

A Two Dimensional Description of Heegaard Splittings

A Dissertation presented

by

Chandrika Sadanand

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Doctor of Philosophy

in

Mathematics

Stony Brook University

May 2017

ProQuest Number: 10283849

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10283849

Published by ProQuest LLC (2017). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

Stony Brook University

The Graduate School

Chandrika Sadanand

We, the dissertation committee for the above candidate for the

Doctor of Philosophy degree, hereby recommend

acceptance of this dissertation

Dennis Sullivan - Dissertation Advisor
Distinguished Professor, Department of Mathematics

Maira Chas - Chairperson of Defense
Associate Professor, Department of Mathematics

Olga Plamenevskaya
Associate Professor, Department of Mathematics

Martin Roček
Professor, Department of Physics

This dissertation is accepted by the Graduate School

Charles Taber
Dean of the Graduate School

Abstract of the Dissertation

A Two Dimensional Description of Heegaard Splittings

by

Chandrika Sadanand

Doctor of Philosophy

in

Mathematics

Stony Brook University

2017

Consider the following two ways to decompose 3-manifolds:

- (i) A compact 3-manifold can be decomposed by a Heegaard splitting into two well-understood, homeomorphic manifolds, glued along their boundary.
- (ii) A compact orientable 3-manifold can be decomposed uniquely as a connect sum of prime 3-manifolds.

Stallings described Heegaard splittings using classes of continuous maps between surfaces and two dimensional complexes. He studied the Poincaré conjecture with these maps using group theory. This dissertation considers these maps more literally, using geometric and topological arguments. One might wonder if these maps fold the surface, or crush handles. We find that minimal genus Heegaard splittings of prime 3-manifolds (with a couple exceptions) can be described by locally injective two dimensional maps. These locally injective maps induce families of conformal structures, and also square complex structures, on the domain surface. The 3-manifold, together with a minimal Heegaard splitting, can be recovered from any member of a family. The construction on primes can be “sewn” together to make a statement for all Heegaard splittings and arbitrary compact orientable 3-manifolds.

For Aee, Appa, and Kala Athai

Contents

List of Figures	vi
1 Introduction	1
2 Background	3
2.1 Heegaard Splittings	3
2.2 Reduction of the Poincaré conjecture to Two Dimensions . . .	10
2.3 Square Complexes	14
3 Main Theorem	19
4 Corollaries	35
4.1 Reducible 3-Manifolds	35
4.2 Square Complexes	38
4.2.1 Handle Slides	39
4.2.2 Relation to Work of Haglund and Wise	41
4.3 Conformal Structures and Quadratic Differentials	45
4.3.1 Tautness	47
4.3.2 Relation to Work of Wolf	49
5 Examples	53
5.1 Lens Spaces	53
5.2 Three Dimensional Torus	58
6 Bibliography	60

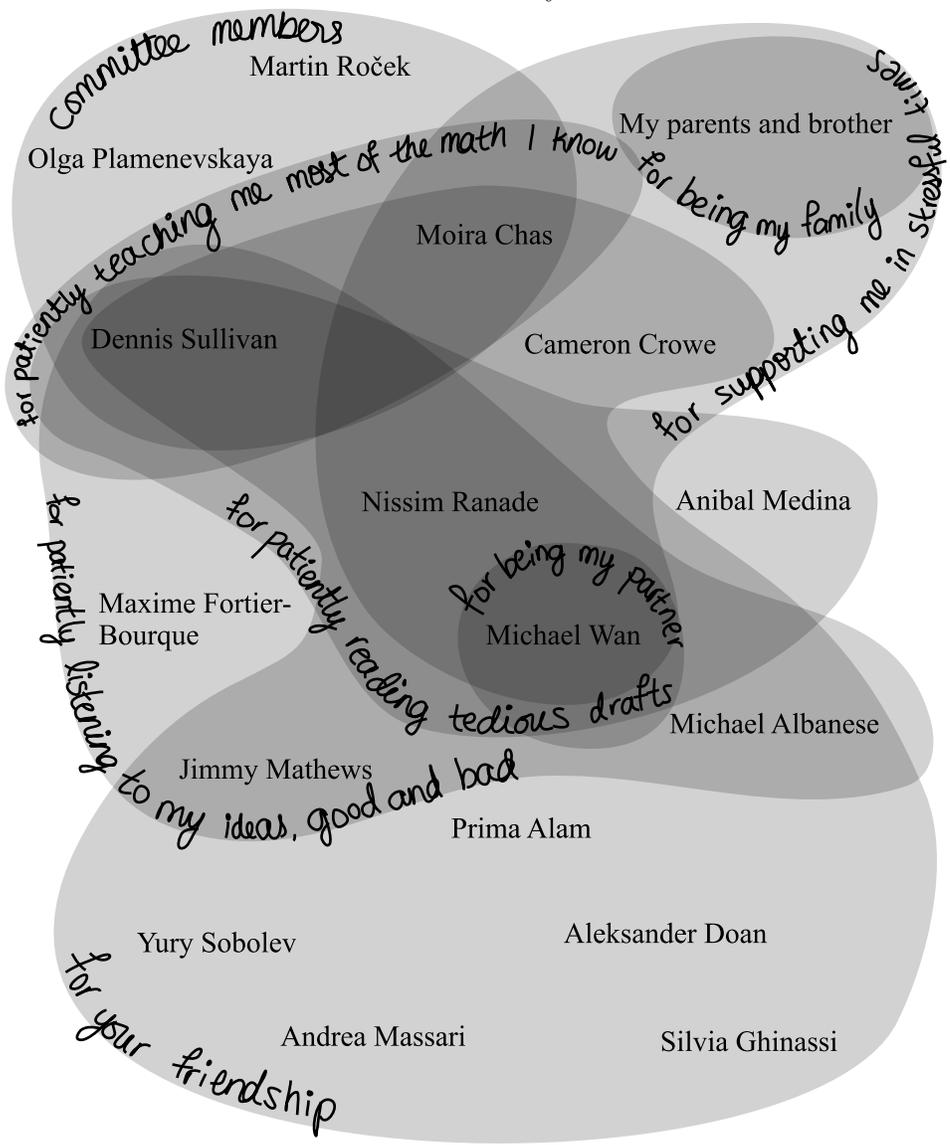
List of Figures

2.1	A handle body of genus four.	4
2.2	Genus four Heegaard splitting of S^3	5
2.3	Genus three Heegaard splitting of S^3	5
2.4	Stabilization.	6
2.5	A handle body of genus four, shown with representatives of a basis for the first homology of the boundary.	7
2.6	A handle body shown as a solid 3-ball with pairs of closed disks on the boundary identified.	8
2.7	A handle slide between two Heegaard diagrams for the genus three Heegaard splitting of S^3	9
2.8	An essential simple closed curve on the splitting surface that is null homotopic in each of the handle bodies, in the genus three Heegaard splitting of S^3	11
2.9	The local topology around various points in $(\vee_g S^1) \times (\vee_g S^1)$	15
3.1	A 3-ball with some closed disks removed on the boundary, obtained by removing disks from a handle body. Letter markings show where the disks used to be.	20
3.2	Half of a Heegaard diagram.	21
3.3	Two regions of Heegaard diagrams. The left is not taut, while the right is taut.	22
3.4	An orange segment that intersects the same side of a diagram curve in two connected components.	23
3.5	Representatives of the free homotopy classes $[\alpha]$ and $[\beta]$, shown in two ways: (i) as a union of segments of the problematic curve and the orange curve, and (ii) as the cuffs of a pair of pants constructed by the problematic curve and the orange segment.	24
3.6	The suggested change to the red half of the Heegaard diagram.	24
3.7	A neighbourhood of the problematic curve before and after the suggested swaps, this time with a few more blue arcs.	25

3.8	The embedding of the pair of pants introduced in Figure 3.5.	26
3.9	S is shown before and after identifications.	26
3.10	A pair of pants decomposition of S	27
3.11	A region of Σ_g is shown with orientations chosen on the Heegaard diagram curves. The two coordinates of $(\vee_g S^1) \times (\vee_g S^1)$ are shown with blue and red points chosen as well as orientations on the edges.	28
3.12	The definition of f on the interior of a square.	29
3.13	The behaviour of edge corners and square corners.	30
3.14	A genus two example of one coordinate of f	32
3.15	The cone structure of half of a solid cylinder.	33
4.1	Two disjoint embedded S^2 in the genus three Heegaard splitting of S^3 , such that the complement is three genus one Heegaard splittings of S^3 (when the boundaries are coned off).	36
4.2	An example of a handle slide with square complexes.	41
4.3	Four pathologies of the 2-skeleton of a cube complex.	41
4.4	f is not a local isometry.	45
4.5	A surface of genus three shown with the fibres of the two coordinates of a map describing a genus three Heegaard splitting of T^3	47
4.6	The foliation in red cannot be induced by an abelian differential.	49
5.1	$S^1 \times S^1$ shown with a blue point and a red point selected on each coordinate circle.	54
5.2	Pre-images of the red and blue points.	55
5.3	Fundamental domains in the universal cover of the splitting surface.	56
5.4	A Heegaard splitting for a lens space with fundamental group $\mathbb{Z}/13\mathbb{Z}$	57
5.5	Connected components of the complement of the Heegaard diagram.	58
5.6	A genus three Heegaard splitting of T^3	59
5.7	Coordinate fibres of a locally injective map describing a genus three Heegaard splitting of T^3	59

Acknowledgements

Thank you



Chapter 1

Introduction

This dissertation investigates a two dimensional geometric equivalent of a group theoretic construction of Stallings. He used this construction to reformulate the Poincaré conjecture [14]. The relationship to the Poincaré conjecture suggests a rich structure in Stallings' construction that I seek to understand in dimension two. This investigation leads to geometric questions about two dimensional mappings that are directly related to three manifolds.

The geometric construction uses Heegaard splittings to associate, to each 3-manifold, a class of continuous maps between a surface and a two dimensional complex from which the 3-manifold can be recovered. This is explained in Chapter 2. We say these maps *describe* the 3-manifold. One might wonder how these maps are related to each other, or if they fold the surface, or crush handles. In Chapter 3, we show folds, crushing and branch covering can be avoided when π_2 of the manifold is trivial. Such 3-manifolds are called *irreducible*.

Theorem. *There is a local embedding describing each irreducible 3-manifold.*

$S^2 \times S^1$ is the only 3-manifold that has non-trivial π_2 that cannot be decomposed non-trivially by connect sum. It is known that every compact orientable 3-manifold that does not have $S^2 \times S^1$ as a connect summand is given by connect sums of irreducible 3-manifolds. The local embeddings in the theorem above can be “sewn” together to give maps representing this larger class of manifolds. These “nice” maps give rise to families of two dimensional structures corresponding to each 3-manifold. Definitions of the terms below are given in Chapter 4, where these corollaries are explained.

Corollary. *For every compact oriented 3-manifold that does not have $S^2 \times S^1$ as a connect summand, there is a family of non-positively curved square*

complexes that have the topology of surfaces that have been glued to each other along a finite set of points.

Corollary. *For every compact oriented 3-manifold that does not have $S^2 \times S^1$ as a connect summand, there is a family of surfaces glued together at finitely many points, equipped with complex structures and quadratic differentials on each surface component.*

Finally, in Chapter 5, some examples are discussed.

Chapter 2

Background

In 1966, John Stallings published a paper titled “How not to prove the Poincaré conjecture” [14]. Here he outlines a failed argument to show that any simply connected 3-manifold must be S^3 . He shares the ideas “in order to deter others from making similar mistakes,” but also to spread the insight gained from this mistake [14]. Indeed, the paper has captured the attention and imagination of topologists over the years.

In the paper, Stallings reduces the Poincaré conjecture to two dimensions and then approaches it with group theory, showing where the argument cannot continue. This dissertation stops to explore the two dimensional scenario in a literal geometric sense. The motivation here is the idea that it must hold interesting structure if it was used to study the Poincaré conjecture.

In what follows, we begin by introducing the main tool in this discussion (Section 2.1). Then we show how this tool was used by Stallings to reduce the Poincaré conjecture to a statement involving continuous maps between two dimensional spaces (Section 2.2). This brings us to the setting of the investigation in this dissertation: continuous maps between two dimensional spaces. Finally we introduce some ideas that will be useful in understanding the main result (Section 2.3).

2.1 Heegaard Splittings

Suppose that M is a compact 3-manifold without boundary.

Definition 1. A *Heegaard splitting* or *Heegaard decomposition* of M is a decomposition $M = H_1 \cup H_2$, where H_1 and H_2 are homeomorphic 3-manifolds

that intersect in their shared boundary, which is a surface of genus g . Further, each H_i is a *handle body*. Namely, H_1 and H_2 are the result of taking a closed 3-ball and attaching g 1-handles, or solid cylinders.

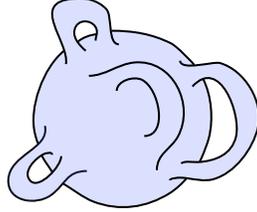


Figure 2.1: A handle body of genus four.

The shared boundary $\partial H_1 = \partial H_2$ will be called the *splitting surface* Σ_g , where g is the genus.

This is an elegant decomposition because each of the parts, the handle bodies and the splitting surface, are well-understood topological spaces. Heegaard splittings were first described by Poul Heegaard in his 1898 thesis [5]. The rest of this section describes his work and ideas.

Theorem 1. (*Heegaard*) *Every compact 3-manifold without boundary has a Heegaard decomposition.*

Proof sketch. One can obtain a Heegaard splitting of a 3-manifold by finitely triangulating it, and then thickening the 1-skeleton and the dual 1-skeleton until they meet along a surface. The genus of the two resulting handle bodies is found to be the same using the Euler characteristic and taking Poincaré duality into account.

Another option is to work in the smooth category, taking a Morse function on the 3-manifold, f , so that $f(\text{critical point of degree } i) = i$. A Heegaard splitting is obtained by taking the pre-image of $\frac{3}{2}$ as a splitting surface, since the two components of the complement of this surface are open handle bodies. \square

Two Heegaard splittings of a 3-manifold are equivalent if there is an automorphism that sends handle bodies to handle bodies and splitting surface to splitting surface. From the above proofs, it follows that each 3-manifold has infinitely many Heegaard splittings (up to equivalence).

Example. Below, S^3 is shown as the one-point compactification of \mathbb{R}^3 and as the union of two handle bodies of genus four. For better viewing, the two handle bodies have been drawn so that there is some space between them. This is schematic only. In a Heegaard splitting, the boundaries of the two handle bodies are the same set.

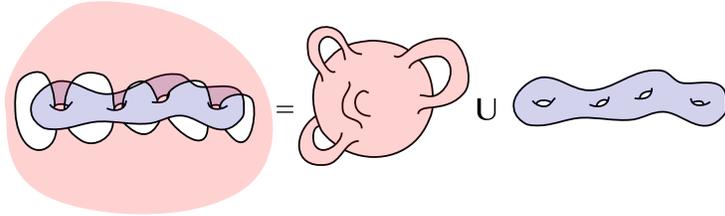


Figure 2.2: Genus four Heegaard splitting of S^3 .

Below, a genus three splitting of S^3 is shown. One can imagine a similar picture to describe a Heegaard splitting of any genus of S^3 .

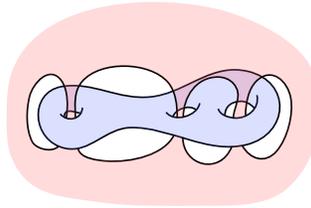
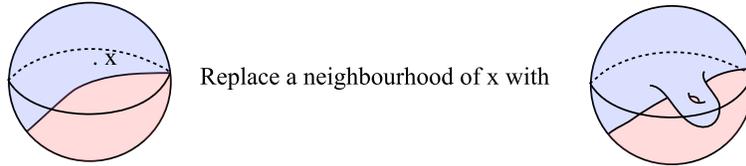


Figure 2.3: Genus three Heegaard splitting of S^3 .

Waldhausen showed in 1968 that all Heegaard splittings of S^3 have the form of the examples above [17].

Given a Heegaard splitting, it is possible to build a new splitting for the same 3-manifold with one more handle. This procedure involves removing a 3-ball from the manifold that intersects the splitting surface in a disk, and replacing it with a 3 ball that contains a torus with one boundary component. In this way, the topology of the manifold is not changed, but the splitting is. The picture below show how this can be done, beginning with any point x on the splitting surface.



