ON THE CLASSIFICATION OF HEEGAARD SPLITTINGS

TOBIAS HOLCK COLDING, DAVID GABAI, AND DANIEL KETOVER

ABSTRACT. The long standing classification problem in the theory of Heegaard splittings of 3-manifolds is to exhibit for each closed 3-manifold a complete list, without duplication, of all its irreducible Heegaard surfaces, up to isotopy. We solve this problem for non Haken hyperbolic 3-manifolds.

0. INTRODUCTION

The main result of this paper is

Theorem 0.1. *Let N be a closed non Haken hyperbolic 3-manifold. There exists an effectively constructible set $S_0, S_1, \cdots, S_n$ such that if S is an irreducible Heegaard splitting, then S is isotopic to exactly one S_i.*

Remarks 0.2. Given $g \in \mathbb{N}$ Tao Li [Li3] shows how to construct a finite list of genus-g Heegaard surfaces such that, up to isotopy, every genus-g Heegaard surface appears in that list. By [CG] there exists an effectively computable $C(N)$ such that one need only consider $g \leq C(N)$, hence there exists an effectively constructible set of Heegaard surfaces that contains every irreducible Heegaard surface. (The methods of [CG] also effectively constructs these surfaces.) However, this list may contain reducible splittings and duplications. The main goal of this paper is to give an effective algorithm that weeds out the duplications and reducible splittings.

Idea of Proof. We first prove the *Thick Isotopy Lemma* which implies that if S_i is isotopic to S_j, then there exists a smooth isotopy by surfaces of area uniformly bounded above and diametric soul uniformly bounded below. (The *diametric soul* of a surface $T \subset N$ is the infimal diameter in N of the essential closed curves in T.) The proof of this lemma uses a 2-parameter sweepout argument that may be of independent interest. We construct a graph $\mathcal{G}$ whose vertices comprise a finite net in the set of genus $\leq C(N)$ embedded surfaces of uniformly bounded area and diametric soul, i.e. up to small perturbations and pinching spheres any such surface is close to a vertex of $\mathcal{G}$. The edges of $\mathcal{G}$ connect vertices that are perturbations of each other up to pinching necks. Thus S_i and S_j are isotopic if and only if they lie in the same component of $\mathcal{G}$. For technical reasons the construction of $\mathcal{G}$ is carried out in the PL category.

The Thick Isotopy Lemma also shows that any reducible Heegaard surface S_i is isotopic to an *obviously* reducible one through surfaces of uniformly bounded area and diametric souls. Thus S_i is reducible if and only if it lies in the same component of $\mathcal{G}$ as an *obviously*

Version 2.00
The first author was partially supported by NSF Grant DM 1404540 and NSF FRG grant DMS 0854774, the second by DMS-1006553, DMS-1607374 and NSF FRG grant DMS-0854969 and the third by an NSF Postdoctoral Fellowship.

reducible one. A surface is obviously reducible if it has small diameter essential curves that compress to each side. The point here is that the surfaces in the isotopy look incompressible at a small scale, but at a somewhat larger scale the isotopy ends at an obviously reducible surface.

By an *effective* algorithm we mean one that produces an output as a function of the initial data. It is interesting to note that our algorithm is elementary and combinatorial, yet the proof that it works requires a 2-parameter sweepout argument and a mulit-parameter min-max argument.

Our main result for manifolds with taut ideal triangulations was earlier obtained by Jesse Johnson [Jo]. His methods inspired some of those used in this paper.

This paper is organized as follows. Basic definitions and facts are given in §1, the Thick Isotopy Lemma is proved in §2 and the graph $\mathcal{G}$ is constructed in §3.

Remark 0.3. We believe that the methods of this paper will have applications to the classification problem for compact 3-manifolds and may also have application to the question of finding the minimal common stabilization to two splittings.

1. Heegaard splittings and Paths of Heegaard foliations

Definition 1.1. A *Heegaard splitting* of a closed orientable 3-manifold M consists of an ordered pair (H_0, H_1) of handlebodies whose union is M and whose intersection is their boundaries. This common boundary S is called a *Heegaard surface.* Two Heegaard splittings (H_0, H_1), (H_0', H_1') are *isotopic* if there exists an ambient isotopy of M taking H_0 to H_0'.

Remark 1.2. The Heegaard splitting (H_0, H_1) may not be isotopic to (H_1, H_0), e.g. see [Bir]. The ordering induces a transverse orientation on the Heegaard surface S, pointing from H_0 to H_1 and an isotopy is required to preserve it. All the results of this paper naturally carry over to the weaker setting where isotopy of Heegaard surfaces need not preserve the transverse orientation.

Definition 1.3. A *Heegaard foliation* $\mathcal{H}$ is a singular foliation of M induced by a submersion to $[0,1]$ such that a Heegaard surface S is a leaf, and if H_0 and H_1 are the handlebodies bounded by S then each H_i has a PL spine E_i such that $\mathcal{H}|(H_i \setminus E_i)$ is a fibration by surfaces that limit to E_i.

Lemma 1.4. *i) Any two Heegaard foliations $\mathcal{H}, \mathcal{H}'$ of the same Heegaard splitting (H_0, H_1) which have the same spines are isotopic, via isotopies fixing S and the spines pointwise.*

ii) Given $\mathcal{H}_0, \mathcal{H}_1$ two Heegaard foliations of the same Heegaard splitting (H_0, H_1) there exists a path $\mathcal{H}_t$ of Heegaard foliations from $\mathcal{H}_0$ to $\mathcal{H}_1$ which varies smoothly away from the spines and fixes the Heegaard surface S throughout.

iii) Given isotopic Heegaard splittings (H_0, H_1) and (H_0', H_1') with Heegaard foliations $\mathcal{H}_0, \mathcal{H}_1$ and Heegaard surfaces S_0, S_1, there exists a path $\mathcal{H}_t$ of Heegaard foliations from $\mathcal{H}_0$ to $\mathcal{H}_1$ which varies smoothly away from the spines and takes the Heegaard surface S_0 to S_1.

Proof of i) It suffices to show that $\mathcal{H}|H_0$ is isotopic to $\mathcal{H}'|H_0$ via an isotopy fixing S. Let $S = S_1, S_2, \cdots,\ \ S = S_1', S_2', \cdots$ be nested sequences of leaves of $\mathcal{H}|H_0$ and $\mathcal{H}'|H_1$ which converge to E_0. Next isotope H_0, fixing S_1 so that S_i is taken to S_i' and if $f : \cup S_i \cup E_0 \to$

$\cup S_i' \cup E_0$, then f is continuous and $f|S_i$ is smooth. By adjusting the isotopy using the normal flow to the leaves we can assume that $\mathcal{H}$ and $\mathcal{H}'$ coincide near each S_i. We abuse notation by denoting $f_*(\mathcal{H})$ by $\mathcal{H}$. Let V_k denote the compact region bounded by S_k and S_{k+1}. By [LB] each V_k can be ambiently isotoped by an isotopy that is fixed near ∂V_k, so that $\mathcal{H}|V_k$ is taken to $\mathcal{H}'|V_k$. With care the union of these isotopies gives a globally defined isotopy, i.e. it is continuous at E_0. □

Proof of ii) It suffices to construct a path from $\mathcal{H}_0|H_0$ to $\mathcal{H}_1|H_0$ which fixes S and is smooth away from the spine. Let E_0 (resp. E_1) denote the spine of H_0 associated to $\mathcal{H}_0$ (resp. $\mathcal{H}_1$). Since handlebodies have unique spines up to sliding of 1-cells, there exists a path $\mathcal{H}_t, t \in [1/2, 1]$ so the spine of H_0 associated to $\mathcal{H}_{1/2}$ is equal to E_0 and the nonsingular leaves vary smoothly. The result now follows from i). □

Proof of iii) An initial ambient isotopy parametrized by $[2/3, 1]$ takes S_0 to S_1, so by the covering isotopy theorem there exists a smoothly varying path of Heegaard foliations $\mathcal{H}_t, t \in [2/3, 1]$ such that $\mathcal{H}_{2/3}$ has S_0 as a Heegaard surface. The result now follows from ii). □

2. Thick Isotopy Lemma

The proof of the Thick Isotopy Lemma relies on a 2-parameter min-max argument to find a path of surfaces joining two isotopic Heegaard splittings with all areas bounded from above.

We first introduce the min-max theory with n parameters. The min-max theory is due to Almgren-Pitts [P] but we will use a refinement by Simon-Smith [SS] that allows one to consider sweepouts of a fixed topology. We also need the optimal genus bounds established in [K].

Throughout this section, M denotes a closed orientable 3-manifold, and $\mathcal{H}^2(\Sigma)$ denotes the 2-dimensional Hausdorff measure of a set $\Sigma \subset M$. Set $I^n = [0,1]^n \subset \mathbb{R}^n$. Let $\{\Sigma_t\}_{t \in I^n}$ be a family of closed subsets of M and $B \subset \partial I^n$. We call the family $\{\Sigma_t\}$ a *(genus-g) sweepout* if

(1) $\mathcal{H}^2(\Sigma_t)$ is a continuous function of $t \in I^n$,
(2) Σ_t converges to Σ_{t_0} in the Hausdorff topology as $t \to t_0$.
(3) For $t_0 \in I^n \setminus B$, Σ_{t_0} is a smooth closed surface of genus g and Σ_t varies smoothly for t near t_0.
(4) For $t \in B$, Σ_t is a 1-complex.

