The Kakimizu complex of a link

Jessica E. Banks
Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2012

This thesis is dedicated to
my family, who do not understand,
and my friends, who pretend they do.

Acknowledgements

I wish to thank Marc Lackenby for being an amazing supervisor.

This research was funded by an EPSRC studentship.

From September to November 2011, I visited the University of California, Davis, supported by the Cecil King Travel Scholarship 2011. I wish to thank the Cecil King Foundation and the London Mathematical Society for this support, and the University of California, Davis, for their hospitality.

Copyrights

- Chapter 2 is based on [2].
- Chapter 3 is based on [4].
- Chapter 4 is based on [3], which is to appear in *Transactions of the American Mathematical Society*.
- Chapter 5 is based on [5], which appeared in *Algebraic and Geometric Topology*.
- Sections 3.3.3, 5.3 and Lemmas 5.2.9, 5.2.10 have not previously appeared in any form.

Abstract

We study Seifert surfaces for links, and in particular the Kakimizu complex MS(L) of a link L, which is a simplicial complex that records the structure of the set of taut Seifert surfaces for L.

First we study a connection between the reduced Alexander polynomial of a link and the uniqueness of taut Seifert surfaces. Specifically, we reprove and extend a particular case of a result of Juhasz, using very different methods, showing that if a non-split homogeneous link has a reduced Alexander polynomial whose constant term has modulus at most 3 then the link has a unique incompressible Seifert surface. More generally we see that this constant term controls the structure of any non-split homogeneous link.

Next we give a complete proof of results stated by Hirasawa and Sakuma, describing explicitly the Kakimizu complex of any non-split, prime, special alternating link.

We then calculate the form of the Kakimizu complex of a connected sum of two non-fibred links in terms of the Kakimizu complex of each of the two links. This has previously been done by Kakimizu when one of the two links is fibred.

Finally, we address the question of when the Kakimizu complex is locally infinite. We show that if all the taut Seifert surfaces are connected then MS(L) can only be locally infinite when L is a satellite of a torus knot, a cable knot or a connected sum. Additionally we give examples of knots that exhibit this behaviour. We finish by showing that this picture is not complete when disconnected taut Seifert surfaces exist.

Contents

Chapter 1

Introduction and background

1.1 Introduction

The Kakimizu complex $\mathrm{MS}(L)$ of a link L is a simplicial complex that records the structure of the set of taut Seifert surfaces for L considered up to ambient isotopy of the complement of L. See Section 1.3 for the definition. This thesis studies the Kakimizu complex, particularly but not exclusively in the context of alternating links. A thorough survey of relevant definitions and known results can be found in [30].

Chapter 2 of this thesis focuses on the relationship between the Alexander polynomial and Seifert surfaces for alternating links. The Alexander polynomial was the first knot polynomial, being defined by Alexander in 1928 ([1]). Crowell and Murasugi independently proved the following result relating the genus of an alternating link L to its reduced Alexander polynomial $\Delta_L^0(t)$.

Theorem 1.1.1 ([8] Theorem 3.5; [28] II Theorem 4.1)**.** *For a non-split, alternating link L with m link components, let R be a Seifert surface given by applying Seifert's algorithm to an alternating diagram for L. Then $\deg \Delta_L^0(t) = 2g(R) + m - 1 = 1 - \chi(R)$, where* deg *denotes degree.*

This theorem has the following corollary.

Theorem 1.1.2 ([13] Theorem 4)**.** *Let L be an alternating link. If R is a surface given by applying Seifert's algorithm to an alternating diagram of L, then R is a taut Seifert surface.*

Here we have chosen to normalise such that $\Delta_L^0(0)$ is defined and positive. In this case, the leading coefficient of $\Delta_L^0(t)$ is $\pm\Delta_L^0(0)$. In [17] Juhasz gives the following relationship between the coefficient $\Delta_L^0(0)$ and the Seifert surfaces for L. He proves this using sutured Floer homology.

Theorem 1.1.3 ([17] Corollary 2.4)**.** *Suppose that K is an alternating knot in $\mathbb{S}^3$ and let $n > 0$. If $\Delta_K^0(0) < 2^{n+1}$ then K can have at most n distinct taut Seifert surfaces that are disjoint in their interiors. In particular, if $\Delta_K^0(0) < 4$ then K has a unique taut Seifert surface.*

We give an alternative proof of this in the case $\Delta_L^0(0) < 4$, using more intuitive geometric techniques, generalised as follows. Homogeneous links are a natural generalisation of both alternating links and positive links (see Definitions 1.2.7, 1.2.8 and 1.6.5).

Theorem 1.1.4. *Let L be a non-split homogeneous link. If $\Delta_L^0(0) < 4$ then L has a unique incompressible Seifert surface.*

If we weaken the condition $\Delta_L^0(0) < 4$ to $\Delta_L^0(0) \leq n$ for some $n \in \mathbb{N}$ we retain some control over the Seifert surfaces for L, as follows.

Theorem 1.1.5. *For fixed $n \in \mathbb{N}$, there is a finite set $\mathcal{S}$ of surfaces embedded in $\mathbb{S}^3$ with the following property. Any non-split, homogeneous link L with $\Delta_L^0(0) \leq n$ has a taut Seifert surface R built from surfaces in $\mathcal{S}$ by reflection, Murasugi sum and plumbing with Hopf bands.*

If $\Delta_L^0(0)$ is prime, R can be formed using only one element of $\mathcal{S}$.

Finally, we show that our methods can be used to give an alternative proof to the following theorem of Riley ([32]).

Theorem 1.1.6. *Choose $g \geq 0$ and $m, n \geq 1$. Then there are only finitely many alternating links L with m link components and genus g such that $\Delta_L^0(0) = n$.*

The aim of Chapter 3 is to give a complete proof of the results stated by Hirasawa and Sakuma in [15], as only 'the idea of the proof' is given in the paper. The main result is an explicit description of $\mathrm{MS}(L)$ for any prime, special alternating link L.

Theorem 1.1.7 ([15] Theorems 1.5, 1.1)**.** *Let L be a prime link with a reduced, special alternating diagram D. Then there is a natural isomorphism between $\mathrm{MS}(L)$ and $\mathcal{K}(D)$. In particular, every taut Seifert surface for L is given by applying Seifert's algorithm to some special alternating diagram for L.*

Here $\mathcal{K}(D)$ is a simplicial complex constructed from D. From this the authors deduce contractibility of $\mathrm{MS}(L)$ for prime, special alternating links. The question of the contractibility of the Kakimizu complex for a general link has since been answered by the following theorem of Przytycki and Schultens.

Theorem 1.1.8 ([31] Theorem 1.1)**.** *The Kakimizu complex of a link is contractible.*

The proof of Theorem 1.1.7 we give makes use of the fact that $\mathrm{MS}(L)$ is simply connected, and in this way differs from that intended by Hirasawa and Sakuma.

We next address the question of the effect on the Kakimizu complex of combining links by taking a connected sum. Kakimizu proved the following, building on work of Eisner in [11].

Theorem 1.1.9 ([19] Theorem B). *Let L_1, L_2 be knots, each not fibred but with a unique incompressible Seifert surface. Then $|\mathrm{MS}(L_1 \# L_2)|$ is homeomorphic to $\mathbb{R}$.*

In Chapter 4 we prove the following more general result, as well as the analogous result for incompressible Seifert surfaces.

Theorem 1.1.10. *Let L_1, L_2 be non-split, non-fibred links in $\mathbb{S}^3$, and let $L = L_1 \# L_2$. Then $|\mathrm{MS}(L)|$ is homeomorphic to $|\mathrm{MS}(L_1)| \times |\mathrm{MS}(L_2)| \times \mathbb{R}$.*

The case that one of the L_i is fibred was dealt with by Kakimizu.

Proposition 1.1.11 ([20] Proposition 2.4). *If L_1 is fibred then $\mathrm{MS}(L_1 \# L_2) \cong \mathrm{MS}(L_2)$.*

In Chapter 5 we consider the question of when the Kakimizu complex is locally infinite. This question was raised by Przytycki and Schultens in [31]. We show the following.

Theorem 1.1.12. *There is a knot whose Kakimizu complex is not locally finite.*

Additionally we prove the following condition on the types of links that might have a locally infinite Kakimizu complex, under the additional assumption that all taut Seifert surfaces for the link are connected.

Theorem 1.1.13. *Let L be an oriented link such that every taut Seifert surface for L is connected. If $\mathrm{MS}(L)$ is locally infinite then L is a satellite of either a torus knot, a cable knot or a connected sum, with winding number 0.*

1.2 Links

Convention 1.2.1. If M is a manifold and $W \subseteq M$, then $\mathcal{N}(W)$ will denote a regular open neighbourhood of W in M, unless otherwise stated.

Definition 1.2.2. A *link* is a piecewise linear embedding $L\colon \coprod_{i=1}^{m} \mathbb{S}^1 \to \mathbb{S}^3$ for some $m \in \mathbb{N}$, considered up to ambient isotopy of $\mathbb{S}^3$. If $m = 1$ we call L a *knot*.

Convention 1.2.3. We will only consider oriented links.

Remark 1.2.4. Intuitively, a link is a finite number of disjoint simple closed curves lying in 3–space, with an arrow pointing one way around each.

Definition 1.2.5. A *diagram* of a link L is a projection D of L from a point in $\mathbb{S}^3 \setminus L$ onto $\mathbb{S}^2 \subset \mathbb{S}^3$ that is a local embedding except at finitely many points, called the *crossings*. In the neighbourhood of a crossing c, the diagram D consists of two simple arcs that meet once transversely at c. We record which of these is the overcrossing and which the undercrossing. We also retain knowledge of the orientation of each link component. Diagrams are considered up to isotopy of $\mathbb{S}^2$.

Remark 1.2.6. We can reconstruct L from D.

Definition 1.2.7. A crossing in a link diagram is defined to be *positive* or *negative* according to the convention show in Figure 1.1. A link L is said to be *positive* (*negative*) if it has a diagram in which all crossings are positive (negative).

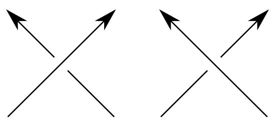

Figure 1.1

Definition 1.2.8. A link diagram D is *alternating* if, when following any link component L_i of L in D, overcrossings and undercrossings are met alternately. A link L is *alternating* if it has an alternating diagram.

Definition 1.2.9. A crossing c in a link diagram D is *nugatory* if there is a simple closed curve ρ that meets D transversely at c and otherwise lies in $\mathbb{S}^2 \setminus D$. The diagram D is *reduced* if it has no nugatory crossings.

Remark 1.2.10. Any link L has a diagram without nugatory crossings, and this can be taken to be alternating if L is alternating.

Definition 1.2.11. A link L is *split* if there is a sphere $S \subset \mathbb{S}^3 \setminus L$ that has components of L on each side.

Convention 1.2.12. We study links that are not split and link diagrams that are connected. If a link L is not split, any diagram of L is connected. Conversely for an alternating link L, Menasco has shown ([24] Theorem 1(a)) that if an alternating diagram D of L is connected then L is not split.

Definition 1.2.13. Given two links L_1 and L_2, a *connected sum* L of L_1 and L_2, denoted $L_1 \# L_2$, is a link formed as follows. Let S be a sphere in $\mathbb{S}^3$, and let V_1, V_2 be the two closed 3–balls into which S divides $\mathbb{S}^3$, so that $\partial V_1 = \partial V_2 = V_1 \cap V_2 = S$. For $i = 1, 2$, position the link L_i in V_i such that $L_i \cap \partial V_i$ is a single embedded arc ρ_i that is not just a point. In doing so, ensure that ρ_1 and ρ_2 coincide but have opposite orientations. Then L is the link $(L_1 \cup L_2) \setminus (\rho_1 \setminus \partial \rho_1)$.

Remark 1.2.14. The resulting link L depends on which component of L_i contains the arc ρ_i for $i = 1, 2$.

Definition 1.2.15. A link L is *prime* if it cannot be expressed as $L_1 \# L_2$ where L_1 and L_2 are links other than the unknot.

Remark 1.2.16. Let D be a diagram of a link L. Suppose there is a simple closed curve ρ in $\mathbb{S}^2$ missing the crossings of D and meeting the edges exactly twice transversely. Then $L = L_1 \# L_2$ for some links L_1, L_2 with diagrams D_1, D_2 respectively. If D is alternating then D_1 and D_2 can be taken to be alternating.

Theorem 1.2.17 ([24] Theorem 1)**.** *Let D be an alternating diagram of a link L with no nugatory crossings. Then L is prime if and only if, whenever ρ is as described above, one component of $\mathbb{S}^2 \setminus \rho$ contains no crossings of D.*

Convention 1.2.18. In other words, a link L with an alternating diagram D is prime if and only if D looks prime. We will therefore use this as the definition of prime for alternating links. That is, we will say D is *prime* if, whenever ρ is as described above, one component of $\mathbb{S}^2 \setminus \rho$ contains only nugatory crossings of D. Note that we do not require D to be reduced.

Definition 1.2.19. A *flype* is a move that changes a link diagram into another diagram of the same link. This move is given by 'turning over' a section of the link corresponding to part of the diagram, untwisting one crossing and creating a new crossing elsewhere, as shown in Figure 1.2.

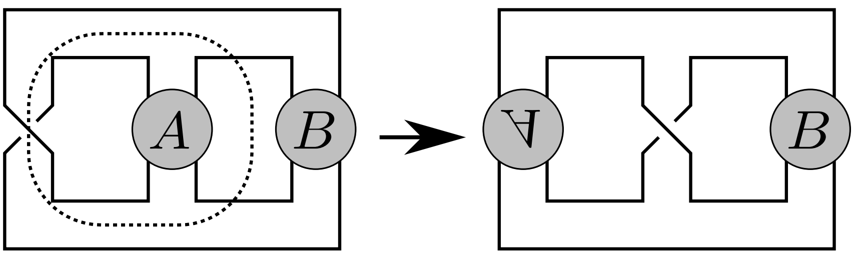

Figure 1.2

Theorem 1.2.20 ([26] Theorem 1; [25] Main Theorem)**.** *Let D and D' be reduced alternating diagrams of a prime link L. Then D and D' differ by a finite sequence of flypes.*

1.3 Seifert surfaces and the Kakimizu complex

Definition 1.3.1. A *Seifert surface* for a link L is a compact, oriented surface R with no closed components embedded in $\mathbb{S}^3$ such that $\partial R = L$ as an oriented link. We consider such surfaces up to ambient isotopy in $\mathbb{S}^3$ keeping L fixed. The surface R can also be viewed as properly embedded in $\mathbb{S}^3 \setminus \mathcal{N}(L)$, up to ambient isotopy of $\mathbb{S}^3 \setminus \mathcal{N}(L)$. We will not explicitly distinguish between these two settings.

Say R is *taut* if it has maximal Euler characteristic among all Seifert surfaces for L.

Lemma 1.3.2 ([13] Lemma 3)**.** *Let R be a Seifert surface for a link L. If R is not taut then there is a Seifert surface R' for L that can be made disjoint from R such that $\chi(R') > \chi(R)$.*

Definition 1.3.3. Given a diagram D of a link L, resolve each crossing of D according to the orientation of the link, as shown in Figure 1.3. This creates a set of disjoint oriented simple closed curves in $\mathbb{S}^2$. A *Seifert circle* is any of these simple closed curves.

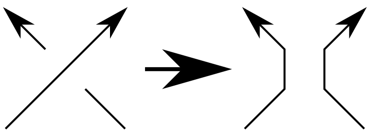

Figure 1.3

Construction 1.3.4 (Seifert's algorithm)**.** Let D be a diagram of a link L. Construct a Seifert surface R for L as follows.

For each Seifert circle C in D, R contains a disc S_C with boundary C. It may not be possible for all these discs to lie disjointly in $\mathbb{S}^2$, but since there are finitely many Seifert circles the discs can be embedded disjointly in $\mathbb{S}^3$ by pushing each slightly above or below $\mathbb{S}^2$. Each disc S_C has an orientation induced by the orientation of its boundary C.

At each crossing c, join the discs that meet there by a half-twisted band, as shown in Figure 1.4. This can be done within $\mathcal{N}(c)$. The orientation of the two discs extends to an orientation of the band, and this orientation agrees with that of L in $\mathcal{N}(c)$. Hence the resulting surface R has the required orientation. As D is connected, so is R.

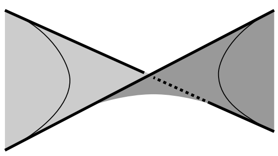

Figure 1.4

Remark 1.3.5. In general some choice is needed when positioning the discs in $\mathbb{S}^3$. This may mean that two or more distinct Seifert surfaces for L can be built in this way.

The Kakimizu complex was first defined as follows.

Definition 1.3.6 ([19] p225)**.** For a link L, the *Kakimizu complex* $\mathrm{MS}(L)$ of L is a simplicial complex, the vertices of which are the ambient isotopy classes of taut Seifert surfaces for L. Distinct vertices $R_0, \ldots, R_n$ span an n–simplex exactly when they can be realised disjointly.

Definition 1.3.7. A metric is defined on the vertices of $\mathrm{MS}(L)$. The distance between two vertices is the distance in the 1–skeleton of $\mathrm{MS}(L)$ when every edge has length 1.

In [19], Kakimizu claimed that the distance between two vertices of $\mathrm{MS}(L)$ can be calculated by considering lifts of the two Seifert surfaces to the infinite cyclic cover of the link complement, as follows.

Proposition 1.3.8 ([19] Proposition 3.1)**.** *Let L be a link, and let $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$. Consider the infinite cyclic cover $\tilde{M}$ of M (that is, the cover corresponding to the kernel of the linking number* $\mathrm{lk}\colon \pi_1(M) \to \mathbb{Z}$*), and let τ be a generator for the group of covering transformations.*

Let R, R' be taut Seifert surfaces for L that represent distinct vertices of $\mathrm{MS}(L)$*. Choose a lift V_0 of $M \setminus R$ to $\tilde{M}$. For $n \in \mathbb{Z}$ let $V_n = \tau^n(V_0)$.*

Take a lift $V_{R'}$ of $M \setminus R'$. Isotope the Seifert surface R' in M so as to minimise the value of $\max\{n : V_{R'} \cap V_n \neq \emptyset\} - \min\{n : V_{R'} \cap V_n \neq \emptyset\}$*. Then* $\mathrm{d}_{\mathrm{MS}(L)}(R, R')$ *is given by* $\max\{n : V_{R'} \cap V_n \neq \emptyset\} - \min\{n : V_{R'} \cap V_n \neq \emptyset\}$.

Przytycki and Schultens pointed out in [31] that this result is not in fact true in full generality. It may fail in the case of a link that bounds a disconnected taut Seifert surface. Figure 1.5 gives an instructive example of this. The link L_α shown consists of two linked copies of the knot 7_4. This knot, which is the $(1, 3, 3)$ pretzel knot, has two taut Seifert surfaces. One is given by applying Seifert's algorithm to an alternating diagram. The other is given by performing a flype on the diagram that moves the single crossing across one of the lines of three crossings, and then applying Seifert's algorithm. Combining two copies of each of these surfaces gives the two disjoint taut Seifert surfaces R_a, R_b for L_α shown in Figure 1.5. The arc ρ shown in Figure 1.6 runs from the top side of R_a to the bottom side, and in doing so passes through R_b in the positive direction twice. This means that if we first take lifts V_n of $M \setminus R_b$ and then a lift V of $M \setminus R_a$ as in Proposition 1.3.8 we find that, for example, V intersects V_0, V_1, V_2. The surfaces R_a, R_b should therefore be distance 2 apart, instead of adjacent.

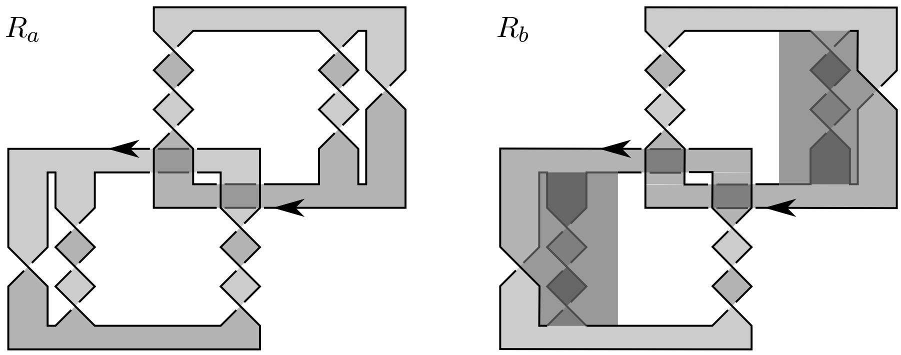

Figure 1.5

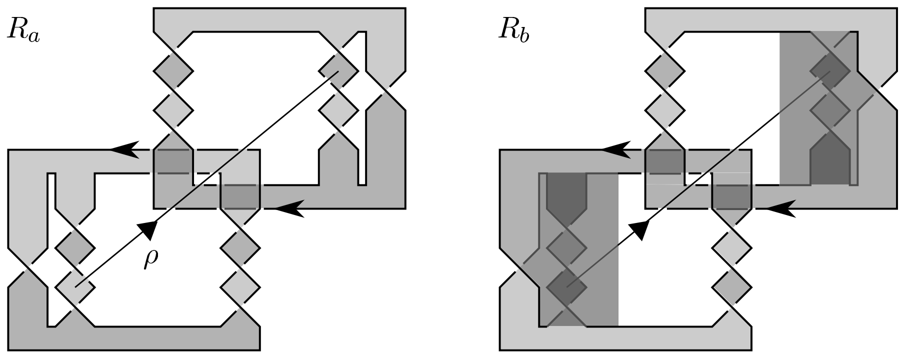

Figure 1.6

To address this difficulty, Przytycki and Schultens redefine the Kakimizu complex as follows. Note that a simplicial complex is said to be *flag* if it contains any simplex whose 1–skeleton is contained in the complex.

Definition 1.3.9 (see [31])**.** Let L be a link, and let $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$. Define the *Kakimizu complex* $\mathrm{MS}(L)$ of L to be the following flag simplicial complex. Its vertices are ambient isotopy classes of taut Seifert surfaces for L. Two distinct vertices span an edge if they have representatives R, R' such that a lift of $M \setminus R'$ to the infinite cyclic cover of M intersects exactly two lifts of $M \setminus R$.

With this definition for $\mathrm{MS}(L)$, Proposition 1.3.8 is true. This is also the definition used in the proof of Theorem 1.1.8. Note that the two definitions are equivalent for knots and for non-split alternating links (see Theorem 1.1.1 and Proposition 2.1.7). In addition, it remains true that if taut Seifert surfaces R and R' are adjacent then they can be realised disjointly.

If we move from studying taut Seifert surfaces to looking at all incompressible Seifert surfaces, a number of results continue to hold, such as the following.

Theorem 1.3.10 ([19] Theorem A; see also [35]). *Let R be an incompressible Seifert surface for a link L. If L has more than one incompressible Seifert surface, up to ambient isotopy, then there is an incompressible Seifert surface R' for L, distinct from R, that can be made disjoint from R in $\mathbb{S}^3 \setminus \mathcal{N}(L)$.*

1.4 Low dimensional manifolds, isotopies and sutured manifolds

Definition 1.4.1 ([31] Section 3). Let M be a connected 3–manifold, and let S, S' be (possibly disconnected) surfaces properly embedded in M.

Call S and S' *almost transverse* if, given a component S_0 of S and a component S'_0 of S', they either coincide or intersect transversely. Call the surfaces *almost disjoint* if, given a component S_0 of S and a component S'_0 of S', they either coincide or are disjoint. Say they are *∂–almost disjoint* if $\partial S = \partial S'$ and, given a component S_0 of S and a component S'_0 of S', they either coincide or have disjoint interiors.

Say S and S' *bound a product region* if the following holds. There is a compact surface T, a finite collection $\rho_T \subseteq \partial T$ of arcs and simple closed curves and a map of $N = (T \times \mathrm{I})/\!\sim$ into M that is an embedding on the interior of N and has the following properties.

- $T \times \{0\} = S \cap N$ and $T \times \{1\} = S' \cap N$.
- $\partial N \setminus (T \times \partial \mathrm{I}) \subseteq \partial M$.

Here $\sim$ collapses $\{x\} \times \mathrm{I}$ to a point for each $x \in \rho_T$. The *horizontal boundary* of N is $(T \times \partial \mathrm{I})/\sim$. Say S and S' have *simplified intersection* if they do not bound a product region.

Proposition 1.4.2 ([36] Corollary 3.2). *Suppose surfaces S_0, S_1 bound a product region N. Let S' be an incompressible surface that is transverse to S_0, S_1. Suppose $S' \cap \mathrm{int}(N) \neq \emptyset$ but $S' \cap (S_1 \cap N) = \emptyset$. Then each component of $S' \cap \mathrm{int}(N)$ bounds a product region with a subsurface of S_0. In particular, if additionally $S' \cap (S_0 \cap N) = \emptyset$ then this section of S' is parallel to those of S_0, S_1 that bound N.*

Proposition 1.4.3 ([34] Proposition 4.8; see also [36] Proposition 5.4 and Corollary 3.2). *Let M be a ∂–irreducible Haken manifold. Let S, S' be incompressible, ∂–incompressible surfaces properly embedded in M in general position.*

1. *If S and S' are isotopic then there is a product region between them.*
2. *Suppose $S \cap S' \neq \emptyset$, but S can be isotoped to be disjoint from S'. Then there is a product region between S and S'.*

Remark 1.4.4. We will usually apply this proposition with $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$ for a link L. If L is neither a split link nor the unknot then M is Haken and ∂–irreducible. Furthermore, if S, S' are taut Seifert surfaces for L then they are properly embedded, incompressible and ∂–incompressible. This remains true if we only consider some components of such surfaces.

The following definitions, which are mostly due to Gabai, are all used by Kobayashi in [21].

Definition 1.4.5. A *sutured manifold* (M, s) is a compact, orientable 3–manifold M, together with a finite set s of disjoint simple closed curves on ∂M, called the *sutures*. The sutures divide ∂M into two (possibly disconnected) compact, oriented surfaces $S_+(M)$ and $S_-(M)$ such that $S_+(M) \cap S_-(M) = s$ and, if ρ is a suture, $S_+(M)$ and $S_-(M)$ meet at ρ with opposite orientations. In addition, for $\rho \in s$ we choose a product neighbourhood $\gamma(\rho) = \rho \times [-1, 1]$ of ρ in ∂M, so $\gamma(s)$ consists of $|s|$ disjoint annuli.

Remark 1.4.6. We could instead first choose suitable annuli $\gamma(s)$, and then take s to be a set of core curves of $\gamma(s)$. A sutured manifold may also have a distinguished set of toral boundary components, but we will not need this. Note also that each suture ρ inherits an orientation from each of $S_+(M)$ and $S_-(M)$, and these orientations agree.

Definition 1.4.7. A *product sutured manifold* is a sutured manifold (M, s) that is homeomorphic to $S_+(M) \times [-1, 1]$ with $s = \partial S_+(M) \times \{0\}$ (and $\gamma(s) = \partial M \times [-1, 1]$).

Remark 1.4.8. A product region as in Definition 1.4.1 is a product sutured manifold, with sutures along the core curves of $(S \cap S') \cup \partial M$, although the orientations of the boundary of the sutured manifold may not agree with those of S and S'.

Definition 1.4.9. Let T be a surface properly embedded in M with $\partial T = s$ as oriented curves. Say T is *parallel* to $S_+(M)$ if there is an embedding $\eta\colon T \times [0, 1] \to M$ such that $\eta(\partial T \times [0, 1]) \subseteq \gamma(s)$, while $\eta(T \times \{0\}) = T$ and $\eta(T \times \{1\}) = S_+(M) \setminus \mathrm{int}_{\partial M}(\gamma(s))$.

Definition 1.4.10. A sutured manifold (M, s) is an *almost product sutured manifold* if every incompressible surface T properly embedded in M with $\partial T = s$ is parallel to $S_+(M)$ or to $S_-(M)$.

Definition 1.4.11. A disc T properly embedded in a sutured manifold (M, s) is a *product disc* if ∂T meets s at exactly two points, where it crosses s transversely. Up to isotopy of T, or of $\gamma(s)$, we may assume $\partial T \cap \gamma(s)$ consists of two simple arcs that are essential in $\gamma(s)$.

Definition 1.4.12. Let (M, s) be a sutured manifold that contains a product disc T. Let ρ be a simple arc on T joining the two points of $\partial T \cap s$ and let $T \times [-1, 1]$ be a product neighbourhood of T in M. The sutured manifold (M', s') obtained from (M, s) by a *product disc decomposition* along T has $M' = M \setminus (T \times (-1, 1))$ and $s' = (s \cap M') \cup (\rho \times \{\pm 1\})$. Figure 1.7 shows what happens in a neighbourhood of T.

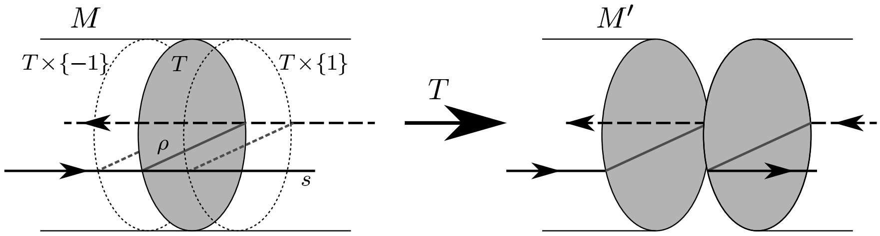

Figure 1.7

Remark 1.4.13. M' is a product sutured manifold if and only if M is.

Remark 1.4.14. Suppose M is irreducible and T' is an incompressible surface properly embedded in M with $\partial T' = s$. By an isotopy we can ensure that $T' \cap (T \times [-1, 1]) = \rho \times [-1, 1]$. Then $T' \setminus (T \times (-1, 1))$ is a surface properly embedded in M' with boundary s'. This leads to the following.

Proposition 1.4.15 (see [21] Lemmas 2.1, 2.2)**.** *Let $(M, s) \xrightarrow{T} (M', s')$ be a product disc decomposition. Suppose that M is irreducible. Then M is an almost product sutured manifold if and only if M' is.*

Definition 1.4.16. Say a product disc is *inessential* if it is separating and the product disc decomposition along it creates a component that is a 3–ball with a single suture. Otherwise, it is *essential*.

Remark 1.4.17. If a product disc T is inessential then $\partial T \cap S^+(M)$ is inessential in $S^+(M)$, and the same is true for $S^-(M)$.

Lemma 1.4.18. *Let (M, s) be a sutured manifold, and let ρ be an essential arc properly embedded in $S^+(M)$. Let T be a product disc in M such that ∂T meets ρ exactly once and such that ρ cannot be isotoped to be disjoint from ∂T. Let (M', s') be the result of the product disc decomposition along T, and let ρ_1, ρ_2 be the two parts of ρ in M'. Suppose there is an essential product disc S in M with $\rho \subset \partial S$. Then there are essential product discs S_1, S_2 in M' with $\rho_1 \subset \partial S_1, \rho_2 \subset \partial S_2$.*

Proof. Notice that we are free to change S, provided we always have $\rho \subset \partial S$. A simple closed curve in $T \cap S$ that is innermost in T bounds a disc both in T and in S. Replacing the subdisc of S with the subdisc of T reduces $|T \cap S|$. Thus we can remove all simple closed curves from $T \cap S$. Exactly one arc of $T \cap S$ has an endpoint on $S^+(M)$, at the point where ∂T crosses ρ, and the other endpoint of this arc lies on $S^-(M)$. Suppose there is at least one other arc of $T \cap S$. Choose such an arc that is outermost in T. This cuts off a subdisc of T disjoint from $S^+(M)$ with interior disjoint from S. It also cuts off a subdisc of S that is disjoint from $S^+(M)$. Replacing the subdisc of S with the subdisc of T again reduces $|T \cap S|$. We may assume, therefore, that $T \cap S$ is a single arc, running from $S^+(M)$ to $S^-(M)$. Since $S \cap S^+(M) = \rho$ has not changed, S is still essential.

Now the disc decomposition along T cuts S into two product discs S_1, S_2. Suppose S_1 is inessential. Then $\rho_1 = \partial S_1 \cap S^+(M')$ is inessential in $S^+(M')$. Let T_1 be a disc in $\partial M'$ between ρ_1 and s'. Then $T_1^* = T_1 \cap \partial M$ is a disc. This disc defines an isotopy of $\rho \cap T_1^*$ in ∂M that makes ρ disjoint from ∂T. This contradicts that no such isotopy exists. Thus S_1 is essential. Similarly, so is S_2. $\square$

Definition 1.4.19. For a sutured manifold (M, s) embedded in $\mathbb{S}^3$, the *complementary sutured manifold* (M', s') is defined by $M' = \mathbb{S}^3 \setminus \mathrm{int}_{\mathbb{S}^3}(M)$ and $s' = s$.

By the *complementary sutured manifold to a Seifert surface* R we mean the complementary sutured manifold to the product sutured manifold given by a product neighbourhood of R. Similarly, by a *product disc decomposition of* R we mean a product disc decomposition of the complementary sutured manifold.

Definition 1.4.20 ([33] 10H1)**.** A link L is *fibred* if there exists a Seifert surface R and a map $F\colon \mathbb{S}^3 \setminus L \to \mathbb{S}^1$ such that the following holds. For $x \in \mathbb{S}^1$ there is a homeomorphism $h_x\colon F^{-1}(\mathcal{N}(x)) \to \mathcal{N}(x) \times R$ such that the restriction of F to $F^{-1}(\mathcal{N}(x))$ is given by h_x followed by projection onto the first factor. The surface R is called a *fibre*.

Remark 1.4.21. A link L with a Seifert surface R is fibred with fibre R if and only if the complementary sutured manifold to R is a product sutured manifold.

In the language of sutured manifolds, Theorem 1.3.10 gives the following (see [21] Corollary 4.2).

Proposition 1.4.22. *Let (M, s) be the complementary sutured manifold to an incompressible Seifert surface R for a link L. If M is an almost product sutured manifold then R is a unique incompressible Seifert surface for L.*

Proof. Suppose R is not unique. By Theorem 1.3.10, there is an incompressible Seifert surface R' for L that is disjoint from R and is not ambient isotopic to it. Then R' can be seen as a surface properly embedded in M with $\partial R' = s$. If R' were parallel to $S^+(M)$ or $S^-(M)$ this would show that R' was ambient isotopic to R in the complement of L, which is not the case. Thus R' shows that M is not an almost product sutured manifold. $\square$

Proposition 1.4.23 ([20] Proposition 1.4)**.** *Let R be a connected taut Seifert surface for a link L. Let R' be the surface obtained from R by a product disc decomposition, and let $L' = \partial R'$. Then the link of the vertex R in* $\mathrm{MS}(L)$ *is isomorphic to the link of R' in* $\mathrm{MS}(L')$*. Furthermore, the isomorphism is induced by the procedure described in Remark 1.4.14.*

1.5 Graphs

Definition 1.5.1. A *graph* $\mathcal{G}$ consists of a set of vertices, denoted $\mathrm{V}(\mathcal{G})$, a set of edges, denoted $\mathrm{E}(\mathcal{G})$, and a function ε that assigns to each edge $e \in \mathrm{E}(\mathcal{G})$ one or two vertices, called the endpoints of e.

Convention 1.5.2. Unless otherwise stated, we assume $\mathrm{V}(G)$ and $\mathrm{E}(G)$ are finite. In general we allow a graph to contain multi-edges (distinct $e, e' \in \mathrm{E}(G)$ with $\varepsilon(e) = \varepsilon(e')$) and loops ($e \in \mathrm{E}(G)$ whose two endpoints are the same). By convention these are usually excluded in the definition of the term 'graph', but we will need them later.