Replace a neighbourhood of x with

Figure 2.4: Stabilization.

Definition 2. This procedure to increase the genus is called *stabilization*.

Sometimes, it is possible to do this process in reverse. More precisely, we mean to find a 3-ball that intersects the splitting surface in a torus with one boundary component, and replace it with a 3-ball that has a disk with the same boundary embedded inside.

Definition 3. The reverse procedure (decreasing genus) is called *destabilization*.

Remark 1. Given a Heegaard splitting, one can always find a 3-ball that intersects the splitting surface in a disk. Such a ball can be found in a small enough neighbourhood of any point on the splitting surface. However, one cannot always find a 3-ball that intersects the splitting surface in a torus with one boundary component. Thus, one can always stabilize, but destabilization is not always possible.

Theorem 2. (*Reidemeister-Singer [12] [13]*) *Beginning with a Heegaard splitting of a 3-manifold, one can obtain any other splitting of this manifold by performing a finite sequence of stabilizations and destabilizations.*

Definition 4. An invariant of a manifold is the smallest number of handles of a Heegaard splitting. This is called the *Heegaard genus*.

The one dimensional homology group of Σ_g is \mathbb{Z}^{2g} . Consider representatives of a standard generating set for this group, shown below on a surface of genus four as blue and black curves. Now consider Σ_g included into a handle body as its boundary. Notice that half of these representatives are now null-homologous. In the example shown below, the handle body is seen shaded and the blue curves are null homologous when included into it. Therefore, given a map describing a Heegaard decomposition, the coordinate projections

each kill half of the first homology. There are instances when they are the same half, for example maps describing the genus one Heegaard decomposition of $S^1 \times S^2$. There are also instances when they are disjoint halves, for example, maps describing the standard genus g Heegaard decomposition of the 3-sphere. Finally there are instances where the halves intersect non-trivially as well, such as maps describing the genus three Heegaard decomposition of the three dimensional torus.

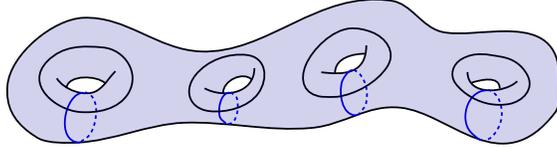


Figure 2.5: A handle body of genus four, shown with representatives of a basis for the first homology of the boundary.

Definition 5. A *Heegaard diagram* describing a Heegaard decomposition is a triple $(\Sigma_g, \mathcal{A}, \mathcal{B})$, where \mathcal{A} and \mathcal{B} are sets of g curves on a surface of genus g , Σ_g , with the following properties.

1. The curves in \mathcal{A} are a maximal set of homologically non-trivial, independent, disjoint representatives each with null homotopic image under ι_1 , the inclusion of Σ_g into the handle body H_1 .
2. the curves in \mathcal{B} are a maximal set of homologically non-trivial, independent disjoint representatives each with null homotopic image under ι_2 , the inclusion of Σ_g into the handle body H_2 .

It is common to call the free homotopy classes in \mathcal{A} *blue diagram curves* and those in \mathcal{B} *red diagram curves*. They will also be depicted in these colours in figures.

Remark 2. Using only the properties (1), (2) above, each of \mathcal{A} and \mathcal{B} must have g elements by the argument in the paragraph preceding this definition.

Definition 6. Two Heegaard diagrams $(\Sigma, \mathcal{A}, \mathcal{B})$ and $(\Sigma', \mathcal{A}', \mathcal{B}')$ are *equivalent* if there is a homeomorphism between the surfaces that takes the free homotopy classes of curves in \mathcal{A} to the free homotopy classes of curves in \mathcal{A}' as sets, and similarly for \mathcal{B} and \mathcal{B}' . They are also equivalent if there is a homeomorphism between the surfaces that takes the free homotopy classes of curves in \mathcal{A} to the free homotopy classes of curves in \mathcal{B}' , and similarly for \mathcal{B} and \mathcal{A}' .

The curves in \mathcal{A} dictate the attachment of H_1 to Σ_g with the following procedure. Cut Σ_g along the curves of \mathcal{A} . This yields a sphere with $2g$ disks removed, because the curves are homologically non-trivial and independent and because there are g of them. Label all the boundary components by the curve in \mathcal{A} they were cut from. Glue $2g$ disks to this space to form a sphere. Attach a 3-ball to the sphere. Finally identify pairs of the glued on disks according to the labeling on their boundary. The picture below shows this last step, where the letters label pairs of disks to be identified.

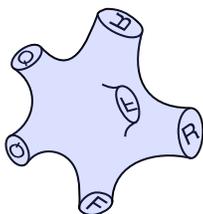


Figure 2.6: A handle body shown as a solid 3-ball with pairs of closed disks on the boundary identified.

Similarly, the curves in \mathcal{B} dictate the attachment of H_2 to Σ_g , and thus one can reconstruct a Heegaard splitting from a diagram.

Remark 3. There are infinitely many Heegaard diagrams describing a single Heegaard decomposition, (up to the equivalence described above).

A simple way to produce several Heegaard diagrams for the same Heegaard decomposition is to perform what is called a handle slide. The definition below is from “An introduction to Heegaard Floer homology,” by Ozsváth and Szabó [10].

Definition 7. Two Heegaard diagrams $(\mathcal{A}, \mathcal{B})$ and $(\mathcal{A}', \mathcal{B}')$ on a surface of genus g differ by a *handle slide* if $\mathcal{B} = \mathcal{B}'$, and $\mathcal{A} \setminus \{\gamma\} = \mathcal{A}' \setminus \{\gamma'\}$ such that for some $\delta \in \mathcal{A}$, δ , γ , and γ' bound a pair of pants. The roles of \mathcal{A} and \mathcal{B} may be reversed as well.

An example of a handle slide is shown below. The curves in \mathcal{A} are shown in blue, and the curves in \mathcal{B} are shown in red.

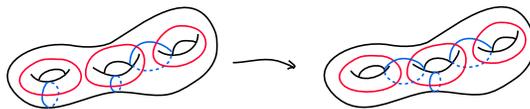


Figure 2.7: A handle slide between two Heegaard diagrams for the genus three Heegaard splitting of S^3 .

To see that performing a handle slide does not change the attachment of a handle body, consider the following. There are disks attached to the splitting surface along γ and δ that are contained in the handle body noted by \mathcal{A} . Notice that there is also an embedded disk in the handle body attached to the splitting surface along γ' . This is best seen by imagining γ in the picture of a 3-ball with pairs of disks labeled for identification (see Figure 2.6). Cutting along the disk attached along γ' , and then gluing along the pair of disks labelled by γ , we arrive at the handle body noted by \mathcal{A}' , but also, we see that it is the same handle body noted by \mathcal{A} .

Theorem 3. (*Reidemeister-Singer [12] [13]*) *Given two Heegaard diagrams for a Heegaard decomposition, one can be obtained from the other by performing a sequence of handle slides.*

When considering a Heegaard diagram, it will be helpful to consider representative curves that intersect minimally. This can be achieved by picking an arbitrary hyperbolic structure on Σ_g , and moving the diagram curves by isotopy to their geodesic representatives. Then the curves intersect transversally and minimally.

Remark 4. This implies that there are finitely many intersection points (since if a pair curves intersected infinitely many times, there would be an accumulation point of these intersections where transversality would be contradicted). There also must be finitely many connected components of the complement of the diagram curves in the surface. We can bound the number of connected components because the number of intersections and number of diagram curves is finite. Every diagram curve segment that is not interrupted by intersections is locally on the boundary of two (not necessarily distinct) connected components of the complement.

2.2 Reduction of the Poincaré conjecture to Two Dimensions

In this section we restrict our attention to orientable 3-manifolds, and explain Stallings' construction relating Heegaard splittings to continuous maps between two dimensional spaces.

The following definition and ideas are directly from Milnor's 1962 paper, "A unique decomposition theorem for 3-manifolds," [8].

Definition 8. A compact, orientable 3-manifold, M , is *prime* if for any 3-manifolds, M_1 and M_2 , $M = M_1 \# M_2$ implies at least one of M_1 or M_2 is the 3-sphere.

Theorem 4. (Kneser [7]) *Every compact, orientable 3-manifold can be written as a finite connect sum of prime 3-manifolds. This decomposition is unique up to reordering, homeomorphism, and adding S^3 as a summand.*

We conclude that if there were an exotic 3-sphere, its decomposition into primes would consist of prime exotic spheres. So it suffices to consider prime 3-manifolds.

Definition 9. A 3-manifold is *irreducible* if every embedded 2-sphere bounds a 3-ball. If a 3-manifold is not irreducible, it is *reducible*.

Remark 5. All primes except $S^2 \times S^1$ are irreducible.

This can be seen by considering a prime 3-manifold, M , that has an embedded S^2 that does not bound a 3-ball. The S^2 cannot separate M because that would contradict M being prime. Therefore there is an embedded S^1 that intersects the S^2 once transversally. We thicken the union of the embedded S^2 and S^1 inside M to give $S^2 \times S^1$ with a 3-ball deleted (this step is an exercise). This implies that $S^2 \times S^1$ is a connect summand of M . M , however is prime, so it must be $S^2 \times S^1$ itself.

Due to this consideration, it suffices to consider only irreducible 3-manifolds (since $S^2 \times S^1$ is not simply connected). There is also a notion of irreducibility of Heegaard splittings.

Definition 10. A Heegaard decomposition is said to be *irreducible* if every essential (i.e., not null homotopic) embedded closed curve on the separating surface Σ_g is not null homotopic when included into at least one of the handle bodies H_1 and H_2 . If a Heegaard splitting is not irreducible is it *reducible*.

Example. The $g = 3$ splitting of S^3 is shown below with an essential simple closed curve that is null homotopic when included into each of the handle bodies. This means that this Heegaard splitting is reducible. In fact, using the theorem of Waldhausen mentioned in the previous section, we can say that all Heegaard splittings of S^3 are reducible except the $g = 0$ and $g = 1$ splittings.

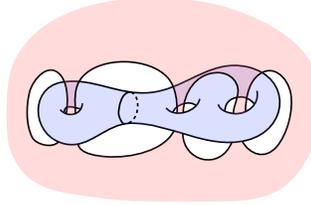


Figure 2.8: An essential simple closed curve on the splitting surface that is null homotopic in each of the handle bodies, in the genus three Heegaard splitting of S^3 .

Naively, the notions of reducibility of 3-manifolds and of Heegaard splittings are related because when a curve on Σ_g is null homotopic when included into both handle bodies, this creates an immersed S^2 in M . However, the relation is much deeper.

Theorem 5. (*Haken's Lemma 1968 [3]*) *Let $M = H_1 \cup H_2$ be a Heegaard splitting of a 3-manifold. Suppose there is an embedded S^2 in M that does not bound a 3-ball. Then there is an (other) embedded S^2 that does not bound a 3-ball in M that intersects the splitting surface in a simple closed curve.*

Remark 6. Every Heegaard splitting of a reducible 3-manifold is reducible. This is due to the preceding theorem.

This statement makes use of the sphere theorem of Papkyriakopoulos.

Theorem 6. (*Papakyriakopoulos 1957*) *If $\pi_2(M)$ is not trivial, there is an embedding of S^2 in M that is not null homotopic.*

Claim 1. *If a Heegaard splitting is reducible, then either the splitting can be destabilized, or the 3-manifold is reducible.*

Proof of Claim 1. Beginning with a reducible Heegaard splitting, there is an essential curve on the splitting surface that is null homotopic when included into both of the handle bodies. Disks can be immersed in each handle body to

realize these null homotopies. Thus, we have an immersed S^2 in M . If it does not represent a trivial element of $\pi_2(M)$, then the 3-manifold is reducible. If it is trivial, first homotope the immersed S^2 to an embedding. Consider the immersed disk that realizes the null homotopy of the reducing curve in one of the handle bodies. If it is already an embedding we are done. If not, lift to the universal cover of the handle body. Here, there are several pre-images of this disk that intersect. The universal cover of a handle body is homeomorphic to a closed 3-ball with a cantor set removed on its boundary. Here, the disks must be immersed with embedded boundaries. In this space, the disks can be moved equivariantly to not intersect each other and so that each one is embedded. Doing this to both disks, we obtain an embedded S^2 . Now since it represents a trivial element, it bounds, and the reducing curve must separate the splitting surface (since the S^2 separates M). Therefore the 3-ball must contain a component of the splitting surface with at least one handle. Finally if we remove this 3-ball and replace it with another that is split with no handle, it would be equivalent to destabilizing, possibly multiple times. \square

Remark 7. Given a reducible Heegaard splitting, the above claim gives way to find a “reducing” embedded sphere in M . Cutting along this sphere and “coning it off” will either destabilize the Heegaard splitting or produce non-trivial connect summands of M .

Remark 8. Consider the genus one Heegaard splitting of S^3 . This can be destabilized because it is possible to find a 3-ball that intersects the splitting surface in a torus with one boundary component. However, this boundary component is null homotopic when included into the splitting surface. In fact, it is not possible to find an essential simple closed curve on the splitting surface that is null homotopic when included into each of the handle bodies. Therefore this is an example of a Heegaard splitting that can be destabilized, but is irreducible.

Begin with an irreducible Heegaard splitting of an irreducible 3-manifold that is simply connected. It is equivalent to the Poincaré conjecture to state that it must be the genus one or genus zero decomposition of S^3 . Now Stallings notes that manifolds with Heegaard genus one (lens spaces) are classified, and none of them is an exotic 3-sphere. So he reformulates the conjecture as follows: If we have a Heegaard splitting of an irreducible 3-manifold that has genus $g \geq 2$, and the 3-manifold is simply connected, then the splitting must be reducible [14]. The simply connected condition is reduced to a group theory condition:

Proposition 1. *Let $M = H_1 \cup H_2$ be a Heegaard splitting, with splitting surface Σ_g , of genus g . Let $\iota_i : \Sigma_g \hookrightarrow H_i$, for $i = 1, 2$ be the inclusions of the*

splitting surface into the handle bodies.

M is simply connected if and only if $\iota_{1*} \times \iota_{2*} : \pi_1(\Sigma_g) \longrightarrow \pi_1(H_1) \times \pi_1(H_2)$ is surjective.

Proof. Consider the Seifert Van Kampen theorem on the Heegaard splitting $M = H_1 \cup H_2$. The fundamental group of M is the push out of $(\pi_1(\Sigma_g), \iota_1, \iota_2)$.

$$\begin{array}{ccc} F_g & \xrightarrow{b} & \pi_1(M) \\ \uparrow \iota_2 & & \uparrow a \\ \pi_1(\Sigma_g) & \xrightarrow{\iota_1} & F_g \end{array}$$

Where a is the homomorphism induced by the inclusion of H_1 into M , and b is induced by the inclusion of H_2 into M . We also have the following commutative diagram, where p_i are the coordinate projections.

$$\begin{array}{ccc} F_g & \xleftarrow{p_2} & F_g \times F_g \\ \uparrow \iota_2 & \nearrow \iota_1 \times \iota_2 & \downarrow p_1 \\ \pi_1(\Sigma_g) & \xrightarrow{\iota_1} & F_g \end{array}$$

If $\iota_1 \times \iota_2$ is surjective, then the above diagrams can be put together in a single commutative diagram.

$$\begin{array}{ccccc} & & F_g & \xrightarrow{b} & \pi_1(M) \\ & & \uparrow p_2 & & \uparrow a \\ & & F_g \times F_g & \xrightarrow{p_1} & F_g \\ \nearrow \iota_2 & & & & \\ \pi_1(\Sigma_g) & \xrightarrow{\iota_1 \times \iota_2} & & & \\ \searrow \iota_1 & & & & \end{array}$$

The push out of a cartesian product is trivial, so M is simply connected.

Now suppose M is simply connected. This implies that for every $\alpha \in \pi_1(\Sigma_g)$ such that $\iota_1(\alpha) \neq \mathbf{1}$, we must have $\iota_2(\alpha) = \mathbf{1}$. The symmetric statement with 1 and 2 switched is true as well. Therefore $\iota_1(\pi_1(\Sigma_g)) \times \mathbf{1}$ and $\mathbf{1} \times \iota_1(\pi_1(\Sigma_g))$ are in the image of $\iota_1 \times \iota_2$. Since ι_i are surjective, these are have $F_g \times \mathbf{1}$ and $\mathbf{1} \times F_g$. This generates $F_g \times F_g$, so $\iota_1 \times \iota_2$ is surjective. \square

So, the Poincaré conjecture is equivalent to: if $M = H_1 \cup H_2$ is a Heegaard splitting of an irreducible 3-manifold with genus $g \geq 2$, and $\iota_{1*} \times \iota_{2*}$ is surjective, the splitting must be reducible.