Given a family of subsets $\{\Sigma_t\}_{t \in \partial I^n}$, we say that $\{\Sigma_t\}_{t \in \partial I^n}$ *extends to a sweepout* if there exists a sweepout $\{\Sigma_t\}_{t \in I^n}$ that restricts to $\{\Sigma_t\}_{t \in \partial I^n}$ at the boundary. A *Heegaard foliation* is a sweepout $\{\Sigma_t\}$ parameterized by $[0,1]$ where Σ_t is a Heegaard surface for $t \neq 1, 0$ and the sets Σ_0 and Σ_1 are 1-complexes in the handlebodies determined by the Heegaard splitting. Additionally, we assume that $\{\Sigma_t\}$ is a singular foliation (with only two singular leaves). If $\mathcal{H}_0(s)_{s=0}^1$ and $\mathcal{H}_1(s)_{s=0}^1$ are two Heegaard foliations with respect to isotopic Heegaard splittings, then we call $\{\Sigma_t\}_{t \in I^2}$ a *Heegaard sweepout joining* $\mathcal{H}_0$ *to* $\mathcal{H}_1$ if it is a sweepout such that for $s \in [0,1]$, $\Sigma_{(0,s)} = \mathcal{H}_0(s)$ and $\Sigma_{(1,s)} = \mathcal{H}_1(s)$ and also for $i \in \{0,1\}$, $\{\Sigma_{(t,i)}\}_{t=0}^1$ is a continuously varying family of 1-complexes.

Beginning with a genus-g sweepout $\{\Sigma_t\}$ we need to construct comparison sweepouts which agree with $\{\Sigma_t\}$ on ∂I^n. We call a collection of sweepouts Π *saturated* if it satisfies the following condition: for any map $\psi \in C^\infty(I^n \times M, M)$ such that for all $t \in I^n$, $\psi(t,.) \in \text{Diff}_0(M)$ and $\psi(t,.) = id$ if $t \in \partial I^n$, and a sweepout $\{\Lambda_t\}_{t\in I^n} \in \Pi$ we have $\{\psi(t,\Lambda_t)\}_{t\in I^n} \in \Pi$. Given a sweepout $\{\Sigma_t\}$, denote by $\Pi = \Pi_{\{\Sigma_t\}}$ the smallest saturated collection of sweepouts containing $\{\Sigma_t\}$. We define the *width* of Π to be

$$W(\Pi, M) = \inf_{\{\Lambda_t\}\in\Pi} \sup_{t\in I^n} \mathcal{H}^2(\Lambda_t). \tag{2.1}$$

A *minimizing sequence* is a sequence of sweepouts $\{\Sigma_t\}^i \in \Pi$ such that

$$\lim_{i\to\infty} \sup_{t\in I^n} \mathcal{H}^2(\Sigma_t^i) = W(\Pi, M). \tag{2.2}$$

Finally, a *min-max sequence* is a sequence of slices $\Sigma_{t_i}^i$, $t_i \in I^n$ taken from a minimizing sequence so that $\mathcal{H}^2(\Sigma_{t_i}^i) \to W(\Pi, M)$. The main point of the Min-Max Theory of Almgren-Pitts [P] is that if the width is bigger than the maximum of the areas of the boundary surfaces, then some min-max sequence converges to a minimal surface in M:

Theorem 2.1. *(Multi-parameter Min-Max Theorem) Given a sweepout $\{\Sigma_t\}_{t\in I^n}$ of genus g surfaces, if*

$$W(\Pi, M) > \sup_{t\in\partial I^n} \mathcal{H}^2(\Sigma_t) \tag{2.3}$$

then there exists a min-max sequence $\Sigma_i := \Sigma_{t_i}^i$ such that

$$\Sigma_i \to \sum_{i=1}^k n_i\Gamma_i \text{ as varifolds,} \tag{2.4}$$

where Γ_i are smooth closed embedded minimal surfaces and n_i are positive integers. Moreover, after performing finitely many compressions on Σ_i and discarding some components, each connected component of Σ_i is isotopic to one of the Γ_i or to a double cover of one of the Γ_i. We have the following genus bounds with multiplicity:

$$\sum_{i\in\mathcal{O}} n_i g(\Gamma_i) + \frac{1}{2}\sum_{i\in\mathcal{N}} n_i(g(\Gamma_i) - 1) \le g, \tag{2.5}$$

where $\mathcal{O}$ denotes the subcollection of Γ_i that are orientable and $\mathcal{N}$ denotes those Γ_i that are non-orientable, and where $g(\Gamma_i)$ denotes the genus of Γ_i if it orientable, and the number of crosscaps that one attaches to a sphere to obtain a homeomorphic surface if Γ_i is non-orientable.

Theorem 2.1 is proved in the Appendix. We can now state the main application of the min-max theory in our setting:

Theorem 2.2 (Bounded area sweepouts)**.** *Let N be a closed hyperbolic 3-manifold and let $\{\Sigma_t\}_{t\in I^n}$ be a genus-g sweepout. Set $C = \max(\sup_{t\in\partial I^n} \mathcal{H}^2(\Sigma_t), 2\pi(2g-2))$. Then for all $\epsilon > 0$, there is a sweepout $\{\Lambda_t\}_{t\in I^n}$ extending $\{\Sigma_t\}_{t\in\partial I^n}$ such that $\sup_{t\in I^n} \mathcal{H}^2(\Lambda_t) \le C + \epsilon$.*

The following is an immediate corollary of Theorem 2.2 and Lemma 1.4.

Corollary 2.3. *Let N be a closed hyperbolic 3-manifold and let $\mathcal{H}_0$ and $\mathcal{H}_1$ be isotopic Heegaard foliations on N representing Heegaard surfaces of genus g. Denote the leaves of $\mathcal{H}_i$ by $\mathcal{H}_i(t), t \in [0,1]$. Let C denote* $\max\{2\pi(2g-2), \max\{\text{area}(\mathcal{H}_s(t))|s \in \{0,1\}, t \in [0,1]\}\}$. *Then for all $\epsilon > 0$, $\mathcal{H}_0, \mathcal{H}_1$ extend to a sweepout $\{\Sigma_t\}_{t\in I^2}$ such that $\sup_{t\in I^2} \mathcal{H}^2(\Sigma_t) \leq C + \epsilon$.*

Proof of Theorem 2.2. Denote by Π the saturation of sweepouts containing $\{\Sigma_t\}_{t\in I^n}$. We argue by contradiction. Thus assume there exists an $\epsilon > 0$ so that all sweepouts in Π have a slice with area greater than $\max(2\pi(g-2), \sup_{t\in\partial I^n} \mathcal{H}^2(\Sigma_t)) + \epsilon$. Then the width $W(\Pi, N)$ of Π satisfies:

$$W(\Pi, N) \geq \max(2\pi(2g-2), \sup_{t\in\partial I^n} \mathcal{H}^2(\Sigma_t)) + \epsilon. \tag{2.6}$$

Since from (2.6) the width $W(\Pi, N)$ is strictly greater than $\sup_{t\in\partial I^n} \mathcal{H}^2(\Sigma_t)$, the Min-Max Theorem 2.1 applies to give a positive integer combination of minimal surfaces $\Gamma = \sum_i n_i\Gamma_i$, so that

$$W(\Pi, N) = \sum n_i \mathcal{H}^2(\Gamma_i). \tag{2.7}$$

The area of a genus g surface in a hyperbolic 3-manifold is at most $2\pi(2g-2)$ and the area of a non-orientable surface is at most $2\pi(k-2)$ (where k denotes the number of cross-caps one must add to a sphere to obtain a homeomorphic surface). Using these bounds, along with the genus bound (2.5) and the fact that the multiplicity n_i of a non-orientable surface occurring among the Γ_i is even we obtain:

$$W(\Pi, N) = \sum n_i \mathcal{H}^2(\Gamma_i) \leq 2\pi \sum_{i\in\mathcal{O}} n_i(2g_i - 2) + 2\pi \sum_{i\in\mathcal{N}} n_i(g_i - 2) \leq 2\pi(2g-2).$$

This contradicts (2.6). □

To prove the main result of this section Lemma 2.10 we need to bound area from above and diametric soul from below. The min-max theory enabled the area bound. The following topological arguments will enable the diametric soul bound.

Definition 2.4. Let M be a Riemannian 3-manifold, S a surface embedded in M. We say that S is δ-compressible if some essential simple closed curve of diameter $< \delta$ bounds an embedded disc in a closed complementary region. Otherwise we say that S is *δ-locally incompressible.*

Let S be a closed separating surface in the 3-manifold M. We say that S is *δ-bicompressible* if there exist essential simple closed curves in S of diameter $< \delta$ that respectively bound discs in each closed complementary region.

Lemma 2.5. *Suppose that $M \neq S^3$. If the Heegaard surface $S \subset M$ is η-bicompressible, where $\eta < \delta_0/2$ and δ_0 is the injectivity radius of M, then S is weakly reducible.*

Proof. Let $\alpha_1, \alpha_2 \subset S$ be essential simple closed curves of diameter $< \eta$ that compress to distinct sides of S. If $\alpha_1 \cap \alpha_2 = \emptyset$, then the result is immediate. Otherwise a slightly perturbed 3-ball E of radius $< \eta$ contains both α_i's and is transverse to S. Since $M \neq S^3$ some essential curve α of $S \cap \partial E$ compresses in one of the handlebodies and that together with one of the α_i's gives a weak reduction. □

Definition 2.6. Let $\mathcal{H}$ be a Heegaard foliation with leaves parametrized by $[0,1]$. If $t \in (0,1)$, then the H_0 *(resp.* H_1*) side* of $\mathcal{H}(t)$ is the component of the closed complement that contains $\mathcal{H}(0)$, (resp. $\mathcal{H}(1)$). More generally if T_t is a 2-dimensional sweepout which coincides on $\partial I \times I$ with a path of Heegaard foliations, then for $v \in [0,1] \times (0,1)$ use the smooth variance to define the H_0 and H_1 sides of T_v. It therefore makes sense to say that a curve in T_v *compresses to the* H_0 *or* H_1 *side.*

Define $\mathrm{mecd}_0(T_v)$ (resp. $\mathrm{mecd}_1(T_v)$) (the *minimal essential compressing diameter*) to be the minimal diameter of a curve in $T_v, v \in I \times (0,1)$ that compresses to the H_0 (resp. H_1) side.