Convention 1.5.3. We will always assume a graph to be connected (although we may consider subgraphs that are disconnected).

Definition 1.5.4. Given a set $A \subseteq \mathrm{V}(\mathcal{G})$, the *induced subgraph* $\mathcal{G}[A]$ is the graph with vertex set A and edge set $\{e \in \mathrm{E}(\mathcal{G}) : \varepsilon(e) \subseteq A\}$.

For $B \subseteq \mathrm{E}(\mathcal{G})$, denote by $\mathcal{G} \setminus B$ the graph obtained by deleting all edges of B from $\mathcal{G}$. That is, $\mathrm{V}(\mathcal{G} \setminus B) = \mathrm{V}(\mathcal{G})$ and $\mathrm{E}(\mathcal{G} \setminus B) = \mathrm{E}(\mathcal{G}) \setminus B$.

Given $e \in \mathrm{E}(\mathcal{G})$, $\mathcal{G}/e$ is the graph obtained by contracting e to a point. This means $\mathrm{V}(\mathcal{G}/e) = (\mathrm{V}(\mathcal{G}) \setminus \varepsilon(e)) \cup \{v_e\}$ where $v_e \notin \mathrm{V}(\mathcal{G})$, while $\mathrm{E}(\mathcal{G}/e) = \mathrm{E}(\mathcal{G}) \setminus e$ and v_e replaces both ends of e in ε. If $B = \{e_1, \ldots, e_n\}$ for some $n \in \mathbb{N}$ and $e_i \in \mathrm{E}(\mathcal{G})$ then $\mathcal{G}/B = ((\cdots(\mathcal{G}/e_1)/e_2)\cdots)/e_n$.

Definition 1.5.5. A vertex $v \in \mathrm{V}(\mathcal{G})$ of a graph $\mathcal{G}$ is a *cut vertex* if $\mathcal{G}[\mathrm{V}(\mathcal{G}) \setminus \{v\}]$ is disconnected.

Definition 1.5.6. A *pointed graph* $(\mathcal{G}, v)$ is a graph $\mathcal{G}$ with a distinguished vertex v.

Definition 1.5.7. For $u, v \in \mathrm{V}(\mathcal{G})$, the distance $\mathrm{d}(u, v)$ is the minimum number of edges in $\mathcal{G}$ between u and v. The *radius* of a pointed graph $(\mathcal{G}, v)$ is $\max\{\mathrm{d}(v, w) : w \in \mathrm{V}(\mathcal{G})\}$.

Convention 1.5.8. For $n \in \mathbb{N}$ and $v \in \mathrm{V}(\mathcal{G})$, denote by $\mathcal{B}(\mathcal{G}, v, n)$ the graph $\mathcal{G}[A_n]$, where $A_n = \{w \in \mathrm{V}(\mathcal{G}) : \mathrm{d}(v, w) \leq n\}$. That is, $\mathcal{B}(\mathcal{G}, v, n)$ is the maximal pointed subgraph of $\mathcal{G}$ with distinguished vertex v and radius n.

Definition 1.5.9. A *digraph* $(\mathcal{G}, \mathcal{O})$ is a graph $\mathcal{G}$ together with an orientation $\mathcal{O}$, which assigns to each $e \in \mathrm{E}(\mathcal{G})$ an initial endpoint $\iota(e)$ and a terminal endpoint $\tau(e)$ such that $\{\iota(e), \tau(e)\} = \varepsilon(e)$. We say that e starts at $\iota(e)$ and ends at $\tau(e)$.

Define the *in-degree* of a vertex $v \in \mathrm{V}(G)$ to be the number of edges $e \in \mathrm{E}(G)$ with $\tau(e) = v$, and define the *out-degree* analogously.

$\mathcal{G}$ is called the *underlying graph* of $(\mathcal{G}, \mathcal{O})$. We will at times consider more than one orientation on the same graph $\mathcal{G}$. Where the choice of orientation is clear, or not important, we will denote $(\mathcal{G}, \mathcal{O})$ by $\mathcal{G}$.

Definition 1.5.10. A *directed path* in G is a path $v_0, e_1, \ldots, e_n, v_n$ such that $\iota(e_i) = v_{i-1}$ and $\tau(e_i) = v_i$ for $1 \leq i \leq n$. A *cycle* is a directed path $v_0, e_1, \ldots, e_n, v_n$ with $v_0 = v_n$.

Two directed paths $v_0, e_1, \ldots, e_n, v_n$ and $v'_0, e'_1, \ldots, e'_m, v'_m$ are said to be *edge-disjoint* if there do not exist n_0, m_0 with $e_{n_0} = e'_{m_0}$.

$\mathcal{G}$ is *$\mathcal{O}$–connected* if for any $u, v \in \mathrm{V}(\mathcal{G})$ there is a directed path in $(\mathcal{G}, \mathcal{O})$ from u to v.

Convention 1.5.11. A graph is *planar* if it has an embedding into $\mathbb{S}^2$. We shall regard this embedding as fixed (some authors call such a graph 'plane').

Definition 1.5.12. Given a planar graph $\mathcal{G}$, we may define the *dual graph* $\mathcal{G}^*$, which is again planar. It has a vertex for each region of $\mathbb{S}^2 \setminus \mathcal{G}$. There is one edge e' in $\mathcal{G}^*$ for each $e \in \mathrm{E}(\mathcal{G})$, joining the vertices corresponding to the regions of $\mathbb{S}^2 \setminus \mathcal{G}$ adjacent to e.

Definition 1.5.13. Given a link L with diagram D, the *underlying graph* $\mathcal{G}$ has a vertex at each crossing in D, and an edge for each arc in D joining two crossings. The *induced orientation* O is that given by the orientation of the link L. We refer to $(\mathcal{G}, O)$ as the *underlying digraph.* We will later put other orientations on $\mathcal{G}$.

Remark 1.5.14. $\mathcal{G}$ is planar. We can reconstruct D from $(\mathcal{G}, O)$, with its embedding into $\mathbb{S}^2$, provided we also know, for each crossing, which arc is the overcrossing and which is the undercrossing.

1.6 Special alternating links and Murasugi sums

Definition 1.6.1. Let D be a diagram of a link L. A Seifert circle C in D is *special* if it bounds a disc in $\mathbb{S}^2 \setminus D$. We will call this disc the *inside* of C (for the diagram of the unknot with no crossings, pick one of the two discs to be the inside). The diagram D is *special* if every Seifert circle in D is special.

Remark 1.6.2. A special alternating link diagram is either positive or negative.

A Seifert circle C may be seen as a cycle in the underlying digraph $(\mathcal{G}, O)$ that turns at every crossing it meets (the direction it turns will always be determined by O). We will not explicitly distinguish between this viewpoint and that in Definition 1.3.3.

Let C be a non-special Seifert circle in a diagram D. We can split D along C to create two new non-trivial link diagrams D_1 and D_2 as follows. Let S be one component of $\mathbb{S}^2 \setminus C$ (so S is an open disc). Let $A = \{v \in V(\mathcal{G}) : v \notin S\}$. Then all vertices of $\mathcal{G}[A]$ are 4–valent except for some of the vertices on C, which are 2–valent. Since C is a cycle, such a vertex has in-degree and out-degree 1. Let $(\mathcal{G}_1, O_1)$ be the digraph obtained by contracting each edge of $\mathcal{G}[A]$ whose terminal vertex is 2–valent. Now take $(\mathcal{G}_1, O_1)$ as the underlying graph of D_1. The choice of undercrossing and overcrossing arcs at a crossing c in D_1 is induced by that at c in D. The diagram D_2 is defined analogously from the other component of $\mathbb{S}^2 \setminus C$.

If D is alternating, then so are D_1 and D_2 (see Figure 1.8).

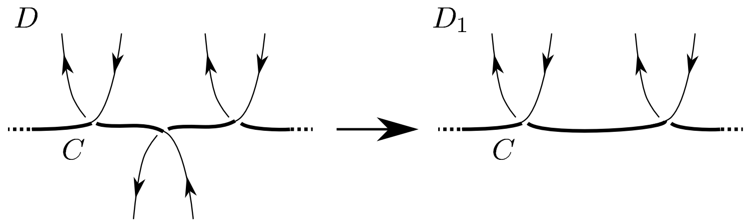

Figure 1.8

Definition 1.6.3 ([29] 3.1)**.** Let L be a link with a diagram D that can be split into diagrams D_1, D_2 as above. Let L_1, L_2 be the links with diagrams D_1, D_2 respectively. We say that D is a $*$*–product* of D_1 and D_2, written $D = D_1 * D_2$. Note that D_1 and D_2 alone do not tell us how to construct D. We also write $L = L_1 * L_2$.

Remark 1.6.4. Let C be the non-special Seifert circle along which D was split. For $i = 1, 2$, let R_i be the Seifert surface for L_i given by applying Seifert's algorithm to D_i, and let S_i be the disc in R_i bounded by C. Let R be a Seifert surface for L given by identifying S_1 and S_2. Then R is given by applying Seifert's algorithm to D.

The interaction of $*$–product with link diagrams makes it a useful tool in [8] and [28], and below.

Definition 1.6.5 ([7] p536)**.** A link diagram D is *homogeneous* if it is formed from special alternating link diagrams by taking connected sums as in Remark 1.2.16 and $*$–products. A link L is *homogeneous* if it has a homogeneous diagram.

Proposition 1.6.6 ([7] Corollary 4.1)**.** *The Seifert surface given by applying Seifert's algorithm to a homogeneous diagram of a link is taut.*

Proposition 1.6.7 ([7] Corollary 3.1, Theorem 8)**.** *A link with a connected homogeneous diagram is not split.*

The $*$–product of Definition 1.6.3 has been generalised as follows (see, for example, [21]).

Definition 1.6.8. Let R, R_1, R_2 be Seifert surfaces for links L, L_1, L_2 respectively. R is a *Murasugi sum* of R_1 and R_2 if the following hold.

- There is a 2–sphere $S \subset \mathbb{S}^3$ dividing $\mathbb{S}^3$ into two closed 3–balls V_1 and V_2.
- $R_1 \subset V_1$, $R_2 \subset V_2$ and $R = R_1 \cup R_2$.
- $T = R_1 \cap S = R_2 \cap S$ is a closed disc.
- T is a $2n$–gon for some $n \in \mathbb{N}$. That is, ∂T consists of $2n$ simple arcs $\rho_1^1, \rho_1^2, \ldots, \rho_n^1, \rho_n^2$ such that, for all i, the arc ρ_i^1 is part of $L_1 = \partial R_1$ and properly embedded in R_2 whereas ρ_i^2 is part of $L_2 = \partial R_2$ and properly embedded in R_1.

When $n = 2$, this operation is known as *plumbing*. A connected sum of two links can be seen as a Murasugi sum, for example by taking $n = 1$.

The behaviour of Seifert surfaces under Murasugi summation depends on whether or not the links involved are fibred.

Theorem 1.6.9 ([21] Theorem 5.1)**.** *Suppose R_1 and R_2 are taut. Then L has a unique taut Seifert surface if and only if L_1, L_2 each have a unique taut Seifert surface and L_i is fibred with fibre R_i for either $i = 1$ or $i = 2$.*

We will wish to apply the 'if' direction of Theorem 1.6.9 in the case of incompressible Seifert surfaces instead of taut ones. The proof in [21] works in this case if we replace [21] Theorem 3.1 with Theorem 1.3.10.

Proposition 1.6.10 ([12] Theorem 3)**.** *L is fibred with fibre R if and only if both L_1 and L_2 are fibred with fibres R_1 and R_2 respectively.*

Let D be a special alternating diagram of a link L, other than a diagram of the unknot with a single crossing. We may colour the regions of D in a checkerboard pattern by making the inside of each Seifert circle black and colouring the remaining regions white. Then the Seifert surface R given by applying Seifert's algorithm to D is formed from the black regions. If a region r of D is a white bigon, it defines a product disc in the complementary sutured manifold to R. The effect on D of the product disc decomposition along this disc is to remove the region r, replacing the two crossings of r with a single crossing. Such a change to D has the effect of pulling off a Hopf band from R, doing the reverse of taking a Murasugi sum (see Figure 1.9). Therefore we have the following.

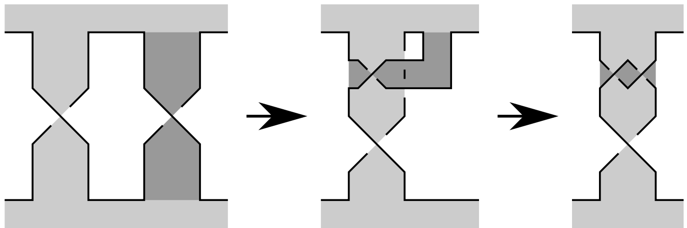

Figure 1.9

Lemma 1.6.11. *Removing a white bigon from a special alternating diagram D corresponds to pulling off a Hopf band from the surface given by applying Seifert's algorithm to D.*

Chapter 2

The Alexander polynomial

2.1 The Alexander polynomial

2.1.1 Alexander's definition

In [1], Alexander defined a link invariant as follows. Take a link L with a reduced diagram D. Let $\mathcal{G}$ be the underlying graph of D with induced orientation O. Each region r_i of $\mathbb{S}^2 \setminus \mathcal{G}$ has a corner $\hat{c}_{ij}$ at each crossing c_j on its boundary. At each crossing, two of the four corners are dotted and each is assigned a value in $\{\pm t, \pm 1\}$ as shown in Figure 2.1.

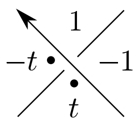

Figure 2.1

Let a_{ij} be the value so assigned to the corner $\hat{c}_{ij}$ if c_j lies on the boundary of r_i, and set $a_{ij} = 0$ otherwise. The matrix $\mathbf{A} = (a_{ij})$ is called the *Alexander matrix.* Since $\mathcal{G}$ is 4–valent, $\mathbf{A}$ is an $(n+2) \times n$ matrix, where n is the number of crossings in D. Choose adjacent regions r_p, r_q, and denote by $\mathbf{A}(p,q)$ the matrix given by deleting rows p and q from $\mathbf{A}$. Let $\Delta_L(t) = \det \mathbf{A}(p,q)$.

Alexander shows that, up to a factor of $\pm t^m$ for some $m \in \mathbb{Z}$, this definition of $\Delta_L(t)$ is independent of the choice of the pair of adjacent regions r_p, r_q and of the choice of the diagram D. We define $\Delta^0_L(t)$ to be $\Delta_L(t)$ normalised such that $\Delta^0_L(0)$ is defined and strictly positive, except when $\Delta_L(t) = 0$. That is, we arrange that the lowest degree term of $\Delta^0_L(0)$ is the constant, and that this is positive.

Definition 2.1.1. The *reduced Alexander polynomial* of L is $\Delta_L(t)$.

Proposition 2.1.2 (see [1] p301). *If $L = L_1 \# L_2$ for some links L_1, L_2 then $\Delta^0_L(t) = \Delta^0_{L_1}(t) \cdot \Delta^0_{L_2}(t)$.*

Lemma 2.1.3 (see [33] 8C4, 7A). *Let D' be the reflection of D, and let L' be the link with diagram D'. Then $\Delta^0_L(t) = \Delta^0_{L'}(t)$.*

Remark 2.1.4. Note that reflection changes positive crossings to negative ones, and vice versa.

Lemma 2.1.5 ([33] 8C7). *Suppose $\Delta^0_L(t) = \sum_{i=0}^{n} a_i t^i$. Then $|a_i| = |a_{n-i}|$ for $0 \leq i \leq n$.*

Proposition 2.1.6 (see [29] (3.7) and the proof of Lemma 3.6). *Suppose $L = L_1 * L_2$ for some links L_1, L_2. Then $\Delta^0_L(0) = \Delta^0_{L_1}(0) \cdot \Delta^0_{L_2}(0)$.*

Proposition 2.1.7 ([23] Proposition 6.14). *Suppose a link L bounds a disconnected oriented surface in $\mathbb{S}^3$. Then $\Delta_L(t) = 0$.*

Lemma 2.1.8 ([33] 10H9). *If a link L is fibred then $\Delta^0_L(0) = 1$.*

2.1.2 Trees

We will follow work of Crowell and Murasugi, which calculates $\Delta^0_L(0)$ for special alternating links by counting trees in graphs. We first give some useful definitions.

Definition 2.1.9. A subgraph $\mathcal{T}$ of a digraph $(\mathcal{G}, \mathcal{O})$ is a *directed tree* if it is connected, it contains no simple closed curves, and any $v \in \mathrm{V}(\mathcal{G})$ is the terminal vertex (with respect to $\mathcal{O}$) of at most one edge of $\mathcal{T}$. There is then one vertex $u \in \mathrm{V}(\mathcal{T})$ that is not the terminal vertex of any edge of $\mathcal{T}$. This vertex is called the *origin* of $\mathcal{T}$. For $v \in \mathrm{V}(\mathcal{T})$, the unique simple path from u to v in $\mathcal{T}$ is a directed path. Any $v \in \mathrm{V}(\mathcal{T})$ that is not the initial vertex of any edge of $\mathcal{T}$ is called a *leaf*.

A directed tree $\mathcal{T}$ is a *directed spanning subtree* of $\mathcal{G}$ if $\mathrm{V}(\mathcal{T}) = \mathrm{V}(\mathcal{G})$. Define $\mathrm{Tr}(\mathcal{G}, v) = \{\mathcal{T} : \mathcal{T}$ is a directed spanning subtree of $\mathcal{G}$ with origin $v\}$. Figure 2.2 shows an example of a digraph $\mathcal{G}$ with a directed spanning subtree $\mathcal{T}$. The origin of $\mathcal{T}$ is u, and v, v' are leaves.

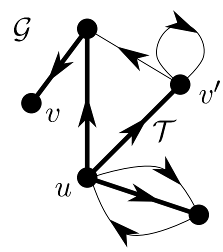

Figure 2.2

Lemma 2.1.10. *Let $v \in \mathrm{V}(\mathcal{G})$ and $e \in \mathrm{E}(\mathcal{G})$ such that $v = \iota(e)$. Then $|\operatorname{Tr}(\mathcal{G}, v)| = |\operatorname{Tr}(\mathcal{G} \setminus e, v)| + |\operatorname{Tr}(\mathcal{G}/e, v)|$.*

Lemma 2.1.11. *Suppose $\mathcal{G}$ is $\mathcal{O}$–connected. Then any directed tree $\mathcal{T}$ in $\mathcal{G}$ can be extended to a directed spanning subtree with the same origin.*

Definition 2.1.12. For a set $\mathcal{F} \subseteq \mathrm{V}(\mathcal{G})$, by a *directed $\mathcal{F}$–spanning subtree* of $\mathcal{G}$ we will mean a directed tree $\mathcal{T}$ such that $\mathcal{F} \subseteq \mathrm{V}(\mathcal{T})$ and every leaf of $\mathcal{T}$ is in $\mathcal{F}$, and such that no proper subtree of $\mathcal{T}$ has the same property.

Lemma 2.1.13. *Suppose $\mathcal{G}$ is $\mathcal{O}$–connected, and that, for some $v \in \mathrm{V}(\mathcal{G})$ and $\mathcal{F} \subseteq \mathrm{V}(\mathcal{G})$, there are n distinct directed $\mathcal{F}$–spanning subtrees of $\mathcal{G}$ with origin v. Then $|\operatorname{Tr}(\mathcal{G}, v)| \geq n$.*

Proof. Since $\mathcal{G}$ is $\mathcal{O}$–connected, Lemma 2.1.11 gives that any directed $\mathcal{F}$–spanning subtree can be extended to a directed spanning subtree with the same origin.

Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be extensions of distinct directed $\mathcal{F}$–spanning subtrees $\mathcal{T}_1^{\mathcal{F}}$ and $\mathcal{T}_2^{\mathcal{F}}$ respectively. Then there exists $w \in \mathcal{F}$ such that the directed path from v to w in $\mathcal{T}_1^{\mathcal{F}}$ is different to that in $\mathcal{T}_2^{\mathcal{F}}$. Then the directed path from v to w in $\mathcal{T}_1$ is different to that in $\mathcal{T}_2$. Thus $\mathcal{T}_1 \neq \mathcal{T}_2$. □

Lemma 2.1.14. *Suppose that $\mathcal{G}$ is planar and that for every region r of $\mathbb{S}^2 \setminus \mathcal{G}$ the boundary of r is a cycle. Then $\mathcal{G}$ is $\mathcal{O}$–connected. That is, if incoming and outgoing edges alternate at every vertex of $\mathcal{G}$ then $\mathcal{G}$ is $\mathcal{O}$–connected.*

Lemma 2.1.15. *Let $e \in \mathrm{E}(\mathcal{G})$ be a loop and let $v \in \mathrm{V}(\mathcal{G})$. Then $|\operatorname{Tr}(\mathcal{G} \setminus e, v)| = |\operatorname{Tr}(\mathcal{G}, v)|$.*

Lemma 2.1.16. *Let $e \in \mathrm{E}(\mathcal{G})$ be such that $\tau(e)$ has in-degree 1. Then $|\operatorname{Tr}(\mathcal{G}/e, v)| = |\operatorname{Tr}(\mathcal{G}, v)|$ for any $v \in \mathrm{V}(\mathcal{G}) \setminus \{\tau(e)\}$.*

Proof. Any directed spanning subtree in $\mathcal{G}$ contains e. □

Definition 2.1.17. Call removing a loop from a digraph $\mathcal{G}$ a *move 1* on $\mathcal{G}$, and collapsing an edge whose terminal vertex has no other incoming edge a *move 2* on $\mathcal{G}$.

2.1.3 Murasugi's proof of Theorem 1.1.1

Murasugi considers Alexander's definition from the following viewpoint in [28].

Once the regions r_p and r_q have been fixed, $\det \mathbf{A}(p, q)$ is formed of terms given by choosing (row, column) pairs $(i_1, j_1), \ldots, (i_n, j_n)$ in such a way that each row and each column is chosen exactly once and then multiplying together the $a_{i_k j_k}$. This is equivalent to choosing a bijection between the crossings c_j and regions r_i other than r_p, r_q, or choosing one corner $\hat{c}_{ij}$ at each crossing c_j provided $\hat{c}_{pj}, \hat{c}_{qj}$ are never chosen and no region is chosen more than once. Call such a bijection an L^s–*correspondence* if the resulting product is $\pm t^s$, or equivalently if s of the chosen corners are dotted.

Murasugi shows ([28] I Lemma 4.2; II Lemmas 6.8, 8.1) that any two L^s–correspondences give terms in the determinant with the same sign. Thus, to find $\Delta_L^0(0)$ for a link L, we need only count ways of choosing a corner for each crossing as above so that as few dotted corners as possible are chosen. Alternatively, by Lemma 2.1.5, we can look for choices where as many dotted corners as possible are chosen.

Now consider a special alternating link diagram D. For a black region r of D, either every corner of r is dotted or every corner of r is undotted. For a white region, corners are alternately dotted and undotted ([28] II Lemma 6.3). Let x be the number of black regions with dotted corners. Then the black regions contribute a constant factor of t^x to $|\prod_k a_{i_k,j_k}|$ and so this factor may be safely ignored for our purposes. This is in fact the power of t we cancel when normalising $\Delta_L(t)$ to $\Delta_L^0(t)$ (see [28] I Lemmas 3.1, 4.1, 5.4). Call an L^{s-x}–correspondence an L_0^s–*correspondence* (this is actually Murasugi's definition of an L^s–correspondence).

Indeed, Murasugi shows that we can forget the black regions altogether by defining a digraph $G^{\mathcal{M}}(D)$ from the diagram D as follows.

$G^{\mathcal{M}}(D)$ has a vertex at the centre of each white region of D, and one edge for each crossing, joining the centres of the white regions meeting at the crossing. $G^{\mathcal{M}}(D)$ is therefore planar. At each crossing, one white corner is dotted, and the other is undotted. Orient each edge from the undotted side to the dotted side (note that this is the reverse of in [28]).

As dotted and undotted corners of white regions alternate, the boundary of any region r of $\mathbb{S}^2 \setminus G^{\mathcal{M}}(D)$ is a cycle with respect to the above orientation o. Thus $G^{\mathcal{M}}(D)$ is o–connected, and in particular $|\operatorname{Tr}(G^{\mathcal{M}}(D), v)| \geq 1$ for $v \in \mathrm{V}(G^{\mathcal{M}}(D))$.

Define a second (unoriented) graph $G_b(D)$ with a vertex at the centre of each black region of D and an edge through each crossing. Then $G_b(D)$ is the dual graph $G^{\mathcal{M}}(D)^*$ of $G^{\mathcal{M}}(D)$.

For a directed spanning subtree $\mathcal{T}$ of $G^{\mathcal{M}}(D)$ let $\mathcal{T}^*$ be the subgraph of $G_b(D)$ consisting of all edges in $\mathrm{E}(G_b(D))$ that do not cross any edge in $\mathcal{T}$. Then $\mathcal{T}^*$ is a tree ([28] I Lemma 5.2). By counting edges and vertices we see that $\mathcal{T}^*$ spans $G_b(D)$.

Let $v_p \in \mathrm{V}(G^{\mathcal{M}}(D))$ and $v_q \in \mathrm{V}(G_b(D))$ be the vertices corresponding to r_p and r_q respectively. Let $\mathcal{T} \in \operatorname{Tr}(G^{\mathcal{M}}(D), v_p)$. There is then a unique way of orienting the edges of $\mathcal{T}^*$ so it becomes a directed tree with origin v_q. Let y be the number of white regions in D. Each region $r \in D$ other than r_p, r_q is the terminal vertex of exactly one edge e_r of $\mathcal{T}$ or $\mathcal{T}^*$. By pairing r with the crossing that e_r corresponds to, we can construct an L_0^y–correspondence. Clearly there is no L_0^s–correspondence for any $s > y$.

Conversely, given an L_0^y–correspondence this process can be reversed to give a directed subgraph $\mathcal{T}$ of $G^{\mathcal{M}}(D)$ with exactly one edge ending at each vertex other than v_p. Suppose $\mathcal{T}$ contains an embedded closed curve, dividing $\mathbb{S}^2$ into two discs. Considering the Euler characteristic χ of the disc that does not contain r_p, r_q gives a contradiction. Hence $\mathcal{T} \in \mathrm{Tr}(G^{\mathcal{M}}(D), v_p)$.

For general s, an L_0^s–correspondence can be used to construct a directed spanning subtree $\mathcal{T}$ of the underlying graph of $G^{\mathcal{M}}(D)$ with an orientation that will not in general agree with that of $G^{\mathcal{M}}(D)$. The number of edges of $\mathcal{T}$ where these two orientations (dis)agree is determined by s. Murasugi and Stoimenow use this in [27] to assign a polynomial to any connected digraph in which the in-degree equals the out-degree at each vertex.

Given a planar digraph $(\mathcal{G}, \mathcal{O})$ in which incoming and outgoing edges alternate at each vertex, we can construct a product sutured manifold $N^{\mathcal{M}}(\mathcal{G})$ embedded in $\mathbb{S}^3$ in such a way that there is an 'obvious' projection of the sutures onto $\mathbb{S}^2$ that gives a link diagram. Roughly speaking, this gives an inverse to $G^{\mathcal{M}}$. We shall examine this in more detail later.

Construction 2.1.18. We first build the 3–manifold $N = N^{\mathcal{M}}(\mathcal{G})$. Centre a 0–handle $\mathbb{D}^3$ on each vertex of $\mathcal{G}^* \subset \mathbb{S}^2 \subset \mathbb{S}^3$. These should be taken to be sufficiently small that they do not intersect. Attach a 1–handle $[0,1] \times \mathbb{D}^2$ for each edge $e \in \mathrm{E}(\mathcal{G}^*)$ with $[0,1] \times \{0\}$ running along e and $\{0,1\} \times \mathbb{D}^2$ glued to the 0–handles corresponding to the endpoints of e.

We now define the sutures s. For a 0–handle V_0, let W_1 be the union of the 1–handles that meet V_0. Then let $V_0 \cap s = (\partial V_0 \setminus \mathrm{int}_N(W_1)) \cap \mathbb{S}^2$. Now let V_1 be a 1–handle. Then $V_1 \cap s$ is made up of two disjoint simple arcs, one running from $\{0\} \times \{1\} \in ([0,1] \times \mathbb{D}^2) \cap \mathbb{S}^2$ to $\{1\} \times \{-1\}$ and the other running from $\{0\} \times \{-1\}$ to $\{1\} \times \{1\}$. The arcs twist around V_1 in the direction shown in Figure 2.3, where the dashed line denotes an arc passing underneath the manifold.

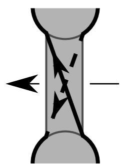

Figure 2.3

Using the orientation $\mathcal{O}$, we can define an orientation on the arcs of s that run along 1–handles, as shown in Figure 2.3. Since incoming and outgoing edges alternate at every vertex of $\mathcal{G}$, this definition of the orientation of s is locally consistent, as shown in Figure 2.4. Therefore the sutures around a 0–handle are either all oriented clockwise or all oriented

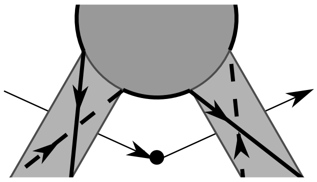

Figure 2.4

anticlockwise, with the orientations on adjacent vertices going in opposite directions. From this we see that (N, s) is a sutured manifold.

Definition 2.1.19. Define $M^{\mathcal{M}}(\mathcal{G}) = \mathbb{S}^3 \setminus \text{int}_{\mathbb{S}^3}(N^{\mathcal{M}}(\mathcal{G}))$ to be the complementary sutured manifold to $N^{\mathcal{M}}(\mathcal{G})$.

Lemma 2.1.20. *Let $\mathcal{G}$ be a planar digraph in which incoming and outgoing edges alternate at each vertex. Let $e \in \mathrm{E}(\mathcal{G})$ be a loop that bounds a disc in $\mathbb{S}^2 \setminus \mathcal{G}$. Then $N^{\mathcal{M}}(\mathcal{G})$ and $N^{\mathcal{M}}(\mathcal{G} \setminus e)$ (and hence also $M^{\mathcal{M}}(\mathcal{G})$ and $M^{\mathcal{M}}(\mathcal{G} \setminus e)$) are equivalent as sutured manifolds embedded in $\mathbb{S}^3$.*

Proof. This can be checked locally, as shown in Figure 2.5. □

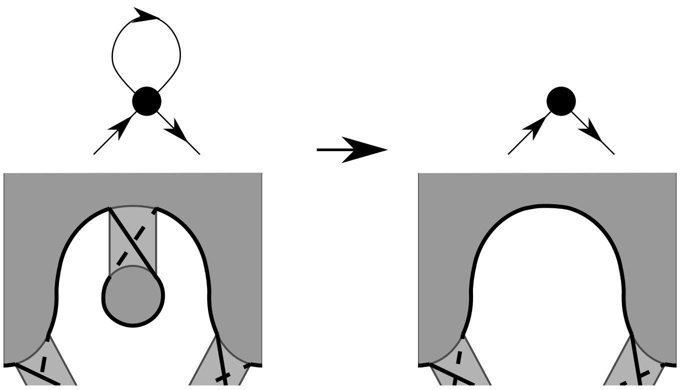

Figure 2.5

Lemma 2.1.21. *Let $\mathcal{G}$ be a planar digraph in which incoming and outgoing edges alternate at each vertex, and let $e \in \mathrm{E}(\mathcal{G})$ be such that $\tau(e)$ is 2–valent. Then $M^{\mathcal{M}}(\mathcal{G}/e)$ is obtained from $M^{\mathcal{M}}(\mathcal{G})$ by a product disc decomposition. In particular, $M^{\mathcal{M}}(\mathcal{G})$ is an almost product sutured manifold if and only if $M^{\mathcal{M}}(\mathcal{G}/e)$ is.*

Proof. See Figure 2.6. □

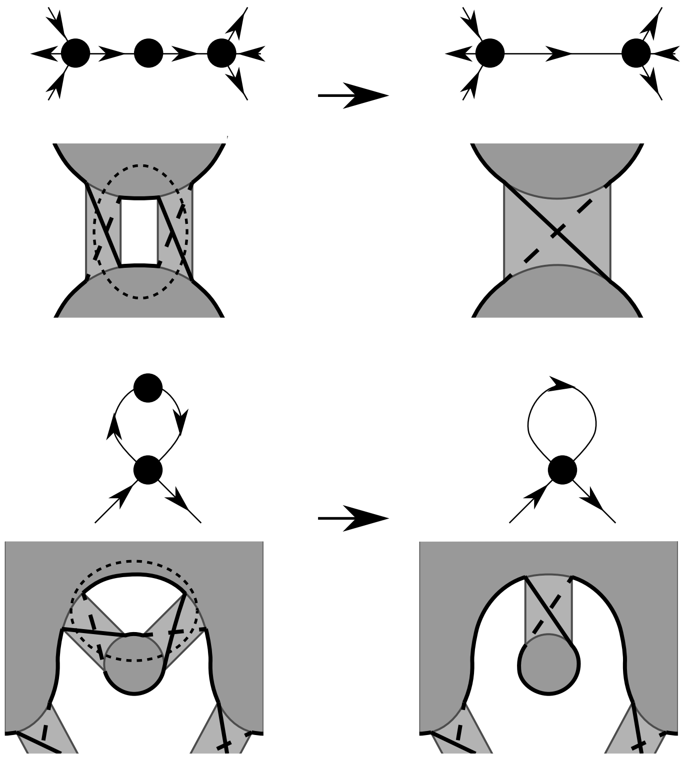

Figure 2.6

2.1.4 Crowell's proof of Theorem 1.1.1

Crowell also made use of directed trees in digraphs to calculate $\Delta_L^0(0)$ for an alternating link L. As we will see, the two methods are similar, but not identical. In the proof of Corollary 2.2.13 we will make use of the difference, as well as the similarity, between the two sets of work.

Let D be an alternating diagram for a link L, and let $(\mathcal{G}, O)$ be the underlying digraph.

Definition 2.1.22. Any edge $e \in \mathrm{E}(\mathcal{G})$ corresponds to an arc ρ_e in D between two crossings. Since D is alternating, one end of ρ_e is part of an undercrossing, and the other end is part of an overcrossing. Let o be the orientation of $\mathcal{G}$ that orients each edge e from the overcrossing to the undercrossing. Near any $v \in \mathrm{V}(\mathcal{G})$, o is as shown in Figure 2.7.

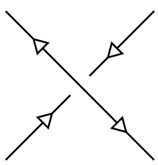

Figure 2.7

Definition 2.1.23. Define subsets $\mathcal{H}, \mathcal{K} \subseteq \mathrm{E}(\mathcal{G})$ by at every vertex putting the incoming edges with respect to o into $\mathcal{H}, \mathcal{K}$ as shown in Figure 2.8.

Define a map $\alpha\colon\ \mathrm{E}(\mathcal{G}) \to \{1, -t\}$ by

$$\alpha(e) = \begin{cases} -t & \text{if } e \in \mathcal{H} \\ 1 & \text{if } e \in \mathcal{K}. \end{cases}$$

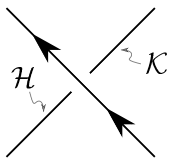

Figure 2.8

Theorem 2.1.24 ([8] Theorem 2.12)**.** *For any $v \in \mathrm{V}(\mathcal{G})$,*

$$\Delta_L(t) = \sum_{\mathcal{T} \in \mathrm{Tr}(\mathcal{G}, v)} \left(\prod_{e \in \mathcal{T}} \alpha(e) \right).$$

Lemma 2.1.25 ([8] 4.7)**.** *A Seifert circle is a cycle with respect to O and o (possibly in opposite directions).*

The special Seifert circles of D are exactly the cycles in $\mathcal{G}$ with respect to o that are contained entirely in $\mathcal{H}$ or entirely in $\mathcal{K}$.