Finally, by definition a Heegaard splitting is reducible if and only if there is an essential embedded closed curve that is null homotopic when included into both of the handle bodies. This is equivalent to a non-trivial element in $\ker(\iota_{1*} \times \iota_{2*})$ having an embedded representative.

The final statement equivalent to the Poincaré conjecture is then: if $M = H_1 \cup H_2$ is a Heegaard splitting of an irreducible 3-manifold with genus $g \geq 2$, and $\iota_{1*} \times \iota_{2*}$ is surjective, there is an element of non-trivial element in $\ker(\iota_{1*} \times \iota_{2*})$ that has an embedded representative.

Stalling's paper is about attempting to prove the following statement. If $\phi : \pi_1(\Sigma_g) \rightarrow F_g \times F_g$ be any homomorphism that is surjective (F_g is the free group on g generators), then there must be an element in $\ker\phi$ that has an embedded representative in Σ_g . Note that this statement implies the Poincaré conjecture.

This dissertation studies continuous maps

$$f : \Sigma_g \rightarrow (\text{bouquet of } g \text{ circles}) \times (\text{bouquet of } g \text{ circles})$$

between the corresponding topological spaces, with the idea that they must have a rich structure if they can be used to study the Poincaré conjecture.

2.3 Square Complexes

The target space in the map described in the previous section, the cartesian product of two bouquets of g circles, will be denoted $(\vee_g S^1) \times (\vee_g S^1)$. Consider its topology. There are g^2 tori (coming from the product of each pair of circles) with identifications. The table below shows the possible topologies of a neighbourhood of a point in this space, organized by the topology of the neighbourhoods of the pairs of points they are indexed by, one in each bouquet of g circles. The image at the bottom right depicts a neighbourhood that cannot be embedded in three dimensions. Instead, I have drawn a subspace of this neighbourhood, and put ellipses to indicate that more ($4g^2$ in total) sheets should approach the centre point. The topology of this neighbourhood is described below.

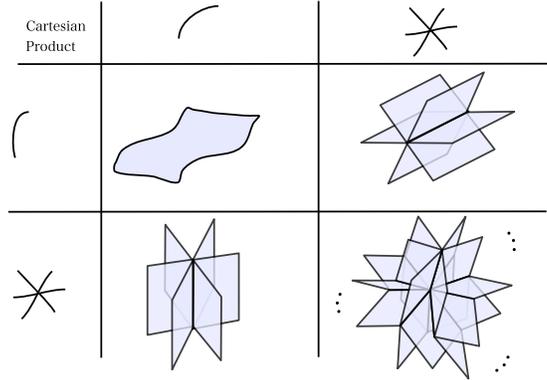


Figure 2.9: The local topology around various points in $(\vee_g S^1) \times (\vee_g S^1)$.

Claim 2. *Let w be the point in $\vee_g S^1$ that has no neighbourhood homeomorphic to an interval (the wedge point). The local topology around $(w, w) \in (\vee_g S^1) \times (\vee_g S^1)$ is the cone on the join of $2g$ points with $2g$ points. This space cannot be embedded in \mathbb{R}^3 when $g > 1$.*

This claim will be proven shortly, but first it will be helpful to have some definitions at our disposal.

Recall the definition of a CW complex [4]: An n -cell is a topological space homeomorphic to a closed n -dimensional ball. A 0-cell is a point. A CW complex, X , is the union $\cup_n X^n$ constructed by beginning with X^0 , a discrete set of 0-cells, and then defining X^n inductively as a quotient of $X^{n-1} \sqcup_\alpha D_\alpha^n$, where D_α^n are n -cells. For each n -cell, there is a continuous attaching map $\phi_\alpha : \partial D_\alpha^n \rightarrow X^{n-1}$, and the quotient is by the equivalence relation $x \sim \phi_\alpha(x)$. X^n is called the n -skeleton of X . In general these spaces can be very complicated. The attaching maps can be restricted to less complicated behaviour by giving the cells the combinatorial structure of cubes, and adding that attaching maps restricted to open faces of cubes be homeomorphisms onto their images. The following definition is from “Special Cube Complexes,” by Haglund and Wise [2].

Definition 11. A 0-cube is a point. An edge (or one dimensional cube) is the interval $I := [0, 1]$. An n -cube is the n -fold cartesian product of the interval I^n taken as a subspace of \mathbb{R}^n . A k -face of an n -cube is a subspace of I^n given by restricting $n - k$ coordinates to a vector in $\{0, 1\}^{n-k}$. An open face is a face with all lower dimensional faces removed.

A cube complex, X , is the union $\cup_n X^n$, where X^0 is a discrete set of 0-cubes, and then X^n is defined inductively as a quotient of $X^{n-1} \sqcup_\alpha C_\alpha^n$, and

C_α^n are n -cubes.

$X^1 := X^0 \sqcup_\alpha C_\alpha^1 / \sim$ where for each edge, C_α^1 , there is an attaching map $\varphi_\alpha : \partial C_\alpha^1 \rightarrow X^0$, and $x \sim \varphi_\alpha(x)$. Thus, as with a CW complex, a one dimensional cube complex is simply a graph. Now we inductively define the equivalence relation \sim , needed to construct higher skeletons. Suppose we have defined \sim for $X^{n-1} := X^{n-2} \sqcup_\alpha C_\alpha^{n-1} / \sim$. The image of the interior of an $(n-1)$ -cube in the quotient space X^{n-1} will be called an *open $(n-1)$ -cube* in X^{n-1} . For example, the open edges of X^1 (a graph) are the maximal open interval subspaces that do not include vertices. For each n -cube, C_α^n in $X^{n-1} \sqcup_\alpha C_\alpha^n$, there is an attaching map $\varphi_\alpha : \partial C_\alpha^n \rightarrow X^{n-1}$ such that φ_α restricted to each open face of I^n is a homeomorphism to an open cell in X^{n-1} .

Given a continuous map between two cube complexes $f : X \rightarrow Y$, it is a *combinatorial map of cube complexes* if each open cube in the domain is sent homeomorphically to an open cube in the target space.

In this dissertation, we will be using two dimensional cube complexes, or square complexes. These are given by identifying the edges of squares (I^2) to edges of graphs with homeomorphisms (such that the attaching maps on the full boundary of each square is continuous).

Definition 12. Consider an open n -cell in a cube complex. Take star-shaped neighbourhoods of all its boundary vertices that do not intersect pairwise, and let T be the union of these open sets. Now intersect T with the open cell. This intersection has 2^n connected components $C_T^1, C_T^2, \dots, C_T^{2^n}$ (because an n cube has 2^n vertices). There was some choice in T . Define an equivalence relation as follows: $C_T^i \sim C_T^j$, if and only if their intersection is non empty. This creates 2^n equivalence classes. Each class is called a *corner* of the n -cell.

A set *intersects a corner* if it intersects all the elements in the equivalence class.

Given a vertex in the cube complex, if every neighbourhood of this vertex intersects a set of corners, we say these corners *meet* at this vertex.

For each square corner, there are exactly two edge corners with the following property: every neighbourhood of every set in the edge corner intersects the square corner. We say that these edge corners are *on the boundary* of the square corner.

Note that $\vee_g S^1$ has a natural CW complex structure given by one vertex and g edges. It is trivially a one dimensional cube complex. The product $(\vee_g S^1) \times (\vee_g S^1)$ has a CW complex structure induced, and it is a square complex. This can be seen by giving the interior of the edges in $\vee_g S^1$ the structure of the interval $[-1, 1]$ and noting that in $(\vee_g S^1) \times (\vee_g S^1)$, the 2-cells are the cartesian products of these edges, which then have the induced structure of I^2 . The attaching map condition is also satisfied. Now we can proceed to the proof of the claim regarding the topology around the point $(w, w) \in (\vee_g S^1) \times (\vee_g S^1)$.

Proof of Claim 2. First we note that for every pair of one dimensional corners that meet at w , there is a two dimensional corner meeting at (w, w) . This gives the $4g^2$ two dimensional corners alluded to earlier, since w is $2g$ -valent. Similarly, for every one dimensional corner meeting w , there are two meeting (w, w) . Therefore there are $4g$ one dimensional corners meeting at (w, w) . We have yet to determine the adjacency of these corners.

Consider the following graph:

$$V = \{1\text{-corners meeting at } (w, w)\}$$

$$E = \{(\alpha, \beta) | \alpha \text{ and } \beta \text{ are the boundary corners of a 2-corner meeting at } (w, w)\}$$

Take the cone on this graph, and consider the cone point. For each edge in the graph, a 2-corner meets the cone point, and for each vertex in the graph, a 1-corner meets the cone point. Further, note that the adjacency of these corners is the same as in the space $(\vee_g S^1) \times (\vee_g S^1)$.

We finish the proof by understanding the structure of this graph. From our observations at the beginning of this proof, $|V| = 4g$ and $|E| = 4g^2$. There are two types of 1-cells in $(\vee_g S^1) \times (\vee_g S^1)$: those that arise as a product of the form 1-cell \times 0-cell, and those that arise as a product of the form 0-cell \times 1-cell. Locally, we can separate the corners of 1-cells meeting at (w, w) into two types depending on the type of the 1-cell to which they belong. This partitions V into two sets. There are $2g$ corners of each type meeting at (w, w) , since w is $2g$ -valent. Note that the attaching maps of 2-cells in $(\vee_g S^1) \times (\vee_g S^1)$ alternate between the two types of edges. This implies that the corner of a 2-cell has one 1-corner of each type in its boundary. Therefore, the graph is bipartite. Now, the first sentence in this proof implies that the graph is in fact complete and bipartite (“for every pair of one dimensional corners that meet at w , there is a two dimensional corner meeting at (w, w) ,”). The join of $2g$ points with $2g$ points is exactly this complete bipartite graph. This shows that the

local topology around (w, w) is the cone on the join of $2g$ points with $2g$ points.

To show that a neighbourhood of (w, w) cannot be embedded in \mathbb{R}^3 , we show that the join of $2g$ points with $2g$ points is not planar. We claim that planarity of the graph is equivalent to the neighbourhood's ability to be embedded in \mathbb{R}^3 . If it is planar, the graph can be embedded on the unit S^2 in \mathbb{R}^3 and the cone on this graph can be embedded by taking the origin as the cone point. If the cone on the graph can be embedded in \mathbb{R}^3 , then we take a small S^2 centred at the cone point, and it must intersect the embedded space in a graph isomorphic to the graph on which we are taking a cone. This implies this graph is planar. The complete bipartite graph on $2g$ points and $2g$ points has $K_{3,3}$ (the complete bipartite graph on three points and three points) as a subgraph, as long as $g > 1$. $K_{3,3}$ is an obstructing subgraph to planarity, thus for $g > 1$, a neighbourhood of (w, w) cannot be embedded in \mathbb{R}^3 . For $g = 1$, $(\vee_g S^1) \times (\vee_g S^1) = T^2$ and can be embedded in 3-space. \square

Note that the graph discussed in this proof is the link of the vertex. If the links of all the vertices in a cube complex satisfy a certain condition, the cube complex exhibits properties of non-positively curved spaces. The following definition is again from Haglund and Wise's paper [2].

Definition 13. A square complex is *non-positively curved* if the cycles in the graphs that are the links of each vertex are of length four or more.

A complete bipartite graph does not have multiple edges, so there are no cycles of length two. The bipartite condition forces cycles to be of even length, so there cannot be cycles of length three. Therefore, $(\vee_g S^1) \times (\vee_g S^1)$ is non-positively curved.

One can begin to imagine how a surface might map to $(\vee_g S^1) \times (\vee_g S^1)$. Several questions arise. How would it wind around the tori? Is it crumpled? Are there branch points and folding? What about if we look only at maps that are related to Heegaard splittings? Can we see the properties of a 3-manifold in these maps? These questions are partially answered in Chapter 4, and motivate the following work.

Chapter 3

Main Theorem

Given a Heegaard splitting $M = H_1 \cup H_2$, with a splitting surface $\Sigma_g = \partial H_1 = \partial H_2$, there are two inclusion maps $\iota_i : \Sigma_g \rightarrow H_i$. Knowing these two maps is sufficient to reconstruct M . The two maps can be expressed as a single map $\iota_1 \times \iota_2 \rightarrow H_1 \times H_2$. We say that $\iota_1 \times \iota_2$ describes the Heegaard splitting $M = H_1 \cup H_2$.

Now, suppose ι_1 and ι_2 were post-composed by a deformation retraction that took H_i to a bouquet of g circles, $\vee_g S^1$. Let ι'_1 and ι'_2 be the respective compositions.

Claim 3. *We can recover the Heegaard splitting $M = H_1 \cup H_2$ from $\iota'_1 \times \iota'_2 : \Sigma_g \rightarrow (\vee_g S^1) \times (\vee_g S^1)$ (for any choice of deformation to a graph in the constructions of ι'_i).*

Proof. We recover a Heegaard diagram for $M = H_1 \cup H_2$ from $\iota'_1 \times \iota'_2$. One can pick g points in $\vee_g S^1$ so that their complement is connected and simply connected. Nudging these points within a small neighbourhood, they can be arranged so that the pre-image of each point by ι'_1 is a union of simple closed curves on Σ_g . Of these curves, some may be null homotopic, in which case they should be ignored. If there is more than one that is essential (there will be at least one), they will all be homotopic to one another. Take this free homotopy class and pick a simple representative α_i . Let α_i be in \mathcal{A} . This gives us g disjoint curves that are homologically non-trivial and independent in \mathcal{A} as desired. The curves in \mathcal{B} are obtained with the same method from ι'_2 .

Now we show that $(\Sigma_g, \mathcal{A}, \mathcal{B})$ is a Heegaard diagram for $M = H_1 \cup H_2$. We will show that \mathcal{A} denotes the attachment of H_1 to Σ_g . The proof is the same

for \mathcal{B} and H_2 .

For each of the curves in \mathcal{A} , consider its image at each moment in time under the chosen deformation retraction from H_1 to a bouquet of g circles. This realizes a map from a disk to H_1 whose boundary is α_i . Using an argument as in the proof of Claim 1, these maps are homotopic to embeddings of disjoint disks in H_1 . Removing the embedded disks, one at a time, must give a handle bodies of decreasing genus. This is because these disks are homotopic to the pre-image of a point in $\vee_g S^1$ by a deformation retraction. Each disk cannot separate H_1 (since the point it is a pre-image of in $\vee_g S^1$, up to homotopy, does not separate). Since there are g disks, removing them results in a 3-ball, as desired. This implies \mathcal{A} denotes the attachment of H_1 .

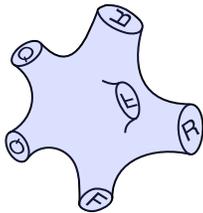


Figure 3.1: A 3-ball with some closed disks removed on the boundary, obtained by removing disks from a handle body. Letter markings show where the disks used to be.

□

From this proof, it is clear that not only can we reconstruct the Heegaard splitting $M = H_1 \cup H_2$ from $\iota'_1 \times \iota'_2$, but that changing this map by a homotopy would not affect our ability to do so. Therefore, we say that all of these types of maps *describe* $M = H_1 \cup H_2$. See the definition below.

Definition 14. A map f from a surface S of genus g to the cartesian product of two bouquets of g circles *describes a Heegaard splitting* $M = H_1 \cup H_2$ if it can be written as $\iota'_1 \times \iota'_2$ up to homotopy. Here ι_i are the inclusions of the separating surface into the handle bodies, followed by some deformation retraction of the handle bodies to a bouquet of circles.

Definition 15. A map $f : S \rightarrow (\vee_g S^1) \times (\vee_g S^1)$ is said to be *irreducible* if it describes an irreducible Heegaard decomposition.

Remark 9. If f is an irreducible map, then $\ker f_*$ has no non-trivial element with a representative that is a simple closed curve.