Since the surfaces T_v vary continuously in v we obtain:

Lemma 2.7. *The functions* mecd_0*,* mecd_1 *are continuous and extend to continuous functions on* $I \times I$*.* □

Lemma 2.8. *Let* $\mathcal{H}$ *be a Heegaard foliation in the Riemannian manifold* M *and let* $t \in (0,1)$*. Suppose that* $\delta < \delta_0/16$ *where* δ_0 *is the injectivity radius. If* $1 > w > t$ *and* $\mathrm{mecd}_0(\mathcal{H}(w)) < 2\delta$*, then either* $\mathrm{mecd}_0(\mathcal{H}(t)) < 4\delta$ *or* $\mathcal{H}(w)$ *is* 4δ*-bicompressible.*

Proof. Let T_s denote $\mathcal{H}(s)$ and hence $\mathrm{mecd}_0(T_w) < 2\delta$. Let B be a smooth ball of diameter $< 4\delta$, transverse to both T_w and T_t, such that some essential curve $\gamma \subset T_w \cap B$ compresses to the H_0 side. Among the components of $T_w \cap \partial B$ that are essential in T_w, choose an innermost one. This curve α has diameter $< 4\delta$ and compresses to either the H_0 or H_1 sides. If it compresses to the H_1 side we are done. Otherwise let D be a compression disc for α that lies the H_0 side of T_w. It can be chosen so that $D \cap T_t \subset \partial B$. Among components of $T_t \cap D$ that are essential in T_t let β be an innermost one. Since T_t and T_w cobound a product such a β exists. This β has diameter $< 4\delta$. □

Definition 2.9. A *C-isotopy* $F : T \times I \to M$ is an isotopy such that for all t, $\mathrm{area}(T_t) \leq C$.

Lemma 2.10. *(**Thick Isotopy Lemma**) Let* N *be a closed non Haken hyperbolic 3-manifold with injectivity radius* δ_0 *and let* $\delta < \delta_0/16$*.*

i) If T_0 *and* T_1 *are isotopic strongly irreducible genus-g Heegaard surfaces that are* 8δ*-locally incompressible, then there exists a computable* $C > 0$ *and there exists a* C*-isotopy* $F : T \times I \to N$ *from* T_0 *to* T_1 *with each* T_t δ*-locally incompressible.*

ii) If T_0 *is weakly reducible and* 8δ*-locally incompressible, then there exists a computable* $C > 0$ *and a* C*-isotopy* $F : T \times I \to N$ *from* T_0 *to a* T_1 *such that each* T_t *is* δ*-locally incompressible and* T_1 *is* 4δ*-bicompressible.*

Remark 2.11. In i), given T_0 and T_1 that look incompressible at a certain intermediate scale, we find an isotopy with each interpolating surface looking incompressible at a certain small scale. In ii) we start with a weakly reducible T_0 that looks incompressible at a certain intermediate scale and *find* a T_1 which is *obviously* weakly reducible at that scale. Further, T_0 is isotopic to T_1 via an isotopy such that each interpolating surface looks incompressible at a certain small scale.

Proof of i). Given a Heegaard surface $T \subset N$, Haken's theory of hierarchies and normal surface theory provides an algorithm for showing that each side of T is a handlebody and hence gives an algorithm for constructing a Heegaard foliation with T as a leaf. Applying these algorithms to T_0 and T_1 produces foliations $\mathcal{H}_0$ and $\mathcal{H}_1$ where for $i=0,1,T_i$ is identified with $\mathcal{H}_i(1/2)$. Let C' denote an upperbound for the area of any leaf of either $\mathcal{H}_0$ or $\mathcal{H}_1$. Let $C=\max\{(7(2g-2),C'\}+1$. Apply Corollary 2.3 to obtain a sweepout between $\mathcal{H}_0$ and $\mathcal{H}_1$. Denote a (possibly singular) surface in this sweepout by either $T_{(s,t)}$ or T_v where $v\in I\times I$.

Let $G_p=\mathrm{mecd}_0^{-1}([0,p))$ and $R_q=\mathrm{mecd}_1^{-1}([0,q))$. Now $G_{1.5\delta}$ is an open set containing $[0,1]\times 0$ and its closure is disjoint from $0\times[1/2,1]$ by Lemmas 2.8, 2.7 and 2.5. A similar statement holds for $R_{1.5\delta}$. Since T_0 is strongly irreducible, it follows from Lemma 2.7 that $\bar{G}_{1.5\delta}\cap\bar{R}_{1.5\delta}=\emptyset$. Hence there exist compact smooth submanifolds G and R of $I\times I$ transverse to $\{0,1\}\times I$ such that $\bar{G}_{1.1\delta}\subset\mathrm{int}(G)\subset G\subset G_{1.5\delta}$ and $\bar{R}_{1.1\delta}\subset\mathrm{int}(R)\subset R\subset R_{1.5\delta}$. Again $R\cap G=\emptyset$ and in particular $G\cap I\times 1=\emptyset$. Elementary topological considerations imply that some component of ∂G has endpoints in both $0\times I$ and $1\times I$. A path in $I\times I$ starting at $(0,1/2)$ giving rise to an isotopy satisfying the conclusion of i) is obtained by concatenating arcs that lie in ∂G and $0\times I$.

Proof of ii). Since T_0 is weakly reducible it's reducible [CaGo] and hence T_0 is isotopic to a stabilization T_1 of some strongly irreducible surface R. Now T_0 is explicitly given and R is isotopic to a surface in a known finite set of surfaces [CG], thus there exist Heegaard foliations $\mathcal{H}_i$ extending T_i, $i=0,1$ such that a computable C' bounds the area of any leaf of $\mathcal{H}_0$ or $\mathcal{H}_1$. Since T_1 is a stabilization we can also assume that for $t\neq 0,1$, $\mathcal{H}_1(t)$ is δ-bicompressible. Parametrize $\mathcal{H}_0$ so that $T_0=\mathcal{H}_0(1/2)$.

Construct the region G as in i), though here $I\times 0\cup 1\times I\subset G$ but $(0,1/2)\notin G$. Let $\sigma\subset I\times I$ be the maximal embedded path starting at $(0,1/2)$ whose interior is disjoint from $0\times[1/2,1]$ such that $\sigma\subset 0\times I\cup\partial G$ and $\sigma\cap\mathrm{int}(G)=\emptyset$. By construction $\mathrm{mecd}_0(T_v)>\delta$ for $v\in\sigma$ and the terminal endpoint of σ lies in $I\times 1\cup 0\times I$. Let λ be the maximal subpath of σ starting at $(0,1/2)$ such that for $v\in\lambda$, $\mathrm{mecd}_1(T_v)\geq\delta$. The desired isotopy is parametrized by λ. Denote the terminal endpoint of λ by ω. If $\omega\in 0\times[0,1/2)$ (resp. $\omega\in 0\times(1/2,1]$), then $\mathrm{mecd}_1(T_\omega)<1.5\delta$ (resp. $\mathrm{mecd}_0(T_\omega)<1.5\delta$) and hence Lemma 2.8 implies that T_ω is 4δ-bicompressible. If $\omega\notin 0\times I$, then T_ω is 1.5δ-bicompressible. $\square$

3. Weeding out duplications and reducibles

Given a triangulated non Haken 3-manifold N and $g\in\mathbb{N}$, Tao Li [Li3] gives an algorithm to construct a set $\{S_0,S_1,\cdots,S_n\}$ of genus-g Heegaard surfaces such that any genus-g Heegaard surface is isotopic to one in this set. In this section, assuming that the manifold N is hyperbolic, we give an algorithm to eliminate all the reducible splittings and duplications. Here is the idea. Suppose that S_0 is weakly reducible. The *Thick Isotopy Lemma* shows that there is an isotopy from S_0 to one that is obviously weakly reducible, i.e there exist reducing curves that lie in small 3-balls, via an isotopy through surfaces of area uniformly bounded above and injectivity radius uniformly bounded below. (The lower bound being smaller than the diameter of those small 3-balls.) Ignoring long fingers and spheres that can be pinched off, one can find a finite net for the totality of such surfaces. Construct a graph $\mathcal{G}$ whose verticies are the points of the net and whose edges correspond to surfaces that differ by small perturbations. Thus if S_0 is reducible, then it is in the same component of $\mathcal{G}$ as an obviously reducible one. This shows how to eliminate reducible splittings from $\{S_0,\cdots,S_n\}$. A similar argument shows that if S_i and S_j are irreducible and isotopic then they are in the same component of $\mathcal{G}$. We carry out this idea in the PL setting. Here smooth isotopies are approximated by PL ones, the relation of normal isotopy gives a net among surfaces transverse to a triangulation (after certain spheres are pinched off), and pinching and elementary isotopies across faces of a triangulation give the edges of the graph.

Definition 3.1. Let Δ be a triangulation of the smooth 3-manifold M. By a surface T *transverse* to Δ we mean that T is transverse to the various skeleta of Δ. A *transverse isotopy* is an isotopy between transverse surfaces through a path of transverse surfaces. An isotopy $F : T \times [n, m] \to M$ between embedded surfaces T_n and T_m is said to be a *generic Δ-isotopy* if for all but finitely many times $n < t_1 < t_2 < \cdots < t_n < m$ each T_t is transverse to Δ. Furthermore, for ϵ sufficiently small the passage from $T_{t_i-\epsilon}$ to $T_{t_i+\epsilon}$ is one of the following four elementary moves or their inverses.

0) Passing through a vertex: See Figure A
1) Passing through an edge: See Figure B
2) Passing through a face (two possibilities): See Figures Ca, Cb

The next result follows from a standard perturbation argument.

Lemma 3.2. *Let Δ be a triangulation of the Riemannian 3-manifold N with metric ρ, $L > 0$ and $\epsilon > 0$, then there exists $K(\Delta, L, \epsilon, \rho) > 1$ such that if T is a closed embedded surface with* $\text{area}(T) < C$, *then T is isotopic to a surface T' such that $|T' \cap \Delta^1| < KC$ and the diameter of the track of any point of the isotopy is at most ϵ.*

If $F : T \times [0, 1] \to M$ is a C-isotopy between surfaces T_0 and T_1 that are transverse to Δ of weight at most L, then there exists a generic $K(C + 1)$-Δ-isotopy G from T_0 to T_1 such that for all $x \in T$ and $t \in [0, 1]$, $d(G(x, t), F(x, t)) < \epsilon$. □

The transverse surface T is said to be obtained by *tubing* the transverse surfaces P and Q if there exists a 3-simplex σ and an embedded $D^2 \times I \subset \text{int}(\sigma)$ such that $D^2 \times 0 \subset P$, $D^2 \times 1 \subset Q$ and $T = (P \cup Q \cup \partial D^2 \times I) \setminus (\text{int}(D^2) \times \{0, 1\})$. Two transverse surfaces T and T' *differ by a pinch* if T' is obtained by tubing T and a 2-sphere or vice versa.