Remark 2.1.26 ([8] 4.2)**.** Neither $\mathcal{H}$ nor $\mathcal{K}$ contains a pair of distinct edges with a common terminal vertex with respect to o.

Definition 2.1.27. For a directed spanning subtree $\mathcal{T}$ of $\mathcal{G}$ with respect to o, define $\mathrm{char}(\mathcal{T})$ to be the number of edges of $\mathcal{T}$ that lie in $\mathcal{K}$. Define an *$\mathcal{H}$–maximal directed spanning subtree* $\mathcal{T}$ to be a directed spanning subtree with $\mathrm{char}(\mathcal{T})$ minimal among such trees with the same origin as $\mathcal{T}$ (that is, $\mathcal{T}$ contains as many edges of $\mathcal{H}$ as possible). For $v \in \mathrm{V}(\mathcal{G})$, let $\mathrm{Tr}^{\mathcal{H}}(\mathcal{G}, v)$ be the set of $\mathcal{H}$–maximal directed spanning subtrees in $\mathcal{G}$ with origin v.

Define a *$\mathcal{K}$–maximal directed spanning subtree* and $\mathrm{Tr}^{\mathcal{K}}(\mathcal{G}, v)$ analogously.

Remark 2.1.28. By Theorem 2.1.24, $\Delta_L^0(0) = |\,\mathrm{Tr}^{\mathcal{K}}(\mathcal{G}, v)| = |\,\mathrm{Tr}^{\mathcal{H}}(\mathcal{G}, v)|$ for any $v \in \mathrm{V}(\mathcal{G})$.

Now suppose D is special. Then every Seifert circle is contained in $\mathcal{H}$ or is contained in $\mathcal{K}$. Further, since no two edges in $\mathcal{H}$ share a terminal vertex, no two Seifert circles in $\mathcal{H}$ share a vertex. We can therefore collapse each such Seifert circle to a point, giving a planar graph $G^{\mathcal{H}}(D)$. Define orientation o on $G^{\mathcal{H}}(D)$ to be that inherited from o on $\mathcal{G}$. Since at each vertex of $\mathcal{G}$ exactly two edges are collapsed, and these are adjacent, incoming and outgoing edges alternate at each vertex of $G^{\mathcal{H}}(D)$. Figure 2.9 shows this process when L is the knot 7_4.

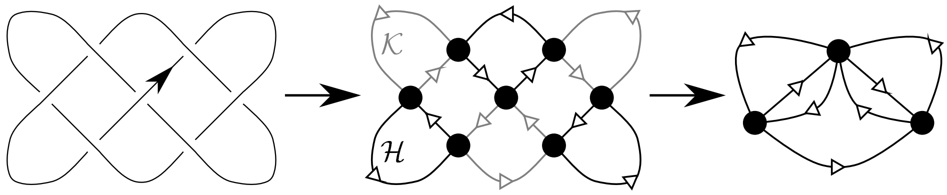

Figure 2.9

Lemma 2.1.29. *Let* $A\colon \mathcal{G} \to G^{\mathcal{H}}(D)$ *be the map that collapses the Seifert circles in* $\mathcal{H}$*. Then, for any* $v \in \mathrm{V}(\mathcal{G})$*,* A *induces a bijection* $A_v\colon \mathrm{Tr}^{\mathcal{H}}(\mathcal{G}, v) \to \mathrm{Tr}(G^{\mathcal{H}}(D), A(v))$*.*

Proof. Define a map $B_v\colon \mathrm{Tr}(G^{\mathcal{H}}(D), A(v)) \to \mathrm{Tr}(\mathcal{G}, v)$ as follows.

Let $\mathcal{T} \in \mathrm{Tr}(G^{\mathcal{H}}(D), A(v))$. Then, for $e \in \mathcal{K} \subseteq \mathrm{E}(\mathcal{G})$, let $e \in B_v(\mathcal{T})$ if and only if $A(e) \in \mathcal{T}$.

Consider a Seifert circle C in $\mathcal{G}$ contained in $\mathcal{H}$. If v lies on C, then no edge of $B_v(\mathcal{T}) \cap \mathcal{K}$ has its terminal vertex on C. Let f_C be the edge of C whose terminal vertex is v. If v instead does not lie on C, then $B_v(\mathcal{T}) \cap \mathcal{K}$ contains exactly one edge e_C whose terminal vertex lies on C. In this case, let f_C be the edge of C that has the same terminal vertex as e_C. In either case, let $B_v(\mathcal{T}) \cap C = C \setminus f_C$ (see Figure 2.10).

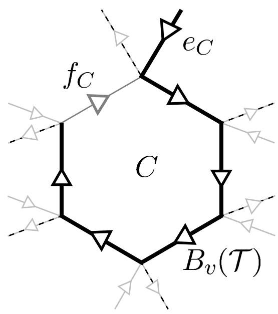

Figure 2.10

Then $A(B_v(\mathcal{T})) = \mathcal{T}$. Thus, since $\mathcal{T}$ contains no simple closed curves, $B_v(\mathcal{T})$ contains no simple closed curves. Let $w \in \mathrm{V}(\mathcal{G})$. Since $\mathcal{T}$ contains a directed path from $A(v)$ to $A(w)$, it is clear that $B_v(\mathcal{T})$ contains a directed path from v to w. Hence $B_v(\mathcal{T}) \in \mathrm{Tr}(\mathcal{G}, v)$.

We can now see that $B_v(\mathcal{T}) \in \mathrm{Tr}^{\mathcal{H}}(\mathcal{G}, v)$. Thus if $\mathcal{T}' \in \mathrm{Tr}^{\mathcal{H}}(\mathcal{G}, v)$ and C is a Seifert circle in $\mathcal{H}$ then $\mathcal{T}'$ contains all but one edge of C, as in Figure 2.10. Therefore, at most one edge of $\mathcal{T}' \cap \mathcal{K}$ has its terminal vertex on C, with no such edge if v lies on C. This means that $A(\mathcal{T}')$ does not contain any simple closed curves, and no edge of $A(\mathcal{T}')$ has terminal vertex $A(v)$. It is now clear that we can define A_v by $A_v(\mathcal{T}) = A(\mathcal{T})$. Knowing this, we see that A_v and B_v are mutual inverses. $\square$

As before, we can construct a product sutured manifold $N^{\mathcal{H}}(\mathcal{G})$ from a digraph $(\mathcal{G}, \mathcal{O})$ that provides an 'inverse' to the map $G^{\mathcal{H}}$, and we define $M^{\mathcal{H}}(\mathcal{G})$ to be the complementary sutured manifold to $N^{\mathcal{H}}(\mathcal{G})$.

Definition 2.1.30. Let $\mathcal{G}'$ be a planar digraph in which incoming and outgoing edges alternate at each vertex. The boundary of any region r of $\mathbb{S}^2 \setminus \mathcal{G}'$ is a cycle. Define a $\mathcal{K}$*–circle* of $\mathcal{G}'$ to be any such cycle that is oriented clockwise around r.

Construction 2.1.31. $N = N^{\mathcal{H}}(\mathcal{G})$ has a 0–handle at each vertex of $\mathcal{G}$, and a 1–handle running along each edge of $\mathcal{G}$. Attach a 2–handle $\mathbb{D}^2 \times [0,1]$ for each $\mathcal{K}$–circle of $\mathcal{G}$. If r is a region of $\mathbb{S}^2 \setminus \mathcal{G}$ whose boundary ∂r is a $\mathcal{K}$–circle, the boundary of the union of the 0–handles and the 1–handles of N meets r in a simple closed curve. Attach the 2–handle along this curve.

For a 1–handle V_1, let $V_1 \cap s = V_1 \cap \partial N \cap \mathbb{S}^2$. This subarc of s is oriented in the same direction as the edge of $\mathcal{G}$ that V_1 runs along. Let V_0 be a 0–handle. Then $\partial V_0 \cap \mathbb{S}^2$ consists of adjacent simple arcs $\rho_1^1, \rho_1^2, \ldots, \rho_n^1, \rho_n^2$ for some $n \in \mathbb{N}$, ordered clockwise around V_0, where ρ_i^1 is properly embedded in N and $\rho_i^2 \subset \partial N$ for each i. For $1 \leq m \leq n$, join the midpoint $\rho_m^2(\frac{1}{2})$ of ρ_m^2 to the far endpoint of ρ_m^1 by a simple arc running over V_0, and to the far endpoint of ρ_{m+1}^1 (where $\rho_{n+1}^1 = \rho_1^1$) by a simple arc running under V_0, as shown in Figure 2.11.

It is clear that N is now a sutured manifold.

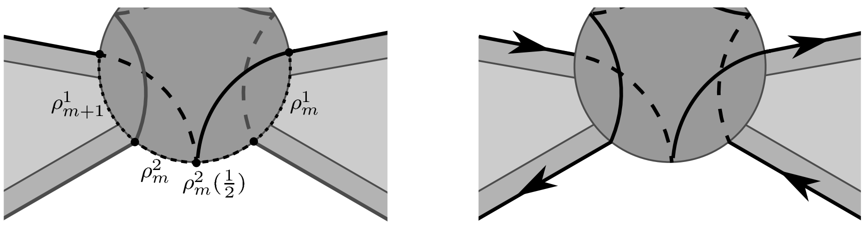

Figure 2.11

Remark 2.1.32. We may similarly define $G^{\mathcal{K}}(D)$, $N^{\mathcal{K}}(\mathcal{G})$ and $M^{\mathcal{K}}(\mathcal{G})$.

2.2 Digraphs constructed from link diagrams

2.2.1 Bounding valency

In both Theorem 1.1.4 and Theorem 1.1.5, we wish to use a bound on $\Delta_L^0(0)$ to control the possibilities for the homogeneous link L. In Section 2.1 we have established the following.

Lemma 2.2.1. *Let L be a link with a reduced, special alternating diagram D. Then $\Delta_L^0(0) = |\operatorname{Tr}(G^{\mathcal{M}}(D), v)| = |\operatorname{Tr}(G^{\mathcal{H}}(D), v)| = |\operatorname{Tr}(G^{\mathcal{K}}(D), v)|$, where in each case v may be any vertex of the relevant digraph.*

Thus we now turn our attention to controlling a digraph using a bound on the number of directed spanning subtrees it contains.

Lemma 2.2.2. *Let $\mathcal{G}$ be a digraph with no loops, and fix $v_0 \in \mathrm{V}(\mathcal{G})$. Suppose there is a directed spanning subtree $\mathcal{T}$ of $\mathcal{G}$ with origin v_0. Let w be any leaf of $\mathcal{T}$ and let n be the in-degree of w in $\mathcal{G}$. Then $\mathcal{G}$ has at least n directed spanning subtrees with origin v_0.*

Proof. Let e be the edge of $\mathcal{T}$ with terminal vertex w. By repeated use of Lemma 2.1.10, $|\operatorname{Tr}(\mathcal{G}, v_0)| \geq |\operatorname{Tr}(\mathcal{G}/(\mathcal{T} \setminus e), v_0)|$. But $\mathcal{G}/(\mathcal{T} \setminus e)$ has two vertices and w still has in-degree n. Since $\mathcal{G}$ contained no loops, $\mathcal{G}/(\mathcal{T} \setminus e)$ has no loops with endpoints at w. Therefore $|\operatorname{Tr}(\mathcal{G}/(\mathcal{T} \setminus e), v_0)| = n$. □

Definition 2.2.3. Define a planar digraph $\mathcal{G}$ to be *prime* if none of the following hold.

- $\mathcal{G}$ contains a loop.
- $\mathcal{G}$ has a cut vertex.
- There is a simple closed curve ρ in $\mathbb{S}^2$ disjoint from the vertices of $\mathcal{G}$ meeting the edges of $\mathcal{G}$ at exactly two points and with at least one vertex of $\mathcal{G}$ on each side of ρ.

Proposition 2.2.4. *Let $\mathcal{G}$ be a prime, planar digraph such that, at every vertex, incoming and outgoing edges alternate. Let $w \in \mathrm{V}(\mathcal{G})$. Then there is a vertex $v_w \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$ and a directed spanning subtree $\mathcal{T}$ of $\mathcal{G}$ with origin v_w such that w is a leaf of $\mathcal{T}$.*

Proof. Assume otherwise. For $v \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$, let $A(v)$ be the set of vertices $v' \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$ such that there is a directed path in $\mathcal{G}$ from v to v' that does not pass through w, and let $B(v) = \mathrm{V}(\mathcal{G}) \setminus (\{w\} \cup A(v))$.

Suppose that $B(v_0) = \emptyset$ for some $v_0 \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$. Then for every $v \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$ there is a directed path $\rho(v)$ in $\mathcal{G}$ from v_0 to v that does not pass through w. Take the union of these paths, and discard edges as necessary to give a tree $\mathcal{T}'$ that includes all vertices in $\mathrm{V}(\mathcal{G}) \setminus \{w\}$. Pick any edge e with terminal vertex w. Then $\mathcal{T} = \mathcal{T}' \cup e$ is a directed spanning subtree of $\mathcal{G}$ with origin v_0 of which w is a leaf. This contradicts the assumption that no such $\mathcal{T}$ exists. Thus, for every $v \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$, the set $B(v)$ is non-empty, as is $A(v)$.

Choose $v_0 \neq w$. By the definition of $B(v_0)$, no edge of $\mathcal{G}$ runs from a vertex in $A(v_0)$ to a vertex in $B(v_0)$. Note that $\mathcal{G}[A(v_0)]$ is connected. On the other hand, $\mathcal{G}[B(v_0)]$ may be disconnected, but every vertex of $B(v_0)$ lies on a path that begins at w and does not meet $A(v_0)$. Thus $\mathcal{G}$ has the form shown in Figure 2.12, where each arrow in the picture may denote multiple edges. Since the boundary of each region of $\mathbb{S}^2 \setminus \mathcal{G}$ is a cycle, e denotes

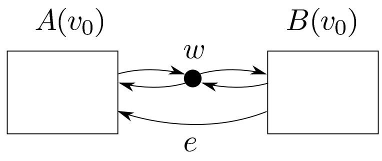

Figure 2.12

at most one edge of $\mathcal{G}$. As $\mathcal{G}$ is prime, e denotes at least one edge. Call this edge $e(v_0)$. Similarly define $e(v)$ for each $v \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$.

For $v, v' \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$, if $v' \in A(v)$ then $A(v') \subseteq A(v)$. This inclusion of sets gives a partial order on $\mathrm{V}(\mathcal{G}) \setminus \{w\}$. Choose v_+ to be maximal with respect to this ordering. Now let v_b be the initial vertex of $e(v_+)$, and v_a the terminal vertex. Then $v_+, v_a \in A(v_+)$ and $v_b \in B(v_+)$. It follows that $A(v_a) \subseteq A(v_+)$, so $v_b \in B(v_a)$. In addition, $v_a \in A(v_a)$, and $v_a \in A(v_b)$, so $A(v_a) \subseteq A(v_b)$. If $v_+ \in A(v_a)$ then $A(v_+) \subseteq A(v_b)$. Since $v_b \in A(v_b) \setminus A(v_+)$, this contradicts maximality of v_+. Thus $v_+ \in B(v_a)$. In summary, $A(v_a) \cap B(v_+) = \emptyset$ while $A(v_a) \cap A(v_+)$, $B(v_a) \cap A(v_+)$ and $B(v_a) \cap B(v_+)$ are non-empty. This gives $\mathcal{G}$ the structure shown in Figure 2.13, where each box is non-empty.

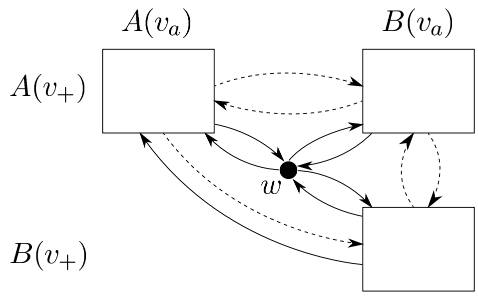

Figure 2.13

Recall that, for any $v \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$, no edge starts in $A(v)$ and ends in $B(v)$, while at most one edge starts in $B(v)$ and ends in $A(v)$. Since $e(v_+)$ starts in $B(v_a) \cap B(v_+)$ and ends in $A(v_a) \cap A(v_+)$, we see that all the dashed arrows in Figure 2.13 are excluded. This contradicts that $\mathcal{G}$ is prime. □

Corollary 2.2.5. *Let $\mathcal{G}$ and w be as in Proposition 2.2.4. Let n be the in-degree of w in $\mathcal{G}$. Then there exists $v \in \mathrm{V}(\mathcal{G}) \setminus \{w\}$ such that $|\operatorname{Tr}(\mathcal{G}, v)| \geq n$. In particular, if $\mathcal{G} \in \{G^{\mathcal{M}}(D), G^{\mathcal{H}}(D), G^{\mathcal{K}}(D)\}$ for some diagram D of a link L then $\Delta^0_L(0) \geq n$.*

2.2.2 Digraph properties

Definition 2.2.6. A diagram D of a link L is *twist-reduced* if, given a simple closed curve ρ in $\mathbb{S}^2$ that passes through two crossings of D transversely and otherwise lies in $\mathbb{S}^2 \setminus D$, the section of L on one side of ρ consists of a line of bigons, like that shown in Figure 2.14.

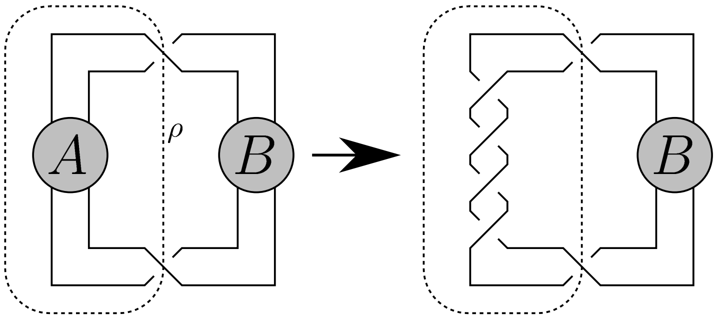

Figure 2.14

If D is special, we say D is *black-twist-reduced* if this holds whenever ρ is contained in the black regions (see [22] p215).

Remark 2.2.7. Any link diagram can be made twist-reduced by a finite sequence of flypes. Given a special alternating, twist-reduced diagram D, the diagram D' formed by deleting white bigons from D (equivalently, removing loops that bound discs from $G^{\mathcal{M}}(D)$) may not be twist-reduced, but will be black-twist-reduced (see [22] p215). These changes preserve the property of being special and alternating.

Definition 2.2.8. Let Λ be the set of link diagrams D with the following properties.

(L1) D is connected.

(L2) D is special alternating and positive.

(L3) D has no nugatory crossings.

(L4) D is prime.

(L5) D is black-twist-reduced.

(L6) D has no white bigons.

Definition 2.2.9. Let Γ be the set of digraphs $\mathcal{G}$ with the following properties.

(G1) $\mathcal{G}$ is connected.

(G2) $\mathcal{G}$ is planar.

(G3) $\mathcal{G}$ is prime (see Definition 2.2.3).

(G4) At every vertex of $\mathcal{G}$, incoming and outgoing edges alternate.

Remark 2.2.10. (G3) and (G4) together imply the following.

(G5) No vertex of $\mathcal{G}$ has in-degree 1.

Lemma 2.2.11. *$G^{\mathcal{M}}$ is a bijection from Λ to Γ.*

Proof. Let $D \in \Lambda$. Recall that $G^{\mathcal{M}}(D)$ consists of a vertex in every white region, and an edge for every crossing, pointing from the undotted side to the dotted side. By (L1) and (L2), $G^{\mathcal{M}}(D)$ is a well-defined digraph with properties (G1), (G2) and (G4). $G^{\mathcal{M}}(D)$ has no loops by (L3), and no cut vertices by (L1) and (L4). By (L5) and (L6), no simple closed curve meets $\mathrm{E}(G^{\mathcal{M}}(D))$ twice and separates $\mathrm{V}(G^{\mathcal{M}}(D))$. Thus $G^{\mathcal{M}}(D)$ is prime. Hence $G^{\mathcal{M}}(D) \in \Gamma$.

Now let $\mathcal{G} \in \Gamma$. By (G2) and (G4), the sutured manifold $N^{\mathcal{M}}(\mathcal{G})$ from Construction 2.1.18 is defined. Let $D_{\mathcal{G}}$ be the diagram constructed from $N^{\mathcal{M}}(\mathcal{G})$ by projecting the sutures onto $\mathbb{S}^2$. We choose the black regions of $\mathbb{S}^2 \setminus D_{\mathcal{G}}$ to be those that correspond to 0–handles in $N^{\mathcal{M}}(\mathcal{G})$. It is clear from the construction of $D_{\mathcal{G}}$ that (L1) and (L2) hold. $D_{\mathcal{G}}$ has no nugatory crossings because there are no loops in $\mathcal{G}$, and no obvious decomposition as a connected sum because $\mathcal{G}$ has no cut vertex. Suppose either (L5) or (L6) does not hold. This means there is a simple closed curve $\rho \subset \mathbb{S}^2$ that meets $D_{\mathcal{G}}$ at exactly two of the crossings of $D_{\mathcal{G}}$ and is otherwise contained in the black regions. Then ρ meets the edges of $\mathcal{G}$ twice and separates $\mathrm{V}(\mathcal{G})$. This contradicts (G3). Hence $D_{\mathcal{G}} \in \Lambda$.

It is now easy to check that, when restricted to Λ and Γ respectively, these two constructions are mutual inverses. $\square$

Lemma 2.2.12. *$G^{\mathcal{H}}$ is a bijection from Λ to Γ.*

Proof. Let $D \in \Lambda$. Recall that $G^{\mathcal{H}}(D)$ is obtained by taking the underlying graph $\mathcal{G}$ of D, with orientations O and o, and collapsing all edges of $\mathcal{H}$. Properties (L1) and (L2) ensure $G^{\mathcal{H}}(D)$ is well defined with properties (G1), (G2) and (G4). Suppose $G^{\mathcal{H}}(D)$ contains a loop e, and consider the copy of e in $\mathcal{G}$. By (L3) there are no loops in $\mathcal{G}$, so e must have both its endpoints on a single Seifert circle contained in $\mathcal{H}$. Since D is special, we can use e to construct a simple closed curve ρ that meets D twice at crossings (the endpoints of e) and otherwise is contained in the black regions of D. This is impossible since (L5) and (L6) hold. Thus $G^{\mathcal{H}}(D)$ contains no loops. If $G^{\mathcal{H}}(D)$ is not prime, this means there is a simple closed curve ρ' crossing the edges of $\mathcal{G}$ twice and dividing the vertices, contradicting (L4). Therefore, $G^{\mathcal{H}}(D) \in \Gamma$.

Conversely, let $\mathcal{G} \in \Gamma$. By (G2) and (G4), we can construct $N^{\mathcal{H}}(\mathcal{G})$. Let $D_{\mathcal{G}}$ be the diagram given by $N^{\mathcal{H}}(\mathcal{G})$, where the black regions are those that correspond to a 0–handle or a 2–handle. $D_{\mathcal{G}}$ is connected, by (G1). By construction, (L2) holds and no black region meets itself. (G3) means no white region meets itself. Thus (L3) holds. Any decomposition of $D_{\mathcal{G}}$ as a connected sum must come either from a similar decomposition of $\mathcal{G}$ or from a cut vertex in $\mathcal{G}$. Neither is possible since $\mathcal{G}$ is prime, so (L4) holds. Finally, suppose there is a simple closed curve ρ that meets $D_{\mathcal{G}}$ at two crossings and otherwise lies in the black regions. This gives a simple closed curve ρ' that meets $\mathcal{G}$ exactly once at a vertex, again contradicting that $\mathcal{G}$ is prime. Thus (L5) and (L6) hold. Hence $D_{\mathcal{G}} \in \Lambda$.

When restricted to Λ and Γ respectively, these two constructions are mutual inverses. □

Corollary 2.2.13. *Let $\mathcal{G} \in \Gamma$. Let $v \in \mathrm{V}(\mathcal{G})$ and let $n = |\operatorname{Tr}(\mathcal{G}, v)|$. Then, for any region r of $\mathbb{S}^2 \setminus \mathcal{G}$, the number of edges of $\mathcal{G}$ in the boundary of r is at most n.*

Proof. Suppose otherwise. Let r be a region of $\mathbb{S}^2 \setminus \mathcal{G}$ with m sides, for some $m > n$. Let $D^{\mathcal{M}} = (G^{\mathcal{M}})^{-1}(\mathcal{G})$. There is a Seifert circle C in $D^{\mathcal{M}}$ corresponding to r, which also has m sides. We may suppose that $C \subseteq \mathcal{H}$ (otherwise, replace $\mathcal{H}$ with $\mathcal{K}$ in the following argument). Let v_C be the vertex of $G^{\mathcal{H}}(D^{\mathcal{M}})$ corresponding to C. Then v_C has in-degree m. By Corollary 2.2.5, $\Delta^0_L(0) \geq m$, where L is the link with diagram $D^{\mathcal{M}}$. This contradicts that $\Delta^0_L(0) = |\operatorname{Tr}(G^{\mathcal{M}}(D^{\mathcal{M}}), v)| = |\operatorname{Tr}(\mathcal{G}, v)| = n < m$. □

2.2.3 Infinite digraphs

Lemma 2.2.14. *Let Φ be an infinite set of planar, pointed digraphs (each with a fixed embedding into $\mathbb{S}^2$) such that every $(\mathcal{G}, v) \in \Phi$ has valence bounded above by $n_1 \in \mathbb{N}$. Then there is a sequence of distinct elements $(\mathcal{G}_i, v_i)_{i=1}^{\infty}$ such that $\mathcal{B}(\mathcal{G}_{m_1}, v_{m_1}, m)$ and $\mathcal{B}(\mathcal{G}_{m_2}, v_{m_2}, m)$ are the same up to ambient isotopy of $\mathbb{S}^2$ whenever $m_1, m_2 \geq m$.*

Proof. See Convention 1.5.8 for the definition of $\mathcal{B}(\mathcal{G}, v, m)$.

Let Θ be the class of all embeddings, up to ambient isotopy of $\mathbb{S}^2$, of planar, pointed digraphs with valence bounded above by n_1. For $m \in \mathbb{N}$, let Θ_m be those elements of Θ with radius at most m. Then $|\Theta_m| < \infty$ for all m, and $\Theta_1 \subset \Theta_2 \subset \cdots \subset \Theta_m \subset \Theta_{m+1} \subset \cdots \subset \Theta$. Choose an enumeration $\theta \colon \Theta \to \mathbb{N}$ such that $\theta(\Theta_m) = \{1, \ldots, |\Theta_m|\}$ for all m.

Let $\Psi_m = \theta(\Theta_m) = \{1, \ldots, |\Theta_m|\}$ for all m, given the discrete topology. Take $\Psi = \prod_{i=1}^{\infty} \Psi_i$, with the Tychonoff product topology. Then Ψ is compact. Define $\phi \colon \Phi \to \Psi$ by $\phi(\mathcal{G}, v) = \big(\theta\big(\mathcal{B}(\mathcal{G}, v, i)\big)\big)_{i=1}^{\infty}$. If $\phi(\mathcal{G}_1, v_1) = \phi(\mathcal{G}_2, v_2)$ then $\mathcal{G}_1, \mathcal{G}_2 \in \Theta_m$ for some $m \in \mathbb{N}$, and $\mathcal{G}_1 = \mathcal{B}(\mathcal{G}_1, v_1, m) = \mathcal{B}(\mathcal{G}_2, v_2, m) = \mathcal{G}_2$. Thus ϕ is injective. In particular, $\phi(\Phi)$ is infinite.

Thus there is a sequence $\big((J_i^j)_{i=1}^{\infty}\big)_{j=1}^{\infty} \subseteq \phi(\Phi)$ of distinct elements that converges to an element $(J_i^{\infty})_{i=1}^{\infty} \in \Psi$. That is, $J_m^j \to J_m^{\infty}$ as $j \to \infty$ for $m \in \mathbb{N}$. As Ψ_m is discrete, there exists $j_m \in \mathbb{N}$ such that $J_m^k = J_m^\infty \in \Psi_m$ for all $k \geq j_m$. By passing to a subsequence of (J_i^j) we may assume that $j_m = m$ for all m. Then $\left(\phi^{-1}(J_i^j)\right)_{j=1}^\infty$ gives the required sequence in Φ. $\square$

Definition 2.2.15. Given a sequence $(\mathcal{G}_i, v_i)_{i=1}^\infty$ as above, define a (not necessarily finite) pointed digraph $(\mathcal{G}_\infty, v_\infty)$ by $(\mathcal{G}_\infty, v_\infty) = \bigcup_{i=1}^\infty \mathcal{B}(\mathcal{G}_i, v_i, i)$. Note that, since $\mathcal{B}(\mathcal{G}_m, v_m, m) = \mathcal{B}(\mathcal{G}_k, v_k, m)$ for $k \geq m$, this is well-defined.

Lemma 2.2.16. *(1) $\mathcal{G}_\infty$ is infinite.*

(2) $\mathcal{G}_\infty$ has valence bounded above by n_1.

(3) There is an injection $\mathcal{G}_\infty \to \mathbb{S}^2$ that is an embedding on any finite subgraph of $\mathcal{G}_\infty$. In particular, given a simple closed curve in $\mathcal{G}_\infty$, there is a well-defined notion of which 'side' of this curve any other point of $\mathcal{G}_\infty$ lies.

Proof. (3) For $n \in \mathbb{N}$, let $f_n \colon \mathcal{G}_n \to \mathbb{S}^2$ be the embedding of $\mathcal{G}_n$ into $\mathbb{S}^2$. Given $n \geq 2$, there is an isotopy $H_n \colon \mathbb{S}^2 \times [0,1] \to \mathbb{S}^2$ from the identity on $\mathbb{S}^2$ to a map $h_n \colon \mathbb{S}^2 \to \mathbb{S}^2$ with $h_n(f_n(x)) = f_{n-1}(x)$ for all $x \in \mathcal{B}(\mathcal{G}_{n-1}, v_{n-1}, n-1)$. Define the inclusion $f_\infty \colon \mathcal{G}_\infty \to \mathbb{S}^2$ by $f_\infty(x) = h_2(h_3(\cdots h_n(f_n(x) \cdots)$ for all $x \in \mathcal{B}(\mathcal{G}_\infty, v_\infty, n)$ for any $n \geq 2$. $\square$

Lemma 2.2.17. *Suppose that, for every $m \in \mathbb{N}$, the boundary of any region r of $\mathbb{S}^2 \setminus \mathcal{G}_m$ is a cycle in $\mathcal{G}_m$ with length at most n_2. Then $\mathcal{G}_\infty$ is $\mathcal{O}$–connected.*

Proof. Let $u, v \in \mathrm{V}(\mathcal{G}_\infty)$. Then $u, v \in \mathcal{B}(\mathcal{G}_\infty, v_\infty, m)$ for some $m \in \mathbb{N}$. As $\mathcal{B}(\mathcal{G}_\infty, v_\infty, m)$ is connected, there is an (unoriented) path from u to v in $\mathcal{B}(\mathcal{G}_\infty, v_\infty, m) = \mathcal{B}(\mathcal{G}_{m+n_2}, v_{m+n_2}, m)$. This can be altered to a directed path from u to v in $\mathcal{B}(\mathcal{G}_{m+n_2}, v_{m+n_2}, m+n_2) \subset \mathcal{G}_\infty$. $\square$

Definition 2.2.18. Let $(\mathcal{G}, \mathcal{O})$ be a (possibly infinite) digraph. Given sets $A, B \subseteq \mathrm{V}(\mathcal{G})$, define an *$(A,B)$–path* ρ to be a simple directed path with respect to $\mathcal{O}$ that begins at a vertex of A and ends at a vertex of B without meeting any vertex of $A \cup B$ in its interior. If ρ contains exactly one edge, call it an *(A,B)–edge*.

Proposition 2.2.19 (Menger's Theorem)**.** *Let $\mathcal{G}$ be a (possibly infinite) digraph. Let $A \subseteq \mathrm{V}(\mathcal{G})$ and $v_B \in \mathrm{V}(\mathcal{G}) \setminus A$. Then either there are edge-disjoint (A, v_B)–paths ρ_1, ρ_2 or there is a partition of $\mathrm{V}(\mathcal{G})$ into sets $\tilde{A}$ and $\tilde{B}$ such that $A \subseteq \tilde{A}$, $v_B \in \tilde{B}$, and there is at most one $(\tilde{A}, \tilde{B})$–edge.*

Proof. Let $A_1 = \{v \in \mathrm{V}(\mathcal{G}) :$ there is an (A, v)–path in $(\mathcal{G}, \mathcal{O})\}$ and let $B_1 = \mathrm{V}(\mathcal{G}) \setminus A_1$. Then $A \subseteq A_1$. If $v_B \in B_1$ then this gives the required partition.

Suppose $v_B \in A_1$. Choose an (A, v_B)–path ρ. Define $\mathcal{O}_\rho$ to be the orientation of $\mathcal{G}$ given by reversing the direction of $\mathcal{O}$ on every edge of ρ. Let $A_2 = \{v \in \mathrm{V}(\mathcal{G}) :$ there is an (A, v)–path in $(\mathcal{G}, \mathcal{O}_\rho)\}$.

Assume $v_B \in A_2$. Let ρ' be an (A, v_B)–path in $(\mathcal{G}, \mathcal{O}_\rho)$, and let $B = \{e \in \mathrm{E}(\mathcal{G}) :$ $e \in \rho$ XOR $e \in \rho'\}$. Consider the edges in B with respect to $\mathcal{O}$. Exactly two such edges start in A while no edge ends in A, and two edges end at v_B. At any vertex $v \notin A \cup \{v_B\}$, the number of edges in B ending at v is equal to the number starting at v. Thus, since B is finite, there are two edge-disjoint (A, v_B)–paths in B.

If instead $v_B \notin A_2$, let $B_2 = \mathrm{V}(\mathcal{G}) \setminus A_2$. By definition, there are no (A_2, B_2)–edges with respect to $\mathcal{O}_\rho$. In particular, no edge of ρ is a (B_2, A_2)–edge with respect to $\mathcal{O}$. This means ρ contains exactly one (A_2, B_2)–edge with respect to $\mathcal{O}$. Hence $\tilde{A} = A_2$ and $\tilde{B} = B_2$ are as required. □

Lemma 2.2.20. *Suppose that $\mathcal{G}_m \in \Gamma$ and the boundary of any region r of $\mathbb{S}^2 \setminus \mathcal{G}_m$ is a cycle in $\mathcal{G}_m$ with length at most n_2 for each $m \in \mathbb{N}$. Then, for any partition of $\mathrm{V}(\mathcal{G}_\infty)$ into non-empty sets A and B, there are at least two (A, B)–edges.*

Proof. By Lemma 2.2.17, $\mathcal{G}_\infty$ is $\mathcal{O}$–connected, so there is at least one (B, A)–edge. Let e be a (B, A)–edge. Choose $m_1 \in \mathbb{N}$ such that $e \in \mathcal{B}(\mathcal{G}_\infty, v_\infty, m_1)$ and set $m_2 = m_1 + n_2$. Let r_1 and r_2 be the regions of $\mathbb{S}^2 \setminus \mathcal{G}_{m_2}$ adjacent to e. The boundaries of these regions are both contained in $\mathcal{B}(\mathcal{G}_{m_2}, v_{m_2}, m_2)$. Let $A_{m_2} = A \cap \mathcal{B}(\mathcal{G}_{m_2}, v_{m_2}, m_2)$ and $B_{m_2} = B \cap \mathcal{B}(\mathcal{G}_{m_2}, v_{m_2}, m_2)$. Then the boundaries of r_1 and r_2 contain (A_{m_2}, B_{m_2})–edges e_1 and e_2 respectively. Since $\mathcal{G}_{m_2}$ is prime, $e_1 \neq e_2$. Thus e_1 and e_2 are distinct (A, B)–edges in $\mathcal{G}_\infty$. □

Corollary 2.2.21. *If $(\mathcal{G}_i)_{i=1}^\infty$ is as in Lemma 2.2.20 and A is a non-empty, proper subset of $\mathrm{V}(\mathcal{G}_\infty)$, there is a vertex $v_B \notin A$ with two (A, v_B)–paths that do not meet except at their endpoints.*

Proof. Choose $v \notin A$. There are edge-disjoint (A, v)–paths ρ_1 and ρ_2. Take v_B to be the first common vertex of ρ_1 and ρ_2 not in A as measured along ρ_1. □

Definition 2.2.22. Call distinct directed paths ρ_1 and ρ_2 with the same endpoints a *spear-pair* if, for some $m \in \mathbb{N} \cup \{0\}$, they run together for the first m vertices, and after that point meet only at their final vertex (see Figure 2.15).