Proposition 2. *Let $M = H_1 \cup H_2$ be an irreducible Heegaard decomposition. The complement of the curves of any associated Heegaard diagram in Σ_g is homeomorphic to a finite union of discs. In other words, the Heegaard diagram fills the surface.*

Proof. Consider the handle body H_i and the g curves of a Heegaard diagram that encode its attachment to Σ_g . Here is a picture of Σ_4 embedded in \mathbb{R}^3 with curves drawn encoding attachment of H_i to the “inside.” Notice that the complement of these curves in Σ_g is sent to a null-homotopic subset of H_i by ι_i .

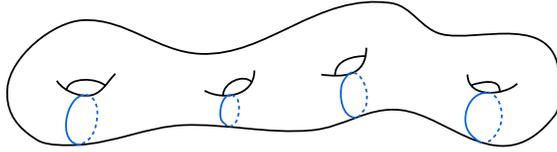


Figure 3.2: Half of a Heegaard diagram.

Now suppose the proposition were false. That is, suppose that there was a Heegaard diagram for $M = H_1 \cup H_2$ that has a complement with a connected component that is not a disk. This implies that there is an essential simple closed curve in the complement of the diagram. Since it does not intersect any curves in \mathcal{A} , it must have null-homotopic image by ι_1 and since it does not intersect any curves in \mathcal{B} , it must have null-homotopic images by ι_2 . This contradicts the irreducibility of $M = H_1 \cup H_2$. \square

The traditional picture of a Heegaard diagram is of red and blue curves. Following this tradition, the discs described above, can be thought of as polygons with alternating red and blue edges made from segments of these curves.

Definition 16. Consider an embedded curve γ on Σ_g . Take a tubular neighbourhood T of γ . Removing the image of γ from this tubular neighbourhood leaves two connected components. One of these components will be called A_T and the other will be called B_T . Now we make an equivalence relation. Any two components of (possibly distinct) deleted tubular neighbourhoods of γ are equivalent if and only if their intersection is non-empty. Notice that there are two equivalence classes and A_T and B_T are in different classes. Call these two

classes the two *sides* of γ .

A set *intersects a side of* γ if it intersects all the sets in this equivalence class. The number of *connected components of such an intersection* is defined as the minimum number of connected components of the intersection with all the sets in the equivalence class.

Alternatively, an irreducible Heegaard diagram is *taut* if every open polygon intersects a side of a diagram curve in at most one connected component.

A region of Σ_g is shown below, with two polygons which are part of two different Heegaard diagrams shown. The first picture is not taut. The open polygon intersects the same side of a curve in two connected components. The second picture is taut. Here, while the open polygon does intersect the sides of a curve in two connected components, they are the opposite sides.

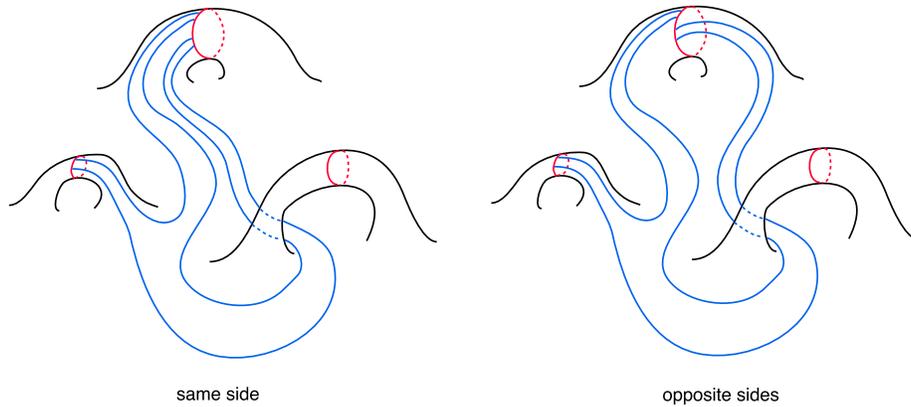


Figure 3.3: Two regions of Heegaard diagrams. The left is not taut, while the right is taut.

Theorem 7. *Every irreducible Heegaard decomposition has at least one associated taut Heegaard diagram.*

Proof. The proof is constructive, beginning with any Heegaard diagram and modifying it until a taut Heegaard diagram is achieved.

If this Heegaard diagram is already taut, we are finished.

If not, there exists a polygon that touches some diagram curve from the same side twice. Such incidences will be called *problems*. Notice that multiple problems may occur on a single polygon, or along a single curve.

Given a problem, without loss of generality, let the problematic curve be red. This is reflected in all the pictures that follow. Consider a curve segment with one end point on each of the two segments shared by the problematic curve and problematic polygon, and with interior inside the problematic polygon. Below, such a segment is shown in orange for the problem shown above.

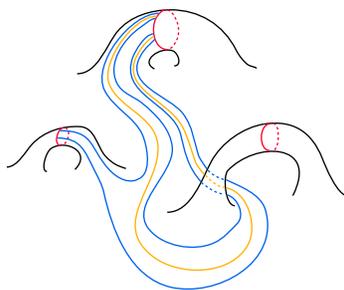


Figure 3.4: An orange segment that intersects the same side of a diagram curve in two connected components.

The problematic curve is broken into two segments by the end points of the orange curve. Consider each of the two closed curves given by the union of each of these segments with the orange curve individually. Isotopic copies of these curves are found as two of the boundaries of the pair of pants obtained by thickening the graph given by the problematic curve and the orange segment in Σ_g . Call these two boundary curves α and β . The third boundary is isotopic to the problematic curve itself.

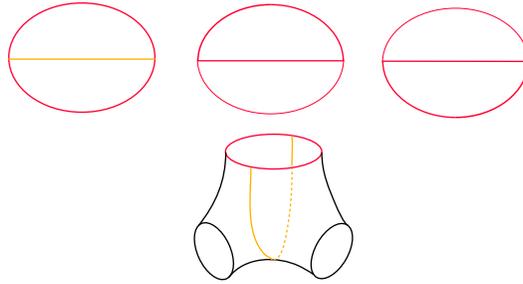


Figure 3.5: Representatives of the free homotopy classes $[\alpha]$ and $[\beta]$, shown in two ways: (i) as a union of segments of the problematic curve and the orange curve, and (ii) as the cuffs of a pair of pants constructed by the problematic curve and the orange segment.

Let us consider replacing the problematic curve (one of the boundaries of these pants) with α or β (one of the other two boundaries on the pants), to produce a new set of red curves. More explicitly, consider erasing the problematic red curve, and colouring either α or β red.

We claim that this swap, if possible, strictly reduces the number of intersections between red and blue curves.

The picture below includes the minimal number of blue arcs possible so that the diagram is filling and does not intersect the orange arc. The elimination of an intersection points is seen.

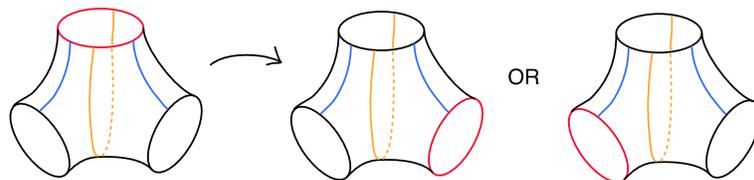


Figure 3.6: The suggested change to the red half of the Heegaard diagram.

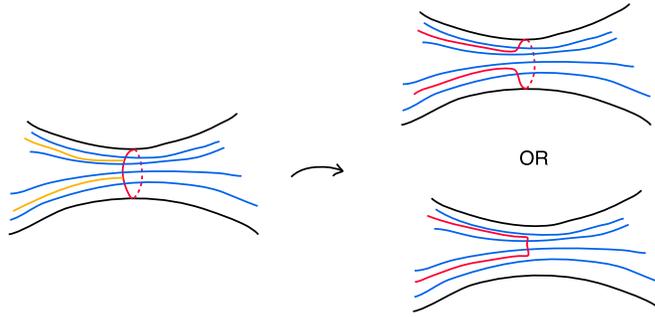


Figure 3.7: A neighbourhood of the problematic curve before and after the suggested swaps, this time with a few more blue arcs.

If one of these swaps were assured to give a new Heegaard diagram for the same Heegaard decomposition, this process could be repeated as long as there is a problem, each time reducing the intersections between red and blue curves by two. The number of intersections is finite and bounded below by zero, as shown in the previous lemma. Therefore, this process would end with a taut configuration.

To show that one of these swaps must give a new Heegaard diagram for the same Heegaard decomposition, several cases must be explored combinatorially.

If α is isotopic to another red curve in the original Heegaard diagram, then β must not be, since the red curves of a Heegaard diagram must be homologically independent. Then, by a single handle slide, the problematic curve can be replaced by β to give a new Heegaard diagram, associated to the same Heegaard decomposition. Similarly, if β is isotopic to a red curve in the original Heegaard diagram, then the problematic curve can be replaced with α .

Assume neither α nor β is isotopic to a red curve in the original Heegaard diagram. In this case, cut Σ_g along all the red curves, and label them to recall the identifications needed to recover Σ_g . This is a sphere with $2g$ discs removed. Note that, by the way it is defined, the orange segment does not intersect any red (or blue for that matter) curves. Therefore, the graph given by the union of the problematic curve and the orange segment can be embedded in the sphere with $2g$ discs removed. When this graph is thickened in this space, the pair of pants bounded by the problematic curve, α , and β are embedded in the sphere with $2g$ discs removed as well.

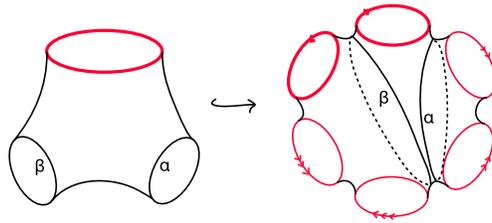


Figure 3.8: The embedding of the pair of pants introduced in Figure 3.5.

Either α or β must separate one copy of the problematic curve from the other. Without loss of generality let it be β . The aim is to show that the problematic curve can be replaced by β to give a new Heegaard diagram associated to the same Heegaard decomposition.

This will be shown by constructing a sequence of embedded pairs of pants that dictate a sequence of handle slides that ultimately make the desired replacement.

Cut along β . This produces two connected components. Pick either one of these to work with. Call this space S . We will embed our pairs of pants that dictate handle slides in this space. Identify as many of the red curves as possible to each other. Each identification creates a handle with a red curve wrapped around it.

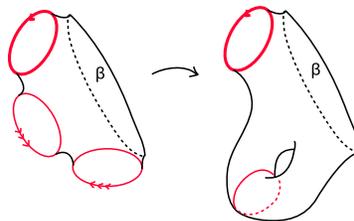


Figure 3.9: S is shown before and after identifications.

In general S is now a surface of genus p with $n + 2$ boundary components (n boundary components besides the problematic curve and β). Further, S has a red curve around each handle, and isotopic to each of the boundary components except for β . Consider the following pair of pants decomposition by curves: the problematic curve, $\gamma_1 \dots, \gamma_{2p}, \dots, \gamma_{2p+n} = \beta$.

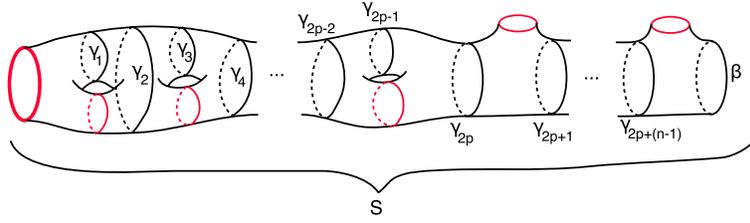


Figure 3.10: A pair of pants decomposition of S .

By a handle slide, the problematic curve can be replaced by γ_1 to give a new Heegaard diagram associated to the same Heegaard decomposition. A second handle slide allows γ_1 to be replaced by γ_2 . Next γ_2 is replaced with γ_3 , and so on until $\gamma_{2p+(n-1)}$ is replaced with β , as desired. \square

Theorem 8. *There is a locally injective map $f : S \longrightarrow (\vee_g S^1) \times (\vee_g S^1)$ describing every irreducible Heegaard decomposition of genus greater than zero.*

Proof. Begin with a taut diagram for an irreducible Heegaard decomposition $M = H_{\text{blue}} \cup H_{\text{red}}$. Let the curves describing the attachment of H_{blue} to Σ_g be blue. Similarly, let the curves describing the attachment of H_{red} to Σ_g be red. Again, we assume that the curves are in a minimally intersecting configuration.

This diagram will be used to give a particular map $f : \Sigma_g \longrightarrow (\vee_g S^1) \times (\vee_g S^1)$ that is locally injective and then we will check that it describes $M = H_{\text{blue}} \cup H_{\text{red}}$.

Defining f : In the target, $(\vee_g S^1) \times (\vee_g S^1)$, consider each coordinate $(\vee_g S^1)$. For simplicity, call one the blue coordinate and the other the red coordinate. After removing the wedge point of $\vee_g S^1$, g open intervals remain. Orient these intervals arbitrarily. In the blue copy of $\vee_g S^1$, Pick g points on the blue $\vee_g S^1$, one at the midpoint of each of these intervals. Call these blue points. Pick a bijection between the set of blue points and the set of blue diagram curves. Repeat the process on the red coordinate, and call the chosen points red points. Now orient the diagram curves on Σ_g arbitrarily.

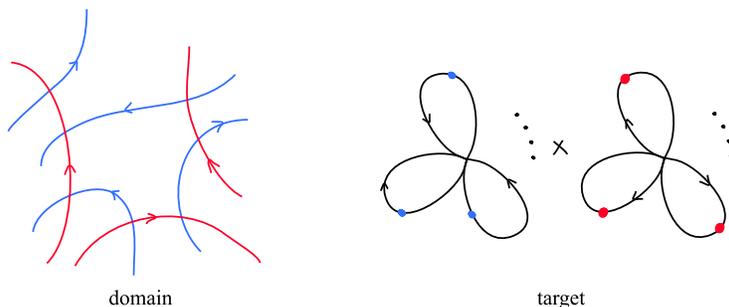


Figure 3.11: A region of Σ_g is shown with orientations chosen on the Heegaard diagram curves. The two coordinates of $(\vee_g S^1) \times (\vee_g S^1)$ are shown with blue and red points chosen as well as orientations on the edges.

Recalling the definition of “side,” there is now a right side (and left side) of each curve using these orientations. There is also a right side of each of the points in the intervals picked out in the preceding paragraph.

Due to Proposition 2, this Heegaard diagram induces a cell decomposition structure on Σ_g . Every vertex in this cell decomposition is quadrivalent. Consider the dual cell decomposition. Every face in this cell decomposition must have four edges (it is a cell decomposition of squares). We take this to be a square complex structure with the diagram curves intersecting the midpoint of the edges of each square. Each square has one blue curve and one red curve running through it, cutting it precisely into quarters. This can be done by moving the diagram by an isotopy. The two diagram curves that run through a given square are associated to two points, one in the blue coordinate $\vee_g S^1$ and one in the red coordinate $\vee_g S^1$. We would like to construct a map $f : \Sigma_g \rightarrow (\vee_g S^1) \times (\vee_g S^1)$. The idea is that f sends the interior of each square homeomorphically to the product of maximal intervals in the two coordinate $\vee_g S^1$ which contain these two points. The interior of the squares are mapped by f such that the right side of the red curve is sent to the right side of the red point and the right side of the blue curve is sent to the right side of the blue point.

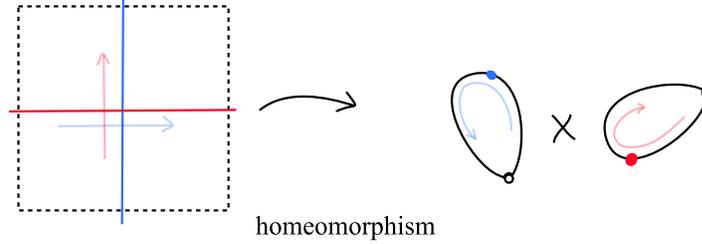


Figure 3.12: The definition of f on the interior of a square.

