Annals of Mathematics Studies, 86, Princeton University Press, Princeton, NJ, 1976).*

[2] Jaco, W., *Lecture Notes on 3-manifold Topology (CBMS Regional Conf. Ser. in Math., No. 43 Amer. Math. Soc., Providence, RI, 1980).*

[3] Freedman, B and Freedman, M, *Knesser-Haken Finiteness for Bounded 3-Manifolds, Locally Free Groups, and Cyclic Covers: preprint, 1996*

[4] Freedman, M, Howards, H, and Wu, Y, *Extension of Incompressible surfaces on the Boundary of 3-Manifolds: preprint 1997*

[5] Howards, H, *Limits of Incompressible Surfaces: to appear Topology and its Applications.*

[6] Sherman, W, *Ph.D Thesis UCLA*