Proposition 2.2.23. *Let $(\mathcal{G}_i)_{i=1}^\infty$ be as in Lemma 2.2.20. Let $n \in \mathbb{N}$. Then there is a set $\mathcal{F}_n \subseteq \mathrm{V}(\mathcal{G}_\infty)$ such that there are at least 2^n (finite) directed $\mathcal{F}_n$–spanning subtrees of $\mathcal{G}_\infty$ with origin v_∞.*

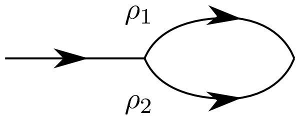

Figure 2.15

Proof. We inductively construct, for $0 \le m \le n$, a vertex $w_m \in \mathrm{V}(\mathcal{G}_\infty)$, a set $A_m \subset \mathrm{V}(\mathcal{G}_\infty)$ and (v_∞, w_m)–paths ρ_m^1, ρ_m^2. We also construct an unoriented (and possibly not simple) closed curve σ_m in $\mathcal{G}_\infty$ and an open disc $\mathbb{D}_m \subset \mathbb{S}^2$. These are all chosen with the following properties.

$(1)_m$ $B_m = \mathrm{V}(\mathcal{G}_\infty) \setminus A_m$ is infinite.

$(2)_m$ $\rho_m^1, \rho_m^2 \subseteq \mathcal{G}_\infty[A_m]$.

$(3)_m$ If $m > 0$ then ρ_m^1, ρ_m^2 are a spear-pair.

$(4)_m$ $\sigma_m \subseteq \bigcup_{i=0}^m (\rho_i^1 \cup \rho_i^2)$.

$(5)_m$ $\mathbb{D}_m$ is a connected component of $\mathbb{S}^2 \setminus \sigma_m$.

$(6)_m$ $B_m \subset \mathbb{D}_m$ and $A_m \cap \mathbb{D}_m = \emptyset$.

$(7)_m$ If $m > 0$ then $w_m \in B_{m-1}$.

$(8)_m$ Let $\mathcal{F}_m = \{w_0, \dots, w_m\}$. For each $(J_i)_{i=1}^m \in \{1,2\}^m$ there is a directed $\mathcal{F}_m$–spanning subtree $\mathcal{T}$ of $\mathcal{G}_\infty$ with origin v_∞ such that $\mathcal{T} \subseteq \mathcal{G}_\infty[A_m]$ and, for $1 \le i \le m$, the edge of $\mathcal{T}$ ending at w_i is the last edge of $\rho_i^{J_i}$.

Let $w_0 = A_0 = \rho_0^1 = \rho_0^2 = \sigma_0 = v_\infty$ and $\mathbb{D}_0 = \mathbb{S}^2 \setminus v_\infty$. Then $(1)_0$–$(8)_0$ hold.

Suppose w_m, A_m, ρ_m^1, ρ_m^2, σ_m and $\mathbb{D}_m$ have been defined for some $m < n$. Then by Corollary 2.2.21 there exists $w_{m+1} \in B_m$ and (A_m, w_{m+1})–paths $\rho_{m+1}^3, \rho_{m+1}^4$ that do not meet except at their endpoints. Since $w_{m+1} \in \mathbb{D}_m$, both ρ_{m+1}^3 and ρ_{m+1}^4 must have their initial vertex on σ_m. By $(4)_m$, these can be extended using directed paths in $\bigcup_{i=0}^m (\rho_i^1 \cup \rho_i^2)$ to a spear-pair of (v_∞, w_{m+1})–paths $\rho_{m+1}^1, \rho_{m+1}^2$.

Now $\rho_{m+1}^3 \cup \rho_{m+1}^4$ forms a finite length simple closed curve or arc contained in $\mathcal{G}_\infty$ with both endpoints on the boundary of $\mathbb{D}_m$. Thus $\mathbb{D}_m \setminus (\rho_{m+1}^3 \cup \rho_{m+1}^4)$ consists of two open discs $\mathbb{D}_{m+1}^1$ and $\mathbb{D}_{m+1}^2$. Further, $B_{m+1}^1 = \mathbb{D}_{m+1}^1 \cap \mathrm{V}(\mathcal{G}_\infty)$ and $B_{m+1}^2 = \mathbb{D}_{m+1}^2 \cap \mathrm{V}(\mathcal{G}_\infty)$ are well defined. At least one of B_{m+1}^1, B_{m+1}^2 is infinite. Suppose B_{m+1}^1 is infinite. Let $B_{m+1} = B_{m+1}^1$ (that is, let $A_{m+1} = \mathrm{V}(\mathcal{G}) \setminus B_{m+1}^1$) and let $\mathbb{D}_{m+1} = \mathbb{D}_{m+1}^1$. Finally, define

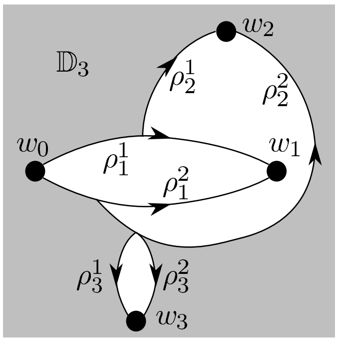

Figure 2.16

σ_{m+1} as the boundary of $\mathbb{D}_{m+1}$. Note that σ_{m+1} consists of $\rho^3_{m+1} \cup \rho^4_{m+1}$ together with all, one section or none of σ_m. Figure 2.16 shows a specific example.

Properties $(1)_{m+1}$–$(7)_{m+1}$ now hold. It remains to check $(8)_{m+1}$. Note that since the final edges of ρ_i^1 and ρ_i^2 differ for $1 \leq i \leq m+1$, $(8)_{m+1}$ implies there are at least 2^{m+1} directed $\mathcal{F}_{m+1}$–spanning subtrees with origin v_∞.

Let $(J_i)_{i=1}^{m+1} \in \{1,2\}^{m+1}$. If $m+1 = 1$, let $\mathcal{T} = \rho_1^{J_1}$. If $m+1 > 1$, choose $\mathcal{T}'$ with origin v_∞ such that $\mathcal{T}' \subseteq \mathcal{G}_\infty[A_m]$ and, for $1 \leq i \leq m$, the edge of $\mathcal{T}'$ ending at w_i is the last edge of $\rho_i^{J_i}$. Since $w_{m+1} \in B_m$, it does not lie on $\mathcal{T}'$. Let v' be the last vertex of $\rho_{m+1}^{J_{m+1}}$ that meets $\mathcal{T}'$, and let ρ' be the section of $\rho_{m+1}^{J_{m+1}}$ from v' to w_{m+1}. Then $\mathcal{T} = \mathcal{T}' \cup \rho'$ has the required properties. □

Theorem 2.2.24. *The set* $\Gamma_n = \{\mathcal{G} \in \Gamma : \exists v \in \mathrm{V}(\mathcal{G}), |\operatorname{Tr}(\mathcal{G}, v)| \leq n\}$ *is finite for any* $n \in \mathbb{N}$.

Proof. Fix $n \in \mathbb{N}$. For each $\mathcal{G} \in \Gamma_n$, fix a vertex $v \in \mathrm{V}(\mathcal{G})$ such that $|\operatorname{Tr}(\mathcal{G}, v)| \leq n$ and fix an embedding of $\mathcal{G}$ into $\mathbb{S}^2$. By Corollary 2.2.5 each $\mathcal{G} \in \Gamma_n$ has valence bounded above by $2n$, and by Corollary 2.2.13 the boundary of any region r of $\mathbb{S}^2 \setminus \mathcal{G}$ is a cycle in $\mathcal{G}$ with length at most n.

Suppose, for a contradiction, that Γ_n is infinite. Then, by Lemma 2.2.14, there is a sequence $(\mathcal{G}_i, v_i)$ from which we can define an infinite pointed digraph $\mathcal{G}_\infty$ as in Definition 2.2.15. By Proposition 2.2.23, there is a set $\mathcal{F} \subseteq \mathrm{V}(\mathcal{G}_\infty)$ and a sequence $(\mathcal{T}_i)_{i=1}^{n+1}$ of finite directed $\mathcal{F}$–spanning subtrees with origin v_∞. Choose $m \in \mathbb{N}$ such that $\mathcal{T}_i \subseteq \mathcal{B}(\mathcal{G}_\infty, v_\infty, m) = \mathcal{B}(\mathcal{G}_m, v_m, m)$ for $1 \leq i \leq n+1$. Since $\mathcal{G}_m$ is $\mathcal{O}$–connected, Lemma 2.1.13 implies that $n+1 \leq |\operatorname{Tr}(\mathcal{B}(\mathcal{G}_m, v_m, m), v_m)| \leq |\operatorname{Tr}(\mathcal{G}_m, v_m)| \leq n$. □

Remark 2.2.25. Theorem 2.2.24 is the key step in the proof of Theorem 1.1.5, which we give in Section 2.3.

2.3 Proofs

2.3.1 Theorem 1.1.4

By Theorem 2.2.24 we now know that the set Γ_3 is finite. In order to prove Theorem 1.1.4, we calculate this set explicitly.

Theorem 2.3.1. *Let $\mathcal{G} \in \Gamma$. Suppose that $|\operatorname{Tr}(\mathcal{G}, v)| < 4$ for any $v \in \operatorname{V}(\mathcal{G})$. Then, up to reflection, $\mathcal{G}$ is one of the digraphs $\mathcal{G}_\alpha, \mathcal{G}_\beta, \mathcal{G}_\gamma, \mathcal{G}_\delta$ shown in Figure 2.17.*

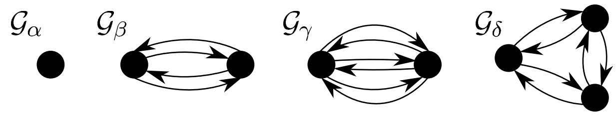

Figure 2.17

Remark 2.3.2. This result for $|\operatorname{Tr}(\mathcal{G}, v)| \in \{1, 2\}$ is already known ([27] Theorem 3(5)).

Proof. First suppose $\mathbb{S}^2 \setminus \mathcal{G}$ is a single region r. Then $\mathcal{G}$ is a tree. Since $\mathcal{G}$ is finite but has no 1–valent vertices, $\mathcal{G} = \mathcal{G}_\alpha$.

Assume there are at least two regions of $\mathbb{S}^2 \setminus \mathcal{G}$. Let r be one such region. Then ∂r is a topological circle, as otherwise r would meet itself at a vertex of $\mathcal{G}$, contradicting that $\mathcal{G}$ is prime. Note that every region of $\mathbb{S}^2 \setminus \mathcal{G}$ has at least two sides. Suppose every region has exactly two sides. Then by considering the Euler characteristic of $\mathbb{S}^2$ and of $\mathcal{G}$ we find that $|\operatorname{V}(\mathcal{G})| = 2$. Since $|\operatorname{Tr}(\mathcal{G}, v)| < 4$ for a vertex v, we see that either $\mathcal{G} = \mathcal{G}_\beta$ or $\mathcal{G} = \mathcal{G}_\gamma$.

This leave the case where $\mathbb{S}^2 \setminus \mathcal{G}$ has a region r_0 with at least 3 sides. Since ∂r_0 is a circle, considering χ shows there is a second region r_1 with at least 3 sides.

We now consider whether certain digraphs can occur as subgraphs of $\mathcal{G}$. Note that, by (G4), the boundary of every region of $\mathbb{S}^2 \setminus \mathcal{G}$ is a cycle, so $\mathcal{G}$ is $\mathcal{O}$–connected. Thus Lemma 2.1.13 applies.

Lemma 2.3.3. *If the graph shown in Figure 2.18 is a subgraph of $\mathcal{G}$, where v and w may coincide or v' and w' may coincide, but not both, then $|\operatorname{Tr}(\mathcal{G}, v)| \geq 4$. Here a dashed line denotes a directed simple path.*

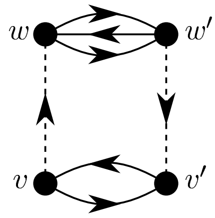

Figure 2.18

Lemma 2.3.4. *If the graph shown in Figure 2.19 is a subgraph of $\mathcal{G}$, then $|\operatorname{Tr}(\mathcal{G}, v)| \geq 4$.*

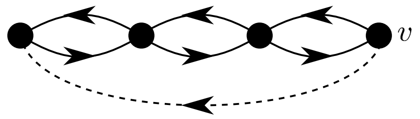

Figure 2.19

Lemma 2.3.5. *Suppose exactly two regions of $\mathbb{S}^2 \setminus \mathcal{G}$ have 3 or more sides. Then, up to reflection, $\mathcal{G} = \mathcal{G}_\delta$.*

Proof. Collapse each bigon region of $\mathbb{S}^2 \setminus \mathcal{G}$ to a line, giving a graph $\mathcal{G}'$. Since this only leaves two regions, $\mathcal{G}'$ is a topological circle. Number the vertices $\mathrm{V}(\mathcal{G}) = \mathrm{V}(\mathcal{G}')$ as $v_1, \ldots, v_n$ around this circle. Then $n \geq 3$. For $1 \leq i \leq n$, let J_i be the number of (unoriented) edges in $\mathcal{G}$ joining v_i to v_{i+1} (where $v_{n+1} = v_1$). By (G4) and (G5), $v_i + v_{i+1}$ is even and at least 4 for each i. Suppose $v_1 \geq 3$. Then, by Lemma 2.3.3, $v_i = 1$ for all $i > 1$. Thus $v_2 + v_3 = 2$, contradicting that $v_2 + v_3 > 2$. Hence $v_i \leq 2$ for all i. By Lemma 2.3.4, this means $n = 3$. ☐

It now suffices to show that at most two regions of $\mathbb{S}^2 \setminus \mathcal{G}$ have three or more sides. We therefore assume otherwise.

Claim 2.3.6. *Up to relabelling the regions of $\mathbb{S}^2 \setminus \mathcal{G}$, r_0 shares an edge with r_1.*

Proof. Suppose otherwise. Then r_0 meets a bigon along every edge. By Lemma 2.3.4, ∂r_0 is a triangle. Since at least three regions of $\mathbb{S}^2 \setminus \mathcal{G}$ have three or more sides and $\mathcal{G}$ is prime, $\mathcal{G}$ contains the graph shown in Figure 2.20a and a directed path disjoint from this graph connecting two distinct vertices of ∂r_0. It must therefore contain one of the graphs in Figure 2.20b. In either case, $|\operatorname{Tr}(\mathcal{G}, v)| \geq 4$. ☐

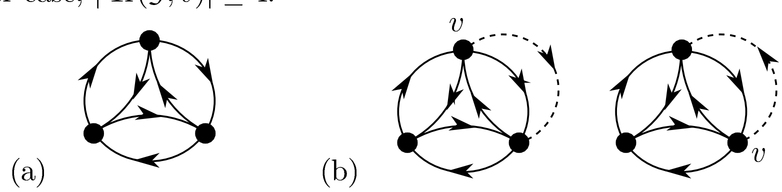

Figure 2.20

Claim 2.3.7. *The regions r_0 and r_1 share exactly one edge e_1 and meet at no vertices other than the endpoints of e_1.*

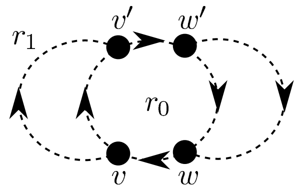

Figure 2.21

Proof. First suppose that $\partial r_0 \cap \partial r_1$ has at least two components. Then $\mathcal{G}$ contains the graph shown in Figure 2.21, where v' and w' may coincide. Then $|\operatorname{Tr}(\mathcal{G}, v)| \geq 4$. Thus $\partial r_0 \cap \partial r_1$ has only one component.

Now suppose r_0 and r_1 share at least two consecutive edges (since ∂r_0 and ∂r_1 are circles, such edges must be consecutive both in ∂r_0 and in ∂r_1). Then the vertex between these edges had in-degree 1. This contradicts (G5). □

We now know $\mathcal{G}$ contains the digraph $\mathcal{G}_1$ shown in Figure 2.22. Let r_2 be the region of

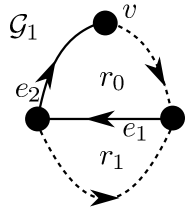

Figure 2.22

$\mathbb{S}^2 \setminus \mathcal{G}$ that meets r_0 along edge e_2. Since ∂r_2 is a cycle, part of ∂r_2 forms a directed path from v to a vertex of $\mathcal{G}_1$. As $\mathcal{G}$ is prime, this path cannot end at v. Thus $\mathcal{G}$ contains one of the digraphs $\mathcal{G}_2$–$\mathcal{G}_5$ shown in Figure 2.23. Note that $|\operatorname{Tr}(\mathcal{G}_2, v')| = |\operatorname{Tr}(\mathcal{G}_4, v')| = 4$, so

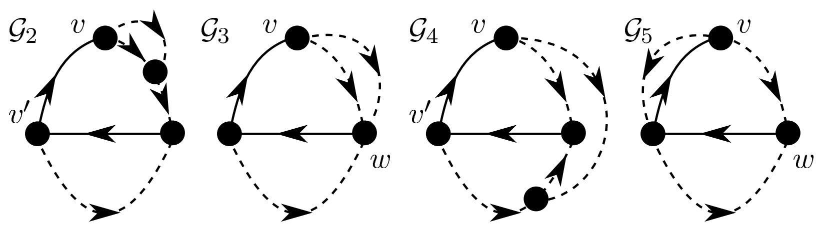

Figure 2.23

neither of these cases can occur. In addition, $\mathcal{G}_3$ cannot arise, as otherwise r_0 and r_2 would contradict Claim 2.3.7.

Thus $\mathcal{G}$ contains $\mathcal{G}_5$. There is a directed path from w, creating one of the digraphs $\mathcal{G}_6$–$\mathcal{G}_{10}$ shown in Figure 2.24. If $i \in \{6, 7, 8, 10\}$ then $|\operatorname{Tr}(\mathcal{G}_i, v)| \geq 4$. This leaves $\mathcal{G}_9$ as the only possibility.

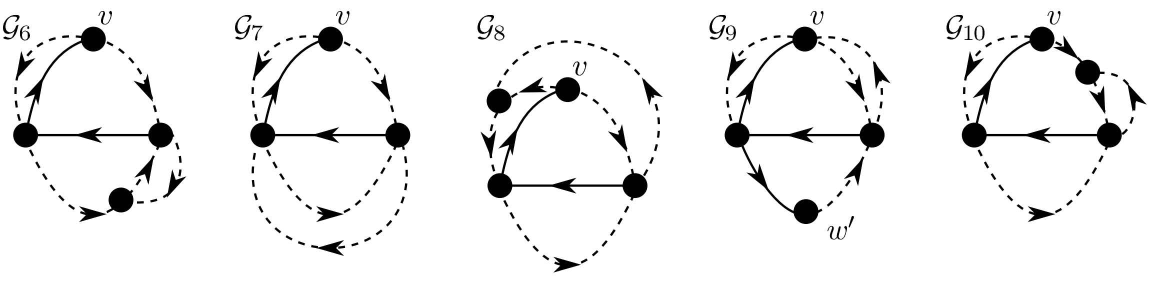

Figure 2.24

There is another directed path beginning at w'. By discarding cases that have already been considered, we see that one of the digraphs $\mathcal{G}_{11}$–$\mathcal{G}_{13}$ shown in Figure 2.25 is contained in $\mathcal{G}$. If $i \in \{11, 12, 13\}$, then $|\operatorname{Tr}(\mathcal{G}_i, w')| > 4$. $\square$

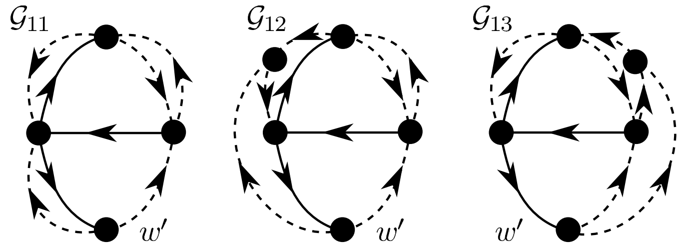

Figure 2.25

Lemma 2.3.8. *Let $\mathcal{G} \in \{\mathcal{G}_\alpha, \mathcal{G}_\beta, \mathcal{G}_\gamma, \mathcal{G}_\delta, \mathcal{G}'_\delta\}$, where $\mathcal{G}'_\delta$ is the reflection of $\mathcal{G}_\delta$. Then $M^{\mathcal{M}}(\mathcal{G})$ (see Definition 2.1.19) is an almost product sutured manifold.*

Proof. $M^{\mathcal{M}}(\mathcal{G}_\alpha)$ is a 3–ball with a single suture, which is a product sutured manifold. The other cases are proved by Kobayashi. $M^{\mathcal{M}}(\mathcal{G}_\beta)$ and $M^{\mathcal{M}}(\mathcal{G}_\gamma)$ give cases of [21] Example 2.8. $N^{\mathcal{M}}(\mathcal{G}_\delta)$ and $N^{\mathcal{M}}(\mathcal{G}'_\delta)$ are both [21] Example 4.5, but with opposite orientations. $\square$

Remark 2.3.9. Considering $M^{\mathcal{H}}$ or $M^{\mathcal{K}}$ would give the same sutured manifolds, but in a different order.

Theorem 2.3.10. *Let L be a prime, special alternating link with $\Delta^0_L(0) < 4$. Then there is a unique incompressible Seifert surface R for L.*

Proof. Let D_0 be a special alternating diagram for L. Then D_0 is either positive or negative. Possibly reflecting D_0 ensures it is positive. By a sequence of flypes we can also ensure that D_0 is twist-reduced. Let R be the surface given by applying Seifert's algorithm to D_0. Then R is a taut Seifert surface for L.

Let $\mathcal{G}_0 = G^{\mathcal{M}}(D_0)$. Apply moves 1 and 2 to $\mathcal{G}_0$ as many times as possible, giving a sequence of digraphs $\mathcal{G}_1, \ldots, \mathcal{G}_n$ where each $\mathcal{G}_{i+1}$ is obtained from $\mathcal{G}_i$ by a move 1 or a move 2 (see Definition 2.1.17). These moves can be chosen so that any move 1 removes an innermost loop e, so that e bounds a disc. Then, for any $v \in \mathrm{V}(\mathcal{G}_n)$, Lemmas 2.1.15 and 2.1.16 give that

$$|\operatorname{Tr}(\mathcal{G}_n, v)| = |\operatorname{Tr}(\mathcal{G}_{n-1}, v)| = \cdots = |\operatorname{Tr}(\mathcal{G}_0, v)| = \Delta_L^0(0) < 4.$$

There are no loops in $\mathcal{G}_n$, since no move 1 is possible. Neither move type can create a cut vertex, and $\mathcal{G}_0$ has no cut vertex because D_0 is connected and prime. Suppose there is a simple closed curve ρ that meets the edges of $\mathcal{G}_n$ twice. We can reverse each move in the complement of ρ, so that ρ meets the edges of $\mathcal{G}_0$ twice. Then ρ meets two crossings of D_0, and otherwise lies in the black regions. As $\mathcal{G}_0$ is twist-reduced, one side of ρ contains only a line of white bigons. This means that, in $\mathcal{G}_0$, that side of ρ contains a single topological arc, composed of edges and 2–valent vertices. Each such vertex can be removed at any stage by a move 2. Thus the points where ρ meets $\mathcal{G}_n$ lie on a single edge of $\mathcal{G}_n$. We conclude that $\mathcal{G}_n \in \Gamma$.

Therefore, up to reflection, $\mathcal{G}_n \in \{\mathcal{G}_\alpha, \mathcal{G}_\beta, \mathcal{G}_\gamma, \mathcal{G}_\delta\}$, so $M^{\mathcal{M}}(\mathcal{G}_n)$ is an almost product sutured manifold by Lemma 2.3.8. Then, from Lemmas 2.1.20 and 2.1.21, we see that $M^{\mathcal{M}}(\mathcal{G}_i)$ is an almost product sutured manifold for each $i \leq n$. In particular, $M^{\mathcal{M}}(\mathcal{G}_0) = M^{\mathcal{M}}(G^{\mathcal{M}}(D_0))$ is an almost product sutured manifold. By Proposition 1.4.22, the incompressible Seifert surface R is unique, as $M^{\mathcal{M}}(G^{\mathcal{M}}(D_0))$ is the complementary sutured manifold to R. □

Remark 2.3.11. The diagram D_0 has properties (L1)–(L5). The graphs $\mathcal{G}_1, \ldots, \mathcal{G}_n$ correspond via $G^{\mathcal{M}}$ to diagrams $D_1, \ldots, D_n$. As we have seen in Lemmas 2.1.20 and 2.1.21, a move 2 corresponds to removing a white bigon, and a move 1 to untwisting a nugatory crossing. This point of view explains why $D_n \in \Lambda$.

Corollary 2.3.12 (see [17] p604)**.** *Let L be a special alternating link. If $\Delta_L^0(0) = 1$ then L is fibred.*

Proof. In this case, $\mathcal{G}_n = \mathcal{G}_\alpha$. Thus $M^{\mathcal{M}}(\mathcal{G}_n)$ is a 3–ball with a single suture, which is the complement of a disc. Hence R is obtained from a disc by plumbing with Hopf bands. By Proposition 1.6.10, as a Hopf band is fibred, L is fibred with fibre R. □

Combining this with Lemma 2.1.8 gives the following.

Corollary 2.3.13. *Let D be a special alternating diagram of a prime link L. Then L is fibred if and only if the graph $G^{\mathcal{M}}(D)$ can be reduced to a single vertex by a sequence of the following moves.*

- *Delete a loop.*
- *Contract an edge, one endpoint of which is at a vertex with valence 2.*

Informally, $G^{\mathcal{M}}(D)$ satisfies this condition exactly when it is a 'tree of loops'.

Theorem 1.1.4. *Let L be a non-split homogeneous link. If $\Delta^0_L(0) < 4$ then L has a unique incompressible Seifert surface.*

Proof. Let D be a homogeneous diagram for L with no nugatory crossings. Break up D along the non-special Seifert circles into special diagrams. Note that this cannot create nugatory crossings. Further break down the diagram into prime pieces. This gives prime, special alternating link diagrams $D_1, \ldots, D_n$ for some $n \in \mathbb{N}$ such that D can be constructed by combining $D_1, \ldots, D_n$ by taking connected sums and $*$–products. Let L_i be the link with diagram D_i for $1 \leq i \leq n$. Then $\Delta^0_L(0) = \prod_{i=1}^n \Delta^0_{L_i}(0)$ by Propositions 2.1.6 and 2.1.2. Since $\Delta^0_L(0) \in \{1, 2, 3\}$, without loss of generality, $\Delta^0_{L_1}(0) = \Delta^0_L(0)$ and $\Delta^0_{L_i}(0) = 1$ for all $i > 1$.

By Theorem 2.3.10, L_i has a unique incompressible Seifert surface R_i for each i, which is given by Seifert's algorithm. For $i > 1$, L_i is fibred with fibre R_i. By repeated use of Theorem 1.6.9, L has a unique incompressible Seifert surface. □

Remark 2.3.14. The link L_δ in Proposition 3.3.13 is an example of a special alternating link with a unique taut Seifert surface but not a unique incompressible Seifert surface.

Remark 2.3.15. It seems unlikely that the methods we have used can be extended to give a complete proof of Theorem 1.1.3. Our proof relies heavily on using the Seifert surface distinguished by applying Seifert's algorithm to a fixed diagram. If a set of distinct, disjoint taut Seifert surfaces with maximal cardinality has cardinality at least 2, it is not known at present whether there always exists such a set containing this particular surface. This remains true if we are allowed to change the alternating diagram to which we apply Seifert's algorithm.

2.3.2 Theorems 1.1.5 and 1.1.6

Theorem 1.1.5. *For fixed $n \in \mathbb{N}$, there is a finite set $\mathcal{S}$ of surfaces embedded in $\mathbb{S}^3$ with the following property. Any non-split, homogeneous link L with $\Delta^0_L(0) \leq n$ has a taut Seifert surface R built from surfaces in $\mathcal{S}$ by reflection, Murasugi sum and plumbing with Hopf bands.*

If $\Delta^0_L(0)$ is prime, R can be formed using only one element of $\mathcal{S}$.

Proof. Let L be a non-split, homogeneous link with $\Delta_L^0(0) \leq n$. Let D be a reduced, homogeneous diagram for L, and let R be the taut Seifert surface given by applying Seifert's algorithm to D. Break up D into prime, special alternating link diagrams $D_1, \ldots, D_m$ for some $m \in \mathbb{N}$ such that D can be constructed by combining $D_1, \ldots, D_m$ by taking connected sums and $*$–products. For $i \leq m$, let R_i be the taut Seifert surface given by applying Seifert's algorithm to D_i. Then R can be built from $\{R_1, \ldots, R_m\}$ by Murasugi summation. In addition, $\Delta_L^0(0) = \prod_{i=1}^m \Delta_{L_i}^0(0)$ by Propositions 2.1.6 and 2.1.2, where L_i is the link with diagram D_i. For $1 \leq i \leq m$ we know that $\Delta_{L_i}^0(0) \geq 1$, so $1 \leq \Delta_L^0(0) \leq n$, and thus $\Delta_{L_i}^0(0) \leq n$ for $1 \leq i \leq m$.

If $\Delta_L^0(0)$ is prime then, without loss of generality, $\Delta_{L_i}^0(0) = 1$ for $i \geq 2$. From the proof of Corollary 2.3.12 we see that each R_i for $i \geq 2$ is obtained from a disc by plumbing with Hopf bands, corresponding to white bigons in the diagram D_i. It is thus possible to plumb Hopf bands on to R_1 in a suitable order to give R. We will explain how to see this at the end of the proof.

We may therefore assume that D is a prime, special alternating diagram. As L is non-split, D is connected. Form a sequence of diagrams $D = D^0, D^1, \ldots, D^l$, where D^{i+1} is obtained from D^i either by removing a white bigon or by untwisting a nugatory crossing, for $0 \leq i < l$, and D^l has no white bigons or nugatory crossings. Note that if a diagram in this sequence is the single crossing diagram of the unknot then the following diagram must be the unknot diagram with no crossings, which is obtained by untwisting the single nugatory crossing. Let R^i be the surface given by applying Seifert's algorithm to D^i. Then R^i can be obtained from R^{i+1} either by an isotopy or by plumbing with a Hopf band.

We wish to show that there are only finitely many possibilities for the diagram D^l and hence also for the surface R^l. We may then take the set $\mathcal{S}$ to consist of these possible surfaces. By reflecting if needed, we can ensure D^l is positive. It is clear that D^l is still special alternating and connected. Any simple closed curve that showed that D^l is not prime could be used to find such a curve in D, which is not possible since D is prime. Thus D^l has all the properties in Definition 2.2.8 except that it may not be black-twist-reduced.

Suppose D^l is not black-twist-reduced. Then there is a simple closed curve ρ_1 that passes through two crossings, c_1 and c_1', of D^l and otherwise lies in the black regions. As D^l is reduced, ρ_1 runs through two distinct black regions r and r'. By a flype (see Definition 1.2.19), c_1' can be moved to be adjacent to c_1 in a new diagram, in the sense that these two crossings are the two crossings of a white bigon.

Let A_1 be the set of all crossings c in D^l where there exists a simple closed curve analogous to ρ_1 through c_1 and c. Also include c_1 in the set A_1. Choose a crossing c_2 that is not in A_1, and similarly form a set A_2 of crossings of D^l. It may be that A_2 consists of only c_2, but for the moment we assume otherwise. First note that A_1 and A_2 are disjoint. Let ρ_2 be a simple closed curve, analogous to ρ_1, through c_2 and another crossing, c_2', of A_2. Isotope ρ_1 and ρ_2 within the black regions to minimise their intersection. Then ρ_1 and ρ_2 can cross at most once in each of the black regions r and r'. Note that $|\rho_1 \cap \rho_2|$ is even. However, if ρ_2 passed through r and r' then c_2 would be in A_1, which is not the case. Hence we see that ρ_1 and ρ_2 are disjoint.

Now suppose that we perform a flype on the crossing c_1' as described above. It is clear that this does not affect the set A_1. Can it affect A_2? Every element of A_2 continues to meet the necessary condition. It is not possible for any new crossings to meet the condition either, as any potential new simple closed curve ρ_2' would have to pass through r and r'. This shows that the flypes we would like to perform using crossings in A_1 are independent of those using crossings in A_2.

Continue to choose crossings c_i and sets A_i until $A_1, \ldots, A_k$ form a partition of the crossings of D^l. For $1 \leq i \leq k$, perform flypes to move to c_i all other crossings of A_i. Note that possibly no flypes are needed for some or all i. Let D^{l+1} be the resulting diagram, and let D^{l+2} be the diagram given by removing from D^{l+1} all of the white bigons created by the flypes. No nugatory crossings are created by this process. Thus D^{l+2} satisfies all the properties in Definition 2.2.8, and so is in Λ. Note that D^{l+2} has k crossings.

Our next step is to show that if D^{l+2} is given, there are only finitely many possibilities for D^l. Suppose A_1 contains at least 3 crossings. Let c_1'' be a third element of A_1, and let ρ_1'' be the simple closed curve through c_1 and c_1''. By an isotopy in the black regions, we may ensure that ρ_1'' lies entirely to one side of ρ_1 except at the crossing c_1. Then $\mathbb{S}^2 \setminus (\rho_1 \cup \rho_1'')$ is three open discs. As D^l contains no white bigons, each of these three discs must contain at least one crossing of D^l, and therefore at least one of $A_2, \ldots, A_k$. This shows that there is a uniform bound on the size of A_i in terms of k. Accordingly, given D^{l+2} there are only finitely many possibilities for D^{l+1}.

Finally, if we know D^{l+1}, counting the number of white bigons tells us how many flypes take place between D^{l+1} and D^l. Since in any link diagram there are only finitely many possible flypes, this gives a finite list of possibilities for D^l, as required. It is now sufficient to show that D^{l+2} can only be one of a finite set of diagrams.

Let $\mathcal{G}_i = G^{\mathcal{M}}(D^i)$ for $1 \leq i \leq l+2$. For $1 \leq i < l$, the digraphs $\mathcal{G}_i$ and $\mathcal{G}_{i+1}$ differ by either a move 1 or a move 2. Getting from $\mathcal{G}_{l+1}$ to $\mathcal{G}_{l+2}$ only requires the use of move 2. Finally, D^l and D^{l+1} are diagrams of the same link, L^l. Therefore, for an appropriate choice of vertex $v_i \in \mathrm{V}(\mathcal{G}_i)$ for each i,

$$\begin{aligned}\Delta_L^0(0) &= |\operatorname{Tr}(\mathcal{G}_0, v_0)| = |\operatorname{Tr}(\mathcal{G}_1, v_1)| = \cdots = |\operatorname{Tr}(\mathcal{G}_l, v_l)| \\ &= \Delta_{L^l}^0(0) = |\operatorname{Tr}(\mathcal{G}_{l+1}, v_{l+1})| = |\operatorname{Tr}(\mathcal{G}_{l+2}, v_{l+2})|.\end{aligned}$$

By Lemma 2.2.11, $\mathcal{G}_{l+2} \in \Gamma$, so $\mathcal{G}_{l+2} \in \Gamma_n$ as $\Delta_L^0(0) \leq n$. Recall that Theorem 2.2.24 tells us that Γ_n is finite. Applying Lemma 2.2.11 again gives the required bound on the number of possibilities for the digraph D^{l+2}.