Now we say precisely how to set this up. Consider f as two coordinates $f = f_{\text{blue}} \times f_{\text{red}}$, where $f_i : \Sigma_g \rightarrow \vee_g S^1$. Each of the 1-cells in the square cell decomposition of Σ_g intersects exactly one diagram curve in one point. If this diagram curve is blue, let f_{blue} send the open 1-cell to the maximal interval subset of $\vee_g S^1$ that contains the blue point assigned to the diagram curve. This should be done in such a way so that the half interval that intersects the right side of the diagram curve is sent to the right of the blue point. Let the boundary of this 1-cell be sent to the wedge point. If the diagram curve is red, let f_{red} send the open 1-cell to the maximal interval subset of $\vee_g S^1$ that contains the red point assigned to the diagram curve, with the boundary of the 1-cell being sent to the wedge point. Now consider a closed 2-cell in the square cell decomposition of Σ_g . f_{blue} and f_{red} are defined on alternating boundary edges of this cell. The rest of the 2-cell is sent by f_{blue} to a straight-line homotopy between the definition(s) of f_{blue} on the boundary. Note that the attaching map of the cell may not be injective, but this will imply that the homotopy is the constant homotopy that does not change with time. f_{red} is defined on the closed 2-cell in the same way. The map f is defined this way on all of Σ_g and it is by definition continuous. It is also clear that f restricted to the interior of a 2-cell is a homeomorphism onto its image.

Checking f is locally injective: Recall that $(\vee_g S^1) \times (\vee_g S^1)$ also has a cell decomposition made up of squares. The map f , as defined, is a combinatorial map of square complexes. Therefore, it suffices to check that f is locally injective at the vertices of the square complex structure on Σ_g .

Claim 4. *No two corners meeting at a vertex are sent to the same image corner.*

Proof. Consider a vertex, v , of the square complex structure on Σ_g . Note that it is contained in the interior of a single polygon of the dual cell decomposition induced by the Heegaard diagram curves.

If the corners of two 1-cells meeting at v have the same image corner, it follows from the definition of f_{blue} and f_{red} on the 1-skeleton of Σ_g that these two 1-cells intersect the same Heegaard diagram curve. (Recall that each 1-cell intersects exactly one diagram curve.) This means that the polygon dual to the vertex v has two edges made up of a single diagram curve. Further, it implies that one side of this diagram curve intersects the open polygon in two connected components, one along each of these edges. Thus the Heegaard diagram is not taut. This is a contradiction.

Suppose there are two square corners, σ and δ , that meet at v and $f(\sigma) = f(\delta)$. The edge corners on the boundary of σ and δ must be sent to the edge corners on the boundary of $f(\sigma)$, because f is a continuous combinatorial map of square complexes. $f(\sigma)$ has two boundary edge corners. The number of edge corners on the boundary of σ or δ is three or four. As the picture below shows, if there were only two, the edge corners would have to be on the boundary of both σ and δ . This would imply the original Heegaard diagram was not in a minimally intersecting configuration, which is a contradiction. In the picture below, the orange edge corners are on the boundary of the shown square corners.

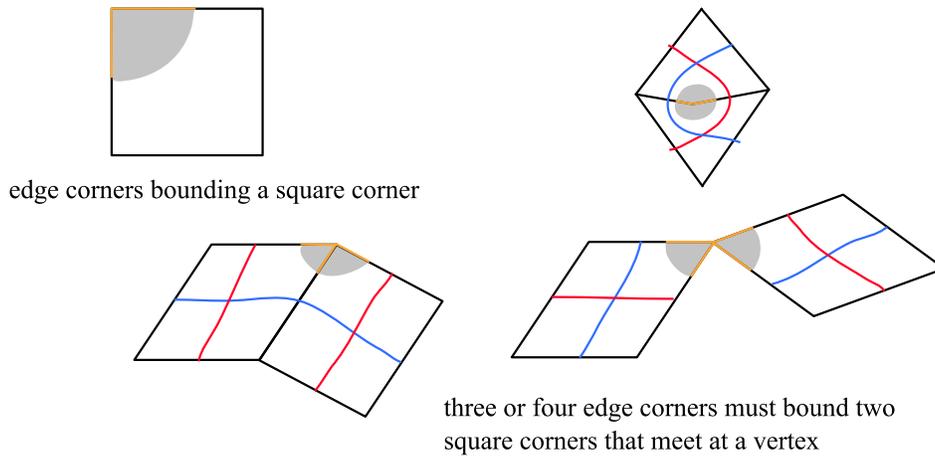


Figure 3.13: The behaviour of edge corners and square corners.

We have already shown that two edge corners meeting at v cannot have the same image. Three or four edge corners being sent to two edge corners by f would force this. Therefore no two square corners meeting at v can have the same image.

□

Consider a small star shaped neighbourhood of a vertex, v , of the square complex structure of Σ_g . There is only one vertex in this neighbourhood, and it is the only point in the neighbourhood sent to a vertex by f (since f is a combinatorial map of square complexes). The rest of the neighbourhood can be decomposed into open intervals and open disks (one in each of the corners that meet at v) that do not intersect pairwise and result from intersecting open cells with the star shaped neighbourhood. The above claim proves that the images of these sets do not intersect pairwise. The map f is defined to be an embedding on each of these sets. Therefore f is injective locally at v .

Checking that f describes the Heegaard decomposition: Recall that f describes $M = H_{\text{blue}} \cup H_{\text{red}}$ if, up to homotopy equivalence of the domain, and product homotopy equivalence of the target, $f = \iota_{\text{blue}} \times \iota_{\text{red}} : \Sigma_g \longrightarrow H_{\text{blue}} \times H_{\text{red}}$, where ι are the inclusion maps from the separating surface to H_{blue} and H_{red} respectively. First, we check that $f_{\text{blue}} = \iota_{\text{blue}}$ up to homotopy equivalence of the target space. Consider the fibres of f_{blue} . An example of what these fibres look like is shown below. Recall that f_{blue} sends the interior of each square in Σ_g to a maximal interval in $\vee_g S^1$. The blue curve running through the square is sent to the blue point in this interval. In fact, the full pre-image of each blue point in Σ_g is a blue diagram curve. The full pre-image of any other point in the same maximal interval as a blue point is an isotopic curve. Again, think of it one square at a time first, and then put them together. The full pre-image of the wedge point is graph on Σ_g given by the 1-skeleton of the square complex structure, with the edges that intersect blue diagram curves removed. This can be seen by thinking about each square first. The pre-image of the wedge point in each closed square is the boundary edges that do not intersect a blue diagram curve. Notice that all the fibres except the pre-image of the wedge point are topological circles. All this follows from the definition of f .

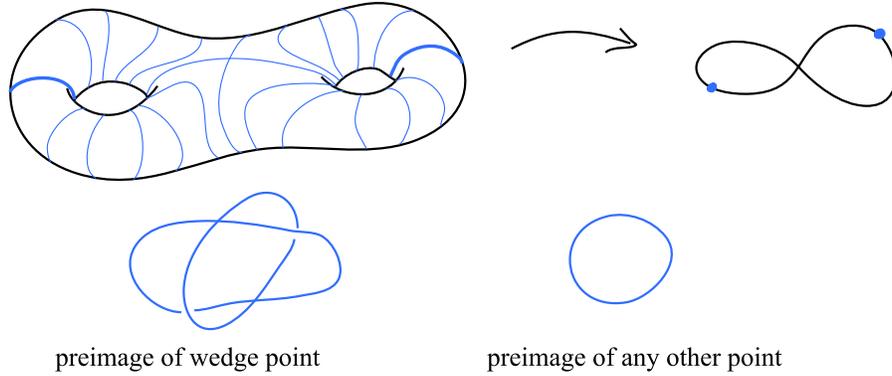


Figure 3.14: A genus two example of one coordinate of f .

Consider including Σ_g into a larger space as follows. Remove the f_{blue} pre-image of the wedge point from Σ_g . Each connected component of the remaining space is made up of the squares through which a given blue diagram curve runs. It is a cylinder of “height” 1. There are g such cylinders remaining, and they have a product structure of $(0, 1) \times S^1$ induced by the fibres of f . Include these cylinders into solid cylinders, $(0, 1) \times D^2$ so that the S^1 fibres are included as boundaries to the D^2 fibres of the solid cylinder. Now glue Σ_g back together, and glue the solid cylinders back together accordingly. We now formalize the preceding argument. First glue two disks to each cylinder to g copies of $[0, 1] \times D^2$. Parametrize the D^2 coordinate of these cylinders as a unit disk using “polar coordinates” (r, θ) , $0 \leq r \leq 1$. Now glue Σ_g back together along the pre-image of the wedge point. If two boundary disks ($\{0\} \times D^2$ or $\{1\} \times D^2$) are identified along an arc on their S^1 boundary, each point in this arc belongs to a ray in each of the two disks. Identify these rays. Note that the centres of all the boundary disks are identified in this process. Call the resulting space H .

Claim 5. H is a handle body and there is a homeomorphism $\phi : H \rightarrow H_{\text{blue}}$ so that the inclusion $\iota : \Sigma_g \hookrightarrow H$ is equal to $\phi \circ \iota_{\text{blue}}$.

Recall we wanted to show that $f_{\text{blue}} = \iota_{\text{blue}}$ up to homotopy equivalence. If the claim is true, the proof is completed by noting that $\vee_g S^1$ is a deformation retract of H . The deformation retraction is realized by taking each point on each disk in the solid cylinders (r, θ) , and continuously shrinking the radius r

to 0. Call this deformation retraction d . Each cylinder becomes an interval, and the way the solid cylinders were glued together leads to these intervals being glued together into a bouquet of circles. This can be seen by considering the glueing locus under d , and noting that it is all sent to a single point in the end. Then note that $f_{\text{blue}} = d \circ \iota$ because we set up the product structure on the hollow cylinders so that each S^1 fibre was the f_{blue} pre-image of a point in $\vee_g S^1$ (not the wedge point). So we have $f_{\text{blue}} = d \circ \iota = d \circ \phi \circ \iota_{\text{blue}}$, and $d \circ \phi$ is a homotopy equivalence as desired (since the composition of a homeomorphism and deformation retraction gives a homotopy equivalence together with inclusion composed with a homeomorphism). This completes the proof of Theorem 8. \square

Now we prove Claim 5.

Proof. Recall how H_{blue} is attached to Σ_g : a disk is attached along each blue diagram curve on Σ_g , and then a 3-ball is attached to the resulting 2-complex. We claim that the method of attaching H is equivalent to this. Consider, in the construction of H , the disk $\{\frac{1}{2}\} \times D^2$ in each solid cylinder. The S^1 boundary of such a disk is a blue diagram curve. So both constructions attach disks along the blue diagram curves. Let A be the space remaining when we remove these disks and Σ_g from H . We will prove that A is a 3-ball that is attached to the rest of H in the same way that the 3-ball is attached in the description of H_{blue} .

The middle disk $\{\frac{1}{2}\} \times D^2$ cuts each solid cylinder in half. Each half is a cone on a disk given by the “bottom” and “sides” of the cylinder. This is shown below.

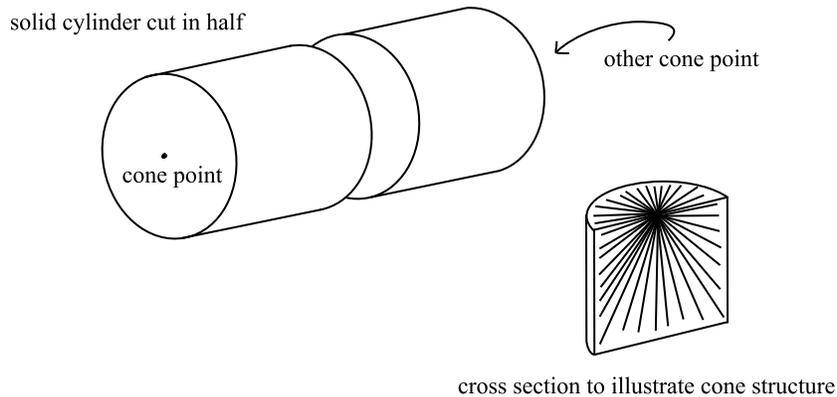


Figure 3.15: The cone structure of half of a solid cylinder.

The two cones that share the disk $\{\frac{1}{2}\} \times D^2$. Note that the cone structure restricted to $\{0\} \times D^2$ and $\{1\} \times D^2$ is compatible with the polar coordinates chosen on these disks. Recall that when the solid cylinders are glued together to form H , rays of one of these disks are identified with rays of another. This implies that A can be expressed as a cone relating to the cone structure we have placed on the half cylinders. Cut Σ_g along the blue diagram curves (cut the cylinders in half) and glue in disks along all boundary components. This is simply S^2 . The cone on this space with S^2 removed is A . This is an open 3-ball. In this way, it is seen that H and H_{blue} are attached to Σ_g in the same way, completing the proof. \square

Chapter 4

Corollaries

4.1 Reducible 3-Manifolds

We can put together the results from Section 2.2 and Chapter 3 to make statements about all compact orientable 3-manifolds (we include reducible 3-manifolds and reducible Heegaard Splittings now).

Definition 17. Begin with a surface Σ with finitely many disjoint planar subsurfaces (surfaces without handles), S_i . Consider the following equivalence classes. Each S_i forms an equivalence class. Points not in any S_i form a singleton equivalence class. The quotient by this equivalence relation collapses each planar subsurface to a point. The result, Σ' will be called a *pinched* surface. Σ' is a *pinching* of Σ .

Corollary 1. *Suppose M is an compact orientable 3-manifold with Heegaard splitting $M = H_1 \cup H_2$ of genus greater than zero. Then it can be described by a map $f : \Sigma_g \rightarrow (\vee_g S^1) \times (\vee_g S^1)$ such that it factors through a map \tilde{f} (shown below)*

$$\begin{array}{ccc} \Sigma_g & \xrightarrow{f} & (\vee_g S^1) \times (\vee_g S^1) \\ \downarrow & \nearrow \tilde{f} & \\ \tilde{\Sigma}_g & & \end{array}$$

such that

- $\tilde{\Sigma}_g$ is a pinching of Σ .
- \tilde{f} restricted to a surface subspace of $\tilde{\Sigma}_g$ describes an irreducible Heegaard splitting $M' = H'_1 \cup H'_2$, where M' is a connect summand of M .

- Let S be a surface subspace of $\tilde{\Sigma}_g$. If $\tilde{f}|_S$ is locally injective unless it describes $S^2 \times S^1$

Proof. Let $M = H_1 \cup H_2$ be a Heegaard splitting of genus greater than zero.

Case 1: $M = S^3$ Then due to Waldhausen, we know that if we model S^3 as the one point compactification of \mathbb{R}^3 , then the splitting surface is the standard embedding of a surface of genus g in \mathbb{R}^3 . Then $g - 1$ disjoint two dimensional spheres can be chosen so that the complement of these spheres, with the boundaries coned off, are each a genus one Heegaard splitting on S^3 .

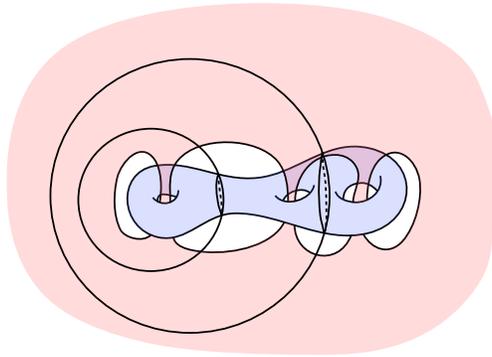


Figure 4.1: Two disjoint embedded S^2 in the genus three Heegaard splitting of S^3 , such that the complement is three genus one Heegaard splittings of S^3 (when the boundaries are coned off).

Each of these spheres intersect Σ_g in a single essential simple closed curve. Call these curves $\{\gamma_i\}_1^{g-1}$. Choose a path δ that has one end point on γ_1 , the other on γ_{g-1} and intersects each γ_i , with $1 < i < g - 1$, once. Consider the union of the spheres with δ in M . the complement of this set is g connected components each with one boundary component homeomorphic to S^2 . The splitting surface is also broken into g connected components, each one homeomorphic to a torus with a disk missing. Cone off the boundary of each remaining 3-manifold component. This gives g copies of the genus one Heegaard splitting of S^3 . (The cone on S^2 is a 3-ball, so this makes the the resulting spaces 3-manifolds. The cone on S^1 is a disk, so this makes the resulting splitting surfaces tori. The resulting spaces must be connect summands of M , so they must be S^3 .)