A *pinched isotopy* F from T_0 to T_1 consists of $0 = n_1 < \cdots < n_k = 1$ and isotopies $f_i : T \times [n_i, n_{i+1}]$, $i = 1, 2, \cdots, n_{k-1}$ such that $f_{n_1}(T, 0) = T_0$, $f_{n_{k-1}}(T, 1) = T_1$ and for all $i < k - 1$ the surfaces $f_i(T, n_{i+1})$ and $f_{i+1}(T, n_{i+1})$ differ by a pinch.

Definition 3.3. Let Δ be a triangulation of the 3-manifold M. We say that T is *crudely almost normal* if T is transverse to Δ and satisfies the following additional properties. If τ is a 2-simplex, then no component of $T \cap \tau$ is a simple closed curve and if σ is a 3-simplex, then each component of $T \cap \sigma$ is either a disc or an unknotted annulus. Also there is at most one annulus component of $T \cap \sigma$. If all the components of intersection with 3-simplices are discs, then T is said to be *crudely normal*. Define the *weight* of T to be $|T \cap \Delta^1|$.

Lemma 3.4. *Let Δ be a triangulation of the 3-manifold M. Given $N \in \mathbb{N}$, there are only finitely many transverse isotopy classes of crudely normal and crudely almost normal surfaces of weight at most N.* □

Definition 3.5. A *generic crudely normal* isotopy between two crudely almost normal surfaces T_0 and T_1 in a manifold with triangulation Δ is a generic Δ-isotopy $F : T \times [0, 1] \to N$ such that away from the finitely many non transverse points, the surfaces $F(T, t)$ are crudely almost normal. A $C_1 - \Delta$-*isotopy* is one whose transverse surfaces all have weight $\leq C_1$ with respect to Δ.

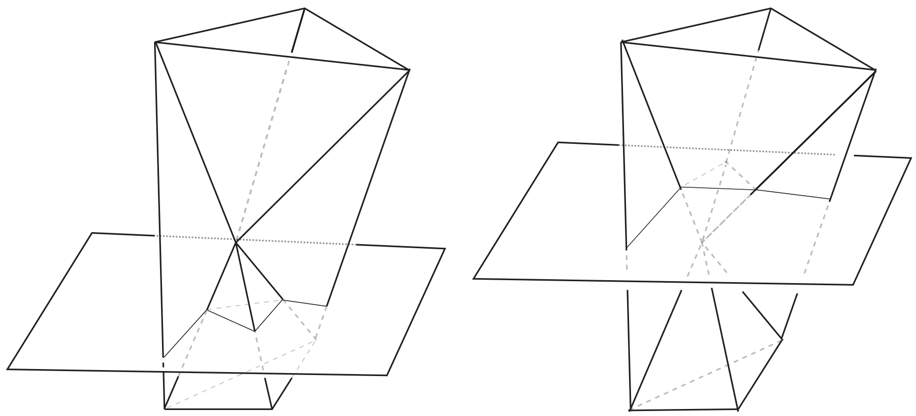

Pushing through a vertex

S_- Figure A S_+

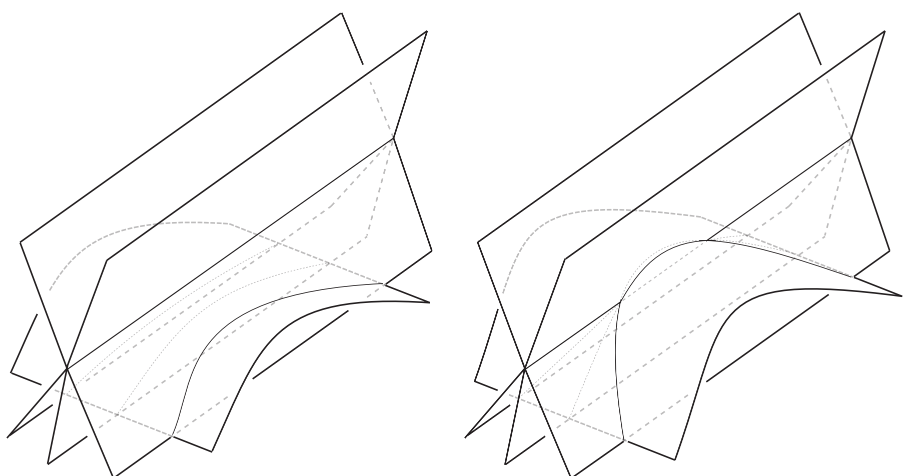

Pushing through a 1-simplex

S_- S_+

Figure B

The following lemma more or less says that provided the diameters of the simplices of Δ are sufficiently small, a generic pinched isotopy between crudely normal surfaces through δ-locally incompressible surfaces can be replaced by one whose generic interpolating surfaces are crudely normal without increasing the uniform upper bound on the weights of the interpolating surfaces. If κ is a subcomplex of the triangulation Δ, then $\mathrm{st}(\kappa)$ denotes the closed star of κ in Δ.

Lemma 3.6. *Let M be a closed irreducible Riemannian 3-manifold with injectivity radius $> \delta_0$. Let $\delta < \delta_0/16$ and Δ a triangulation on M such that if $\kappa \in \Delta$, then $st^2(\kappa)$ and $st(\kappa)$ are 3-balls of diameter $< \delta$. Let G be a generic pinched C_1-Δ-isotopy between the crudely normal genus-g surfaces T_0 to T_1 such that for each t, $G(T,t) := S_t$ is δ-locally incompressible. (At*

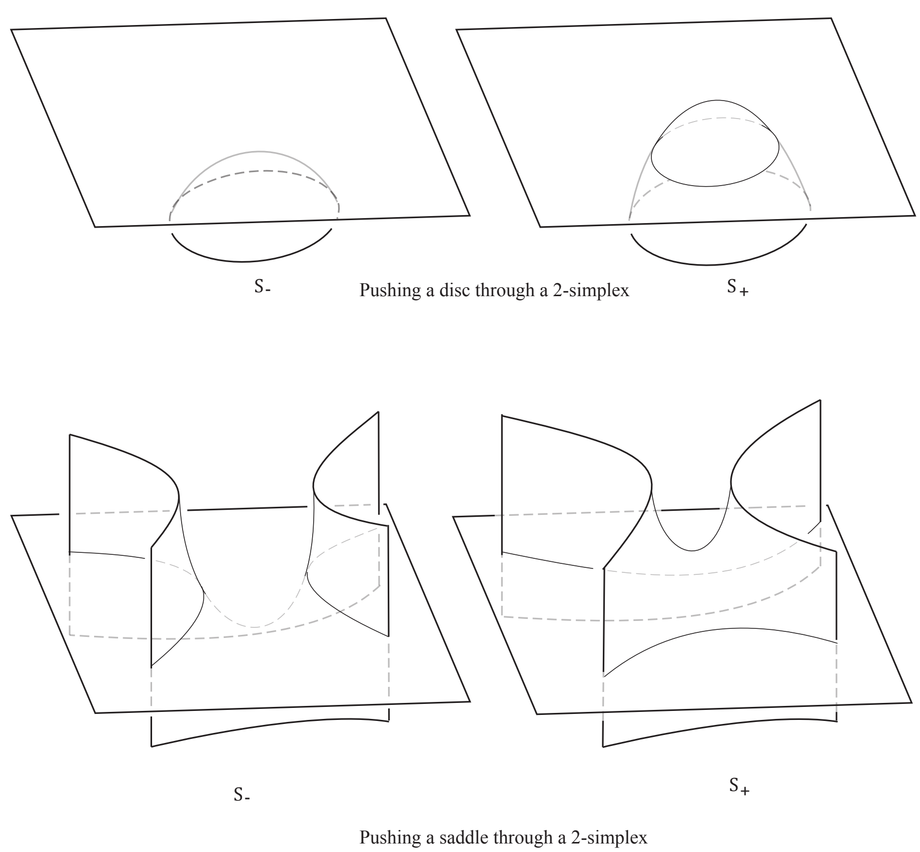

Figure C

the pinch times there are two possibilities for S_t.) Then there exists a generic pinched crudely normal C_1-Δ-isotopy H from T_0 to T_1.

Proof. If S_t is transverse to Δ, then let T_t be the crudely normal surface obtained as follows. First, let S'_t be the surface obtained by compressing S_t near each component of $S_t \cap \partial\sigma$ for every 3-simplex σ and then deleting the components disjoint from Δ^2. I.e. if $N(\Delta^2)$ is a small regular neighborhood of Δ^2, then $S'_t \cap N(\Delta^2) = S_t \cap N(\Delta^2)$ and for all every 3-simplex σ each component of $S'_t \cap \sigma$ is a disc. Since S_t is δ-locally incompressible, each component of $S_t \cap \partial\sigma$ is inessential in S_t. It follows that S'_t is a union of 2-spheres and a single component T_t of genus-g. Note that T_t is crudely normal and up to transverse isotopy T_t is locally constant in t. If t is a pinch time, the result of this construction is independent of the two possible choices for S_t. Finally the local incompressibility condition implies that if σ is a 3-simplex, then any closed curve in $S'_t \cap \sigma$ is inessential in S'_t.

Away from a small neighborhood of the non transverse times, we let $H(S_t) = T_t$. To complete the proof, it suffices to show that if S_s is not transverse to Δ and ϵ is sufficiently small, then there is a generic pinched crudely normal $C_1 - \Delta$-isotopy from $T_{s-\epsilon}$ to $T_{s+\epsilon}$. We abuse notation by letting $S_{s-\epsilon}$ (resp. $S'_{s-\epsilon}$, $T_{s-\epsilon}$) be denoted S_- (resp. S'_-, T_-) with analogous notation for $S_{s+\epsilon}$ (resp. $S'_{s+\epsilon}$, $T_{s+\epsilon}$).