We now return to the case where D is not prime or not special but $\Delta_L^0(0)$ is a prime number.

First suppose that D is special but not prime. Without loss of generality, D is formed from D_1 by taking a connected sum with the diagrams $D_2, \ldots, D_m$ in order. It is sufficient to show that we can remove the surface R_m from R by removing Hopf bands. Let ρ be a simple closed curve that crosses the diagram D transversely at two points, showing that D is a connected sum of D_m and $D_1 \# \cdots \# D_{m-1}$. As L_m is fibred, D_m contains a white bigon r. There is a region r' of D corresponding to r, which is either a white bigon or the white region through which ρ passes. If r' is a white bigon, then deleting r' from D corresponds to pulling a Hopf band off the section of the surface R that corresponds to R_m, by Lemma 1.6.11. The new diagram produced, D', is special and has fewer crossings than D. It may not be reduced, but removing any nugatory crossings will not change the surface given by applying Seifert's algorithm.

If instead r' is the white region through which ρ passes then we could have made a different choice for the bigon r. To see this, recall that R_m can be reduced to a disc by repeatedly removing Hopf bands, and consider how the diagram D_m changes as these Hopf bands are removed. As D_m has no nugatory crossings, the first move is to delete a white bigon. This may create a single nugatory crossing, which is then untwisted. This process as a whole does not increase the total number of white bigons. However, by repeating it, we will reach the point where only one Hopf band remains, and the diagram is the standard two-crossing diagram of the Hopf link. This contains two white bigons. Thus the original diagram D_m must contain at least two white bigons.

Finally we need to consider the effect of taking a $*$–product when one of the links involved is fibred. For this purpose we will change notation slightly, and assume that D is formed from D_1 by taking $*$–products with $D_2, \ldots, D_m$ in order, where each D_i is special but not necessarily prime.

Let C_m be the Seifert circle along which D_m is added. Then C_m is a non-special Seifert circle in D. Since D is reduced, D_m has at least two crossings. As L_m is fibred, D_m contains a white bigon r. Let r' be the corresponding region of D. If the boundary of r' does not meet C_m then deleting r' again corresponds to pulling a Hopf band off R and reduces the number of crossings in D_m. This leaves the case where both crossings of the bigon r lie on the copy of C_m in D_m. It may be that the section of C_m that forms part of the boundary of r' contains other crossings, coming from $D_1 * \cdots * D_{m-1}$. It is nevertheless possible to perform a change on the section of D corresponding to D_m that reduces the number of crossings and equates to pulling off a Hopf band from the section of R corresponding to R_m. This is shown in Figure 2.26. Note that in general there is some choice needed when

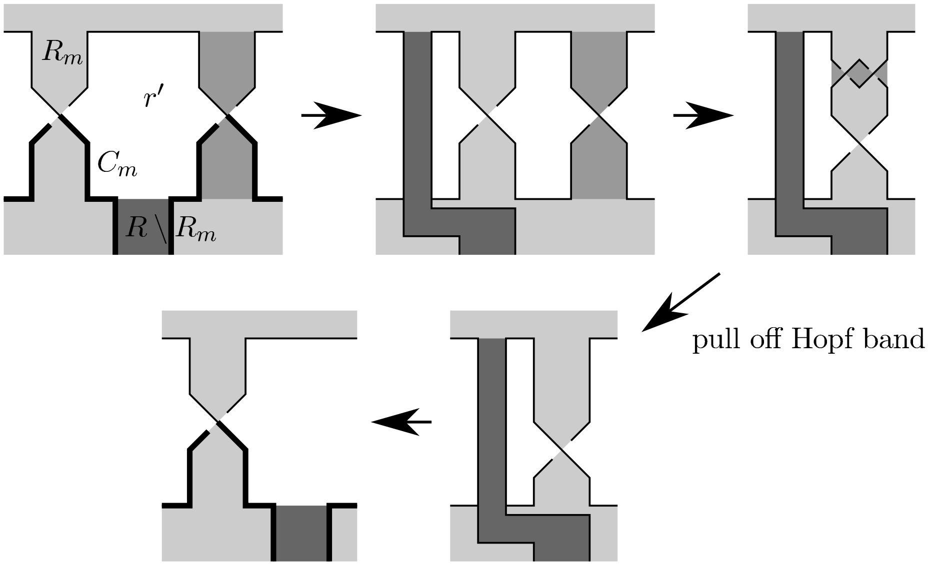

Figure 2.26

applying Seifert's algorithm to a non-special diagram. To build the surface R depicted in Figure 2.26, take the surface R_m formed of the black regions of D_m, and then position $R \setminus R_m$ (which is given the darkest shading) above the disc bounded by C_m.

It is therefore possible to construct the surface R from R_1 by plumbing with Hopf bands, as required. □

In [32], Riley proves the following using Hermitian forms. He notes that an alternative proof based on [8] is possible.

Theorem 1.1.6. *Choose $g \geq 0$ and $m, n \geq 1$. Then there are only finitely many alternating links L with m link components and genus g such that $\Delta_L^0(0) = n$.*

Proof. Let L be such a link with a reduced, alternating diagram D. Let L_1, L_2 be non-trivial alternating links with reduced, alternating diagrams D_1, D_2. Suppose $D = D_1 * D_2$. Let R, R_1, R_2 be the taut Seifert surfaces given by applying Seifert's algorithm to D, D_1, D_2 respectively. Then $\chi(R_1), \chi(R_2) > \chi(R)$, and $\Delta_{L_1}^0(0), \Delta_{L_2}^0(0) \leq \Delta_L^0(0)$. In addition, there are only finitely many ways of combining D_1 and D_2 by a $*$–product. The same is true if $L = L_1 \# L_2$. These facts allow us to reduce to the case that D is prime and special.

Starting from $G^{\mathcal{M}}(D)$, form a sequence of graphs as in the proof of Theorem 2.3.10, with final graph $\mathcal{G}$. Then $\mathcal{G} \in \Gamma_n$, so by Theorem 2.2.24 there are only finitely many possibilities for $\mathcal{G}$. We know that $G^{\mathcal{M}}(D)$ is obtained from $\mathcal{G}$ by reversing a finite sequence of moves 1 and 2. It therefore remains to bound the length of this sequence. Note that no move 1 can be performed on $G^{\mathcal{M}}(D)$ as D is reduced, and each move 2 creates at most one loop. It thus suffices to bound the number of times a move 2 occurs. Since a move 2 increases the Euler characteristic of the corresponding surface, such a bound exists. $\square$

Chapter 3

Special alternating links

3.1 The paper of Hirasawa–Sakuma

The purpose of this chapter is to prove the results stated in [15]. There are a number of definitions we will need from [15], which we now give. We will change some of the notation and terminology.

Let L be a prime link with a reduced, special alternating diagram D. Consider the planar graph $\mathcal{G}$ that has a vertex in each black region and an edge through each crossing of D. Note that $\mathcal{G}$ contains no loops. It may be that distinct edges $e, e' \in \mathrm{E}(\mathcal{G})$ bound a bigon region of $\mathbb{S}^2 \setminus \mathcal{G}$. If this occurs, remove one of e, e'. Repeat this until $\mathbb{S}^2 \setminus \mathcal{G}$ has no bigon regions.

Suppose there exists a simple closed curve ρ in $\mathbb{S}^2$ such that ρ consists of an edge e of $\mathcal{G}$ together with an arc ρ' that only meets $\mathcal{G}$ at its endpoints and such that $\mathbb{S}^2 \setminus (\mathcal{G} \cup \rho')$ has no bigon regions. Add a new edge to $\mathcal{G}$ along ρ'. Repeat this as many times as possible. Since any two such arcs can be made disjoint on their interiors, the result of this process is well-defined given D.

Definition 3.1.1. Define $G^{\mathcal{F}}(D)$ to be the graph that results from this process.

Figure 3.1 shows a digram D_β of a link L_β together with the graph $G^{\mathcal{F}}(D_\beta)$.

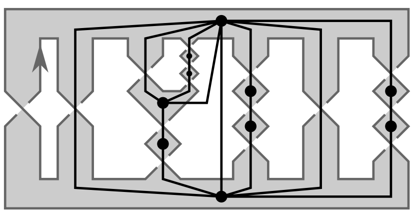

Figure 3.1

Definition 3.1.2. A *flype circle* ϕ is a simple closed curve in $\mathbb{S}^2$ that meets the link diagram D as shown in Figure 3.2a, where the tangles A and B each contain at least one crossing. The flype circle ϕ determines a flype that changes D to another special alternating diagram of L, as shown in Figure 3.2b. This change is realised by an isotopy of $\mathbb{S}^3$.

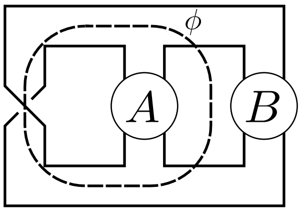

(a)

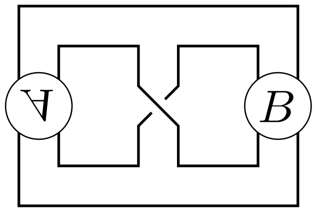

(b)

Figure 3.2

Let ϕ be a flype circle in D. Let D_ϕ be the diagram given by performing the flype defined by ϕ on D. Let R, R_ϕ be the surfaces given by applying Seifert's algorithm to D, D_ϕ respectively. Finally, let R' be the preimage of R_ϕ under the flype.

Definition 3.1.3. The flype circle ϕ is called *inessential* if R' is ambient isotopic to R, and *essential* otherwise.

If ϕ lies in the white regions of the diagram where it meets the crossing of D then ϕ is inessential. From now on we will generally ignore all such flype circles. That is, we assume that ϕ lies on the surface R where it meets the crossing.

Definition 3.1.4. We call the crossing through which ϕ passes the *flype crossing*, and the arc of ϕ disjoint from R the *flype arc.*

Recall that white bigons in a special alternating diagram signify plumbed on Hopf bands in the Seifert surface (see Lemma 1.6.11). Because the Hopf link is fibred, a flype that interchanges the two crossings of a white bigon is inessential. Thus ϕ is inessential if the flype crossing and flype arc of ϕ differ only by a line of white bigons. We will see later that the converse is also true.

If ϕ is essential, the flype crossing and flype arc correspond to distinct edges e, e' of $G^{\mathcal{F}}(D)$ with the same endpoints. We will also refer to the simple closed curve $e \cup e'$ as the flype circle ϕ, provided we retain knowledge of which edge is the flype crossing and which the flype arc. Note that distinct flype circles ϕ and ϕ' in D may give the same flype circle in $G^{\mathcal{F}}(D)$. In this case, the flype crossings of ϕ and ϕ' differ by at most a line of white bigons, as do the flype arcs. Hence the diagrams D_ϕ and $D_{\phi'}$ are the same.

Definition 3.1.5. Call such flype circles *equivalent.*

Definition 3.1.6. Let $v, v' \in \mathrm{V}(G^{\mathcal{F}}(D))$ such that v, v' are joined by at least two edges of $G^{\mathcal{F}}(D)$. Then the subgraph $G^{\mathcal{F}}(D)[\{v, v'\}]$ is called a *θ–graph*. For an edge e of a θ–graph, denote by e^{θ} the θ–graph containing e. Define $\theta(D)$ to be the subgraph of $G^{\mathcal{F}}(D)$ that is the union of all θ–graphs in $G^{\mathcal{F}}(D)$.

Definition 3.1.7. Each vertex of $G^{\mathcal{F}}(D)$ inherits an orientation (clockwise or anticlockwise) from the Seifert circle it lies inside. This orientation extends to a transverse orientation of each edge of $G^{\mathcal{F}}(D)$. We define the *positive* and *negative sides* of each edge, such that the normal points from the negative side to the positive side. See Figure 3.3. We similarly define the *positive* and *negative sides* of each crossing in D.

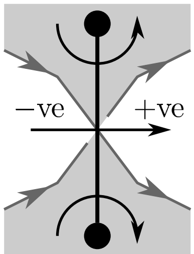

Figure 3.3

Definition 3.1.8. Given an essential flype circle ϕ in D, define the *positive side of* ϕ to be the component of $\mathbb{S}^2 \setminus \phi$ that meets the positive side of the flype crossing of ϕ. The other component is called the *negative side*.

Each edge $e \in \mathrm{E}(\theta(D))$ inherits a weight $w(e) \in \mathbb{N} \cup \{0\}$, given by the number of crossings in D that correspond to e. Define $w^{\theta}(e^{\theta}) = \sum_{e' \in e^{\theta}} w(e')$. Number the edges of $\theta(D)$ as $e_1, \ldots, e_n$. These weights are used to define a simplicial complex $\mathcal{K}(D)$. First, if $n \geq 1$, set

$$\mathrm{V}(\mathcal{K}(D)) = \left\{ (w_1, \ldots, w_n) \in \mathbb{Z}^n_{\geq 0} : \sum_{\{j : e_j \in e_i^{\theta}\}} w_j = w^{\theta}(e_i^{\theta}) \text{ for } 1 \leq i \leq n \right\}.$$

If D contains no θ–graphs, and so $\theta(D)$ is empty, take $\mathcal{K}(D)$ to consist of a single vertex.

Definition 3.1.9. Let $\tilde{\theta}(D)$ be the planar graph obtained from $\theta(D)$ by cutting apart vertices to make the θ–graphs disjoint.

Convention 3.1.10. Note that a region of $\mathbb{S}^2 \setminus \tilde{\theta}(D)$ corresponds to a union of regions of $\mathbb{S}^2 \setminus \theta(D)$. We will refer to these as the *regions of* $\theta(D)$.

Definition 3.1.11. Let r be a region of $\theta(D)$. Define the *positive boundary* $\partial^+ r$ of r to be the edges of $\theta(D)$ which r meets exactly on the negative side, and the *negative boundary* $\partial^- r$ to be those it meets exactly on the positive side.

Using this, the region r defines a map r^θ from a subset of $\mathrm{V}(\mathcal{K}(D))$ to $\mathrm{V}(\mathcal{K}(D))$, given by $r^\theta(w_1,\ldots,w_n) = (w'_1,\ldots,w'_n)$, where

$$w'_i = \begin{cases} w_i + 1 & \text{if } e_i \in \partial^+ r \\ w_i - 1 & \text{if } e_i \in \partial^- r \\ w_i & \text{else.} \end{cases}$$

Thus $r^\theta(w_1,\ldots,w_n)$ is defined when $w_i > 0$ for all $e_i \in \partial^- r$.

We describe the process of applying r^θ to $v \in \mathrm{V}(\mathcal{K}(D))$ as *adding the region r to v*.

The higher-dimensional simplices of $\mathcal{K}(D)$ are defined as follows. A set of distinct vertices $v_0,\ldots,v_k$ spans a k–simplex if there is an ordering $r_1,\ldots,r_m$ of the regions of $\theta(D)$ such that $r^\theta_j(\cdots r^\theta_1(v_0)\cdots)$ is defined for $1 \le j \le m$ and each v_i occurs as $r^\theta_j(\cdots r^\theta_1(v_0)\cdots)$ for some j. That is, every simplex can be extended to an $(m-1)$–simplex, and we can cycle through the vertices of an $(m-1)$–simplex by adding each region once, in some order. Note that $r^\theta_m(\cdots r^\theta_1(v_0)\cdots) = v_0$.

A metric is defined on $\mathrm{V}(\mathcal{K}(D))$ in the same way as that defined on $\mathrm{MS}(L)$. That is, the distance between two vertices is the distance in the 1–skeleton of $\mathcal{K}(D)$ when every edge has length 1.

Note that $\mathcal{K}(D)$ depends only on $\theta(D)$. A vertex v of $\mathcal{K}(D)$ specifies a reduced, special alternating diagram D_v of L, to which Seifert's algorithm can be applied. D_v is obtained from D by a canonical set of flypes. A flype on D can result in either of two diagrams, which differ by turning over the whole diagram. However, the surface given by Seifert's algorithm is independent of this choice. This gives a map $\mathcal{A}$ from $\mathrm{V}(\mathcal{K}(D))$ to $\mathrm{V}(\mathrm{MS}(L))$.

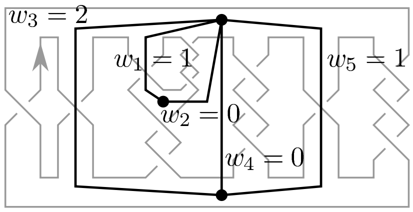

Figure 3.4

Figure 3.4 shows the graph $\theta(D_\beta)$ for the diagram D_β in Figure 3.1, marked with the weights that come from D_β. The complex $\mathcal{K}(D_\beta)$ is 3–dimensional, and this vertex lies in four 3–simplices, corresponding to the cycles

$$(1,0,2,0,1) \to (1,0,3,0,0) \to (0,1,3,0,0) \to (1,0,2,1,0) \to (1,0,2,0,1),$$

$$(1,0,2,0,1) \to (0,1,2,0,1) \to (0,1,3,0,0) \to (1,0,2,1,0) \to (1,0,2,0,1),$$

$$(1,0,2,0,1) \to (0,1,2,0,1) \to (1,0,1,1,1) \to (1,0,2,1,0) \to (1,0,2,0,1),$$

$$(1,0,2,0,1) \to (0,1,2,0,1) \to (1,0,1,1,1) \to (1,0,1,0,2) \to (1,0,2,0,1).$$

3.2 Special form

3.2.1 Definition

Let L be a special alternating link with a special alternating diagram $D \subset \mathbb{S}^2 \subset \mathbb{S}^3$.

Definition 3.2.1. Let $\mathbb{B}_a$ be the 3–ball lying above $\mathbb{S}^2$ in $\mathbb{S}^3$, $\mathbb{B}_b$ the 3–ball lying below $\mathbb{S}^2$, and P the set of midpoints of edges of D. We call a section of L between two consecutive points of P an *arc of* L.

By an isotopy, L can be arranged such that $L \cap \mathbb{S}^2 = P$, with overcrossing arcs of L lying in $\mathbb{B}_a$ and undercrossing arcs lying in $\mathbb{B}_b$ (as in [13]). Let R be the Seifert surface given by applying Seifert's algorithm to D.

Consider an incompressible Seifert surface R' for L that is disjoint from R. By an isotopy, $\partial R'$ can be made to run along the opposite side of $\partial \mathcal{N}(L)$ to ∂R (again as in [13]). In a neighbourhood of a crossing, ∂R and $\partial R'$ are as shown in Figure 3.5.

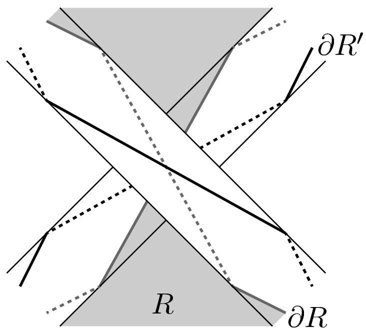

Figure 3.5

Now isotope R', keeping $\partial R'$ fixed and keeping R' disjoint from R, to minimise $|R' \cap \mathbb{S}^2|$. Note that $R' \cap \mathbb{S}^2$ is disjoint from the black regions of D. Since R' is incompressible and $\mathbb{S}^3 \setminus \mathcal{N}(R)$ is irreducible, $R' \cap \mathbb{S}^2$ contains no closed components, so it consists of arcs with their endpoints on $\partial R' \cap \mathbb{S}^2$. Similarly, no such arc has both endpoints at the same point of $R' \cap \mathbb{S}^2$. On the other hand, every point of $R' \cap \mathbb{S}^2$ is an endpoint of at least one arc. We will identify these points with P.

Suppose there exists $p \in P$ such that at least two arcs of $R' \cap \mathbb{S}^2$ have an endpoint at p. Then R' can be isotoped to reduce the number of endpoints of arcs at p by 2. This reduces the number of arcs of $R' \cap \mathbb{S}^2$ and any closed curve created can be removed, which contradicts the minimality of $|R' \cap \mathbb{S}^2|$. Hence each point $p \in P$ is the endpoint of exactly one arc.

Definition 3.2.2. Call a set of disjoint arcs lying in the white regions of D with exactly one arc endpoint at each point of P a *set of P–arcs*.

Consider the (probably disconnected) surface $R'_a = R' \cap \mathbb{B}_a$. Its boundary $\partial R'_a$ projects to disjoint simple closed curves in $\mathbb{S}^2$, together composed exactly of the set of P–arcs and the overcrossing arcs of L. Unless L is the unknot, each component includes both overcrossings and P–arcs, alternating around the curve.

We wish to show that every component of R'_a is a disc. To see this, first isotope each component downwards so that its boundary lies in $\mathbb{S}^2$. Now choose a closed curve C in $\partial R'_a$ that is innermost in $\mathbb{S}^2$. It bounds a disc in $\mathbb{B}_a$ that is otherwise disjoint from R'_a, and so bounds a disc in R'. This disc cannot lie below C in R' since part of C lies on $\partial R'$, so it lies above C. If the interior of the disc meets $\mathbb{S}^2$, then it also meets $\partial R'$, which cannot be the case. Thus the disc lies in $\mathbb{B}_a$. By isotoping it down towards $\mathbb{S}^2$, we see that for our current argument we may discard this component of R'_a. Thus inducting on $|\partial R'_a|$ gives the required result.

A similar argument holds for $R'_b = R' \cap \mathbb{B}_b = R' \setminus \text{int}_{R'}(R'_a)$. Hence R' is completely specified by the set of P–arcs $R' \cap \mathbb{S}^2$ added to the diagram D.

Definition 3.2.3. We will say that a surface defined in this way by a set of P–arcs is in *special form.*

Suppose D has n crossings. Then there are n P–arcs. Let n_a be the number of discs in R'_a, and n_b the number in R'_b. Then $\chi(R') = 2n - (2n + n) + (n_a + n_b) = -n + n_a + n_b$.

Note that $n_a = |R' \cap \mathbb{B}_a| = |\partial(R' \cap \mathbb{B}_a)|$. Hence R' is taut if and only if $|\partial(R' \cap \mathbb{B}_a)| + |\partial(R' \cap \mathbb{B}_b)|$ (that is, the number of simple closed curves formed from the P–arcs and the arcs of L) is maximised.

3.2.2 Some surfaces in special form

Fix a special alternating diagram D of a link L. Let R be the Seifert surface given by applying Seifert's algorithm to D. We can define a Seifert surface R' in special form by putting a P–arc across the negative side of each crossing, as show in Figure 3.6. By pushing

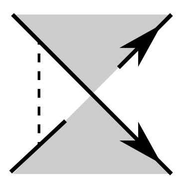

Figure 3.6

each P–arc close to the corresponding crossing, it is easy to see that $R' = R$. Placing a P–arc on the positive side of each crossing again gives a surface equivalent to R. It is given by pushing R' through the parallel surface R to the other side.

Figure 3.7 shows two diagrams D_γ and D_γ^* of a link L_γ that differ by a flype along the flype circle ϕ shown. Let R_γ, R_γ^* be the surfaces given by applying Seifert's algorithm

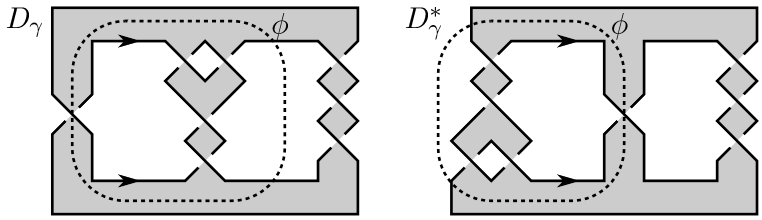

Figure 3.7

to D_γ, D_γ^* respectively. Consider the effect on R_γ^* of an isotopy that takes D_γ^* to D_γ. By inspection, R_γ^* can be put into special form with respect to D_γ as shown in Figure 3.8a. Figure 3.8b shows $\partial(R_\gamma^* \cap \mathbb{B}_a)$.

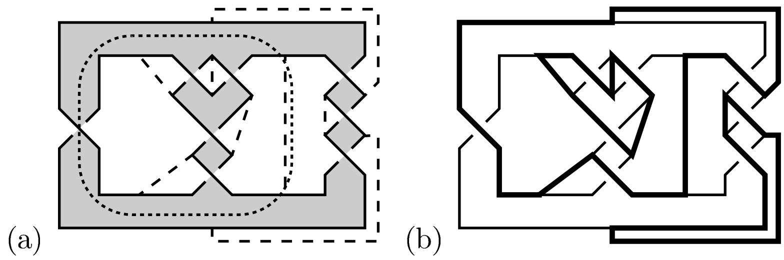

Figure 3.8

The special form of R_γ^* has the following description. The flype arc of ϕ is a P–arc. Every crossing $c \neq c_\phi$ (where c_ϕ is the flype crossing of ϕ) has a P–arc across it. This P–arc crosses c on its negative side if c is on the positive side of ϕ, and on its positive side if c lies on the negative side of ϕ.

We can extend this description to the case of more than one flype circle on a link diagram D, provided we can consistently define the notion of being on the positive/negative side of the flype circles.

Definition 3.2.4. Call a set of distinct flype circles with this property *coherent*. Otherwise call them *incoherent*.

Remark 3.2.5. Coherent flype circles must be pairwise disjoint. Figure 3.9 shows a set of incoherent disjoint flype circles on the diagram D_β.

Using the above description, a set of coherent flype circles gives a set of arcs in the white regions of D. We need to check whether they form a set of P–arcs.

Lemma 3.2.6. *The arcs defined by a coherent set of flype circles form a set of P–arcs.*

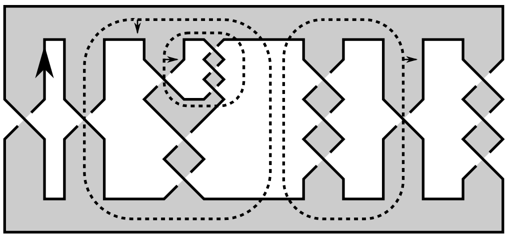

Figure 3.9

Proof. This will be true if and only if exactly one endpoint of an arc lies on any given edge ε of D. By counting the arcs, it is enough to check that at least one arc has an endpoint on ε.

If a flype circle crosses ε, then an arc has an endpoint on ε. Assume otherwise. Let c_1, c_2 be the crossings at the ends of ε such that ε lies on the positive side of c_1 and on the negative side of c_2. Then both, one or neither of c_1, c_2 is a flype crossing. Since ε is not crossed by any flype circle, c_1 and c_2 cannot both be flype crossings, as otherwise the flype circles would be incoherent (see Figure 3.10). Suppose neither is a flype crossing. They

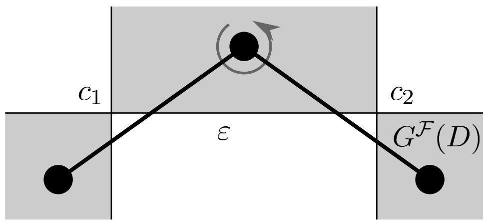

Figure 3.10

then lie on the same side of the flype circles, say the negative side. Then the arc across the positive side of c_1 has an endpoint on ε. Finally suppose, without loss of generality, that c_1 is a flype crossing, while c_2 is not. Then c_2 lies on the positive side of the flype circles, so has an arc across it on the negative side. This arc has an endpoint on ε. ☐

Definition 3.2.7. Given a special alternating link diagram D, say that a Seifert surface R for D is *admissible at* D if R can be put into special form with the following description.

- There is a set of coherent flype circles in D.
- The flype arc of each flype circle is a P–arc.
- Every crossing of D that is not a flype crossing has a P–arc across it, on the negative side if it lies on the positive side of the flype circles, and on the positive side if it lies on the negative side of the flype circles.

Call such a description of R *admissible*.

Remark 3.2.8. We allow the case where there are no flype circles. We can then take the crossings of D to all lie on the positive side of the flype circles, or all lie on the negative side. As we have seen, both these special forms describe the surface given by applying Seifert's algorithm to D.

To construct a set of P–arcs, it is in fact sufficient to specify a coherent set of flype circles in $\theta(D)$, provided the edge of each such flype circle chosen as a flype crossing has at least one crossing in D corresponding to it. To see this, notice by inspection that equivalent flype circles create the same set of P–arcs (see Figure 3.11).

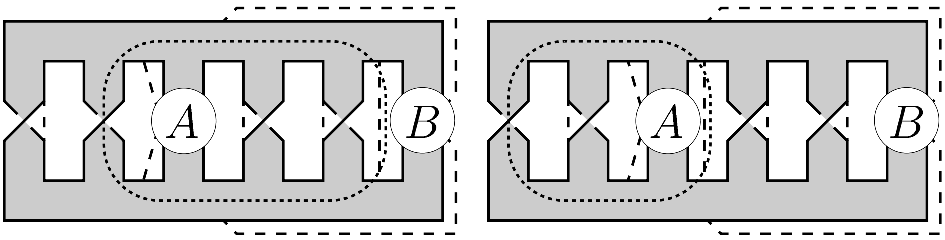

Figure 3.11

Lemma 3.2.9. *Let R' be a surface in admissible special form at a special alternating diagram D. Let D' be the diagram given by applying all the flypes indicated by the flype circles. Then R' is the surface given by applying Seifert's algorithm to D'.*

Proof. We prove this by induction on the number of flype circles. If there are no flype circles in the admissible special form then the result holds.

Suppose there is at least one flype circle, and choose an innermost flype circle ϕ. Then D has the form shown in Figure 3.12a, where A contains no flype circles. The flype ϕ changes the diagram to that in Figure 3.12b. This also gives R' in admissible form. By induction, the result follows. ☐

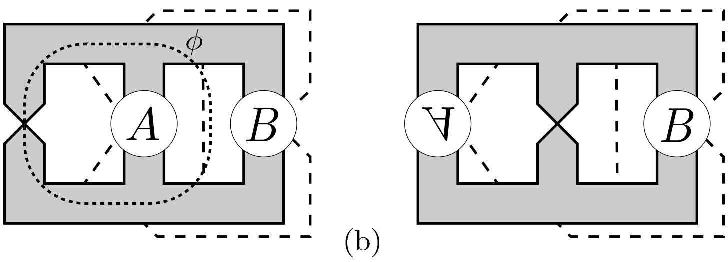

(a) (b)

Figure 3.12

3.2.3 Relating $\mathcal{K}(D)$ and $\mathrm{MS}(L)$

Our aim is to relate $\mathcal{K}(D)$ to $\mathrm{MS}(L)$. We have a map $\mathcal{A}\colon \mathrm{V}(\mathcal{K}(D)) \to \mathrm{V}(\mathrm{MS}(L))$, as described in Section 3.1. To study the properties of this map, we will establish a local description at each vertex of $\mathcal{K}(D)$.

Proposition 3.2.10. *Let $v_0 \in \mathrm{V}(\mathcal{K}(D))$. As above, v_0 corresponds to a prime, reduced, special alternating diagram D_{v_0}, and $R_0 = \mathcal{A}(v_0)$ is given by applying Seifert's algorithm to D_{v_0}. Without loss of generality, we may assume $D = D_{v_0}$.*

Let $v_1 \in \mathrm{V}(\mathcal{K}(D))$ with $\mathrm{d}_{\mathcal{K}(D)}(v_0, v_1) = 1$. Then $\mathrm{d}_{\mathrm{MS}(L)}(R_0, \mathcal{A}(v_1)) \leq 1$.

Proof. Recall, for $R, R' \in \mathrm{V}(\mathrm{MS}(L))$, that $\mathrm{d}_{\mathrm{MS}(L)}(R, R') \leq 1$ if and only if R and R' can be made disjoint.

Since $\mathrm{d}_{\mathcal{K}(D)}(v_0, v_1) = 1$, vertex v_1 is obtained from v_0 by adding a sequence of distinct regions of $\theta(D)$. Let Λ be the union of these regions. Then $\partial\Lambda$ is split into its positive boundary $\partial^+\Lambda$ and its negative boundary $\partial^-\Lambda$. The change from v_0 to v_1 is achieved by subtracting 1 from the weight of all edges in $\partial^-\Lambda$ and adding one to the weight of all edges in $\partial^+\Lambda$.

Consider a θ–graph in $\theta(D)$, with each edge labelled by the effect of adding Λ (see, for example, Figure 3.13). The total of these labels is 0, and the 1s and -1s must alternate.

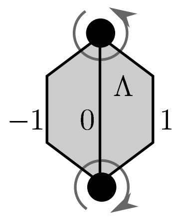

Figure 3.13

Pair each -1 with the 1 on its positive side. Each such pair defines a flype circle, whose crossing circle is the edge with label -1, and whose flype arc is the edge with label 1. Note that this is always possible, because an edge with label -1 must have at least one crossing in D corresponding to it. Doing this for each θ–graph gives a set of disjoint flype circles in D. Note also that these flype circles are coherent.

$\mathcal{A}(v_1)$ is given by applying these flypes to D, yielding a diagram D_1, and then applying Seifert's algorithm to D_1. By Lemma 3.2.9, $\mathcal{A}(v_1)$ is the surface described in admissible special form by this set of flype circles on D. Hence $\mathcal{A}(v_1)$ can be made disjoint from R_0. ☐

Proposition 3.2.11. *Let v_0 and R_0 be as in Proposition 3.2.10. Let R_1 be a surface given in admissible special form at D. Then there exists $v_1 \in \mathrm{V}(\mathcal{K}(D))$ such that $\mathrm{d}_{\mathcal{K}(D)}(v_0, v_1) \leq 1$ and $\mathcal{A}(v_1) = R_1$.*

Proof. Choose a set of flype circles showing that R_1 is in admissible special form. If there are no flype circles then $R_1 = R_0$ and $v_1 = v_0$.

Assume there is at least one flype circle. Each flype circle gives a weight of -1 to one edge of $\theta(D)$ and a weight of 1 to another edge. Modifying v_0 using these values gives another vector. Call this vector v_1. We aim to show that v_1 has the required properties.

Consider the set of flype circles in $\theta(D)$. Note that a flype circle cannot run over an edge of $\theta(D)$ twice. Suppose that two or more flype circles run over an edge $e \in \mathrm{E}(\theta(D))$. Choose two such flype circles ϕ, ϕ' that are adjacent in e. Since the flype circles are coherent, one must give e a weighting of 1 while the other gives it a weighting of -1. That is, e is the flype crossing for ϕ, say, and the flype arc for ϕ'. By combining ϕ and ϕ' as shown in Figure 3.14 to give a new flype circle ϕ'', we can reduce the number of flype circles running over e without changing the admissible form or the vector v_1. Hence we may assume that at most one flype circle runs over any edge of $\theta(D)$.

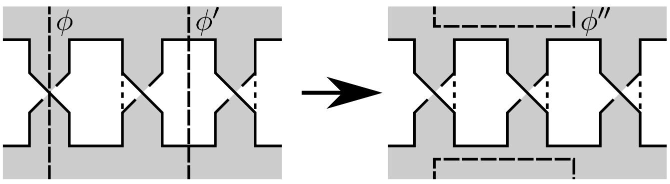

Figure 3.14

Let $\Lambda \subseteq \mathbb{S}^2$ be the positive side of the flype circles, and $\lambda \subseteq \mathbb{S}^2$ the negative side. If e receives a weighting of 1 then a flype arc runs along e, and Λ lies on the negative side of e. If an edge e receives a weighting of -1 then e corresponds to a flype crossing, and Λ lies on the positive side of e. If e receives a weighting of 0, no flype circle runs along it, and it lies in the interior of either Λ or λ.