To construct f , we let f take the union $\delta \cup_i \gamma_i$ to the vertex of $(\vee_g S^1) \times (\vee_g S^1)$. Thus f factors through $\tilde{\Sigma}_g =$ a bouquet of g tori. Note that $\delta \cup_i \gamma_i$, when thickened in Σ_g , is a planar surface. On each torus, we let \tilde{f} be a locally injective map describing the genus one decomposition of S^3 . This should be done in a way that makes \tilde{f} well defined on the entire bouquet (only the wedge point of $\tilde{\Sigma}_g$ is sent to the vertex of $(\vee_g S^1) \times (\vee_g S^1)$ by \tilde{f}). f is then defined as the composition of the quotient to the bouquet of tori followed by \tilde{f} . One can see that f describes the Heegaard splitting by noting that the intersection of the two dimensional spheres with H_i together with δ form a contractible subset of H_i . Therefore, a deformation retraction of H_i to a graph can be given by retracting this set to a point, to get a bouquet of solid tori, and then retracting these solid tori to circles, to get the desired bouquet of circles. The definition of \tilde{f} implies that each coordinate of f is such a deformation retraction up to homotopy.

Case 2: $M \neq S^3$ Then let $M = P_1 \# P_2 \dots P_n$ be the unique decomposition of M into primes, where $P_i \neq S^3$.

If M is reducible, using Theorem 5 and Theorem 6 an embedded sphere can be found in M that realizes a non trivial connect sum of M and intersects the splitting surface in a single essential simple closed curve. Consider the complement of this sphere, and cone off the boundaries to give two Heegaard splittings whose genera sum to g . For each of these, if $\pi_2 = \mathbb{1}$, repeat the process. Continue this until all P_i are constructed. By pushing each sphere off of the 3-ball formed by coning off the boundary in the previous step, there is a natural embedding of the spheres in M . This process constructs a set of disjoint embedded spheres in M that realize the connect sum $M = P_1 \# P_2 \dots P_n$, and that each intersect the splitting surface in a single essential simple closed curve.

If M is irreducible we can skip the previous paragraph and begin here. If the Heegaard splittings of P_i can be destabilized, then continue choosing two dimensional spheres that realize the destabilizations until the complement of all the spheres in M is a number of minimal genus Heegaard decompositions with 3-balls missing.

Proceed as in the previous case. Again, call the family of curves that is the intersection of the embedded two dimensional spheres with the splitting

surface $\{\gamma_i\}_1^n$. Again, pick a path, δ , that has one end point on γ_1 and the other on γ_n and intersects γ_i once transversally, for $1 < i < n$. Thickening $\delta \cup_i \gamma_i$ give a planar subsurface of Σ_g . Collapsing this subsurface to a point gives $\tilde{\Sigma}_g$, which is a bouquet of surfaces whose genera add to g . \tilde{f} takes the wedge point to the vertex of $(\vee_g S^1) \times (\vee_g S^1)$. Each surface subspace of $\tilde{\Sigma}_g$, S_j , is the image of a component of the complement of $\delta \cup_i \gamma_i$ in Σ_g . Each such component is embedded in a component of the complement of the two dimensional spheres in M . Each of these three dimensional components is either a prime P_i with a 3-ball missing, or a 3-sphere with a 3-ball missing. The map \tilde{f} is defined on S_j to be a locally injective map describing P_i or S^3 accordingly. That is unless $P_i = S^2 \times S^1$, in which case a locally injective map is not possible, and any map describing the genus one Heegaard splitting of $S^2 \times S^1$ should be chosen. Again, this should be done so that \tilde{f} is well defined on the entire bouquet of surfaces $\tilde{\Sigma}_g$. \square

Remark 10. There are reducible Heegaard splittings that can be described by locally injective maps. Osborne found examples of genus two Heegaard diagrams that are filling and taut. These genus two Heegaard splittings are reducible, but they can be described by locally injective maps using. This can be shown using the methods of the proof of Theorem 8. An example like this can be seen in Section 5.1.

In the the remainder of the corollaries (Chapter 4) we by and large do not discuss the reducible case discussed in this section. However the reader should keep in mind that analogous arguments to the proof of Corollary 1 can be used to address reducible Heegaard splittings using a pinched surface in the discussion that follows.

4.2 Square Complexes

Note that in the proof of Theorem 8, we saw that the splitting surface in an irreducible Heegaard splitting has a natural family of square complex structures. We also saw that that there are maps, $f : \Sigma_g \rightarrow (\vee_g S^1) \times (\vee_g S^1)$, describing the Heegaard splitting that are combinatorial maps of cube complexes with respect to these structures. In fact, the square complex structures on the surface must be non-positively curved. This is because the links of vertices on the surface are sent to the link of the vertex of $(\vee_g S^1) \times (\vee_g S^1)$. Recall that the links are each graphs, and the map between links must send vertices to vertices and edges to edges, because f is a combinatorial map of cube complexes. Therefore, if the minimum length of cycles in the domain is bounded

below by the minimum length of cycles in the target. Since $(\vee_g S^1) \times (\vee_g S^1)$ is non-positively curved, this implies the square complex structures on the splitting surfaces are also non-positively curved.

Corollary 2. *For every irreducible Heegaard splitting, there is a natural family of surfaces with non-positively curved square complex structures.*

For every compact orientable 3-manifold with Heegaard splitting $M = H_1 \cup H_2$, there is a natural family of pinched surfaces with non-positively curved square complexes.

In both cases, the 3-manifold (with Heegaard splitting) can be recovered from a single member of this family. Members of the family do not have self osculating hyperplanes.

The last sentence of this corollary will be explained later, but for now, without explanation, we state that not having self osculating hyperplanes is equivalent to tautness (Definition 16).

Corollary 2 places Heegaard decompositions of 3-manifolds into the category whose objects are two dimensional cube complexes and whose morphisms are combinatorial maps.

Consider a combinatorial map between two square complexes that describe Heegaard splittings. One way in which this is possible is a n -sheeted branch covering from a surface to another with a pair of branch points of degree n (the branch points must occur at vertices). In this case, the number of Heegaard diagram curves and the genus are both divided by n (taking Euler characteristic into account). This can be extended to a branch covering between the 3-manifolds they describe. By taking an arc inside each handle body that connects the two ramification points, and does not intersect any of the disks whose boundaries are Heegaard diagram curves (dual to the square complex structure), an embedded circle in the 3-manifold is formed. A branch map can be created with the circle being the branch points of degree n .

Question 1. Which combinatorial maps between splitting surfaces with square complex structures can be extended to maps between 3-manifolds?

4.2.1 Handle Slides

Handle slides are realized in the square complexes as follows. Dual to the square complex structure, is a Heegaard diagram describing the 3-manifold. A

blue handle slide of this Heegaard diagram is determined by choosing a simple path γ in the complement of the blue curves on the splitting surface, with end points on distinct blue curves, α and α' . This is because the space given by the union of the path and the two blue curves can be thickened to give a pair of pants, P with two blue cuffs. Then by Definition 7, a handle slide is performed by replacing one of the blue curves on the cuffs of these pants with the remaining cuff that is not blue. This path, γ , can be changed by a homotopy that keeps its end points on the blue curves, and keep the rest of the path in the complement of the blue curves so that it is in the 1-skeleton of the square complex structure. From here, we modify the square complex structure as follows. Consider the cuff of P that is not blue. It is homotopic to a closed curve ρ in the one skeleton of the square complex given by concatenating four paths:

1. Without loss of generality, assume γ intersects the right side of α and the left side of α' . Beginning at the point x where γ intersects the boundary of a square containing a segment of α' , travel “parallel” to α' (not intersecting α') along the boundaries of squares intersecting α' , ending at x again.
2. Beginning at x , travel along $\bar{\gamma}$ ($= \gamma(1 - t)$) to the point y , where γ intersects the boundary of a square containing a segment of α .
3. Beginning at y , travel “parallel” to $\bar{\alpha}$ (not intersecting α) along the boundaries of squares intersecting α , ending at y .
4. Beginning at y , travel along γ , ending at x .

Suppose the curve ρ is n edges long. Cut along this curve and glue in a cylinder of height one square and circumference n squares. After this, remove the cylinder consisting of squares that α intersects. Glue together the exposed boundaries by identifying pairs of edges that were on the boundary of a square that intersected α . This will give a new square complex that describes the same 3-manifold and this operation is dual to a handle slide. Below, on the top left, α and α' are shown in blue and γ is shown in orange. On the top right, the curve ρ is shown in orange. On the bottom left, the added cylinder is shown in green. Finally, on the bottom right, the cylinder consisting of squares that intersect α is removed. Notice the difference in the shown blue diagram curves.

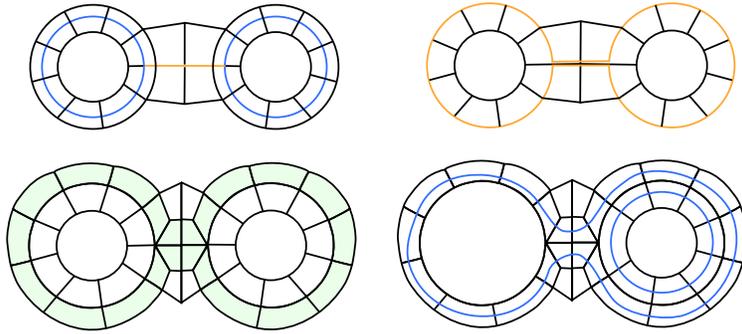


Figure 4.2: An example of a handle slide with square complexes.

4.2.2 Relation to Work of Haglund and Wise

Corollary 2 relates to the main theorem (Theorem 1.1) in the aforementioned paper [2], which says that if a cube complex does not have any of four types of pathologies, then there is a local isometry from the cube complex to a standard cube complex of a finitely generated right-angled Artin group. The various parts of this statement are explained, as it relates to this work, as follows.

The pathologies are illustrated below [2].

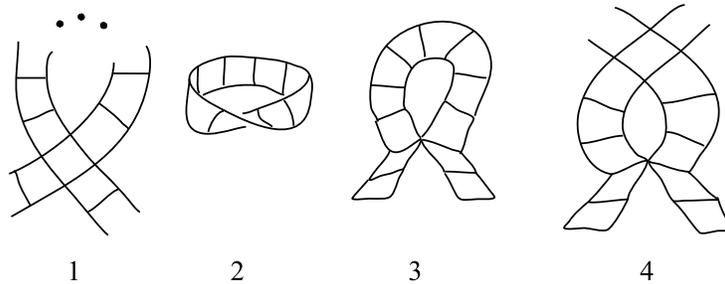


Figure 4.3: Four pathologies of the 2-skeleton of a cube complex.

They are defined in terms of *strips* of the surface Σ .

Definition 18. Consider an alternating sequence of squares, σ_i , and edges e_i , $\{\sigma_0, e_0, \sigma_1, e_1, \sigma_2, e_2, \dots, \sigma_n, e_n\}$, such that:

- $\sigma_0 = \sigma_n$ and $e_0 = e_n$
- e_i and e_{i+1} are edges of σ_{i+1} that are not adjacent.

Let S be the closure of the union of all the edges and squares in the above sequence. S is a *strip* of Σ .

Each pathology is then defined as follows. The numbers match Figure 4.3.

1. A strip for which there is i, j such that $\sigma_i = \sigma_j$ and $e_i \neq e_j$.
2. A strip that is not an orientable subspace of Σ .
3. The third pathology is called a *self osculating hyperplane*. Orient e_0 arbitrarily. Orient the remaining e_i so that their formal sum is a cocycle. Pathology (3) is the existence of a vertex that contains e_i and e_j in its coboundary with the same sign.
4. The fourth pathology is called *inter osculating hyperplanes*. Suppose there are two strips S and S' such that $\sigma_i = \sigma'_j$ but e_i and e'_j are consecutive edges in the boundary of the shared square σ_i . Orient e_i arbitrarily, and then orient the remaining edges of S so that they form a cocycle. Orient e'_j so that e_i and e'_j appear with the same sign in the coboundary of the vertex at which they meet. Then orient the remaining edges in S' so that they form a cocycle. Pathology (4) occurs when there is a vertex, v disjoint from σ_i that is in the intersection of S and S' such that an oriented edge of S and an oriented edge of S' appear with the same sign in the coboundary of v .

In our setting, which begins with a cube complex structure on Σ that describes a Heegaard splitting, each pathology can be described in terms of the Heegaard diagram dual to the square complex structure. Again, the numbers correspond to the picture above.

1. A Heegaard diagram curve intersects itself.
2. A neighbourhood of a diagram curve is a Möbius band (instead of a cylinder).

3. (self osculating hyperplane) A Heegaard diagram is not taut. This can be seen by considering the vertex in Figure 4.3, (3), where four squares are shown meeting. This vertex is dual to some polygon. This polygon intersects one side of the following curve in two connected components. A single Heegaard diagram curve cuts all the shown squares in (3) consecutively in half. By definition this is not taut.
4. The fourth pathology is described by two Heegaard diagram curves, α and β . The pathology occurs when there are two polygons that intersect the same side of α and they also both intersect the same side of β . One of the polygons has a segment of α and a segment of β consecutively on its boundary, while for the other polygon, the segment of α and the segment of β are not consecutive on its boundary.

A square complex can be built where X^0 is a single point. There is an edge for each Heegaard diagram curve, so that X^1 is a bouquet of $2g$ circles. Then a square is attached to a pair of edges so as to form a torus, if the two corresponding diagram curves intersect. Again, this is described in our setting, and the more general definitions can be found in [2]. Call this square complex A .

Definition 19. A *finitely generated right angled Artin group* is a group that can be presented with finitely many generators and relations that are commutators of pairs of these generators.

Note that A is the presentation complex of such a group.

Definition 20. A is called the *standard cube complex* of its fundamental group (which is a finitely generated right angled Artin group).

A is a subspace of $(\vee_g S^1) \times (\vee_g S^1)$ (the target space of a map describing the Heegaard splitting). Dual to each square in Σ_g , there are two diagram curves intersecting. For each pair of diagram curves that intersect, there is a square in A that forms a torus together with the edges that correspond to these two curves. Thus there is a natural combinatorial map of cube complexes $\Sigma_g \rightarrow A \subset (\vee_g S^1) \times (\vee_g S^1)$ given by sending each square, σ , in Σ_g to the square in A which forms a torus together with the two edges that correspond to diagram curves that intersect in the interior of σ .

In fact, this map is exactly the map we have been considering. To be precise, the results of Chapter 3 could be stated with this vocabulary as follows.

Corollary 3. *If M is an irreducible 3-manifold, there exist Heegaard diagrams for M such that dual to the diagram curves on the splitting surface, Σ_g , is a square complex that does not exhibit pathologies (1), (2), (trivially) or (3) (tautness), and therefore, the natural combinatorial map of square complexes $\Sigma_g \rightarrow A \subset (\vee_g S^1) \times (\vee_g S^1)$ is locally injective.*

Now the theorem of Haglund and Wise supposes a metric on the square complexes by taking each square to have the standard euclidean metric of a 1×1 square (and each edge is of length one). Their theorem restricted to surfaces says that a square complex on a surface, Σ , that does not exhibit pathologies (1), (2), (3), or (4) will map locally isometrically by the natural combinatorial map of square complexes $\Sigma \rightarrow A$.

Remark 11. The square complexes in this dissertation have “3/4” of the hypothesis of this Haglund and Wise’s theorem, and result in a local embedding instead of a local isometry to the same target space, which is a subspace of the product of bouquets.

To see that the map constructed in Theorem 8 is not a local isometry with the metric defined above, consider the following. In the case $g > 1$, any square complex structure on Σ_g must have a vertex where more than four squares meet. This is due to Euler characteristic considerations. Now suppose we have a square complex structure on Σ_g that describes an irreducible Heegaard splitting. Let f be the combinatorial map of square complexes $\Sigma_g \rightarrow (\vee_g S^1) \times (\vee_g S^1)$ constructed in Theorem 8. Take a vertex, v , that is of valence greater than four. Consider the map from the link of v to the link of $f(v)$ in the target space. Recall the link of $f(v)$ is the complete bipartite graph on $(2g, 2g)$. The link of v is circular graph with number of vertices and edges equal to the valence of v . The domain link contains vertices that are three edges apart because the valence of v is greater than four. However, in the target link, their image is only one edge apart (the two vertices are not in the same partition of the bipartite graph because a path of three edges must begin and end in different partitions). The existence of an edge that connects these two points implies there is another square corner meeting $f(v)$. This corner contains a right angled triangle in $(\vee_g S^1) \times (\vee_g S^1)$ whose orthogonal edges are the image of the shortest path between two points in the domain, and the hypotenuse is the shortest path between their images. Thus f is not a local isometry at v . Below, v is shown. The link of v is shown in grey. Two points in the link that are three edges apart and the shortest path between them are shown in orange. The image of all this is shown in the same colours. As proven previously, a neighbourhood of the $f(v)$ cannot be embedded in \mathbb{R}^3 ,

and so all of it cannot be simultaneously pictured. Nevertheless, due to the argument above, there is a square meeting $f(v)$ as shown in green. It is clear that the orange path in the green square is not the shortest between its endpoints. The shortest path is shown in green.