Case 0: S_s passes through a vertex.

Proof of Case 0. Let v denote the vertex. Let σ be a 3-simplex having v as a vertex. We say an edge of σ with vertex v is *up* if it lies to the S_+ side of S and *down* otherwise. If σ has three down (resp. up) edges, then $S_+ \cap \sigma$ differs from $S_- \cap \sigma$ by the loss (resp. gain) of a normal triangle. If it has one or two down edges, then exactly one component is non transversely isotoped in the passage from S_- to S_+. It follows that S'_+ differs from S'_- by an elementary move of type 0) and either T_+ is transversely isotopic to T_- or they differ by a move of type 0). □

Case 1: S_s is tangent to a 1-simplex.

Proof of Case 1. Let η denote the 1-simplex. We can assume that the isotopy from S_- to S_+ is as in Figure 4 as opposed to the inverse operation. Let $\kappa_1, \cdots, \kappa_m$ denote the 2-simplices containing η. We say that κ_i is a *down* simplex if some component of $\kappa_i \cap S_-$ splits into two components of $\kappa_i \cap S_+$, otherwise it is an *up simplex*. By perturbing the isotopy slightly, at a cost of creating moves of type 2b), we can assume that there exists a unique down simplex. It follows that S'_+ is obtained from S'_- by a type 1) move and that either T_+ is transversely isotopic to T_- or they differ by a type 1) move. □

Case 2a: S_s intersects a 2-simplex locally at a point.

Proof of Case 2a. Let τ denote the 2-simplex. Again we can assume that the isotopy is as in Figure 5 as opposed to the inverse one. Up to transverse isotopy S'_+ is equal to the union of S'_- and a 2-sphere that intersects Δ^2 in a single circle and $T_- = T_+$. □

Case 2b: S_s intersects a 2-simplex at a saddle.

Proof of Case 2b. Let τ denote the 2-simplex, $x \in \tau$ the point of tangency and σ_1 and σ_2 the 3-simplices that contain τ. We can assume that σ_1 (resp. σ_2) lies *below* (resp. *above*) x, i.e. if a normal vector to S_s based at x points from σ_1 to σ_2, then the isotopy locally moves S_s in that direction. Let γ be the component of $S_s \cap \tau$ that contains x. We have various subcases.

Case 2bi: $\gamma \cap \Delta^1 = \emptyset$.

Proof. By replacing the isotopy by the reverse if necessary, it suffices to consider the case that two S^1 components of $S_- \cap \tau$ transform to one component of $S_+ \cap \tau$. Up to transverse isotopy S'_- is the union of S'_+ and a S^2 component that intersects Δ^2 in a single circle and $T_- = T_+$. □

Case 2bii: $|\gamma \cap \Delta^1| = 2$. We can assume that an arc and an S^1 component of $S_- \cap \tau$ coalesce to an arc component of S_+ and therefore have the same conclusion as Case 2bi. □

Case 2biii: $|\gamma \cap \Delta^1| = 4$. Here two arc components p_1, q_1 of $S_- \cap \tau$ transform to two other arc components p_2, q_2 of $S_+ \cap \tau$. Observe that p_1, q_1 belong to the same component of $S_- \cap \partial\sigma_1$ if and only if p_2, q_2 belong to different components of $S_+ \cap \partial\sigma_1$. A similar fact holds for $\partial\sigma_2$.

We next show that it suffices to consider the case that p_1, q_1 are contained in different components of both $S_- \cap \partial\sigma_1$ and $S_- \cap \partial\sigma_2$. If they belong to the same components, then p_2, q_2 are contained in different components of both $S_+ \cap \partial\sigma_1$ and $S_+ \cap \partial\sigma_2$, so the reverse isotopy from S_+ to S_- has the desired feature. If p_1, q_1 belong to different components p, q of $S_- \cap \partial\sigma_1$ but the same component of $S_- \cap \partial\sigma_2$, then $p \subset \sigma$ and is homologically essential, contradicting our local incompressibility condition.

Now assume that p_1, q_1 are contained in different components of both $S_- \cap \partial\sigma_1$ and $S_- \cap \partial\sigma_2$. Let P be the component of S'_- containing p_1 and Q the component containing q_1. Again if $P = Q$, then p is homologically essential in S_- contradicting δ-local incompressibility. Therefore one of them, say Q is a 2-sphere. If P is also a 2-sphere, then T_- is transversely isotopic to T_+. If $P = T_-$, then let T' be the crudely almost normal surface obtained by tubing P and Q, i.e. T' is pinch equivalent to T_-. Finally T_+ is obtained from T' by a type 2b) move and the proof of the lemma is complete. □

Proof of Theorem 0.1 The initial data for N is a totally geodesic triangulation Δ_1. (E.g. see [Br].) Use it to find a lower bound δ_1 for the injectivity radius of N. By [CG] there exists an effectively computable $C(N)$ bounding the genus of any irreducible Heegaard surface. Now fix $g \leq C(N)$ and apply Tao Li's algorithm [Li3] (or [CG]), to effectively construct a set $\{S_1, \cdots, S_n\}$ of Heegaard surfaces such that any Heegaard surface of genus-g is isotopic to one of these surfaces. By construction the surfaces from [Li3] are almost normal with respect to Δ_1. Next, pass to the second barycentric subdivision Δ_2 of Δ_1. It follows that after a small isotopy of Δ_2, each S_i is normal to Δ_2 and $st(\kappa)$ and $st^2(\kappa)$ are closed 3-balls for each simplex $\kappa \in \Delta_2$. Let $\delta_2 \leq \delta_1$ be such that any essential curve of any S_i has diameter $> \delta_2$. Let $\delta_3 \leq \delta_2$ be such that if $X \subset N$ and $\operatorname{diam}(X) \leq \delta_3$, then $X \subset st(\kappa)$ some $\kappa \in \Delta_2$. Let $\delta_4 \leq \delta_3/100$ and subdivide Δ_2 to Δ_3 so that $\operatorname{diam}(\kappa) \leq \delta_4$ for all $\kappa \in \Delta_3$ and so that each S_i is normal. Next let L be the maximal Δ_3 weight of all the S_i's, $C = 7(2g-2)$ and $K = K(\Delta_3, \delta_4/2, L)$ as in Lemma 3.2. Note that $K \geq L$ and that C is greater than the hyperbolic area of all the S_i's that arise from [CG], by Gauss - Bonnet since such surfaces are carried by $-1+\eta$-negatively curved branched surfaces, where η is very close to 0. If using [Li3] it is also easy to effectively bound the hyperbolic area these S_i's are surfaces carried by a set of almost normal surfaces to Δ_1.

Construct a graph $\mathcal{G}$ whose vertices are transverse isotopy classes of Δ_3 crudely almost normal surfaces of weight at most $K(C+1)$. Connect two vertices by an edge if they differ either by a pinch or a move of type 0), 1), 2a) or 2b) or their inverses. Call a vertex of $\mathcal{G}$ corresponding to the surface T *obviously weakly reducible* with respect to Δ_2 if there exists (possibly equal) $\kappa_0, \kappa_1 \in \Delta_2$ and essential curves $\alpha_0 \subset T \cap st(\kappa_0)$ and $\alpha_1 \subset T \cap st(\kappa_1)$ such that α_0 and α_1 respectively compress to opposite sides of T. Note that by Haken's normal surface algorithm it is effectively decidable if these subsurfaces of T compress to one side or the other. Indeed, one can pass to a further subdivision Δ_4 such that if such a subsurface compresses to one side, then there exists an essential compressing disc normal with respect to Δ_4 which is a fundamental solution to the appropriate normal surface equations.

By construction each S_i is $16\delta_4$-locally incompressible and of Δ_3 weight $\leq K(C+1)$. Since weakly reducible implies reducible [CaGo] in non Haken 3-manifolds, it suffices to apply the following Lemma 3.7, to weed out the reducible S_i's. Simply remove those S_i's lying in the same component of $\mathcal{G}$ as an obviously weakly reducible splitting.

Once reducible splittings are eliminated from the list, duplications are eliminated by applying Lemma 3.8. Again S_i and S_j are isotopic if and only if they are in the same component of $\mathcal{G}$. □

Lemma 3.7. *i) If T is obviously weakly reducible with respect to Δ_2, then it is weakly reducible.*

ii) If T is transverse to Δ_2 and δ_3-bicompressible, then T is obviously weakly reducible with respect to Δ_2.

iii) If T is $16\delta_4$-locally incompressible and $\text{weight}_{\Delta_2}(T) \leq L$, then T is weakly reducible if and only if it lies in the same component of $\mathcal{G}$ as a vertex that is obviously weakly reducible with respect to Δ_2. Here L is the maximal Δ_3 weight of all the S_i's.

Proof. i) This is immediate if $\text{int}(st(\kappa_0)) \cap \text{int}(st(\kappa_1)) = \emptyset$. Otherwise these stars have a simplex λ in common and hence $st(\kappa_0) \cup st(\kappa_1) \subset st^2(\lambda)$. By construction $st^2(\lambda)$ is a 3-ball, so some component α_3 of $\partial(st^2(\lambda)) \cap T$ is essential in T (else N is the 3-sphere) and compresses to one side or the other. Thus α_3 and one of α_0 or α_1 provide a weak reduction.

ii) Since T is δ_3-bicompressible there exist essential compressing curves α_0, $\alpha_1 \subset T$ that compress to opposite sides and have diameter at most δ_3 and hence lie in stars of simplices of Δ_2.

iii) If T is weakly reducible, then by Lemma 2.10 there exists a C-isotopy from $T = T_0$ to a $16\delta_4$-bicompressible surface T_1'' such that each interpolating surface is $2\delta_4$-locally incompressible.