Each of Λ and λ is a non-empty union of regions of $\theta(D)$. We wish to show that those in Λ can be ordered such that each partial composition of the corresponding maps is defined for v_0. By symmetry, those in λ can then also be suitably ordered so that the regions in Λ followed by those in λ define a simplex in $\mathcal{K}(D)$. Then $v_1 \in \mathcal{K}(D)$ and $\mathrm{d}_{\mathcal{K}(D)}(v_0, v_1) = 1$. It is sufficient to find a single region that can be added to v_0. Induction on the number of regions in Λ then completes the proof.

Suppose no such region exists. Then each region r in Λ has a boundary edge $e(r)$ such that r lies on the positive side of $e(r)$ and $w(e(r)) = 0$ (in v_0). This means $e(r)$ cannot lie on $\partial\Lambda$, so the region on the other side of $e(r)$ from r is also in Λ. Construct a digraph $\mathcal{G}_\Lambda$ with an edge from r across $e(r)$ for each r in Λ. There are only finitely many regions in Λ, so $\mathcal{G}_\Lambda$ contains a simple closed curve ρ. Then there is a θ–graph in $\theta(D)$ such that ρ

runs through every edge of the θ–graph. Hence every edge of this θ–graph has value 0 in v_0. This contradicts the construction of $\theta(D)$. □

Lemma 3.2.12. *The complex $\mathcal{K}(D)$ is flag.*

Proof. Let $v_0, \ldots, v_m$ be distinct vertices of $\mathcal{K}(D)$ that are pairwise adjacent, for some $m \geq 2$. For $1 \leq i \leq m$, there is a unique set A_i of regions of $\theta(D)$ such that v_i is obtained from v_0 by adding the regions in A_i in some order. Let $A_0 = \emptyset$. Then there is a partial order on the vertices $v_0, \ldots, v_m$ given by inclusion of the sets $A_0, \ldots, A_m$. We wish to show that $v_0, \ldots, v_m$ span a simplex in $\mathcal{K}(D)$. From the proof of Proposition 3.2.11 we see that it is sufficient to prove that this partial ordering of the vertices is actually a total order.

Suppose otherwise. Then there are two vertices that are not comparable. Without loss of generality, these are v_1 and v_2. Thus $A_1 \setminus A_2$ and $A_2 \setminus A_1$ are non-empty. Let Λ_1 be the union of the regions in $A_1 \setminus A_2$, and Λ_2 the union of the regions in $A_2 \setminus A_1$. Then Λ_1 and Λ_2 have disjoint interiors.

Suppose an edge e of $\theta(D)$ lies on the boundary of Λ_1 and on the boundary of Λ_2. Then it lies on the positive boundary of one of Λ_1, Λ_2 and on the negative boundary of the other. Thus the coordinates of v_1 and v_2 corresponding to the edge e differ by two. This contradicts that v_1, v_2 are adjacent.

As v_1 is adjacent to v_2 in $\mathcal{K}(D)$, there is a unique set B of regions of $\theta(D)$ such that v_2 is obtained from v_1 by adding the regions in B in some order. Let Λ be the union of the regions in B. The boundary of Λ is the union of the boundaries of Λ_1 and Λ_2, as these are the edges of $\theta(D)$ at which the coordinates of v_1 and v_2 differ. However, Λ lies on the outside of $\partial\Lambda_1$ (that is, on the side away from Λ_1) and on the inside of $\partial\Lambda_2$, which is impossible. This gives the required contradiction. □

Remark 3.2.13. Since MS(L) is also flag (by definition), each of the two complexes is defined by its 1–skeleton, so in relating them we may restrict our attention to comparing vertices and edges.

3.3 Putting surfaces into admissible special form

Consider a surface in special form. Pick a white region of the diagram. At least one P–arc in this region will run across a crossing. Following Gabai ([13]), we can cut both the diagram and the surface along this crossing, giving a Seifert surface for a new link with fewer crossings than the original. The effect on the diagram and P–arcs is as shown in Figure 3.15. This gives the basis for an inductive argument.

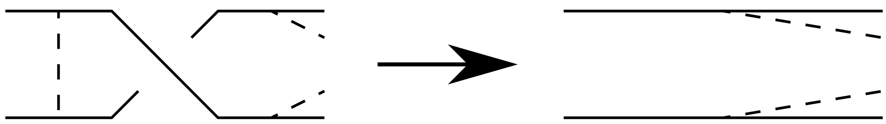

Figure 3.15

Our inductive hypothesis in the proof of Proposition 3.3.8 will relate to prime links. It is therefore useful to find a white region such that cutting along any crossing of that region will result in a prime link diagram.

Definition 3.3.1. Say that such a region is *cuttable*.

Lemma 3.3.2. *If D is a prime, reduced, special alternating diagram of a link L, then there is a cuttable white region of D.*

Proof. Suppose otherwise. Choose a white region of D. Then this region has a crossing c such that cutting along c gives a connected sum. Since the original diagram was prime, D has the form shown in Figure 3.16a, where neither A nor B is trivial or consists of a single

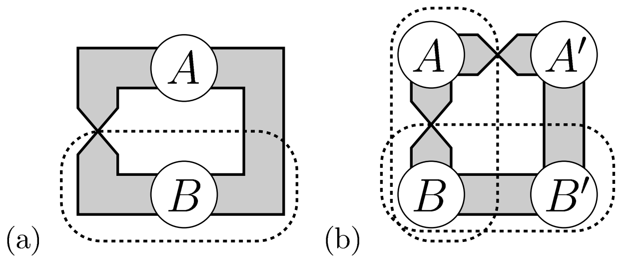

Figure 3.16

line of black bigons. In particular, A and B must each contain a white region. In this way, we can choose a simple closed curve ρ_r for each white region r. We will allow two white regions to share one suitable curve. Suppose, for regions r, r', the curves $\rho_r, \rho_{r'}$ are distinct and cannot be isotoped to be disjoint. Then D has the form shown in Figure 3.16b. In this case, we could have chosen ρ_r for both the regions r and r'. Thus the simple closed curves can be chosen to be disjoint. Choose a region r such that ρ_r is innermost. Then no white region r' lies entirely inside ρ_r. This contradicts the choice of ρ_r. □

3.3.1 Reducing a diagram

In this section we consider some specific situations that will arise in the proof of Proposition 3.3.8.

Let D, D^c be special alternating link diagrams such that D^c is obtained from D by removing a crossing c by a single type I Reidemeister move. Let R be an incompressible

Seifert surface given in special form at D. What changes need to be made to give R in special form at D^c?

Any P–arcs that are distant from c can be copied to D^c. We consider those close to c. The P–arcs around c are as shown in Figure 3.17a. Let p be the point of D^c corresponding to c in D. One way to remove c is to pull the undercrossing arc upwards into $\mathbb{B}_a$ and then

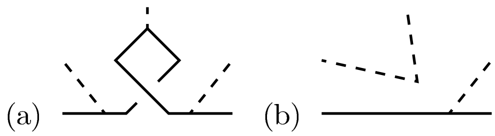

Figure 3.17

untwist. This leaves $R \cap \mathbb{S}^2$ as shown in Figure 3.17b, giving a set of P–arcs on D^c, possibly with a single simple closed curve lying within the white region adjacent to p. Does this give a special form for R? One of the discs of $R \cap \mathbb{B}_b$ has been moved, with a subdisc S along part of its boundary being pulled up into $\mathbb{B}_a$. Clearly this still leaves a disc in $\mathbb{B}_b$. The disc S is glued onto discs of $R \cap \mathbb{B}_a$ along two disjoint subarcs of its boundary. If these two arcs are glued to distinct discs in $\mathbb{B}_a$, this again gives a disc, and the P–arcs give a special form of R. Alternatively, the two arcs may be glued to the same disc, in which case an annulus T is formed. The core curve of T bounds a disc in $\mathbb{B}_a$ disjoint from R. As R is incompressible, one component of ∂T bounds a disc in R disjoint from the other. This component cannot run along the link, so must be a simple closed curve ρ of $R \cap \mathbb{S}^2$. In this case, the P–arcs around c in D must be as shown in Figure 3.18. The P–arcs in D^c, without ρ, give a special form for R.

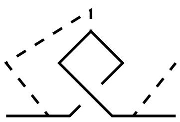

Figure 3.18

We could instead have removed c by first pushing the overcrossing arc downwards into $\mathbb{B}_b$. Similar reasoning applies in this case. If no annulus is formed in either case then around c the simple closed curves bounding the discs of R in $\mathbb{B}_a$ and $\mathbb{B}_b$, which are made up of P–arcs and arcs of D, must be connected as shown in Figure 3.19.

Definition 3.3.3. Call the change of D to D^c given by pushing c into $\mathbb{B}_a$ and untwisting an *a–reduction at c*, and say that we *a–reduce* D. Similarly define a *b–reduction* and *b–reducing.*

Remark 3.3.4. If there is a P–arc across c in D then a–reducing D at c and b–reducing it give the same result, and do not affect any of the other P–arcs.

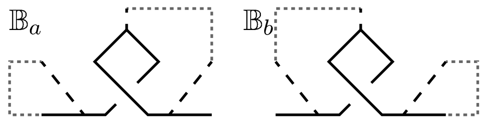

Figure 3.19

For the rest of this section we assume D^c is the a–reduction of D.

Lemma 3.3.5. *Suppose D and D^c are prime, and D^c is reduced. Suppose further that the special form for R at D^c is admissible (and hence R is taut). Then there is a P–arc in D across c.*

Proof. Consider the flype circles in the admissible special form. Let D_0^c be the diagram obtained from D^c by performing all the flypes so as to leave p fixed. By Lemma 3.2.9, the surface given by the P–arcs on D_0^c is R, and in particular is taut.

This process of changing the diagram by a flype ϕ moves a P–arc only if it lies strictly inside ϕ on the side that is moved. Since p was not moved by any flype, the P–arcs in D^c that came from those adjacent to c in D were also not moved. This means we can perform the same changes to the diagram D with its P–arcs as were made to D^c, to give a diagram D_0. These changes will correspond to performing flypes on D, but these flypes may not come from flype circles in the special form on D. However, we see that D_0^c is obtained from D_0 by untwisting c, and the P–arcs in D_0^c are those given by a–reducing D_0 at c. Hence we may assume that the admissible special form of R at D^c has no flype circles.

If either a–reducing or b–reducing D at c creates an annulus as described above, there is a P–arc across c in D. Hence we may assume no annulus is formed. We use this to build up a picture of part of D^c, and so derive a contradiction.

From above (Figure 3.19), in $\mathbb{B}_a$ the P–arcs on D^c near p are connected as shown in Figure 3.20. Since there is no P–arc across c in D, σ_1 cannot consist entirely of P–arcs.

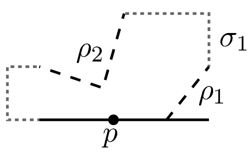

Figure 3.20

Then, since there are no flype circles in the admissible special form on D^c, ρ_1 and ρ_2 each run across a single crossing. From this and the fact that D^c is reduced we have Figure 3.21. There is a P–arc across c_1, so a neighbourhood of c_1 is as shown in Figure 3.22a. Consider the next crossing c_2 around the black region r that σ_2 meets. Again, there is a P–arc across c_2, so we have the set-up shown in Figure 3.22b. Inductively we see that the arc σ_2 will

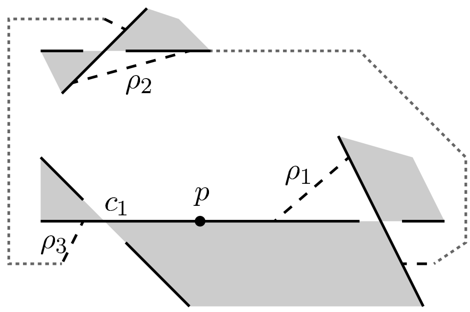

Figure 3.21

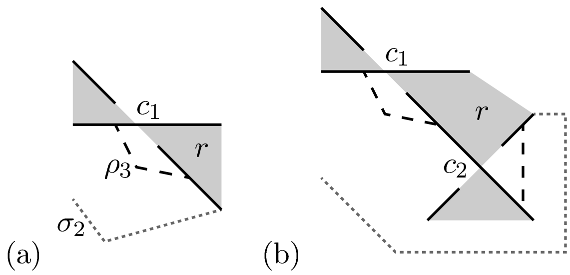

Figure 3.22

return to ∂r after every crossing. Thus D^c has the structure shown in Figure 3.23. Because

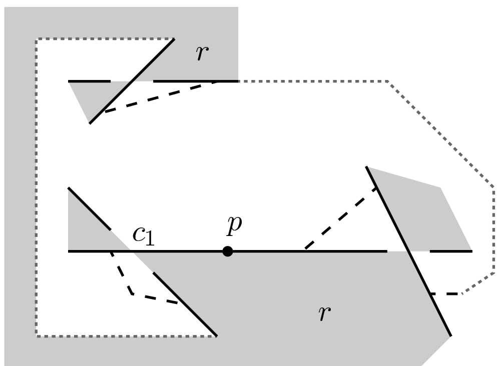

Figure 3.23

D^c is reduced and alternating, this contradicts that it is prime. □

Lemma 3.3.6. *Suppose there is a reduced diagram $D^{c,c'}$ given by a–reducing D^c at a crossing c', where p does not lie on the edge of D^c connecting c' to itself. Suppose further that all the diagrams are prime and the special form of R at $D^{c,c'}$ is admissible. Then in D there are P–arcs across each of c and c'.*

Proof. By Lemma 3.3.5 there is a P–arc across c' in D^c. Suppose there is no P–arc across c' in D. Then c and c' lie on the same white region of D, and the P–arcs around them are connected in one of the patterns shown in Figure 3.24. In the first case, the P–arcs around

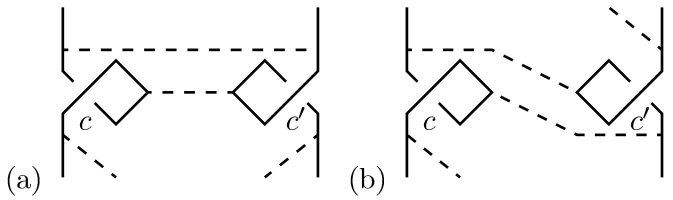

Figure 3.24

c' are not connected as in Figure 3.19. Therefore the situation in Figure 3.24b occurs.

Consider the set of flype circles in $D^{c,c'}$ that show that the special form of R is admissible. None of these can separate c and c' in D. Thus, as in the proof of Lemma 3.3.5, we may assume that there are no flype circles needed and every crossing in D has a P–arc across it. Further following the proof of Lemma 3.3.5 shows that D has the structure shown in Figure 3.25, contradicting that it is prime.

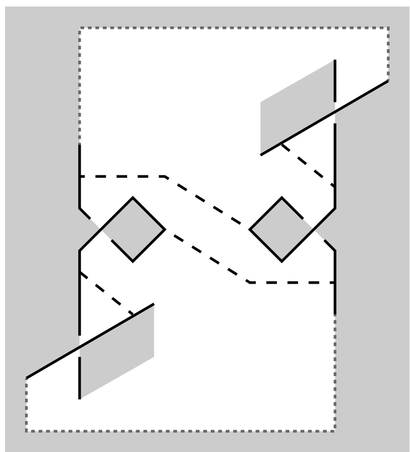

Figure 3.25

Thus there is a P–arc across c' in D. This means that untwisting c' has no impact on the rest of the picture, so Lemma 3.3.5 gives that there is a P–arc across c in D. □

Lemma 3.3.7. *Suppose p lies on an edge of D^c connecting a crossing c' to itself, so c' can be removed by a type I Reidemeister move. Suppose the special form of R at D^c has a P–arc across c'. Then there is a P–arc across c in D.*

Proof. Suppose otherwise. From above (Figure 3.19) the arcs of $R \cap \mathbb{B}_a$ on D must connect as in Figure 3.26a, where σ does not consist entirely of P–arcs. This contradicts that there is a P–arc across c' in D^c (see Figure 3.26b). □

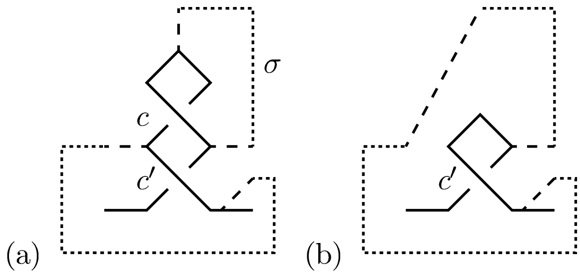

Figure 3.26

3.3.2 Adjacent surfaces can be put into admissible special form

Proposition 3.3.8. *Let D be a prime, reduced, special alternating diagram of a link L. Let R be the (taut) Seifert surface for L given by applying Seifert's algorithm to D. Let R' be a taut Seifert surface for D disjoint from R, given in special form. Then this special form is admissible.*

Proof. First suppose D has no crossings. Then L is the unknot, $R' = R$ and the special form of R' is admissible with no P–arcs or flype circles.

Now suppose the result holds for any diagram with at most $n-1$ crossings, and that D has n crossings. By Lemma 3.3.2, there is a cuttable region r of D. Choose a crossing c of r with a P–arc across it. Cut along this P–arc and c as in Figure 3.15 to give a new diagram D_c and a new taut Seifert surface R'_c in special form at D_c. Since r is cuttable, D_c is prime. It is also special and alternating, and has $n-1$ crossings. However, it may not be reduced.

Suppose D_c is reduced. Then the inductive hypothesis holds, so the special form of R'_c is admissible. Take a flype circle ϕ in D_c, and consider ϕ in D. The flype arc of ϕ is a P–arc in D_c, so also is in D. This means that c and the P–arc across it lie on one side of ϕ, so ϕ is a flype circle in D. Thus the flype circles in D_c give a set of flype circles in D. These flype circles inherit coherence from D_c. Suppose c does not lie between a crossing c' and the P–arc that crosses it in D_c. Then each crossing in D that is not a flype crossing or c has a P–arc across it on the required side. By counting the number of endpoints of P–arcs in each white region of D we find that the P–arc across c also lies on the required side. Thus the special form of R' is admissible. Instead suppose c does lie between a crossing c' and the P–arc that crosses it in D_c. We see that c and c' are as show in Figure 3.27. Again, the special form of R' at D is admissible.

Now suppose D_c is not reduced. Take a simple closed curve ρ that shows that D_c is not reduced. That is, ρ passes through a single crossing of D_c and otherwise lies in a white regions of D_c. Since D is reduced, ρ must pass through c in D, giving D the form shown

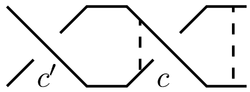

Figure 3.27

in Figure 3.28. As D is reduced and D_c is prime, either A or B consists of a single line

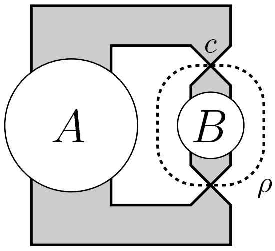

Figure 3.28

of crossings. Thus cutting along c leaves one or two lines of nugatory crossings joined by black bigons in a single white region of what is otherwise a reduced diagram (for example as shown in Figure 3.29).

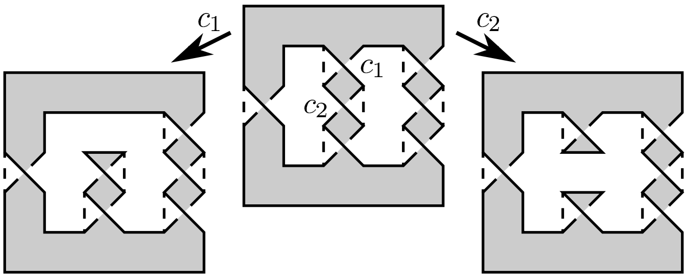

Figure 3.29

Repeatedly a–reduce the diagram until only one crossing from each of these lines remain, giving a new diagram D_1. Let D_0 be the diagram given by a–reducing the final one or two nugatory crossings of D_1. Then, by the inductive hypothesis, the special form of R' at D_0 is admissible, so the hypotheses of either Lemma 3.3.5 or Lemma 3.3.6 hold in relation to D_1. One of these lemmas followed by repeated application of Lemma 3.3.7 gives that the special form of R'_c at D_c has a P–arc across every nugatory crossing. Thus a–reducing D_c does not affect the other P–arcs, and every flype circle in D_0 is a flype circle in D.

Let c' be a crossing of D_0. If c' is a flype crossing in D_0 and in D then the flype arc of the flype circle is a P–arc in D. Suppose c' is not a flype crossing. Then there is a P–arc ρ across it on the required side in D_0. Suppose that ρ is not across c' in D. Then at least one of the nugatory crossings of D_1 lies between c' and ρ. This then means that in fact c

and all of the nugatory crossings of D_c lie between c' and ρ. Hence there is at most one crossing in D_0 for which this occurs.

If no such crossing exists, we may again check that, for each crossing of D that is not in D_0, the P–arc across it lies on the required side. Suppose there is such a crossing c_ϕ. Then c_ϕ and the P–arc ρ_ϕ across it in D_0 form a new flype circle ϕ in D. Without loss of generality, c lies on the positive side of ϕ. It is then easily checked that every crossing on the positive side of ϕ in D has a P–arc across it on the negative side. Since we know that ρ_ϕ lies on the positive side of c_ϕ in D_0, the crossing c_ϕ (and hence also the flype circle ϕ) lies on the negative side of the flype circles in D_0. Hence these together with ϕ give a set of flype circles in D proving that the special form of R' at D is admissible. □

Corollary 3.3.9 ([15] Theorem 1.1)**.** *Let L be a prime, special alternating link with a reduced, special alternating digram D. Let R be a taut Seifert surface for L. Then R is given by doing a finite sequence of flypes on D and then applying Seifert's algorithm to the resulting diagram.*

Proof. This is a combination of Proposition 3.3.8, Lemma 3.2.9 and the fact that MS(L) is connected. □

3.3.3 Incompressible Seifert surfaces that are not taut

In this section we show that the 'taut' hypothesis in Corollary 3.3.9 is crucial to our method of proof.

Let D be a special alternating diagram of a link L, and R the surface given by Seifert's algorithm. Let R' be a Seifert surface for L given in special form. That is, let R' be a surface defined by a set of P–arcs. Suppose R' is compressible, and let T be a compression disc for R', isotoped to have minimal intersection with $\mathbb{S}^2$. Since R is incompressible, T can also be made disjoint from R. Note that $T \cap \mathbb{S}^2 \neq \emptyset$ by the construction of R'. Then $T \cap \mathbb{S}^2$ consists of arcs in the white regions of D. As ∂T does not meet L, the endpoints of these arcs lie on the P–arcs. The interiors of these arcs are all disjoint from each other and from the P–arcs. Suppose one of these arcs has both endpoints on the same P–arc. An innermost such arc cobounds a disc in $\mathbb{S}^2$ with the P–arc. Using this, we can split T into two discs each with fewer arcs of intersection with $\mathbb{S}^2$. At least one of these will still be a compression disc for R'. We may therefore assume this does not occur.

Definition 3.3.10. Call an arc in a white region of the diagram, joining two distinct P–arcs, a *Q–arc*. Let ρ be a Q–arc. Let P_ρ be the operation on the special form of R' that changes the P–arcs along ρ, as shown in Figure 3.30.

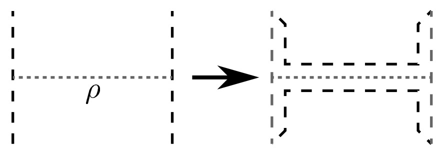

Figure 3.30

Lemma 3.3.11. *If $|T \cap \mathbb{S}^2| = 1$ then there is a Q–arc ρ with endpoints that lie on the boundary of the same disc in $\mathbb{B}_a$ and on the boundary of the same disc in $\mathbb{B}_b$, and T is defined by ρ.*

Proof. The arc $\rho = T \cap \mathbb{S}^2$ divides T into two subdiscs, $T_a = T \cap \mathbb{B}_a$ and $T_b = T \cap \mathbb{B}_b$. Then $\partial T_a \setminus \rho$ is an arc in the disjoint union of discs $R' \cap \mathbb{B}_a$. It is thus contained in a single such disc, and is unique given ρ, up to isotopy keeping its endpoints fixed. In particular, the ends of ρ lie on the same disc in $\mathbb{B}_a$. The same holds for T_b. Since $\mathbb{S}^3 \setminus \mathcal{N}(L)$ is irreducible, T is defined up to isotopy by ∂T. □

Lemma 3.3.12. *There is a finite sequence $\rho_1, \ldots, \rho_n$ of disjoint arcs that satisfy the following.*

- *The arc ρ_1 is a Q–arc on D.*
- *For $1 < m \leq n$, ρ_m is a Q–arc on $P_{\rho_{m-1}}(\cdots P_{\rho_1}(D)\cdots)$.*
- *For $1 \leq m < n$, the endpoints of ρ_m lie on the boundary of the same disc in exactly one of $\mathbb{B}_a$ and $\mathbb{B}_b$.*
- *The endpoints of ρ_n lie on the boundary of the same disc in $\mathbb{B}_a$ and on the boundary of the same disc in $\mathbb{B}_b$.*

Proof. We prove this by induction on $|T \cap \mathbb{S}^2|$, with the additional hypothesis that each ρ_m is an arc of $T \cap \mathbb{S}^2$. This ensures that the ρ_m are disjoint. The base case $|T \cap \mathbb{S}^2| = 1$ is given by Lemma 3.3.11. Suppose $|T \cap \mathbb{S}^2| > 1$. Let ρ_1 be an arc of $T \cap \mathbb{S}^2$ that is outermost in T. Without loss of generality, the disc T_a cut off by ρ_1 lies in $\mathbb{B}_a$. As above, the endpoints of ρ_1 lie on the boundary of the same disc of R' in $\mathbb{B}_a$.

The disc T_a defines an isotopy of R' in a neighbourhood of T_a that removes ρ_1 from $T \cap \mathbb{S}^2$. The map P_{ρ_1} records the change of $R' \cap \mathbb{S}^2$ that results from this isotopy. $P_{\rho_1}(D)$ is a special form for R' if and only if before the isotopy the endpoints of ρ_1 lie on the boundaries of different discs in $\mathbb{B}_b$. If this is not the case, the proof is complete directly. Otherwise, it is complete by induction. □

Proposition 3.3.13. *Let L_δ be the link shown in Figure 3.31, and let R'_δ be the Seifert surface for L_δ given by the special form shown. Then R'_δ is incompressible but not taut. In particular, there is no admissible special form for R'_δ.*

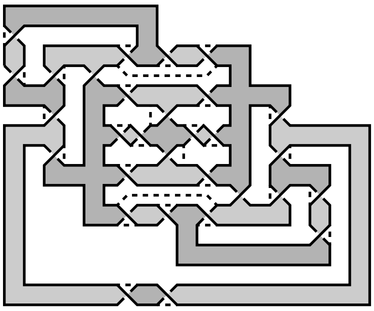

Figure 3.31

Proof. Figure 3.32 shows the boundaries of the discs $R'_\delta \cap \mathbb{B}_a$ and $R'_\delta \cap \mathbb{B}_b$. From this we see

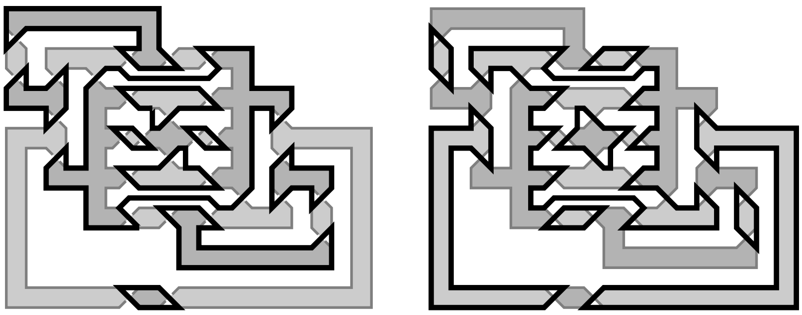

Figure 3.32

that $\chi(R'_\delta) = -12$ whereas $\chi(R_\delta) = -10$, where R_δ is the Seifert surface given by Seifert's algorithm. We also see that there are only two Q–arcs in this diagram that have endpoints on the same disc in either $\mathbb{B}_a$ or $\mathbb{B}_b$. By inspection, these do not satisfy the conclusions of Lemma 3.3.12. Hence R'_δ is incompressible. □

Remark 3.3.14. R'_δ can be constructed from the surface shown in Figure 3.33a by first tubing together two points as shown in Figure 3.33b, and then attaching another piece of surface through the middle of the tube. If we attempted to apply an induction argument to this surface, at some point cutting open a crossing of an incompressible surface would yield a compressible surface, to which we could not immediately apply an inductive hypothesis.

Lemma 3.3.15. *There is a special alternating diagram D_ε and a surface R_ε given in special form at D_ε with no Q–arcs whose endpoints lie on the same disc in either $\mathbb{B}_a$ or $\mathbb{B}_b$.*

Proof. The diagram D_ε is shown in Figure 3.34, and Figure 3.35 shows the boundaries of the discs in $\mathbb{B}_a$ and $\mathbb{B}_b$. □

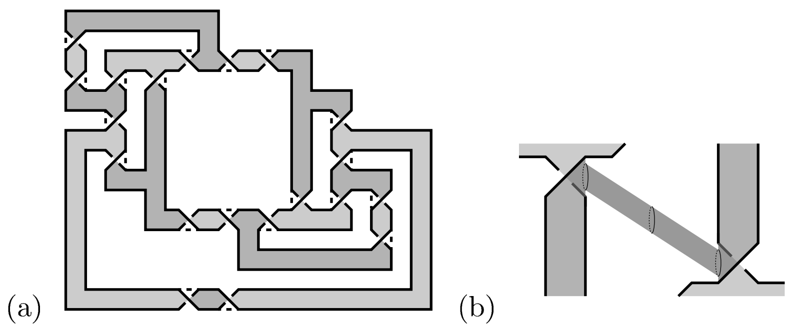

Figure 3.33

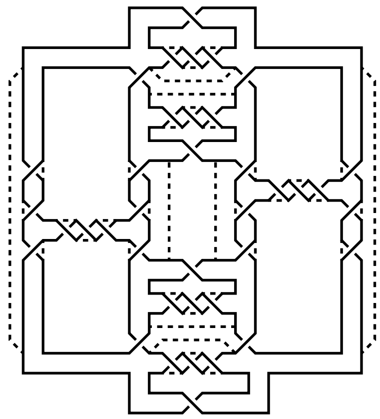

Figure 3.34

3.4 Distinguishing surfaces in special form

Our current aim is to prove Theorem 1.1.7. To do so, we wish to show that the map $\mathcal{A}\colon \mathrm{V}(\mathcal{K}(D)) \to \mathrm{V}(\mathrm{MS}(L))$ extends to an isomorphism of the simplicial complexes. By Proposition 3.2.10 we know that $\mathcal{A}$ can be extended to a map on the edges of the complexes. Since $\mathrm{MS}(L)$ is flag, this gives a simplicial map $\mathcal{A}\colon \mathcal{K}(D) \to \mathrm{MS}(L)$. We will show that $\mathcal{A}$ is a local isomorphism at each point, and so is a covering map. As $\mathrm{MS}(L)$ is simply connected, $\mathcal{A}$ is therefore an isomorphism.

Because both $\mathrm{MS}(L)$ and $\mathcal{K}(D)$ are flag, we need only show that $\mathcal{A}$ acts as required on the 1–skeleton of $\mathcal{K}(D)$. That is, we will consider the 1–skeleton of the closure of the star of a vertex v in $\mathcal{K}(D)$, and show that this is mapped bijectively under $\mathcal{A}$ to the corresponding subgraph of $\mathrm{MS}(L)$. From Proposition 3.3.8, together with Proposition 3.2.11, we already know that this is a surjection. In this section we will show that the distances between the vertices are not reduced. This will complete the proof of Theorem 1.1.7.

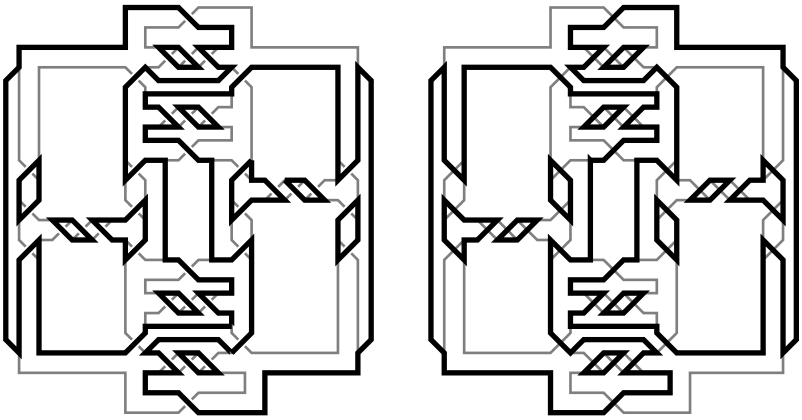

Figure 3.35

3.4.1 Adjacent surfaces are distinct

Convention 3.4.1. Given a special alternating diagram D, let R be the Seifert surface given by applying Seifert's algorithm. Recall that R is oriented. Divide the black regions of D into *a–regions* and *b–regions* according to whether the normal to R in a given region points into $\mathbb{B}_a$ or into $\mathbb{B}_b$.

Proposition 3.4.2. *Let D be a reduced, special alternating diagram for a prime link L. Let $v, v' \in \mathrm{V}(\mathcal{K}(D))$ such that $\mathrm{d}_{\mathcal{K}(D)}(v, v') = 1$ and v is given by D. Then $R = \mathcal{A}(v)$ is given by applying Seifert's algorithm to D. Let $R' = \mathcal{A}(v')$. Then $\mathrm{d}_{\mathrm{MS}(L)}(R, R') = 1$.*

Proof. Since $\mathrm{d}_{\mathcal{K}(D)}(v, v') = 1$, the proof of Proposition 3.2.10 gives an admissible special form for R' at D. This special form has at least one flype circle. Choose a set of flype circles that minimises the number of flype circles (as in the proof of Proposition 3.2.11). Together R and R' cut $\mathbb{S}^3 \setminus \mathcal{N}(L)$ into two sutured manifolds. Fix an a–region r of D. Call the two manifolds M^+ and M^-, where M^+ lies above r and M^- lies below it. We will show that neither of M^+ and M^- is a product sutured manifold. By symmetry, it is sufficient to prove this for M^+. The result will then follow by Proposition 1.4.3.

Consider a flype circle ϕ. The arcs shown in bold in Figure 3.36 form part of the

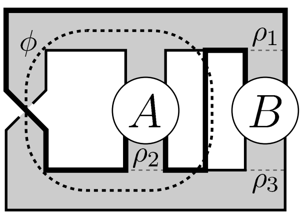

Figure 3.36

boundary of a single disc of $R' \cap \mathbb{B}_a$. To see this, note that only two overcrossing arcs of L cross the boundary of A. Hence one of ρ_1 and ρ_2 defines a product disc in M^+ contained in $\mathbb{B}_a$. Assume, without loss of generality, this is ρ_1. Then the product disc in M^+ defined by ρ_3 is contained in $\mathbb{B}_b$.

Up to isotopy, $M^+ \cap \mathbb{B}_a$ and $M^+ \cap \mathbb{B}_b$ are defined by where they meet $\mathbb{S}^2$. We now consider where M^+ lies relative to D and the special form of R'. As M^+ always meets the same side of R, it lies above all a–regions of D, and below all b–regions. In the white regions, M^+ lies on exactly one side of each P–arc. Let ρ be a P–arc. First suppose ρ is the flype arc of a flype circle ϕ. If ϕ and ρ are as in Figure 3.37, where the black region r'

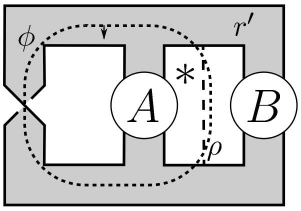

Figure 3.37

shown is an a–region, then M^+ lies on the positive side of ϕ, as marked $*$. Instead suppose that ρ lies across a crossing c. If c lies on the negative side of the flype circles, ρ lies on the positive side of c and M^+ lies between ρ and c. Otherwise, ρ lies on the positive side of the flype circles. In this case, M^+ does not lie between c and ρ. It lies on the far side of c from ρ, and on the far side of ρ from c.