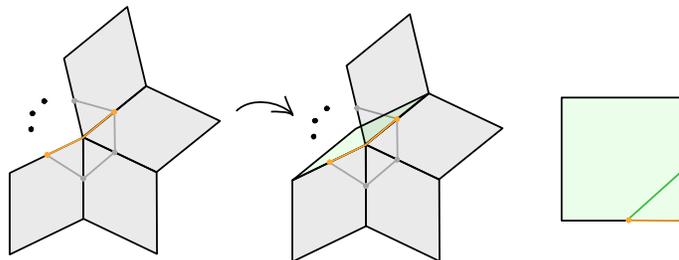


Figure 4.4: f is not a local isometry.

Question 2. What Heegaard splittings can be described by square complex structures that do not inter osculating hyperplanes? What Heegaard splittings can be described by square complexes structures that do not have any of the described pathologies?

4.3 Conformal Structures and Quadratic Differentials

We can also consider the conformal structure induced by these squares to get a family of Riemann surface structures on the splitting surface. However, in general a Riemann surface can be broken into flat squares in many ways, so a conformal structure alone does not capture the information of the Heegaard splitting or 3-manifold. An exception to this is genus one splittings of lens spaces, described in Section 5.1. For other Heegaard splittings, additional information is required.

The fibres of the two coordinates of a combinatorial map of square complexes describing a Heegaard splitting are transversal foliations with singularities. This can be seen by considering the construction of these maps on the

interior of each square. This information can be captured with a quadratic differential.

Definition 21. A *quadratic differential*, on a Riemann surface Σ , is a holomorphic section of $T^*\Sigma \otimes T^*\Sigma$.

Definition 22. Locally, a vector field of unit length on which a quadratic differential, ϕ , is positive can be integrated to give a global foliation of Σ with singularities. This is called the *horizontal foliation of ϕ* . Another foliation with singularities can be constructed such that ϕ on tangent directions is negative. This is called *the vertical foliation of ϕ* .

Remark 12. These two foliations with singularities are orthogonal (Note that if $\phi = f(z)dz \otimes dz$ and $\phi(a\partial x + b\partial y) > 0$, then $\phi(-b\partial x + a\partial y) < 0$, where $a, b \in \mathbb{R}$). In particular $dz \otimes dz$ has positive foliation made up of horizontal lines in \mathbb{C} , and negative foliation made of vertical lines.

A quadratic differential with horizontal and vertical foliations equal to the fibres of the two coordinates of the map describing the Heegaard splitting can be created as follows. Make coordinate patches for each square to a unit square in \mathbb{C} with quadratic differential $dz \otimes dz$, and making small disk shaped patches at the vertices that map to the unit disk in \mathbb{C} with quadratic differential $z^n dz \otimes dz$ where $2n + 4$ is the valence of the vertex. The square patches should be chosen to overlap each other but not include the vertices.

In the picture below, which represents the splitting surface of a genus three decomposition of the three dimensional torus, these fibres are shown. It should be imagined embedded in a cube with opposite faces identified. The entire cube with opposite faces identified is T^3 , and the two components of the complement of the surface in this space are genus three handle bodies. Points on the same blue leaf have the same image by the first coordinate, and points on the same red leaf have the same image by the second coordinate. This example is explained in more detail in Section 5.2.

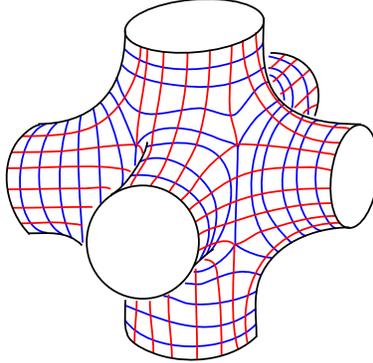


Figure 4.5: A surface of genus three shown with the fibres of the two coordinates of a map describing a genus three Heegaard splitting of T^3 .

Corollary 4. *For every irreducible Heegaard splitting, there is a natural family of pairs (Riemann surface, quadratic differential).*

For every compact orientable 3-manifold with Heegaard splitting $M = H_1 \cup H_2$, there is a natural family of pinched surfaces with each surface subspace having a conformal structure and a quadratic differential.

Again, in both cases the 3-manifold and Heegaard splitting can be recovered from any member of this family.

One can also ask about maps between two pairs $(\Sigma, \phi) \rightarrow (\Sigma', \phi')$ inducing maps between the 3-manifolds they describe, as in Question 1. The effect of handle slides on these structures is also similar to Section 4.2.1, and is described in terms of strips as before.

It is interesting to observe the change to a 3-manifold and Heegaard splitting induced by changing the conformal structure and quadratic differential on a surface describing it.

Remark 13. Rational fractional Dehn twists along certain simple closed curves can change a pair (Σ, ϕ) , describing one 3-manifold, into a pair (Σ', ϕ') , describing another 3-manifold. This is described in detail in the case $g = 1$ in Section 5.1.

4.3.1 Tautness

The tautness of Definition 16 that is crucial for Theorem 8 is not captured in Corollary 4. This can be described this in the language of quadratic differentials by contrasting *abelian differentials* with quadratic differentials. The

following follows the exposition in Papadopoulos' Handbook of Teichmüller Theory [11].

Definition 23. An *abelian differential* is a holomorphic 1-form on a Riemann surface.

A natural question is whether a given quadratic differential, ϕ can be written as $\omega \otimes \omega$, where ω is an abelian differential.

Definition 24. Two orthogonal foliations arise from an abelian differential ω by integrating ω in simply connected neighbourhoods to give holomorphic coordinate patches. Consider two such overlapping patches U and V . Let $p \in U$, $q \in V$ and $a \in U \cap V$. Then the coordinate maps are $u(a) = \int_p^a \omega$ and $v(a) = \int_q^a \omega$. Note that $u(a) - v(a)$ is the constant $\int_p^q \omega$. Thus the transition map is $u = v + \int_p^q \omega$, which is a translation in \mathbb{C} . Now, one can pull back the foliation by horizontal lines and the foliation by vertical lines in \mathbb{C} to the Riemann surface consistently over coordinate patches. These define the *horizontal foliation* and *vertical foliation* of ω .

Remark 14. Note that the horizontal and vertical foliations of ω are the horizontal and vertical foliations of $\omega \otimes \omega$ (away from the singularities).

One can construct a branched double cover of a Riemann surface with a quadratic differential (Σ, ϕ) . If locally $\phi = f(z)dz \otimes dz$, a simply connected neighbourhood of each point is covered by two conformally equivalent neighbourhoods, one with the abelian differential $\sqrt{f(z)}dz$ and one with the abelian differential $-\sqrt{f(z)}dz$. If this double cover is disconnected, then ϕ can globally be written as the square of one of its roots. The transition map between the coordinate map $u(a) = \int_p^a \sqrt{f(z)}dz$ and $v(a) = \int_p^a -\sqrt{f(z)}dz$ is $u = -v$. This is a rotation by π . Thus if we have a quadratic differential that is not the square of an abelian differential, we

Restrict attention to a strip (Definition 18) of Σ . What is the obstruction to finding an ω here such that $\phi = \omega \otimes \omega$? We can think of a strip as a cylinder (with one foliation by circles and a transverse foliation whose leaves have end points on opposite boundary components) that is immersed in Σ . On the cylinder, it is possible to choose coordinate charts to rectangles in \mathbb{C} such that the transition maps are translations. When the strip is immersed in Σ , if a segment of a boundary component of the cylinder is identified with itself, transition maps that are rotations will also be required. In the picture

below, consider the shown coordinate patches. If each are sent to \mathbb{C} so that the red foliation is horizontal, all the transition maps cannot be translations.

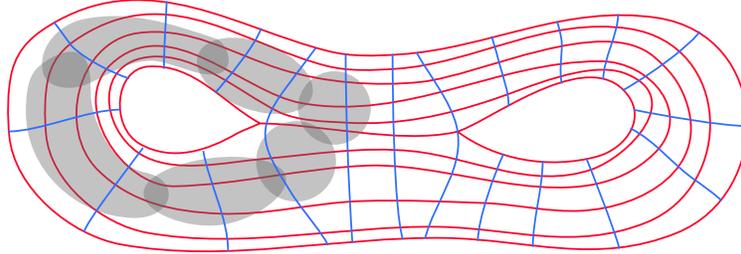


Figure 4.6: The foliation in red cannot be induced by an abelian differential.

Corollary 5. *A taut Heegaard diagram of an irreducible 3-manifold, M , constructs a pair (Σ, ϕ) describing M such that ϕ is the square of an abelian differential when restricted to any strip of Σ .*

Unfortunately, we cannot say that ϕ is the square of an abelian differential globally on Σ because the foliations may have odd ordered singularities.

Question 3. What Heegaard splittings are described by pairs (Σ, ϕ) such that $\phi = \omega \otimes \omega$?

4.3.2 Relation to Work of Wolf

Wolf had the idea of studying quadratic differentials on Riemann surface, Σ , using equivariant maps from the universal cover of a Riemann surface to trees whose edges have lengths in \mathbb{R} (there is a $\pi_1(\Sigma)$ action on both domain and target) [18]. Given such an equivariant map, a foliation with transverse measure is given on the universal cover by the fibres of the map. A path transverse to the foliation is measured by taking the length of its image in the tree. This induces a foliation with transverse measure on the Riemann surface because of the equivariance. Classes of foliations can be constructed by looking at homotopy classes of these equivariant maps. Conversely, beginning with a quadratic differential on a Riemann surface, it can be lifted to the universal cover. Taking a quotient to the vertical leaf space, we get a $\pi_1(\Sigma)$ equivariant map to a tree. Each edge on the tree corresponds to a one parameter family of vertical leaves. The length of this edge is transverse width of this family, obtained by integrating the real value of the local square root along a path

transverse to this family, and taking its absolute value.

Wolf proved that given a Riemann surface, Σ , and a class of measured foliations, \mathcal{F} , there is a unique quadratic differential, ϕ , such that the vertical leaves of ϕ are \mathcal{F} . This is a classic theorem of Hubbard and Masur [6], to which Wolf's proof brings a new perspective.

His approach is to show the existence of a harmonic representative in a class of maps corresponding to a measured foliation \mathcal{F} . He then goes on to show that one can construct a quadratic differential from derivatives of this map, and that its vertical leaves are the fibres of the map and that this quadratic differential is unique with this property [18].

Definition 25. A *harmonic map* $h : \Sigma \rightarrow T$ between a Riemann surface and a tree with real edge lengths is a map such that except for at a discrete set of points, it is locally a harmonic map onto its image. That is for each point in Σ , there is a neighbourhood around it on which h mapping to its image is the real part of a holomorphic function.

It is immediate that the maps $\Sigma_g \rightarrow (\vee_g S^1) \times (\vee_g S^1)$ constructed in this dissertation have harmonic coordinates with respect to the conformal structure on the splitting surface discussed in this section.

Corollary 6. Consider a combinatorial map of square complexes $f : \Sigma_g \rightarrow (\vee_g S^1) \times (\vee_g S^1)$ describing an irreducible 3-manifold. Apply a conformal structure on Σ_g such that each square isomorphic to a unit square in \mathbb{C} . The coordinate maps f_1 and f_2 (where $f = f_1 \times f_2$) are harmonic.

Recall that harmonic maps are also energy minimizing maps.

Definition 26. Given a map between Σ with a Riemannian metric and \mathbb{R} , $f : \Sigma \rightarrow \mathbb{R}$, the *energy* is $E(f) = \int_A \langle df, df \rangle dA$, where dA is the area form, and the inner product is induced by the Riemannian metric.

Remark 15. While the definition of energy requires a conformal metric to be chosen on Σ , energy does not vary depending on this choice.

Remark 16. Consider a harmonic coordinate of a map describing a 3-manifold, $f_1 : \Sigma_g \rightarrow \vee_g S^1$. Using the flat metric on each square of Σ_g , there exists a vector field v on the square, such that at each point, $df_1(w) = \langle v, w \rangle$ for all tangent vectors w . Then $\langle df, df \rangle := \langle v, v \rangle$. Consider that we have charts that take each square to the square $\{a + ib | a, b \in (0, 1)\} \subset \mathbb{C}$ and with these coordinates, f_1 is a (the real part) of each point. Therefore $df(w) = \langle \partial_y, w \rangle$,

and $\langle df, df \rangle = \langle \partial y, \partial y \rangle = 1$. Therefore the energy of each coordinate on each square is 1 (since the area of each square is 1).

Consider an arbitrary Heegaard splitting of a 3-manifold, M . Consider a map $f : \Sigma_g \rightarrow (\mathbb{V}_g S^1) \times (\mathbb{V}_g S^1)$ describing it. Let $f = f_1 \times f_2$. For every conformal structure on Σ_g , we can calculate $E(f_1) + E(f_2)$ using a conformal metric. If M is irreducible, a conformal structure that minimizes this function is found using the methods previously discussed in the first paragraph of this section (Section 4.3). If M is not irreducible, there are two possibilities.

Case 1 $E(f_1) + E(f_2)$ does not have a minimum in the Teichmüller space of Σ_g . In this case, the infimum of the total energy $E(f_1) + E(f_2)$ must still exist since it is bounded below by zero. Therefore there is a sequence of conformal structures on Σ_g that does not converge in the Teichmüller space with total energy converging to this infimum. In the augmented Teichmüller space, this sequence converges to a marked Riemann surface with nodes. The 3-manifolds described by points in the augmented Teichmüller space are discussed in Section 4.1, Corollary 1. This occurs when there is a simple closed curve in Σ_g that is sent by f (and therefore by f_1 and f_2) to a point. Then both $E(f_1)$ and $E(f_2)$ decrease when the length of this curve decreases. Decreasing the length of this curve limits to collapsing this curve to a point, giving a Riemann surface with nodes. Recall the diagram of Corollary 1.

$$\begin{array}{ccc}
 \Sigma_g & \xrightarrow{f} & (\mathbb{V}_g S^1) \times (\mathbb{V}_g S^1) \\
 q_1 \downarrow & & \nearrow \tilde{f} \\
 \Sigma_g / \gamma_i & & \\
 q_2 \downarrow & & \\
 \tilde{\Sigma}_g & &
 \end{array}$$

We have added the Riemann surface with nodes Σ_g / γ_i (γ_i are disjoint simple closed curves that are sent to distinct points by q_1). One can think of Riemann surfaces with nodes as several Riemann surfaces with pairs of points identified. Then q_2 sends any Riemann surface components that are spheres to points to create the pinched surface $\tilde{\Sigma}_g$. The composition $\tilde{f} \circ q_2$ has harmonic coordinates on Riemann surface components that are not spheres and do not describe $S^2 \times S^1$.

Case 2 $E(f_1) + E(f_2)$ has a minimum in the Teichmüller space of Σ_g . In this case, we have transverse foliations (if they were not transverse, it would be possible to decrease the energy which moving to infinity in the Teichmüller space, as discussed in in Case 1). This case occurs when a filling Heegaard diagram can be recovered from f . The examples of Osborne described in Section 5.1 and also in Remark 10 are a case of this. He gives Heegaard diagrams of lens spaces of genus two which are filling and taut. This results in a locally injective map, which results in a Riemann surface and quadratic differential describing these splittings. However, the intersection number of the diagrams he provides are lower than the intersection number of the genus one diagram of the same space. This implies there are fewer squares on the genus two splitting surface than on the genus one splitting surface. Therefore the total energy of the harmonic maps describing these genus two splittings is less than the energy of harmonic maps describing the genus one splittings. It is perhaps surprising that a reducible Heegaard splitting could be described by a lower energy map than an irreducible splitting for the same manifold.