By applying Lemma 3.2, with $\Delta = \Delta_3$ and $\epsilon = \delta_4/2$ we can assume that there exists a generic $K(C+1) - \Delta_3$-isotopy G from T_0 to T_1', where T_1' is transverse to Δ_3, is $17\delta_4$-bicompressible and each interpolating surface is δ_4-locally incompressible. Finally apply Lemma 3.6 to find a generic pinched crudely normal $K(C+1) - \Delta_3$-isotopy H from T_0 to T_1 where T_1 is crudely normal of weight $\leq K(C+1)$ and $19\delta_4$-bicompressible. Since $19\delta_4 < \delta_3$ it follows T_1 is obviously weakly reducible with respect to Δ_2 and that H determines a path within $\mathcal{G}$ from T_0 to T_1. □

Lemma 3.8. *If S_i and S_j are strongly irreducible then they are isotopic if and only if they are in the same component of $\mathcal{G}$.*

Proof. If they are in the same component of $\mathcal{G}$, then it is immediate that they are isotopic. Conversely by Lemma 2.10 there exists a C-isotopy F from S_i to S_j such that each interpolating surface is $2\delta_4$-locally incompressible. By Lemma 3.2 there exists a generic $K(C+1) - \Delta_3$-isotopy from S_i to S_j such that each interpolating surface is δ_4-locally incompressible. By Lemma 3.6 there exists a generic pinched crudely normal $K(C+1) - \Delta_3$-isotopy H from S_i to S_j. This H determines a path in $\mathcal{G}$ from S_i to S_j. □

4. Three Fundamental Problems

We record three long-standing problems related to the work of this paper.

Problem 4.1. *Given two irreducible Heegaard splittings of a Haken manifold or a Seifert fibered spaces, find an algorithm to decide whether or not they are isotopic.*

Problem 4.2. *Given a Heegaard splitting of a Haken manifold or Seifert fibered space, give an algorithm to decide whether or not it is reducible.*

Problem 4.3. *Given a closed non Haken 3-manifold, construct the tree of Heegaard splittings. In particular if F_1 and F_2 are irreducible, how many stabilizations are needed before they become isotopic.*

Remarks 4.4. i) The Reidemeister - Singer theorem [Re], [Si] implies that any two Heegaard splittings have a common stabilization, thus by Li [Li2] the tree is finite.

ii) Waldhausen [Wa2] proved that S^3 has a unique Heegaard splitting up to stabilization and hence it's tree consists of a single point. More generally Bonahon and Otal [BO] showed that lens spaces have unique Heegaard *surfaces* up to stabilization. Viewing Heegaard splittings as ordered pairs (H_1, H_2), reversing the orientation of an oriented Heegaard surface may give rise to distinct Heegaard splittings, [Bir].

iii) The Heegaard tree has been determined for infinitely many non Haken, non lens space Seifert fibered spaces by the results of [Mor], [Sed], [ML], [MS] and [Sch]. As of this writing, the general case is still open.

iv) Schultens [Sch] proved that for Seifert Fibered spaces with orientable base, given two Heegaard splittings, one stabilization on the larger genus splitting is a common stabilization for both. For general 3-manifolds more than one stabilization may be needed [HTT], [Ba], [Jon2]. Johannson [Jo] gave a polynomial upper bound on the number of stabilizations needed for splittings of Haken manifolds to become equivalent and Rubinstein and Scharlemann [RS] gave linear bounds for pairs of irreducible splittings of non Haken manifolds.

5. Appendix: Proof of Theorem 2.1

The proof of the Min-Max Theorem 2.1 is similar to that of the one-parameter case handled in [CD]. The following two things must be established:

(1) Starting from any given minimizing sequence $\{\Sigma_t\}^i$, one can "pull-tight" to produce a new minimizing sequence $\{\Gamma_t\}^i$ so that any min-max sequence obtained from $\{\Gamma_t\}^i$ converges to a stationary varifold.
(2) At least one min-max sequence obtained from $\{\Gamma_t\}^i$ is *almost minimizing (a.m.) in sufficiently small annuli.* (see Definition 3.2 in [CD] for the relevant definition).

Given (1) and (2) it follows from [CD] that the min-max limit is regular. The genus bound (2.5) then follows from [K] since only the almost minimizing property is used there. Thus (1) and (2) together imply the Min-Max Theorem 2.1. Item (1) follows from straightforward modifications of [CD]. We will give a proof of (2) since the combinatorial argument is more involved than the one-parameter case handled in [CD].

Proof of (2). The idea of the proof is that if *no* min-max sequence were almost minimizing, we can glue together several isotopies to produce a competitor sweepout with all areas strictly below $W(\Pi, M)$ and thus violate the definition of $W(\Pi, M)$. To accomplish this, we will need to find many disjoint annuli on which to pull down the slices, which requires a combinatorial argument due to Almgren-Pitts [P].

Given $x \in M$, for $r, s > 0$, denote by $An(x, r, s)$ the open annulus centered around x with outer radius s and inner radius r. Given sequences $\{r_i\}_{i=1}^k$ and $\{s_i\}_{i=1}^k$ with $s_i < r_{i+1}$, consider the family of annuli $\{An(x, r_i, s_i)\}_{i=1}^k$. Such a family is called *admissible* if $r_{i+1} > 2s_i$ for each i. A family of admissible annuli containing precisely L annuli is called L-*admissible*. Given a surface $\Sigma \subset M$ and an L-admissible family $\mathcal{F}$, we will say that Σ is ϵ-almost minimizing in an $\mathcal{F}$ is it ϵ-almost minimizing in at least one of the annuli comprising $\mathcal{F}$.

Lemma 5.1. *There exists an integer $L = L(n)$ and a min-max sequence Σ_j converging to a stationary varifold such that Σ_j is almost minimizing in every L-admissible family of annuli.*

Proof. Much of the proof of Lemma 5.1 follows [CD], and thus we focus only on the parts that are different. The proof is by contradiction. If it failed, we would obtain a set $S \subset [0,1]^n$ of "large slices" and an open covering $\{\mathcal{O}_i\}_{i=1}^M$ of S so that to each open set $\mathcal{O}_i$ is associated an L-admissible family $An_{\mathcal{O}_i}$. It is enough to prove that there is another cover $\{\mathcal{U}_i\}$ of S refining $\{\mathcal{O}_i\}_{i=1}^M$ (in the sense that each element in the collection $\{\mathcal{U}_i\}$ is contained in at least one of the $\mathcal{O}_i$) such that:

1. To each $\mathcal{U}_i$ we can fix one annulus An_i which is one of the annuli comprising $An_{\mathcal{O}_j}$ for some $\mathcal{O}_j$ containing $\mathcal{U}_i$.
2. Each $\mathcal{U}_i$ intersects at most $d = d(n)$ other elements in the collection $\{\mathcal{U}_i\}$
3. To each $\mathcal{U}_i$ we can associate a compactly supported C^∞ function $\phi_i : \mathcal{U}_i \to [0,1]$. Moreover, for any $x \in S$ we have that $\phi_i(x) = 1$ for some i.
4. If $\phi_i(x)$ and $\phi_j(x)$ are both nonzero for some i and j, then $An_i \cap An_j = \emptyset$

To achieve this, it will help to have some notation (following Pitts). For each $j \in \mathbb{N}$, denote by $I(1, j)$ be the cell complex on the interval $I^1 = [0,1]$ whose 1-cells are $[0, 3^{-j}], [3^{-j}, 2 \times 3^{-j}], ...[1 - 3^{-j}, 1]$ and whose 0-cells are $[0], [3^{-j}], ...[1]$. We are considering the n-dimensional cell complex I^n, which we can write as a tensor product:

$$I(n, j) = I(1, j) \otimes I(1, j) \otimes ... \otimes I(1, j) \text{ (n times).}$$

For $0 \leq p \leq n$ we define a p-cell as an element $\alpha = \alpha_1 \otimes \alpha_2 ... \otimes \alpha_n$ where $\sum_{i=1}^n \dim(\alpha_i) = p$. The sets $\mathcal{U}_i$ will be "thickened" p-cells for suitably large j. Precisely, for each p cell $\alpha \in I(n, j)$ expressed as $\alpha = \alpha_1 \otimes \alpha_2 ... \otimes \alpha_n$, the *thickening* of α denoted $T(\alpha)$ is the open set obtained by replacing each α_i in its tensor expansion that is a 0-cell $[x]$ by $[x - 3^{-j-1}, x + 3^{-j-1}]$. Note that for an n-cell β, $T(\beta) = \beta$. Denote the collection of thickened cells by $T(I(n, j))$. For some j_0 large enough we have that any element in $T(I(n, j_0))$ has diameter smaller than the Lebesgue number of the covering $\{\mathcal{O}_i\}$ and thus each element of $T(I(n, j_0))$ is contained in some $\mathcal{O}_i$. We set our refinement $\mathcal{U}_i$ to then be the collection $T(I(n, j_0))$. Because each thickened cell $T(\alpha)$ only intersects those faces β such that α and β are faces of a common cell γ, we easily obtain (2). For (1), (3), and (4), we can now invoke the following combinatorial lemma of Pitts:

Lemma 5.2. *(Proposition 4.9 in* [P]*) For each $\sigma \in I(n, j)$, let $A(\sigma)$ be an $(3^n)^{3^n}$-admissible family of annuli. Then there exists a function F which assigns to each $\sigma \in I(n, j)$ an element $F(\sigma) \in A(\sigma)$ such that $A(\sigma)$ and $A(\tau)$ are disjoint whenever σ and τ are faces (possibly of different dimensions) of some common cell $\gamma \in I(n, j)$.*

To apply Lemma 5.2 to our setting, set $L = (3^n)^{3^n}$ and for each cell $\alpha \in I(n, j_0)$ we choose some L-admissible family $A(\alpha)$ from one of the $\mathcal{O}_i$ containing it (it does not matter which).