From this we see that on the negative side of the flype circles R and R' are parallel, with M^+ as the product region between them. This product structure extends to the product discs described above where M^+ meets each flype circle. Thus we only need to consider the pieces of M^+ left by removing these product regions. Let Φ be the union of the flype circles. The remaining pieces of M^+ correspond to the components of $\mathbb{S}^2 \setminus \Phi$ that lie on the positive side of the flype circles. Let Λ be one such component, and M^+_Λ the corresponding piece of M^+. Create a new diagram D_Λ from D by changing it as shown in Figure 3.38 along each

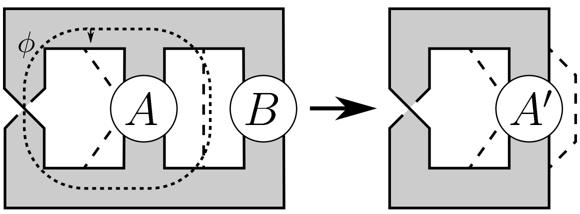

Figure 3.38

flype circle that bounds Λ. Let R_Λ be the surface given by applying Seifert's algorithm to

D_Λ, and R'_Λ that given by the P–arcs. Since the special form of R'_Λ has no flype circles, R_Λ and R'_Λ are parallel through a product sutured manifold N. The complementary sutured manifold to N is isotopic to M^+_Λ. Hence M^+ is a product sutured manifold if and only if every diagram $D_{\Lambda'}$ constructed in this way is of a fibred link with fibre given by the surface $R_{\Lambda'}$. Again by symmetry, we need only consider D_Λ.

Consider $G^{\mathcal{M}}(D)$ and $G^{\mathcal{M}}(D_\Lambda)$, and a fixed flype circle ϕ_Λ that bounds Λ. As can be seen from in Figure 3.39, $G^{\mathcal{M}}(D_\Lambda)$ is obtained from $G^{\mathcal{M}}(D)$ by collapsing everything on the

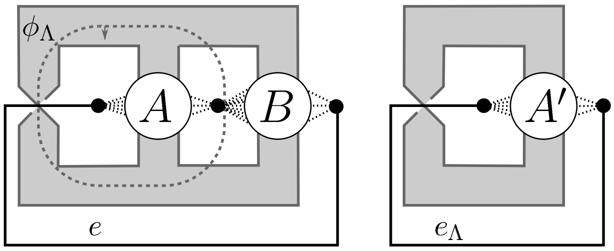

Figure 3.39

negative side of the flype circles. There may be other flype circles in A that are collapsed to give A'. Assume D_Λ is fibred. Then by Corollary 2.3.13 the edge e_Λ forms part of a subdivision of a loop.

Suppose that $G^{\mathcal{M}}(D_\Lambda)$ is not homeomorphic to $\mathbb{S}^1$. We can then divide up D_Λ and $G^{\mathcal{M}}(D_\Lambda)$ as shown in Figure 3.40, where at least three edges of $G^{\mathcal{M}}(D_\Lambda)$ run from each of

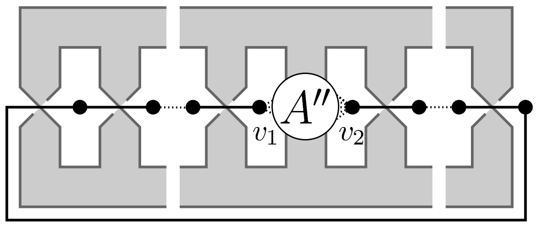

Figure 3.40

v_1 and v_2 into A''. Because of the 'tree-like' structure of $G^{\mathcal{M}}(D_\Lambda)$, either v_1 is a cut vertex or there is a loop attached to v_1. The collapsing of $G^{\mathcal{M}}(D)$ to $G^{\mathcal{M}}(D_\Lambda)$ does not create either of these. Hence v_1 is also a cut vertex or has a loop attached to it in $G^{\mathcal{M}}(D)$. This contradicts that D is prime and reduced.

Therefore, $G^{\mathcal{M}}(D_\Lambda)$ is homeomorphic to $\mathbb{S}^1$. Since the flype ϕ_Λ is essential, at least one vertex of $G^{\mathcal{M}}(D_\Lambda)$ in A' comes from collapsing a flype circle in A. By 'uncollapsing' all such vertices in $G^{\mathcal{M}}(D_\Lambda)$ we see that $G^{\mathcal{M}}(D)$ decomposes as shown in Figure 3.41. We can then combine two flype circles as in Figure 3.14, contradicting our choice of flype circles. Thus D_Λ is not fibred. □

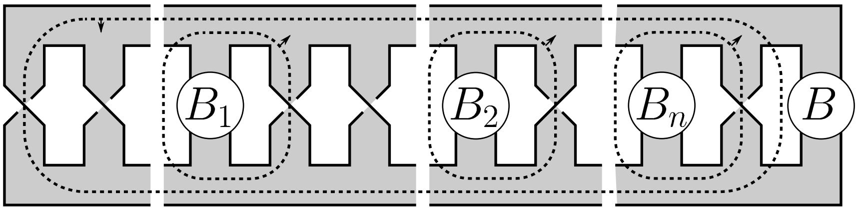

Figure 3.41

Proposition 3.4.3. *Let D be a reduced, special alternating diagram for a prime link L, and let $v_0 \in \mathrm{V}(\mathcal{K}(D))$ be given by D. Let $v_{-1}, v_1 \in \mathrm{V}(\mathcal{K}(D))$ such that $\mathrm{d}_{\mathcal{K}(D)}(v_{-1}, v_0) = \mathrm{d}_{\mathcal{K}(D)}(v_0, v_1) = 1$ and $\mathrm{d}_{\mathcal{K}(D)}(v_{-1}, v_1) = 2$. Let $R_i = \mathcal{A}(v_i)$. Then $\mathrm{d}_{\mathrm{MS}(L)}(R_{-1}, R_1) = 2$.*

Remark 3.4.4. This proof requires careful consideration of the relative positions of the Seifert surfaces in $\mathbb{S}^3$ and their interaction with the projection plane $\mathbb{S}^2$. As such, it is more suited to being drawn in a specific case than described in words for a general link. Therefore, a worked example is given in Section 3.4.2.

Proof. Following the proof of Proposition 3.2.10, construct an admissible special form at D for R_{-1} and for R_1. These special forms will each have at least one flype circle. Position each flype circle with the distance from the positive side of the flype crossing to the negative side of the flype arc as small as possible. Let Φ denote the set of flype circles for R_{-1} and R_1, where we do not require that each passes through the vertices of $\theta(D)$ and instead make them disjoint. For $i = \pm 1$, let Λ_i be the positive side of the flype circles for R_i, and λ_i the negative side. Note that each of these is a union of regions of $\theta(D)$.

Suppose that at least one region of $\theta(D)$ lies in $\Lambda_{-1} \cap \Lambda_1$. As in the proof of Proposition 3.2.11, there is at least one region r of $\theta(D)$ in $\Lambda_{-1} \cap \Lambda_1$ such that r^θ is defined at D. Adding r to v_0 gives a new vertex v_0' in $\mathcal{K}(D)$ with $\mathrm{d}_{\mathcal{K}(D)}(v_{-1}, v_0') \leq 1$ and $\mathrm{d}_{\mathcal{K}(D)}(v_0', v_1) \leq 1$. Thus $\mathrm{d}_{\mathcal{K}(D)}(v_{-1}, v_0') = \mathrm{d}_{\mathcal{K}(D)}(v_0', v_1) = 1$ since $\mathrm{d}_{\mathcal{K}(D)}(v_{-1}, v_1) = 2$. Hence, without loss of generality, we may assume $\Lambda_{-1} \cap \Lambda_1$ does not contain any region of $\theta(D)$. In particular this means that no two flype arcs lie on any one edge of $\theta(D)$, and neither do two flype crossings.

Let λ be a component of λ_{-1}. Create a new diagram D_λ by, for every flype circle ϕ of R_{-1} on the boundary of λ, changing D as shown in Figure 3.42a. All P–arcs in λ can be copied to D_λ. If D_λ together with the P–arcs is not connected, choose a simple closed curve ψ around each component, separating it from the other components. We can choose these

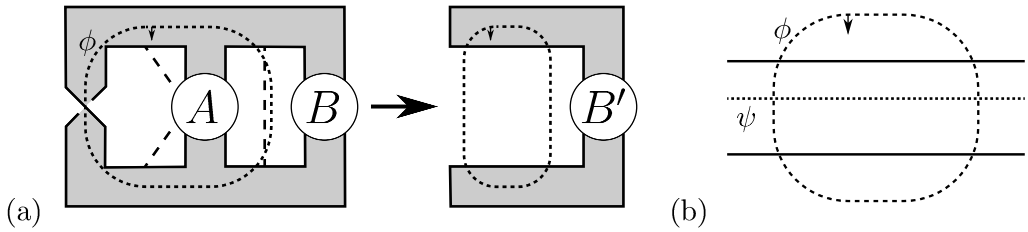

Figure 3.42

curves to be disjoint and to have minimal intersection with the images of the flype circles. As D is connected, each curve ψ must run through the image of at least one flype circle of R_{-1} as shown in Figure 3.42b. Repeat this process for each component of λ_{-1}. Let Ψ be the union of the flype circles in D together with those sections of each curve ψ contained in λ_{-1}.

Colour the P–arcs of R_{-1} red, and those of R_1 blue. Also colour the components of $\mathbb{S}^2 \setminus \Psi$. Colour Λ_{-1} red and Λ_1 blue. Colour a component in $\lambda_{-1} \cap \lambda_1$ blue if it meets Λ_1 anywhere along its boundary (that is, if it meets a component already coloured blue), and colour it red otherwise. Thus a point x on D in $\lambda_{-1} \cap \lambda_1$ is in a blue component of $\mathbb{S}^2 \setminus \Psi$ if there is a path ρ in D_λ from x to a flype circle of R_1, where λ is the component of λ_{-1} containing x, and is in a red component if no such path exists.

We wish to apply Proposition 1.4.3 to R_{-1} and R_1 to show that they cannot be made disjoint. To do so, we must specify precisely how to position them in $\mathbb{S}^3 \setminus \mathcal{N}(L)$. We will build up a description of the position of each disc in the constructions of R_{-1} and R_1 using the flype circles and the P–arcs, working inwards from the edge of the disc. Figure 3.46 gives an example of the type of picture we will build up.

First note that $\partial R_{-1} = \partial R_1$. Let $\overline{\rho}$ denote the closure of $(R_{-1} \cap R_1) \setminus \partial R_1$. Since R_{-1}, R_1 are incompressible and $\mathbb{S}^3 \setminus \mathcal{N}(L)$ is irreducible, $\overline{\rho}$ can be made to consist only of properly embedded arcs. We will arrange that each such arc is contained in a single disc from the construction of each of R_{-1}, R_1, with its endpoints on L. Let $\partial\overline{\rho}$ denote the (as yet undefined) set of endpoints of these arcs.

We now arrange the P–arcs of each surface. Let c be a crossing in D. Then c lies in the interior of λ_{-1} or λ_1, and so has at least one P–arc across it on the positive side.

Definition 3.4.5. Given two special forms, and a crossing c' such that the P–arcs around c' are disjoint, call the P–arc closest to c' on the positive side *inside* at c', and the other P–arc(s) on the positive side of c' *outside*.

If c only has one P–arc across it on the positive side, this P–arc must be positioned inside to avoid any intersections of P–arcs. There are three ways that c could come to have two P–arcs across it. One is that c lies in the interior of $\lambda_{-1} \cap \lambda_1$. If the corresponding component of $\mathbb{S}^2 \setminus \Psi$ has been coloured blue, put the blue P–arc outside, and if it is coloured red then put the red P–arc outside. The second possibility is that one, but not both, of the P–arcs across c is a flype arc of a flype circle ϕ. Without loss of generality, c is in the interior of λ_{-1} and ϕ is a flype circle of R_1. Then c lies on the positive side of ϕ, which contradicts the choice of the position of ϕ. Hence this case does not occur. The third possibility is that the white region of D on the positive side of c is a bigon, and one P–arc has been placed on the positive side of c, while the other P–arc has been placed on the negative side of the other crossing. In this case, put the P–arc that is paired with c on the inside.

Consider a flype circle ϕ with flype crossing c, and the pattern of simple closed curves in D given by $\partial(R_{-1} \cap \mathbb{B}_a)$. Let C_{-1} be the curve that runs along the overcrossing arc at c. After pushing the endpoints of P–arcs that lie on ϕ to the positive side of ϕ, as shown in Figure 3.43, we see that C_{-1} also runs along ρ_0. This is because each edge of $\theta(D)$ is only crossed by one overcrossing arc. The same is true of the analogous curve C_1 in $\partial(R_1 \cap \mathbb{B}_a)$.

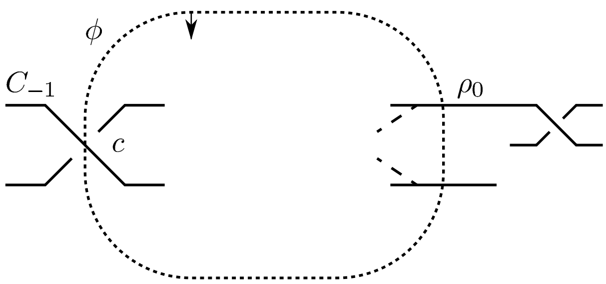

Figure 3.43

Suppose ϕ is a flype circle of R_{-1}, so that the positive side of ϕ is coloured red. Then the two points where C_{-1} crosses ϕ lie on the boundary of the same component of $\mathbb{S}^2 \setminus \Psi$ in λ_{-1}. Let A be the set of all such points where this component of $\mathbb{S}^2 \setminus \Psi$ is coloured blue, and define an involution $\widehat{\cdot}\colon A \to A$ such that if $a_1 \in A$ then $\widehat{a_1}$ is the other point of A on the same flype circle as a_1. Note that if points $a_1, a_2 \in A$ lie on the same simple closed curve C_{-1} of $\partial(R_{-1} \cap \mathbb{B}_a)$ then the pairs $a_1, \widehat{a_1}$ and $a_2, \widehat{a_2}$ do not interleave on C_{-1}.

We would like it to be the case that no two points of A lie on the same overcrossing arc of D. However, this is not in general true (see, for example, Figure 3.44). We therefore create a subset a of A with a new involution where this does not occur, as follows. Given points a_1 and a_2 of A that lie on a single overcrossing arc of D, remove a_1 and a_2 and change the involution to pair $\widehat{a_1}$ with $\widehat{a_2}$. Repeat this as many times as possible. Note that it is possible, as in Figure 3.44, to reach a pair where $\widehat{a_1} = a_2$. In such a case the two points are simply deleted from A. Since A is finite, this process will terminate. Notice that the

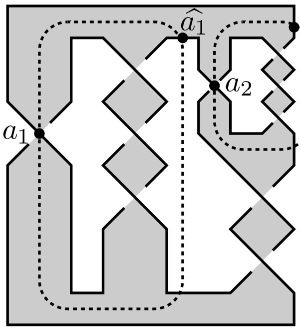

Figure 3.44

final result will not depend on the order in which pairs of points are chosen. Let a be the resulting set of points, and $\bar{\cdot}\colon a \to a$ the resulting involution. The following three properties of A are also true in a. If $a_1 \in a$ then a_1 and $\overline{a_1}$ lie on the same curve C_{-1} of $\partial(R_{-1} \cap \mathbb{B}_a)$ and on the same curve C_1 of $\partial(R_1 \cap \mathbb{B}_a)$. Additionally, a_1 and $\overline{a_1}$ lie on the boundary of the same component of $\mathbb{S}^2 \setminus \Psi$ in $\lambda_{-1} \cap \lambda_1$. Furthermore, if $a_2 \in a$ then the pairs $a_1, \overline{a_1}$ and $a_2, \overline{a_2}$ do not interleave on C_{-1} or on C_1.

The same process in $\mathbb{B}_b$ gives another set b with involution $\bar{\cdot}\colon b \to b$. We will end up with $\partial\overline{\rho} = a \cup b$.

Armed with these, we next connect the ends of the P–arcs to give the position of the neighbourhood of the boundary of every disc in the construction of R_{-1} and R_1. For each disc, this neighbourhood is an annulus, one boundary component of which we have already positioned (as ∂R_{-1} and ∂R_1 are the same fixed curve on $\partial\mathcal{N}(L)$). We describe the relative positions of the annuli by drawing on D the other boundary curve of each annulus. We will describe this process in $\mathbb{B}_a$. That in $\mathbb{B}_b$ is analogous.

Consider two crossings c and c' of D that are adjacent in D. Suppose the same colour P–arc lies inside at each, as shown in Figure 3.45a. We can then connect the P–arcs along the overcrossing arc at c' without them crossing (see Figure 3.45a). Suppose instead that

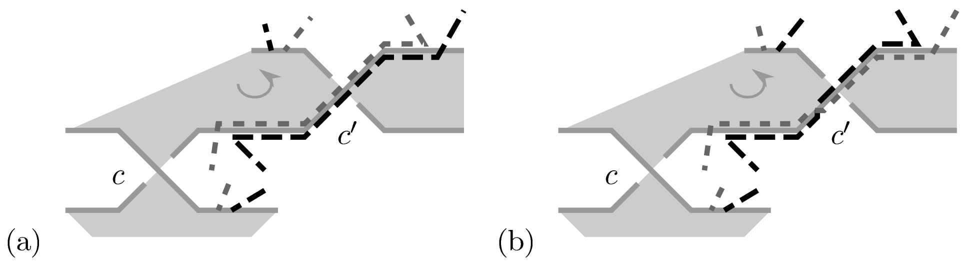

Figure 3.45

the P–arcs inside at c and c' are not the same colour. Then one of the crossings lies in Λ_{-1} whereas the other is either in Λ_1 or in a blue region of $\lambda_{-1} \cup \lambda_1$. Hence exactly one point

a_1 of a lies between c and c'. This time we connect the P–arcs so that the curves created cross once at a_1, as shown in Figure 3.45b. In this way we can connect all the P–arcs to form two sets of simple closed curves, as required. Figure 3.46 is an example of the result of this process.

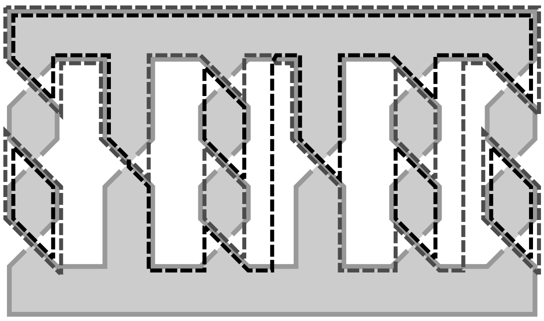

Figure 3.46

Finally we position the interior of each disc of R_{-1} and R_1. We may view those of R_{-1} as fixed, and vary those of R_1. In $\mathbb{B}_a$ we wish to arrange the discs such that an arc of $\overline{\rho}$ connects points $a_1, a_2 \in a$ if and only if $\overline{a_1} = a_2$. Recall that a_1 and $\overline{a_1}$ lie on the same curves of $\partial(R_{-1} \cap \mathbb{B}_a)$ and $\partial(R_{-1} \cap \mathbb{B}_a)$ for any $a_1 \in a$, and for $a_3 \in a$ the pairs $a_1, \overline{a_1}$ and $a_3, \overline{a_3}$ do not interleave along these curves. The same holds in $\mathbb{B}_b$. Thus there is no obstruction to our choice of $\overline{\rho}$. This completes the construction of $R_{-1} \cup R_1$.

It now remains to check that R_{-1} and R_1 have simplified intersection. Then, since R_{-1} and R_1 are not disjoint, Proposition 1.4.3 shows that $\mathrm{d}_{\mathrm{MS}(L)}(R_{-1}, R_1) > 1$. Let $M = \mathbb{S}^3 \setminus \mathcal{N}(L \cup R_{-1} \cup R_1)$.

Suppose a component M' of M does not meet $\mathbb{S}^2$. Then $\partial M'$ is composed of sections of the discs in the construction of R_{-1} and R_1, with at least one from each surface. Consider a disc S of R_1 that meets $\partial M'$. Then $\partial(S \cap \partial M')$ is a collection of simple closed curves in $R_{-1} \cap R_1$. Without loss of generality, assume $M' \subset \mathbb{B}_a$. Then $\partial(S \cap \partial M')$ is made up of overcrossing arcs of L and arcs in $\overline{\rho}$. As L is not the unknot, any simple closed curve of $\partial(S \cap \partial M')$ includes both types of arc. This is not possible, since no overcrossing arc includes more than one point of $\partial\overline{\rho}$. Hence every component of M meets $\mathbb{S}^2$. This means that, in checking whether any component is a product region, it is enough to consider those that meet each region of D.

Consider a point of $\partial\overline{\rho}$ on $\partial\mathcal{N}(L)$. Near this point, M has three components, only one of which, M_L, meets $\partial\mathcal{N}(L)$. On ∂M_L we see the pattern shown in Figure 3.47. Suppose M_L is a product region between R_{-1} and R_1. Then $\mathcal{N}_{\partial M_L}(\overline{\rho} \cup \partial\mathcal{N}(L))$ is of the form $\partial S_L \times \mathrm{I}$ for some surface S_L, and in particular is a collection of annuli. Since any component of

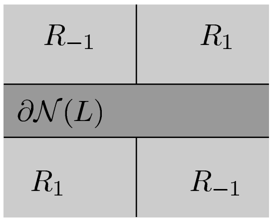

Figure 3.47

$\partial\mathcal{N}(L) \cap \partial M_L$ is an annulus, and each endpoint of an arc of $\overline{\rho}$ lies on $\partial\mathcal{N}(L)$, this cannot be the case. Hence M_L is not a product region.

Let r be a black region of D, and let M_r be the component of M that meets r. Then M_r meets the Seifert surface R_0. Since R_0 is disjoint from $R_{-1} \cup R_1$ and connected, R_0 is entirely contained in M_r. The manifold M_L also meets R_0. This shows that $M_r = M_L$, so M_r is not a product region. The same is true if r is instead a section of a white region of D that meets L. We are therefore left to consider those sections of white regions of D that are entirely bounded by P–arcs. These components of M form sutured manifolds, where the sutures are $R_{-1} \cap R_1$.

Let r be a section of a white region of D that lies between P–arcs in $\Lambda_{-1} \cup \Lambda_1$, and let M_r be the component of M that meets it. Let Λ be the component of $\mathbb{S}^2 \setminus \Psi$ containing r. Inside Λ, the discs of R_{-1} and R_1 are parallel to R_0, and M_r is isotopic to the complement of R_0 there. As in the proof of Proposition 3.4.2, we aim to decompose M_r along product discs, and to edit the diagram D to find a special alternating link diagram D' such that one piece of M_r is the complement of the surface given by applying Seifert's algorithm to D'. Again we will see that the link with diagram D' is not fibred, and hence that M_r is not a product sutured manifold.

Let c be the flype crossing of a flype circle ϕ in the boundary of Λ. Let C_{-1} be the curve of $\partial(R_{-1} \cap \mathbb{B}_a)$ that crosses c and C_1 the corresponding curve of $\partial(R_1 \cap \mathbb{B}_a)$. Then C_{-1} and C_1 also cross the flype arc of ϕ together. Let S_i be the disc of $R_i \cap \mathbb{B}_a$ with boundary C_i for $i = \pm 1$. Choose a point x on $C_{-1} \cap C_1$ slightly inside Λ from c, and a similar point x' near the flype arc of ϕ.

Suppose that there exists a point a_1 of a on C_{-1} or on C_1 such that $a_1, \overline{a_1}$ interleave with x, x'. Then there is a point $a_2 \in A$ such that $a_2, \widehat{a_2}$ interleave with x, x'. This contradicts that all the flype circles in Φ are disjoint. Thus no such point exists.

This means there is a product disc in M_r between S_{-1} and S_1 meeting $R_{-1} \cap R_1$ at x and x'. Similarly there is a corresponding product disc in $M_r \cap \mathbb{B}_b$. Cut M_r along these product discs, and retain the component that meets r. The corresponding change in D is to collapse the negative side of ϕ to a single crossing, as in Figure 3.38. After so collapsing all flype circles that bound Λ, we arrive at a diagram D' as described above. By the same reasoning as in the proof of Proposition 3.4.2, the remaining section of M_r is not a product region, and so M_r is not.

Finally, let r be a section of a white region of D that lies between a red P–arc and a blue P–arc parallel to it, and let M_r be the component of M that meets it. In this case, we show by contradiction that M_r also meets $\mathbb{S}^2$ in a white region between P–arcs in $\Lambda_{-1} \cup \Lambda_1$, and hence it has already been shown to not be a product region.

Suppose M_r does not meet $\mathbb{S}^2$ in any such white region. Let λ be the component of $\mathbb{S}^2 \setminus \Psi$ containing r. Note that every crossing in the interior of λ has both a red P–arc and a blue P–arc across it on the positive side. Since D is connected, there is a path σ from r to a flype circle ϕ on the boundary of λ that is contained in D and the P–arcs. If σ passes between an overcrossing and an undercrossing at a crossing c in the interior of λ, change it to instead switch between the two via the P–arcs across the positive side of c. After such changes, σ gives a path in ∂M_r.

Given that M_r does not meet the white region between the P–arcs on the other side of the flype crossing or flype arc, the endpoint of σ must be a point of $\partial\overline{\rho} = a \cup b$. From this we see that λ is coloured blue, which means there is a path σ in $D_{\lambda'}$ with the P–arcs from r to a flype circle of R_1, where λ' is the component of λ_{-1} containing λ. Again we can make this path run along P–arcs rather than moving directly between overcrossings and undercrossings. Our aim is to show that σ gives a path in ∂M_r. Then M_r will meet a white region of D between the P–arcs in Λ_1, contradicting our assumption.

We know that σ gives a path in ∂M_r until it first reaches a flype circle of R_{-1}, as D and $D_{\lambda'}$ are the same in the interior of λ_{-1}. Without loss of generality, suppose σ meets a flype circle ϕ of R_{-1} on an overcrossing arc of D. There it meets a point $a_1 \in a$. Then ∂M_r contains the arc of $\overline{\rho}$ that runs from a_1 to $\overline{a_1}$. Note that the path σ must also run through the point of $D_{\lambda'}$ that corresponds to $\overline{a_1}$.

In this way, we can see that σ defines a path in ∂M_r as required. Therefore M_r is not a product region. □

3.4.2 Proof by picture: a worked example of Proposition 3.4.3

Let L_ζ be the link with taut Seifert surfaces R_{-1} and R_1 given in admissible special form in Figure 3.48a and b respectively. We show the steps of the proof of Proposition 3.4.3 in this case, as follows.

i. Construct a minimal set of flype circles. Perform flypes to remove $\Lambda_{-1} \cap \Lambda_1$. Change flype circles to equivalent ones where needed. This gives Figure 3.48c.

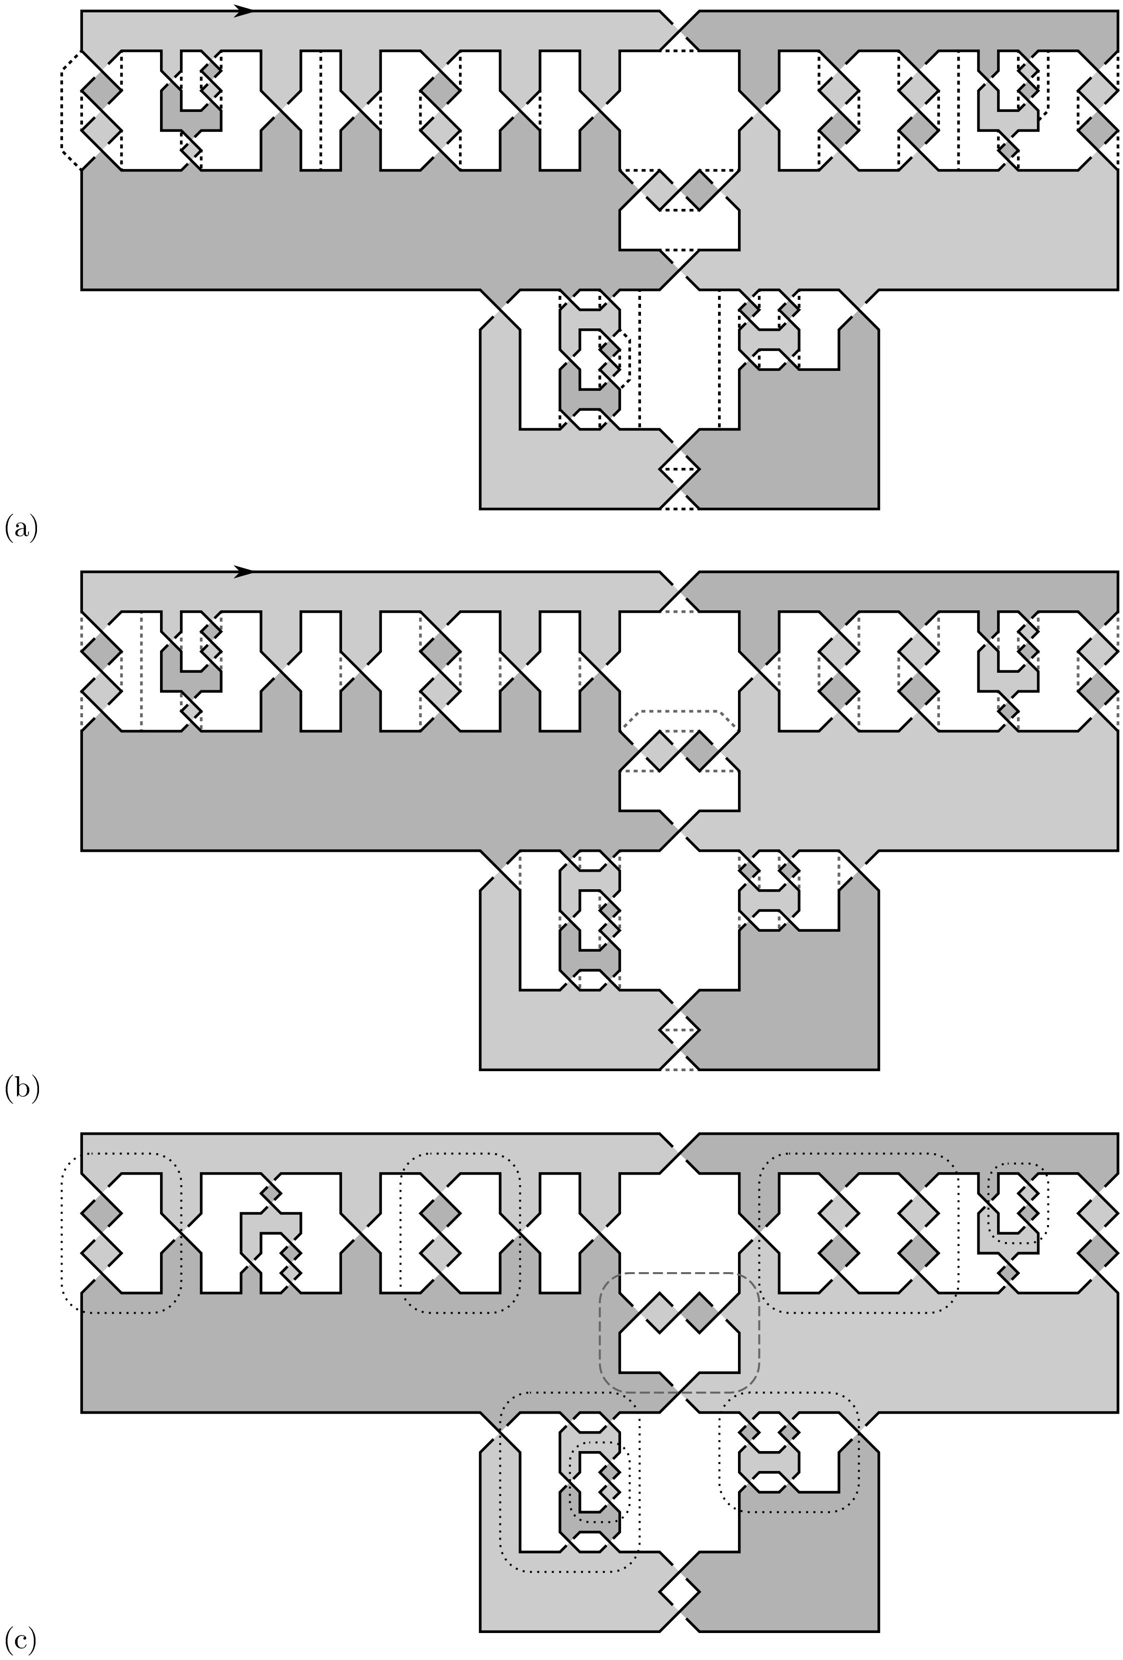

Figure 3.48

ii. For each region λ of λ_{-1}, construct the diagram D_λ and, if it has more than one component, choose a curve ψ around each component. The two new diagrams are shown in Figure 3.49a and b.

iii. Colour the components of $\mathbb{S}^2 \setminus \Psi$. Define a and b, and the corresponding involutions. See Figure 3.49c (the pairs of points in a are marked with stars).

iv. Position the P–arcs and connect them. Figure 3.50 shows the resulting picture in $\mathbb{B}_a$.

3.5 Non-special alternating links

Theorem 1.2 of [15] gives two families of links that illustrate the importance of assuming, in Theorem 1.1.7, that the link L has a special alternating diagram, and not just an alternating diagram. These families of links are the subjects of the next two propositions.

Proposition 3.5.1. *If L_{η_n} is the link shown in Figure 3.51 then* $\mathrm{MS}(L_{\eta_n})$ *contains an $(n-1)$–simplex in which exactly one vertex is given by applying Seifert's algorithm to an alternating diagram.*

Proof. There is only one non-trivial flype possible on the diagram given, and it is not an essential flype. Hence only one taut Seifert surface, R, for L_{η_n} comes from an alternating diagram.

There is a product disc decomposition of R that removes the plumbed on Hopf band. The resulting surface, R', is given by applying Seifert's algorithm to the special alternating link L'_{η_n} shown in Figure 3.52. By Theorem 1.1.7 the link of the vertex R' in $\mathrm{MS}(L'_{\eta_n})$ is an $(n-2)$–simplex, and so the same is true of R in $\mathrm{MS}(L_{\eta_n})$ by Proposition 1.4.23. □

Proposition 3.5.2. *If L_{θ_n} is the link shown in Figure 3.53 then* $\mathrm{MS}(L_{\theta_n})$ *is as shown in Figure 3.54, where a white circle represents a surface given by applying Seifert's surface to an alternating diagram, and a black circle represents a surface that cannot be so constructed.*

Proof. Again there are no essential flypes of this diagram, so only one taut Seifert surface comes from an alternating diagram.

The surface R_m, for $0 \leq m \leq n$, is the surface given by applying Seifert's algorithm to the special (but not alternating) diagram D_m shown in Figure 3.55. In Figure 3.53 we may see each of the horizontal twisted bands (which are plumbed on Hopf bands) as lying either above or below the rest of the surface. This does not change the ambient isotopy class of the Seifert surface because the Hopf link is fibred. From this we can see that R_0 is the surface shown in Figure 3.53 (see Figure 3.56).