Chapter 5

Examples

5.1 Lens Spaces

Consider 3-manifolds that have a Heegaard splitting with a single handle. These are called lens spaces. They were introduced by Tietze in 1908 [16]. They are all prime, and $S^2 \times S^1$ is the only lens space that is reducible. Excluding $S^2 \times S^1$, all the genus one Heegaard splittings of lens spaces are irreducible. This is because any essential simple closed curve on the two dimensional torus does not separate the surface. Therefore the existence of such a curve that bounds a disk in both handle bodies implies the existence of a two dimensional sphere embedded in the lens space that does not separate. As we have seen earlier, in the discussion after Remark 5, this implies that $S^2 \times S^1$ is a connect summand of the lens space. This is a contradiction, as we have excluded $S^2 \times S^1$ in our hypothesis and lens spaces are prime. Therefore, there cannot be a reducing curve and the Heegaard splitting is irreducible. We restrict attention to irreducible lens spaces in what follows (we do not consider $S^2 \times S^1$).

A lens space can be described by two integer parameters p and q , and is denoted $L(p, q)$. Recall they have genus one Heegaard splittings, so Heegaard diagrams for these are pairs of simple essential curves on a two dimensional torus. For simplicity we define the p and the q in terms of Heegaard diagrams. This is not how it is usually done, but this view ties in best with this dissertation. The first integer, p is equal to the number of intersections of a Heegaard diagram for $L(p, q)$ (notice there are no handle slides for the single genus case). The second integer, q is coprime to p , and is found by starting at a point on a diagram curve and numbering the intersection points as one

travels along the curve (pick a direction). Then starting at a point on the other curve, read off the numbers while travelling along it (pick a direction). The consecutive numbers on the list obtained will differ by a constant, modulo p . This constant is q , and $L(p, q) = L(p, -q(\text{mod } p)) = L(p, q^{-1}(\text{mod } p))$. It turns out the fundamental group of the lens space $L(p, q)$ is \mathbb{Z}/\mathbb{Z}_p .

Consider the domain and the target for maps that describe genus one Heegaard splittings. The splitting surface is the two dimensional torus, and the target space is the product of two bouquets of a single circle. This is also the two dimensional torus. Now because of this thesis, we know there is a locally injective representative map that describes the Heegaard splitting. Locally injective endomorphisms of T^2 are covering maps. And so, the irreducible lens spaces can be described by finite sheeted covering maps of the torus.

Claim 6. *The degree of the covering map is p .*

Proof. Begin with a covering map $f : T^2 \rightarrow S^1 \times S^1$ that describes $L(p, q)$. Recall, that one can obtain a Heegaard diagram for the splitting described by taking the curves that are the coordinate pre-images of a point on each circle coordinate of $S^1 \times S^1$. Below we see two points whose pre-image should be considered to obtain a diagram. The pre-image of the blue points should be coloured blue and the pre-image of the red points should be coloured red.

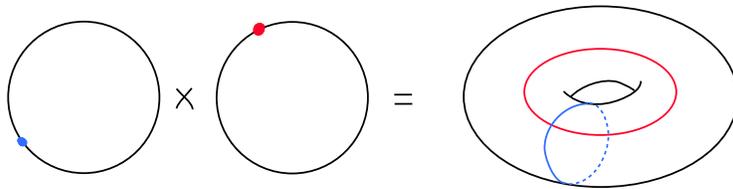


Figure 5.1: $S^1 \times S^1$ shown with a blue point and a red point selected on each coordinate circle.

Note that the blue and the red curves intersect once above. Call the intersection point y . Now, in the domain T^2 , the the diagram we obtained has p intersection points. Therefore there are p pre-images of y . This shows that the degree of f is p . \square

In the example below, using the notation of the proof above, f is the map between tori and it describes $L(5, 2)$.

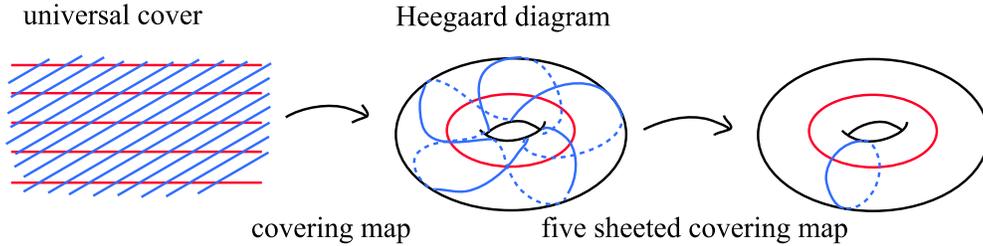


Figure 5.2: Pre-images of the red and blue points.

The other integer, q , can be seen in the square complex structures and conformal structures. Taking each S^1 to be of unit length, we have a metric on $S^1 \times S^1$, which imparts a conformal structure of a single square on the target torus. This induces a conformal structure made of p squares on the domain torus. Recall that the squares can be combinatorially drawn as the dual to the Heegaard diagram. This shows how the p squares are attached together to give a torus.

Proposition 3. *The conformal structure of the domain T^2 is given by attaching p squares to make a cylinder of circumference p and height one, and then attaching the remaining boundary components with a $\frac{q}{p}$ twist. There is a distinct conformal structure for each lens space.*

Proof. The first sentence becomes clear when one draws the squares dual to the diagram, and observes how they are attached to each other. This is best seen in the universal cover. Above, we saw part of the universal cover of the Heegaard diagram of $L(5, 2)$. One can imagine the corresponding picture for any lens space. If we lift a cell decomposition dual to the Heegaard diagram to the universal cover, we obtain a tiling by squares dual to the previous picture. The orange cell decompositions below are this dual decomposition. Looking back at the picture above, one may consider what the fundamental domains of the infinite sheeted covering map on the left look like. These fundamental domains are drawn in black on the picture below. From here, it can be seen that the top of an orange square is not glued to its own bottom but to the bottom of the square q squares to the right. This shows that a $\frac{q}{p}$ twist is

present, since there are p orange squares in one fundamental domain. Below $L(5, 1)$ is shown for comparison.

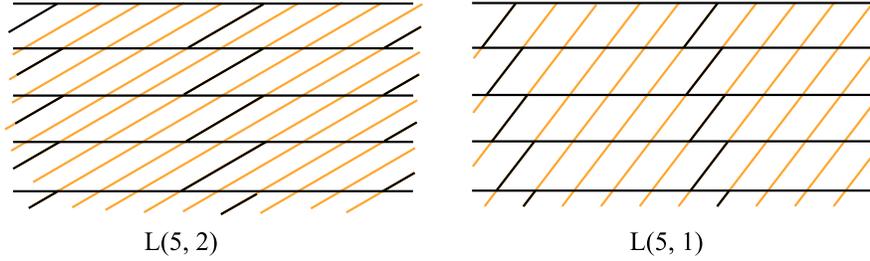


Figure 5.3: Fundamental domains in the universal cover of the splitting surface.

The second sentence in this claim comes from the following observations.

1. The two spaces obtained by gluing a flat cylinder of circumference p and height one with a $\frac{q}{p}$ twist and with a $\frac{q+np}{p}$ twist are conformally equivalent for any integer n .
2. The two spaces obtained by gluing a flat cylinder of circumference p and height one with a $\frac{q}{p}$ twist and with a $\frac{-q(\text{mod } p)}{p}$ twist are conformally equivalent (there is an orientation reversing conformal homeomorphism).
3. After the squares are glued together, “tops and bottoms” of the squares form a simple closed curve on the torus, while the two other sides of all the squares form another simple closed curve. A fractional Dehn twist of $\frac{q}{p}$ along one of these curves gives the same conformal structure as performing a fractional Dehn twist of $\frac{q^{-1}(\text{mod } p)}{p}$ along the other curve. This can be seen with a calculation in modular arithmetic.

These are exactly the considerations when one is taking a quotient of the upper half plane to obtain the moduli space of conformal structures on the torus. Therefore the conformal structure on T^2 induced by $L(p, q)$ and $L(p', q')$ are equivalent if and only if $L(p, q)$ is homeomorphic to $L(p', q')$. \square

Example of Osborne: A Genus Two Splitting of a Lens Space

The author understood these examples through conversation with John Hempel, who referenced the work of Osborne [9], who cites the work of Stevens [15]. Consider the following Heegaard diagram.

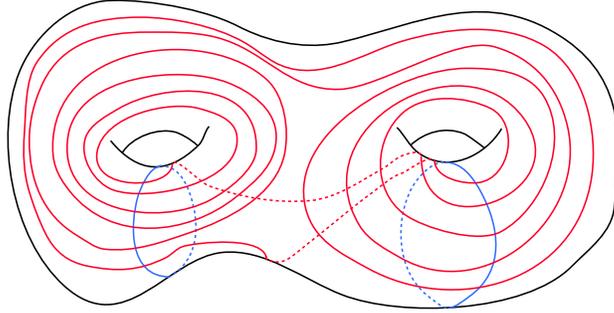


Figure 5.4: A Heegaard splitting for a lens space with fundamental group $\mathbb{Z}/13\mathbb{Z}$.

Let a, b, c, d be non zero integers such that $|ac - bd| = p$, where p is a prime. Then $\mathbb{Z}/p\mathbb{Z} = \langle x, y, |x^a y^b = 1, x^c y^d = 1 \rangle$. Further, this presentation can be realized by a Heegaard diagram [9], [15]. This means that there is a genus two Heegaard diagram such that if the blue curves are labelled x and y , the two relations describe the intersection of two red curves with these blue curves. The 3-manifold this describes is a lens space with fundamental group $\mathbb{Z}/p\mathbb{Z}$. The example above is such a diagram with $a = 5, b = 2, c = 1, d = 3$. It describes a lens space with fundamental group $\mathbb{Z}/13\mathbb{Z}$.

By construction it has $a + b + c + d = 11$ intersection points, which is fewer than the 13 intersection points needed for a diagram for this manifold of genus. Any Heegaard decomposition for genus greater than one is reducible [1], so in particular this one must be reducible. Still, this diagram is filling and taut. This can be seen by careful inspection of the connected components of the complement. Most components are rectangles, and can be seen at a glance. The complements that might be more difficult to follow as they wind around the surface are shown below. after some inspection, it can be seen that they are polygons. They are all intersect each side of each curve in at most one component (the diagram is taut).

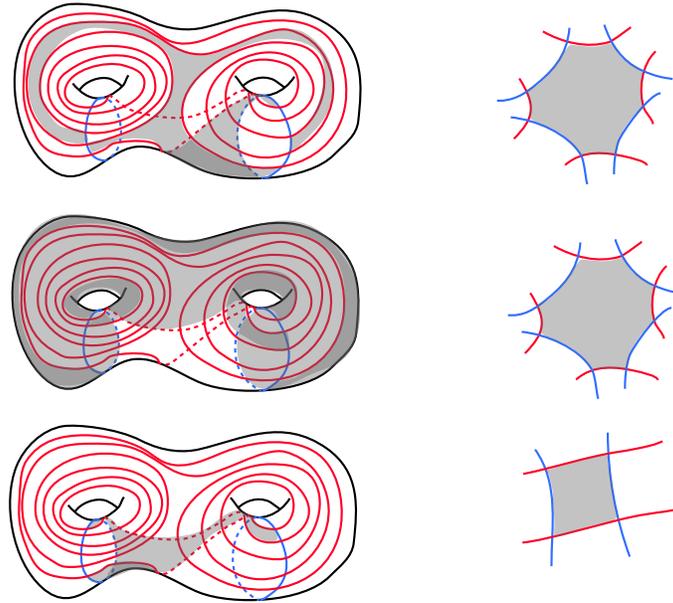


Figure 5.5: Connected components of the complement of the Heegaard diagram.

5.2 Three Dimensional Torus

Consider the three dimensional torus, T^3 , as a solid cube with opposite faces identified. A Heegaard splitting can be found by considering two embedded graphs. Give the solid cube the structure of a cell complex with eight vertices, twelve edges, six faces and one 3-cell. After identifying opposite faces of the cube, the 1-skeleton of the resulting cell complex has one vertex and three edges. The 1-skeleton of the dual cell complex has one vertex (because the original decomposition had one 3-cell), and three edges (because the original decomposition had three faces). These two graphs are disjointly embedded as shown below.

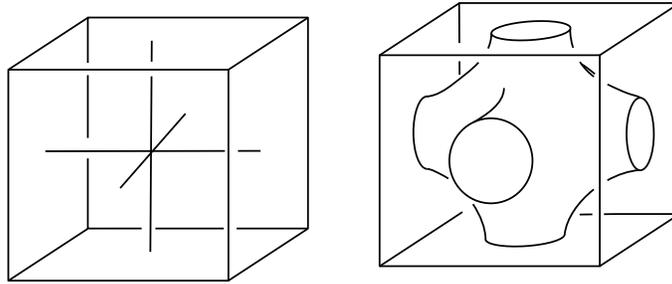


Figure 5.6: A genus three Heegaard splitting of T^3 .

Thickening these two graphs until T^3 is exhausted results in a genus three Heegaard splitting. The picture above shows the splitting surface. This Heegaard splitting can be described by the map induced on the splitting surface by retracting each handle body to the shown 1-skeleton it contains. In the picture below (also seen in the previous section), the fibres of the coordinate taking the splitting surface to the edges of the cube are shown in red. The fibres of the coordinate taking the splitting surface to the dual 1-skeleton are shown in blue. This Heegaard splitting is irreducible and the described map is locally injective.

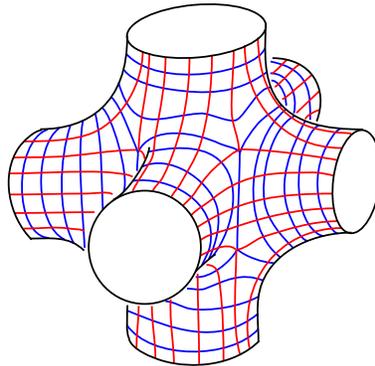


Figure 5.7: Coordinate fibres of a locally injective map describing a genus three Heegaard splitting of T^3 .

Chapter 6

Bibliography

- [1] F. Bonahon and J. P. Otal. Scindements de heegaard des espaces lenticulaires. *Annales scientifiques de l'École Normale Supérieure*, 16(3):451–466, 1983.
- [2] F. Haglund and D. Wise. Special cube complexes. *Geometric and Functional Analysis*, 17:1551–1620, 2008.
- [3] W. Haken. Some results on surfaces in 3-manifolds. *Studies in Modern Topology*, pages 39–98, 1968.
- [4] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [5] P. Heegaard. *Forstudier til en topologisk Teori for de algebraiske Fladers Sammenhang*. PhD thesis, University of Copenhagen, 1898.
- [6] J. Hubbard and H. Masur. Quadratic differentials and foliations. *Acta Math.*, 142:221–274, 1979.
- [7] H. Kneser. Geschlossen flaschen in dreidimensionalen mannigfaltigkeiten. *Jahresbericht der Deutschen Mathematiker Vereinigung*, 38:252–255, 1929.
- [8] J. Milnor. A unique decomposition theorem for 3-manifolds. *American Journal of Mathematics*, 84:1–7, 1962.
- [9] R. P. Osborne. Heegaard diagrams of lens spaces. *Proceedings of the American Mathematical Society*, 84(3):412–414, 1982.
- [10] P. Ozsváth and Z. Szabó. An introduction to Heegaard Floer homology. In *Floer homology, gauge theory, and low-dimensional topology*, volume 5 of *Clay Math. Proc.*, pages 3–27. Amer. Math. Soc., Providence, RI, 2006.

- [11] A. Papadopoulos. *Handbook of Teichmüller theory, Volume II*. European Mathematical Society, 2009.
- [12] K. Reidemeister. Zur dreidimensionalen topologie. *Abh. Math. Sem. Univ. Hamburg*, 11:189–194, 1933.
- [13] J. Singer. Three-dimensional manifolds and their heegaard diagrams. *Trans. Amer. Math. Soc.*, 35:88–111, 1933.
- [14] J. R. Stallings. How not to prove the Poincaré Conjecture. *Annals of Mathematics Studies*, 60:83–88, 1966.
- [15] R. Stevens. Classification of 3-manifolds with certain spines. *Transactions of the American Mathematical Society*, 205:151–166, 1975.
- [16] H. Tietze. Über die topologischen Invarianten mehrdimensionaler Mannigfaltigkeiten. *Monatsh. Math. Phys.*, 19:1–118, 1908.
- [17] F. Waldhausen. Heegaard-zerlegungen der 3-sphäre. *Topology*, 7(2):195 – 203, 1968.
- [18] M. Wolf. On realizing measured foliations via quadratic differentials of harmonic maps to R-trees. *Journal d'Analyse Mathématique*, pages 107–120, 1996.