This then determines a map A and the lemma applies to give an annulus associated to each cell α and thus to each to thickened cell $T(\alpha)$ satisfying (4). Item (3) then follows easily. □

While Lemma 5.1 gives an almost minimizing property in many annuli about each point, we need the property for all annuli sufficiently small:

Lemma 5.3. *Some subsequence of the min-max sequence produced by Lemma 5.1 is almost minimizing in sufficiently small annuli.*

Proof. We first show the following claim:

(1*) Given a point $p \in M$, and a min-max sequence Σ_j that is a.m. in all L-admissible familes of annuli about p, one can find a positive number $r(p)$ and a subsequence of Σ_j that is a.m. in sufficiently small annuli of outer radius at most $r(p)$ about p.

To prove (1*), first fix an L-admissible family $\mathcal{F}$ of annuli about p. By the pigeonhole priniciple Σ_j is a.m. in one of the annuli comprising $\mathcal{F}$ for infinitely many j. Fix the outermost such annulus $A(r,s)$ and a pass to the subsequence (not relabeled) of the Σ_j that are a.m. in $A(r,s)$. Now consider the family $\mathcal{F}'$ consisting of the annuli in $\mathcal{F}$ exterior to $A(r,s)$, the annulus $A(r/2,s)$ as well as the annuli in $\mathcal{F}$ interior to $A(r,s)$ whose inner and outer radii are multiplied by $1/2$ (such a family is still admissible). By the a.m. property infinitely many of the Σ_j must be a.m. in one of the annuli in the family $\mathcal{F}'$. Again choose the outermost such annulus (which can be no further out than $A(r/2,s)$ by construction) and pass to a subsequence a.m. in this new annulus. Iterating this procedure and a diagonal argument gives a min-max sequence a.m. in *all* annuli sufficiently small about p, proving (1*).

To prove Lemma 5.3, we combine (1*) with a Vitali-type covering argument. For each $p \in M$ and given a min-max sequence Σ_j, let $m(p)$ denote the supremum of positive numbers η so that some subsequence of Σ_j is a.m. in annuli with outer radius η. By (1*), $m(p) > 0$. Set $r_m(p) := m(p)/2$. Choose some $p_1 \in M$ with $r_m(p_1) > \frac{1}{2}\sup_{q\in M} r_m(q)$. Then pass to a subsequence of Σ_j (not relabeled) that is a.m. in annuli with outer radius at most $r_m(p_1)$. Choose then $p_2 \in M \setminus B_{r_m(p_1)}(p_1)$ so that

$$r_m(p_2) > \frac{1}{2} \sup_{q\in M\setminus B_{r_m(p_1)}(p_1)} r_m(q) \tag{5.1}$$

and pass to a further subsequence of Σ_j a.m. in annuli about p_2 and outer radius at most $r_m(p_2)$. Iterating this procedure gives rise in finitely many steps to a cover of the entire manifold with the desired properties since by the maximality of the construction and (1*) it cannot happen that $r_m(p_i) \to 0$. □

References

[Ba] D. Bachman , *Stabilizing and destabilizing Heegaard splittings of sufficiently complicated 3-manifolds*, Math. Ann.**355** (2013), no. 2, 697–728.

[Bir] J. Birman, *On the equivalence of Heegaard splittings of closed, orientable 3-manifolds*, Ann. of Math. Studies **84** (1975), 137–164.

[BO] F. Bonahon & J. Otal, *Scindements de Heegaard des espaces lenticulaires*, Ann. Sci. Ecole Norm. Sup. (4) **16** (1983) 451–466.

[Br] W. Breslin, *Curvature bounds for surfaces in hyperbolic 3-manifolds*, Canad. J. Math. **62** (2010), 994–1010

[CaGo] A. Casson & C. Gordon, *Reducing Heegaard splittings*, Topol. Appl. **27** (1987), 275–283.

[CD] T.H. Colding & C. De Lellis. *The min-max construction of minimal surfaces*, Surveys in Differential geometry VIII, Intl. Press, Somerville, 2013.

[CG] T.H. Colding & D. Gabai, *Effective Finiteness of irreducible Heegaard splittings of non Haken 3-manifolds*, arXiv:1411.2509.

[HTT] J. Hass, A. Thompson, W. Thurston, *Stabilization of Heegaard splittings*, Geom. Topol. **13** (2009), 2029–2050.

[Jo] K. Johannson *Topology and combinatorics of 3-manifolds*, Lecture Notes in Mathematics, **1599** (1995), Springer-Verlag, Berlin.

[Jon] J. Johnson, *Calculating isotopy classes of Heegaard splittings*, arXiv:1004.4669.

[Jon2] J. Johnson, *Bounding the stable genera of Heegaard splittings from below*, J. Topol. **3** (2010), no. 3, 668–690.

[K] D. Ketover, *Degeneration of min-max sequences in* 3*-manifolds*, arXiv:1312.2666.

[LB] F. Laudenbach & S. Blank, *Isotopie de formes fermes en dimension trois.*, Invent. Math. **54** (1979), no. 2, 103–177.

[Li1] T. Li, *Heegaard surfaces and measured laminations I: The Waldhausen conjecture*, Invent. Math. **167** (2007), 135-177.

[Li2] T. Li, *Heegaard surfaces and measured laminations, II: non-Haken 3-manifolds*, J. AMS **19** (2006), 625-657.

[Li3] T. Li, *An algorithm to determine the Heegaard genus of a 3-manifold*, Geometry & Topology **15** (2011), 1029-1106.

[Mor] Y. Moriah, *Heegaard splittings of Seifert fibered spaces*, Invent. Math. **91** (1988), 465–481.

[ML] Y. Moriah & M. Lustig, *Nielsen Equivalence in Fuchsian groups and Seifert fibered spaces*, Topology, **30** (1991), 191–204.

[MS] Y. Moriah & J. Schultens, *Irreducible Heegaard splittings of Seifert fibered spaces are either vertical or horizontal*, Topology, **37** (1998), 1089–1112.

[P] J. Pitts, *Existence and Regularity of minimal surfaces on Riemannian manifolds*, Mathematical Notes 27, Princeton University Press, Princeton, 1981.

[PR] J. Pitts & J. Rubinstein, *Applications of minimax to minimal surfaces and the topology of 3-manifolds*, Proc. Center Math. Applic., Australian National University **12**, (1987), 137–170.

[Re] K. Reidemeister, *Zur dreidimensionalen Topologie*, Abh. Math. Sem. Univ. Hamburg **9** (1933), 189–194.

[RS] H. Rubinstein and M. Scharlemann, *Comparing Heegaard splittings of non-Haken 3-manifolds* Topology **35** (1996), 1005–1026.

[Sc] M. Scharlemann, *Local detection of strongly irreducible Heegaard splittings* Topol. App. **90** (1998), 135–147.

[Sch] J Schultens, *The stabilization problem for Heegaard splittings of Seifert fibered spaces*, Topology Appl. **73** (1996) 133–139.

[Sed] E. Sedgwick, *The irreducibility of Heegaard splittings of Seifert fibered spaces*, Pac. J. Math., **190** (1990), 173–190.

[Si] J. Singer, *Three-dimensional manifolds and their Heegaard-diagrams*, Trans. Amer. Math. Soc. **35** (1933), 88–111.

[SS] F. Smith. *On the existence of embedded minimal 2-spheres in the 3-sphere, endowed with an arbitrary Riemannian metric*, PhD thesis, University of Melbourne (1982).

[Wa] F. Waldhausen, *Some problems on 3-manifolds*, Proc. Symp. Pure Math. **32** (1978), 313–322.

[Wa2] F. Waldhausen, *Heegaard Zerlegungen der 3-sphare*, Topology **7** (1968), 195–203.

Department of Mathematics, MIT, Cambridge, MA 02139-4307

Department of Mathematics, Princeton University, Princeton, NJ 08544

E-mail address: colding@math.mit.edu, gabai@math.princeton.edu and dketover@math.princeton.edu

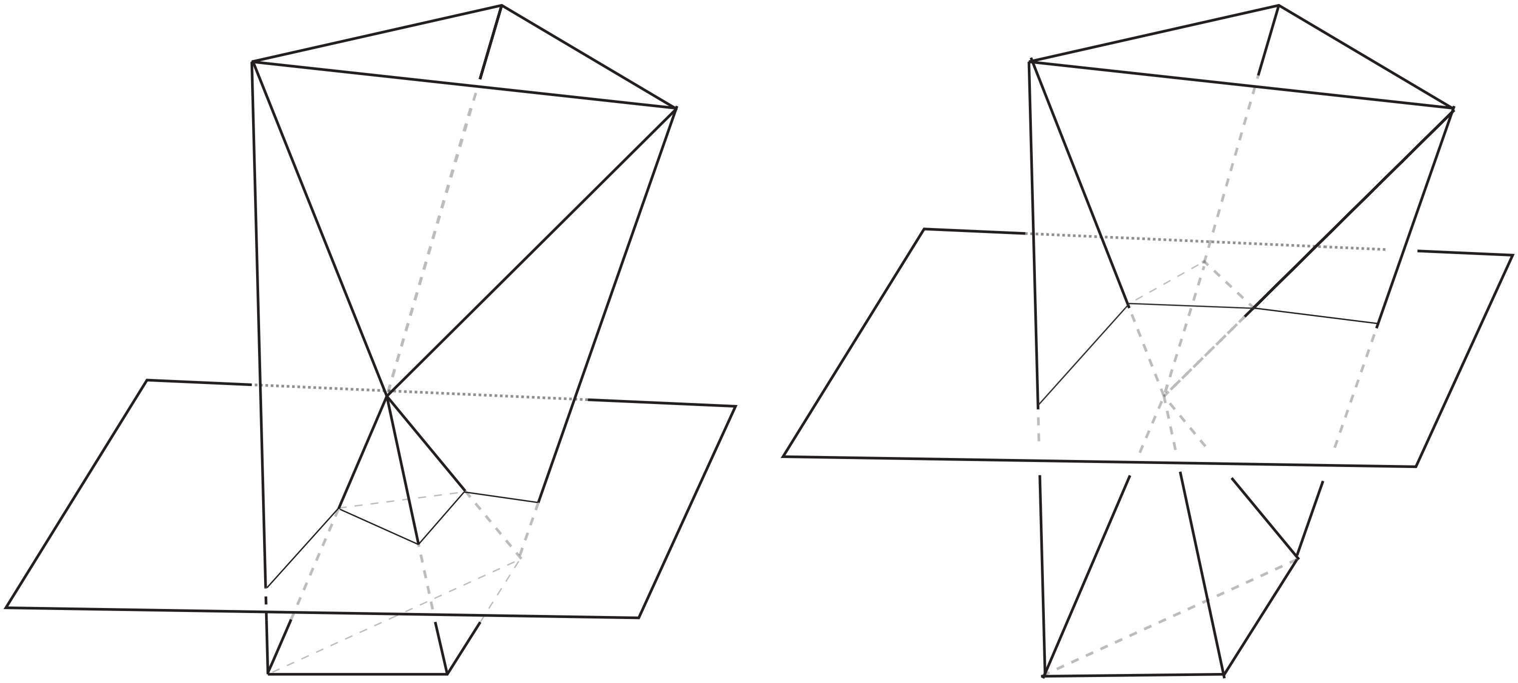

Pushing through a vertex

S_- Figure A S_+

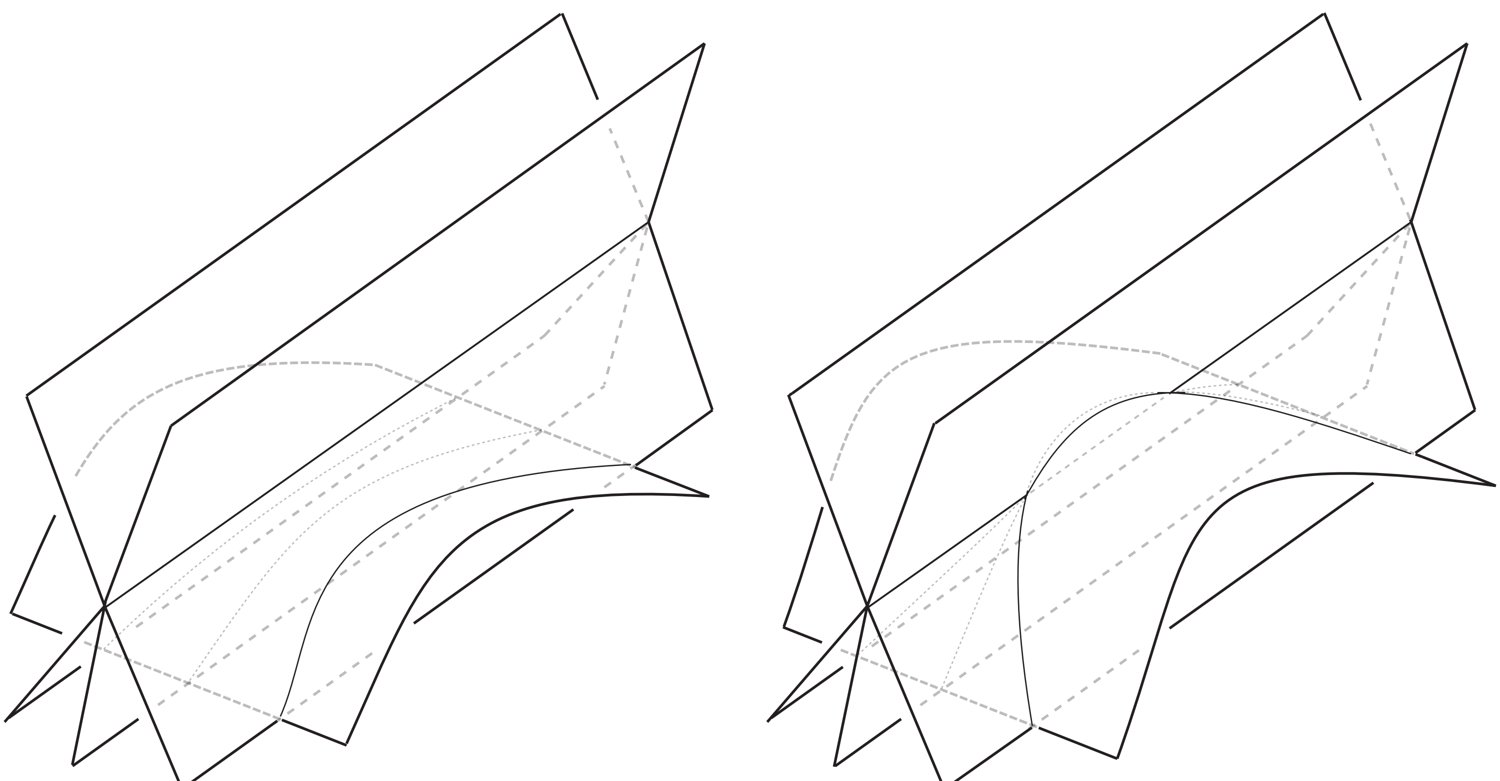

Pushing through a 1-simplex

S_- S_+

Figure B

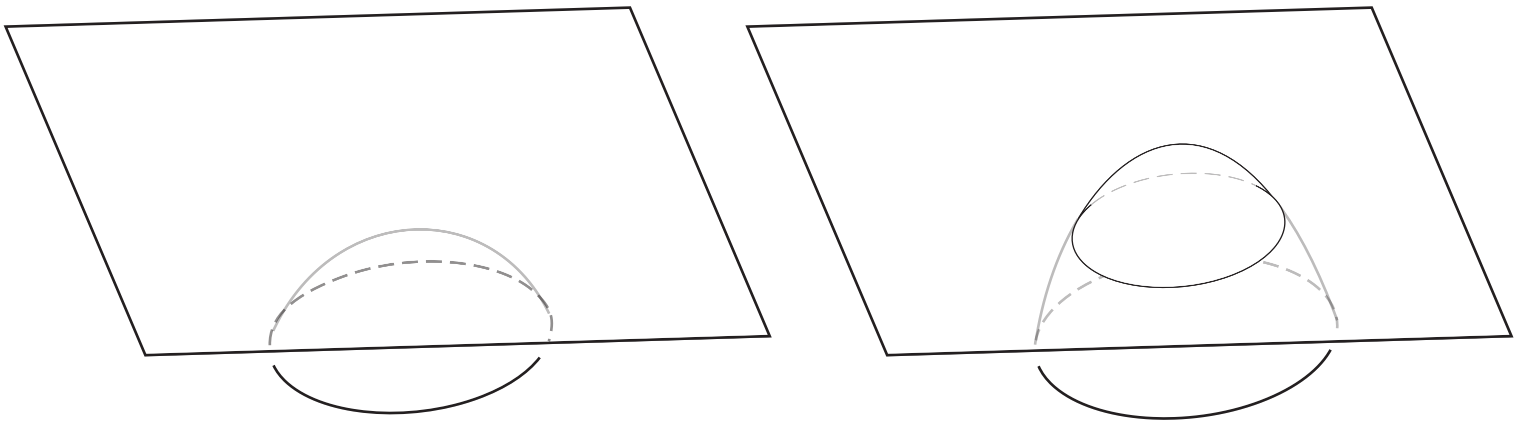

S₋ Pushing a disc through a 2-simplex S₊

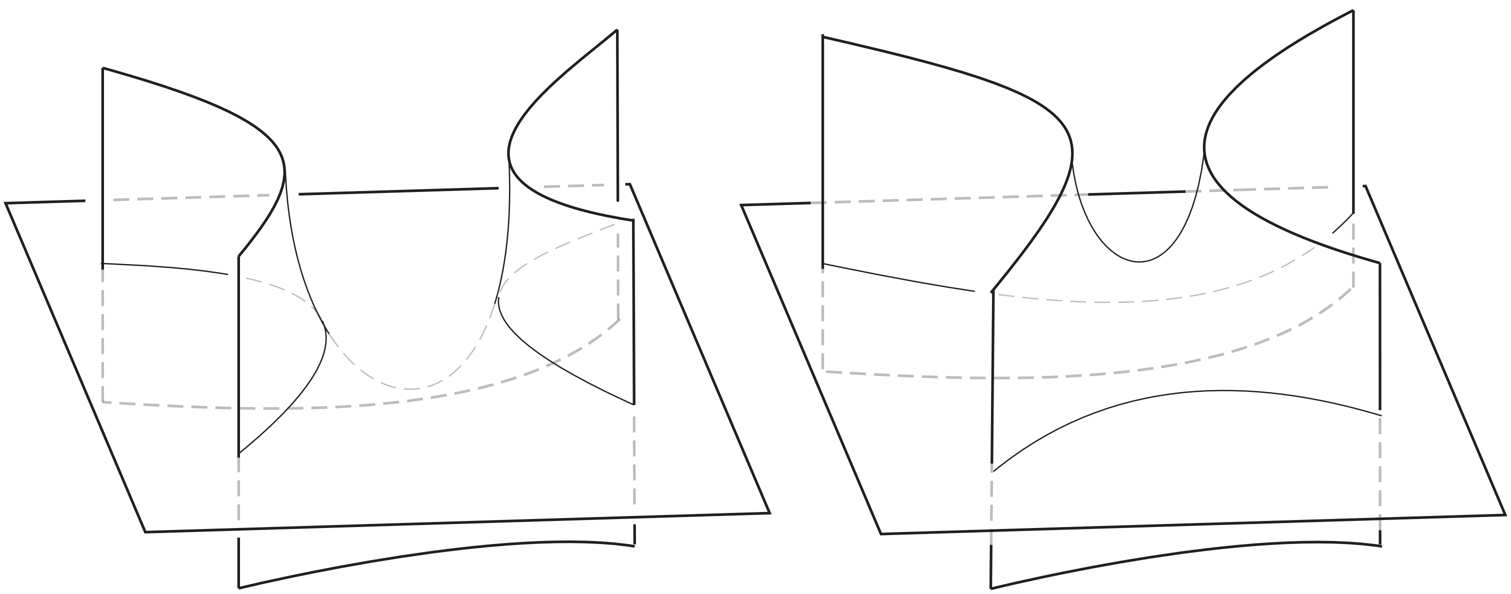

S₋ S₊

Pushing a saddle through a 2-simplex

Figure C