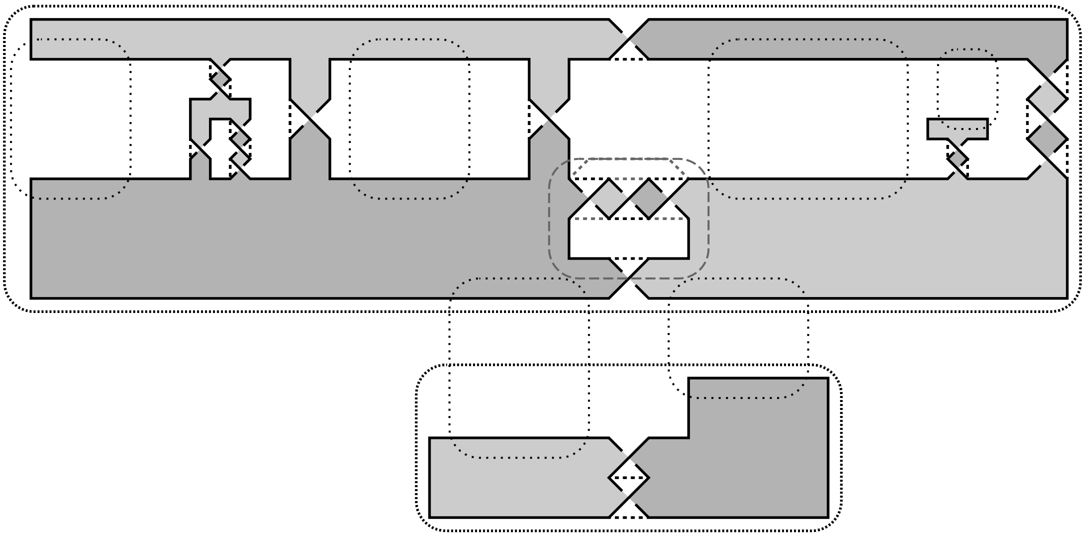

(a)

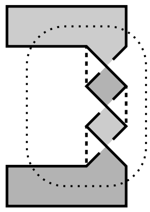

(b)

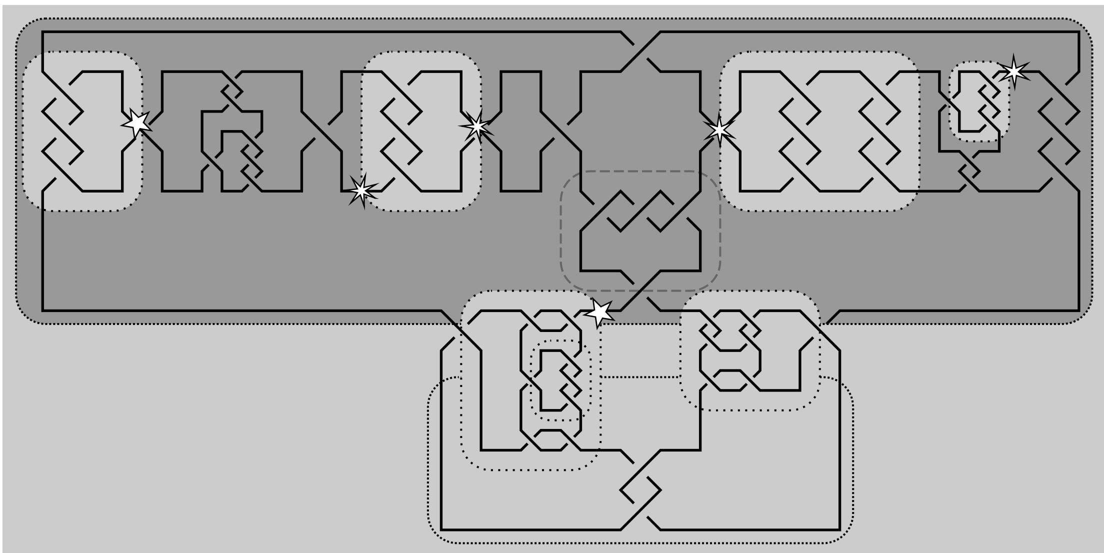

(c) red blue

Figure 3.49

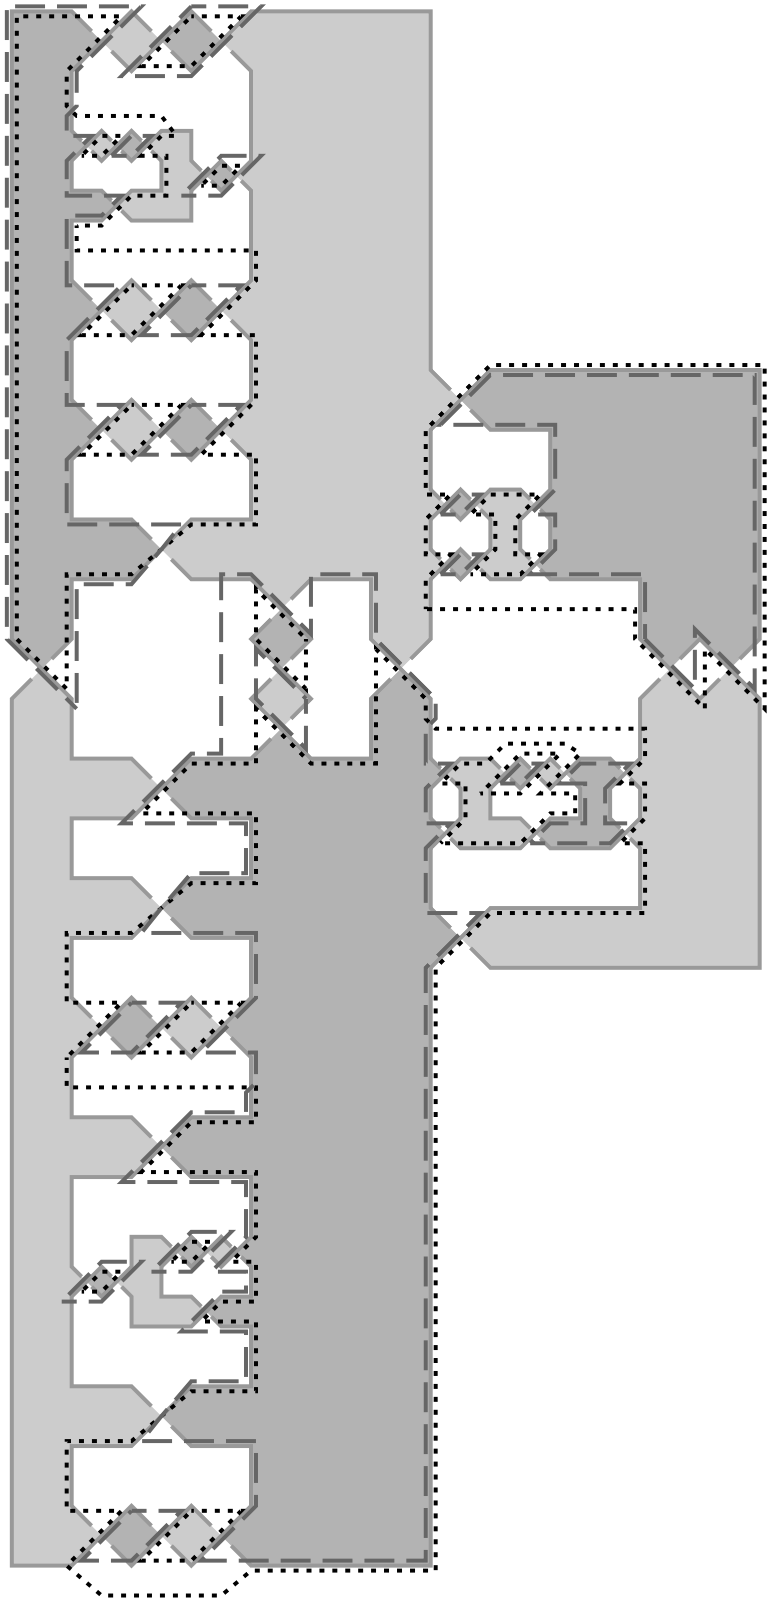

Figure 3.50

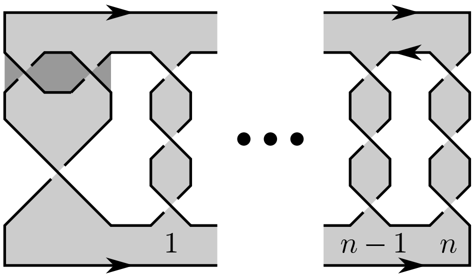

Figure 3.51

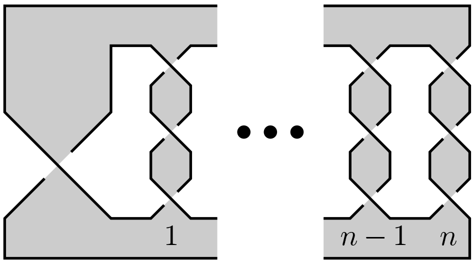

Figure 3.52

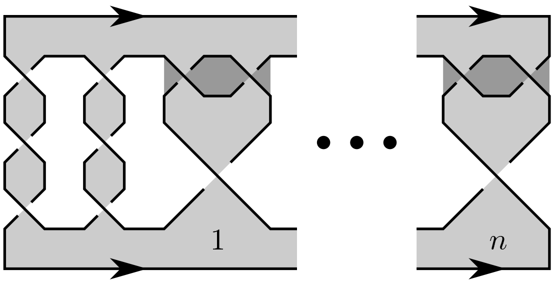

Figure 3.53

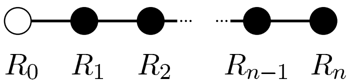

Figure 3.54

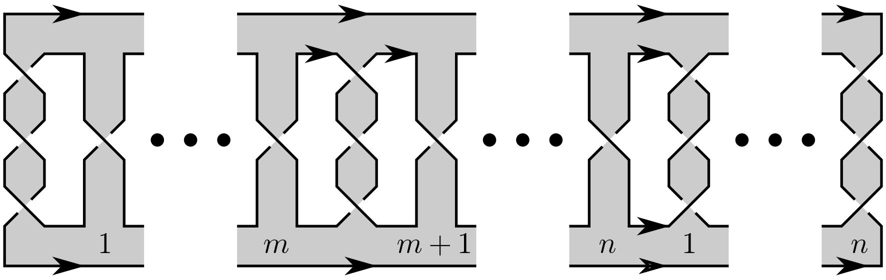

Figure 3.55

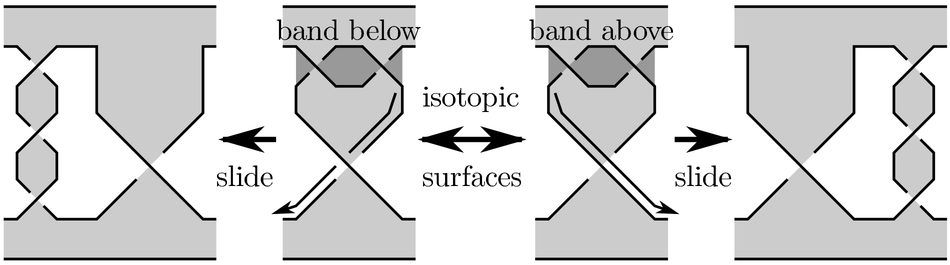

Figure 3.56

Consider the diagram D_0. We can express R_1 on this diagram as a surface in special form on an alternating diagram, with another surface attached in the obvious way as shown in Figure 3.57. Thus we see that R_1 is disjoint from R_0. We want to show that $R_0 \neq R_1$,

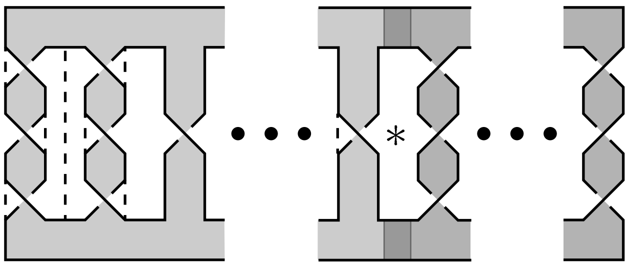

Figure 3.57

and that no other taut Seifert surface is disjoint from R_0. To do this, we show that the link of R_0 in $\mathrm{MS}(L_{\theta_n})$ consists of only one vertex, and this vertex is R_1.

The white region marked $*$ in Figure 3.57 defines a product disc in the complementary sutured manifold to R_0. Doing the product disc decomposition gives the same diagram, but with one fewer twisted bands in the right-hand surface. After n such decompositions, we arrive at a surface $R_0^{(n)}$ given by just the alternating section of the diagram. By Theorem 1.1.7 the link of $R_0^{(n)}$ is a single vertex. Hence, by repeated application of Proposition 1.4.23, the link of R_0 is also a single vertex, as required. By following the constructions in Theorem 1.1.7 and Proposition 1.4.23 we find that this single vertex is R_1.

We now show, for $1 \leq m \leq n-1$, that the link of R_m is $\{R_{m-1}, R_{m+1}\}$. Again we find that these surfaces can be expressed as a surface in special form on an alternating diagram together with a second surface attached in the obvious way, as shown in Figure 3.58a and b respectively. Also as before, after n product disc decompositions we may apply Theorem 1.1.7, giving that the link of $R_m^{(n)}$ contains exactly two vertices and no edges. This shows that the link of R_m is as required.

It remains to consider the surface R_n. We now know that R_{n-1} is in the link of R_n. Thus we need only show that the link of R_n contains only one element. Each white bigon

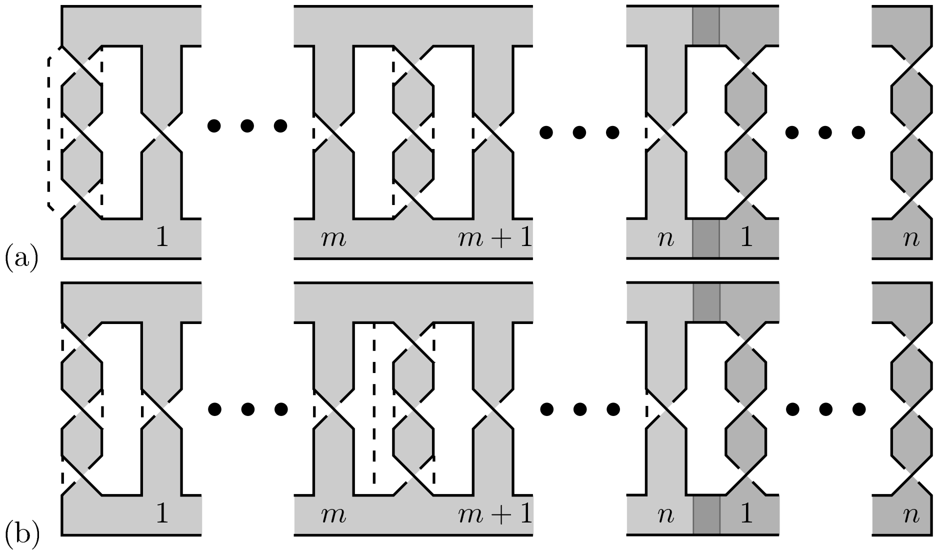

Figure 3.58

region of D_n defines a product disc in the complement of R_n. Using $n-1$ such product disc decompositions, removing all but one of the once-twisted bands, and an isotopy, we arrive at the diagram and surface shown in Figure 3.59, where the horizontal twisted band lies

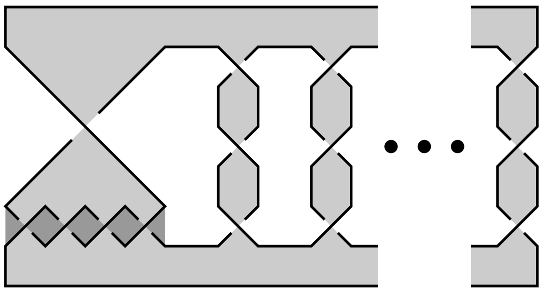

Figure 3.59

above the rest of the surface. Thus we are interested in a plumbing of two surfaces, S_a and S_b, shown in Figure 3.60. Call the plumbed surface S_{ab}.

S_a is an annulus with two full twists, bounding a torus link L_a. After n further disc decompositions, removing the twisted bands from the right-hand side of the diagram, S_b is also reduced to an annulus with two full twists. From this we see that both S_a and S_b are unique taut Seifert surfaces for L_a and $L_b = \partial S_b$ respectively, and also that neither link is fibred. By Theorem 1.6.9, this is enough to show there is at least one taut Seifert surface for ∂S_{ab} in the link of S_{ab}. However, we already knew that such a surface exists (the image of S_{n-1} under the product disc decompositions), and we have not yet shown there are no others.

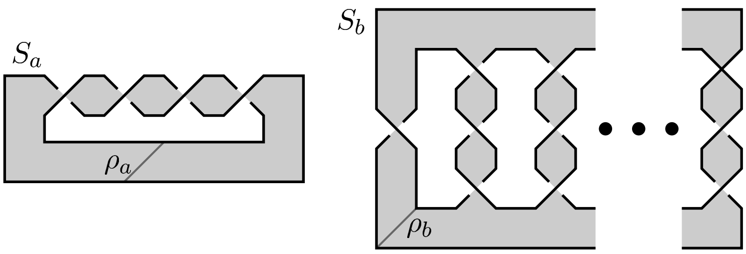

Figure 3.60

To do this, we consider the arcs ρ_a and ρ_b shown in Figure 3.60. These should be seen as lying in the boundary of the complementary sutured manifolds to the surfaces S_a, S_b. Note that ρ_a, ρ_b both lie on the upper side of the surfaces as shown. By [20] Proposition 3.4, the proof will be complete if we show that there is no product disc with either ρ_a or ρ_b as part of its boundary. Since ρ_a and ρ_b are essential, any such product disc will be essential. This result is clear for ρ_a as there are no essential product discs in the complement of S_a.

As we reduce S_b to an annulus by product disc decompositions, what is the effect on ρ_b? Figure 3.61 shows the result of the first decomposition. This leaves two copies of the

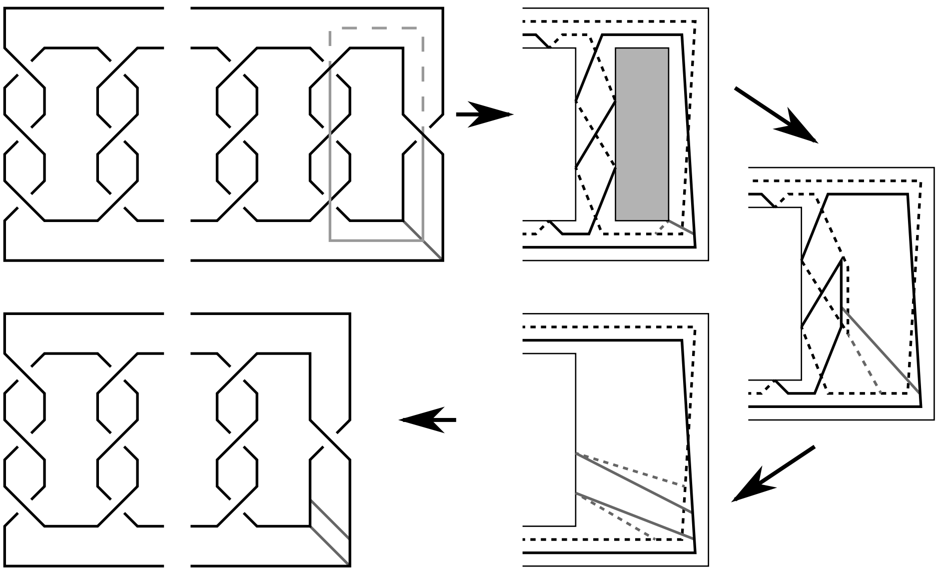

Figure 3.61

arc ρ_b', the curve in S_b' analogous to ρ_b in S_b. Inductively, we are therefore only concerned with the curve $\rho_b^{(m)}$ in $S_b^{(m)}$. The final pair $(S_b^{(n)}, \rho_b^{(n)})$ is the same as (S_a, ρ_a), so there is no essential product disc with $\rho_b^{(n)}$ contained in its boundary. By n applications of Lemma 1.4.18, there is no product disc T with $\rho_b \subset \partial T$. □

Remark 3.5.3. These two propositions together complete the proof of [15] Theorem 1.2.

3.6 The realisation of MS(L)

Let L be a prime link with a reduced, special alternating diagram D. In this section we turn our attention to Theorem 1.6 of [15], which asserts that MS(L) is homeomorphic to a ball in $\mathbb{R}^n$ for some n. The results we will use to prove this refer to simplicial complexes with a partial ordering on the vertices. Thus we will need to give such an order to the vertices of $\mathcal{K}(D)$.

Note that the simplices of $\mathcal{K}(D)$ already have a cyclic ordering on their vertices. We need to break this circle to give a linear order. To do this, pick a region r of $\theta(D)$. We break each circle in the ordering at the region r. That is, if v_1 is obtained from v_0 by adding the region r then we remove the relation $v_0 < v_1$. Note that the order this gives to the vertices of a face of a simplex matches that induced by the ordering of the vertices of the whole simplex. This method also orders the edges of each θ–graph in $\theta(D)$.

Definition 3.6.1 ([9])**.** Fix $m, n \in \mathbb{N}$, and let $v_0, \ldots, v_n$ be the (ordered) vertices of an n–simplex $[v_0, \ldots, v_n]$.

A *colour scheme* is an $m \times (l+1)$ matrix $\mathbf{X} = (x_{ij})$ for some $l \leq n$ with $x_{ij} \in \{0, \ldots, n\}$. In addition, $\mathbf{X}$ must have pairwise distinct columns, and its entries must satisfy $x_{10} \leq x_{11} \leq \cdots \leq x_{1l} \leq x_{20} \leq \cdots \leq x_{ml}$. For $k \leq l$, column k of $\mathbf{X}$ defines a point v'_k of $[v_0, \ldots, v_n]$ by $v'_k = (v_{x_{1k}} + \cdots + v_{x_{mk}})/m$. Thus $\mathbf{X}$ defines an l–simplex $[v'_0, \ldots, v'_l]$.

The *edgewise subdivision* $\mathrm{Esd}_m([v_0, \ldots, v_n])$ of $[v_0, \ldots, v_n]$ is made up of all simplices given by colour schemes.

Remark 3.6.2. In [9] Section 3 it is proved that $\mathrm{Esd}_m([v_0, \ldots, v_n])$ is indeed a subdivision of $[v_0, \ldots, v_n]$.

Proposition 3.6.3 ([15] Theorem 1.6(1))**.** *Suppose $\theta(D)$ consists of a single θ–graph with $n+1$ edges and total weight m. Then $\mathcal{K}(D)$ is the edgewise subdivision of an n–simplex.*

Proof. Fix a region r_0 of $\theta(D)$ to give an ordering on the vertices of $\mathcal{K}(D)$ as above (note that in this case $\tilde{\theta}(D)$ is $\theta(D)$). Label the edges of $\theta(D)$ as $e_0, \ldots, e_n$ where $e_0 < e_1 < \cdots < e_n$. Label the other regions of $\theta(D)$ in order around the θ–graph in the positive direction, so that $\partial^- r_i = e_{i-1}$ and $\partial^+ r_i = e_i$.

Let $v_0, \ldots, v_n$ be the (ordered) vertices of an n–simplex $[v_0, \ldots, v_n]$. Given a vertex $u = (w_0, \ldots, w_n)$ of $\mathcal{K}(D)$, we may define an $m \times 1$ matrix $\mathbf{X} = (x_{ij})$ by taking $x_{i0} = k$ where $\sum_{j<k} w_j < i \leq \sum_{j \leq k} w_j$. That is, the number of times k occurs in $\mathbf{X}$ is the weight that u assigns to e_k, and the entries of $\mathbf{X}$ are arranged so as to be non-decreasing. The matrix $\mathbf{X}$ is then a colour scheme, and so defines a vertex of $\mathrm{Esd}_m([v_0, \ldots, v_n])$. Thus there is a map B from the vertices of $\mathcal{K}(D)$ to the vertices of $\mathrm{Esd}_m([v_0, \ldots, v_n])$. It is clear that

B is a bijection. We will show that B induces a bijection on the n–simplices of the two complexes. This will complete the proof as in each of the complexes every simplex is a face of an n–simplex.

Let $u_0, \ldots, u_n$ be the vertices of an n–simplex in $\mathcal{K}(D)$ with $u_0 < u_1 < \cdots < u_n$. Then there is a permutation $\pi \colon \{1, \ldots, n\} \to \{1, \ldots, n\}$ such that u_i is obtained from u_{i-1} by adding the region $r_{\pi(i)}$. Let $\mathbf{X}$ be an $m \times (n+1)$ matrix, where column k of $\mathbf{X}$ is the column vector defined by u_k in the construction of the map B. The column vectors from u_k and u_{k+1} differ exactly in one coordinate, which changes from $\pi(k) - 1$ to $\pi(k)$. Because $u_0, \ldots, u_n$ is an n–simplex, we also know that u_0 is obtained from u_n by adding r_0. This has the effect of dropping the last coordinate of the column vector, moving the remaining entries down one place, and inserting a 0 in the top coordinate. Hence $\mathbf{X}$ is a colour scheme defining an n–simplex of $\mathrm{Esd}_m([v_0, \ldots, v_n])$.

Conversely, choose an n–simplex of $\mathrm{Esd}_m([v_0, \ldots, v_n])$ and let $\mathbf{X}$ be the colour scheme defining it. Let u_k be the vertex of $\mathcal{K}(D)$ given by column k of $\mathbf{X}$. That is, u_k assigns a weight of 1 to e_i for each time i occurs in the column vector. As the columns of $\mathbf{X}$ are pairwise distinct, at least one element changes between u_k and u_{k+1}. The ordering of the values of the elements of $\mathbf{X}$ ensure this is an increase in each case, and that the sum of the sizes of these increases is at most n. Thus exactly one coordinate increases between u_k and u_{k+1} and this coordinate increases by 1. Moreover, for each $k \leq n$, exactly one of these changes is from $k-1$ to k, and $x_{(k-1)n} = x_{k0}$ for $1 < k \leq m$ while $x_{10} = 0$ and $x_{mn} = n$. Define a permutation $\pi \colon \{1, \ldots, n\} \to \{1, \ldots, n\}$ by $\pi(j) = k$ if the move from u_{j-1} to u_j sees a coordinate change from $k-1$ to k. Then u_j is obtained from u_{j-1} by adding the region $r_{\pi(j)}$ for $j \leq n$ and adding r_0 to u_n gives u_0. This shows that $u_0, \ldots, u_n$ span an n–simplex in $\mathcal{K}(D)$.

Each of these two maps is an injection, as simplices are defined by their vertices. As there are only finitely many n–simplices in either complex, both maps are therefore bijections. □

Definition 3.6.4 ([10] Chapter II Definition 8.7)**.** A simplicial complex $\mathcal{X}$ is *ordered* if there is a binary relation $\leq$ on the vertices of $\mathcal{X}$ with the following properties.

(P1) $(u \leq v \text{ and } v \leq u) \Rightarrow u = v$.

(P2) If u, v are distinct, $(u \leq v \text{ or } v \leq u) \Leftrightarrow$ u and v are adjacent.

(P3) If u, v, w are vertices of a 2–simplex then $(u \leq v \text{ and } v \leq w) \Rightarrow u \leq w$.

Remark 3.6.5. It is clear that in searching for such a relation we may use the following weaker version of (P2).

(P2)$'$ If u, v are adjacent then ($u \leq v$ or $v \leq u$).

To see this, note that we can remove all relationships between non-adjacent vertices.

Definition 3.6.6 ([10] Chapter II Definition 8.8)**.** Let $\mathcal{X}_1, \mathcal{X}_2$ be ordered simplicial complexes. We define the simplicial complex $\mathcal{X}_1 \times \mathcal{X}_2$. Its vertices are given by the set $\mathrm{V}(\mathcal{X}_1) \times \mathrm{V}(\mathcal{X}_2)$. Vertices $(u_0, v_0), \ldots, (u_n, v_n)$ span an n–simplex if the following hold.

- $\{u_0, \ldots, u_n\}$ is an m–simplex of $\mathcal{X}_1$ for some $m \leq n$.
- $\{v_0, \ldots, v_n\}$ is an m–simplex of $\mathcal{X}_2$ for some $m \leq n$.
- The relation defined by $(u, v) \leq (u', v') \Leftrightarrow (u \leq u'$ and $v \leq v')$ gives a total linear order on $(u_0, v_0), \ldots, (u_n, v_n)$.

Remark 3.6.7. The projection maps on the vertices extend to simplicial maps of the complexes.

Theorem 3.6.8 ([10] Chapter II Lemma 8.9)**.** *The map $|\mathcal{X}_1 \times \mathcal{X}_2| \to |\mathcal{X}_1| \times |\mathcal{X}_2|$ induced by projection is a homeomorphism.*

Remark 3.6.9. With the ordering on the vertices defined above, $\mathcal{K}(D)$ is an ordered simplicial complex.

Theorem 3.6.10 ([15] Theorem 1.6(2))**.** *The realisation $|\mathcal{K}(D)|$ of $\mathcal{K}(D)$ is homeomorphic to a ball of dimension $\sum (n^\theta(e^\theta) - 1)$, where the sum is taken over all θ–graphs e^θ in $\theta(D)$ and $n^\theta(e^\theta)$ is the number of edges in e^θ.*

Proof. We proceed by induction on the number of θ–graphs in $\theta(D)$. If there are no θ–graphs then $\mathcal{K}(D)$ is a single vertex. The case of one θ–graph is covered by Proposition 3.6.3.

The construction of the simplicial complex $\mathcal{K}(D)$ is dependent only on $\tilde{\theta}(D)$ together with a choice of positive direction and total weight on each θ–graph. Let Θ be the set of such graphs. That is, each element of Θ is a finite collection of disjoint θ–graphs in $\mathbb{S}^2$, with a choice of positive direction and total weight on each. Then for every such graph θ we can construct a simplicial complex $\mathcal{K}(\theta)$. It is on this set of complexes that we will induct. Note that the base cases hold in this more general setting.

Now suppose that the result holds for all elements of Θ with at most $(m-1)$ θ–graphs. Let θ_0 be one with m θ–graphs. Choose a region r of $\mathbb{S}^2 \backslash \theta_0$ that meets at least two θ–graphs, and pick a simple closed curve ρ contained in r that separates the θ–graphs of θ_0. This gives two new elements of Θ, each with at most $(m-1)$ θ–graphs. Call these θ_1, θ_2. Use the region r to order the vertices of $\mathcal{K}(\theta), \mathcal{K}(\theta_1), \mathcal{K}(\theta_2)$ as above. Then, by Theorem 3.6.8, $|\mathcal{K}(\theta_1) \times \mathcal{K}(\theta_2)| \cong |\mathcal{K}(\theta_1)| \times |\mathcal{K}(\theta_2)|$. The inductive hypothesis gives that $|\mathcal{K}(\theta_1)|, |\mathcal{K}(\theta_2)|$ are balls of the relevant dimension, and so the result holds for $|\mathcal{K}(\theta_1) \times \mathcal{K}(\theta_2)|$.

It remains to show that $\mathcal{K}(\theta_1) \times \mathcal{K}(\theta_2)$ is $\mathcal{K}(\theta_0)$. Clearly the obvious map on the vertices is a bijection. Again we will check that it induces a bijection on the top-dimensional simplices.

Consider an ordered n–simplex $[w_0, \ldots, w_n]$ in $\mathcal{K}(\theta_0)$. Then w_0 assigns a weight to each edge of θ_0. There is an ordering of the regions of $\mathbb{S}^2 \setminus \theta_0$ with r last such that the vertices of $[w_0, \ldots, w_n]$ are given by adding these regions in turn in order. This induces similar orderings of the regions of $\mathbb{S}^2 \setminus \theta_i$. For $0 \leq i \leq n$, write $w_i = (u_i, v_i)$, where u_i gives the weights on the edges of θ_1 and v_i gives the weights on the edges in θ_2. If $i < j$ then $u_i \leq u_j$ and $v_i \leq v_j$. Thus $[(u_0, v_0), \ldots, (u_n, v_n)]$ is a simplex in $\mathcal{K}(\theta_1) \times \mathcal{K}(\theta_2)$.

Conversely, consider an n–simplex $[(u_0, v_0), \ldots, (u_n, v_n)]$ in $\mathcal{K}(\theta_1) \times \mathcal{K}(\theta_2)$, where the u_i are vertices of $\mathcal{K}(\theta_1)$, the v_i are vertices of $\mathcal{K}(\theta_2)$ and $(u_0, v_0) < (u_1, v_1) < \cdots < (u_n, v_n)$. For $0 < i \leq n$, as $u_{i-1} \leq u_i$, either $u_{i-1} = u_i$ or u_i is given by adding some regions of $\mathbb{S}^2 \setminus \theta_1$ other than r to u_{i-1}. Similarly, either $v_{i-1} = v_i$ or v_i is given by adding some regions of $\mathbb{S}^2 \setminus \theta_2$ other than r to v_{i-1}. Since there are only n regions of $\mathbb{S}^2 \setminus \theta_0$ other than r, for $0 < i \leq n$ either $u_i = u_{i-1}$ and v_i is given by adding a single region to v_{i-1} or instead $v_i = v_{i-1}$ and u_i is given by adding a single region to u_{i-1}. Because $[u_0, \ldots, u_n]$ is a simplex in $\mathcal{K}(\theta_1)$, adding r to u_n gives u_0. Likewise, adding r to v_n gives v_0. Thus there is an ordering on the regions of $\mathbb{S}^2 \setminus \theta_0$ with r last that shows that $[(u_0, v_0), \ldots, (u_n, v_n)]$ is an n–simplex in $\mathcal{K}(\theta_0)$. □

Chapter 4

Connected sums

The aim of this chapter is to prove Theorem 1.1.10, describing the Kakimizu complex of a connected sum $L = L_1 \# L_2$ in terms of those for L_1 and L_2. The idea of the proof is to first define a triangulation of $|\operatorname{MS}(L_1)| \times |\operatorname{MS}(L_2)| \times \mathbb{R}$, then choose a representative of each element of $\operatorname{V}(\operatorname{MS}(L_i))$, which we use to define a map between the vertices of this triangulation and those of $\operatorname{MS}(L)$, and finally show that this map is an isomorphism.

4.1 Product regions and detecting adjacency

In this section we recall some results from 3–manifold theory that we will need. These include a number of related results; we include full proofs as later proofs are sensitive to the details. The concepts covered in this section are generally well-known, although some of the details are specific to the case at hand.

Let L be a non-split link other than the unknot, and let $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$. Let S and S' be incompressible, ∂–incompressible surfaces properly embedded in M in general position, so that Proposition 1.4.3 applies.

Remark 4.1.1. If ∂S and $\partial S'$ are transverse, a product region N given by Proposition 1.4.3 is always embedded in M. However, we will want to apply the proposition when S and S' are Seifert surfaces for L that may have components that either coincide or have boundaries that coincide. It continues to hold in this situation, but may result in a product region that is not embedded. There are two ways that this can occur.

One option is that the components S_0 of S and S'_0 of S' that bound N coincide. Then N is the whole of M, the surfaces S and S' are connected, and the link L is fibred with fibre S. As we are currently only interested in non-fibred links, this case will not arise.

The second possibility is that S_0 and S'_0 do not coincide but their boundaries do. Then ∂N covers an entire component of $\partial \mathcal{N}(L)$, and $T \times \{0\}$ meets $T \times \{1\}$ along at least one component of ∂T. See Figure 4.4.

Corollary 4.1.2. *Suppose L is not fibred. If S, S' are isotopic by an isotopy fixing their boundaries then there is a product region $N = (T \times \mathrm{I})/\sim$ between them with $\rho_T = \partial T$ in the notation of Definition 1.4.1.*

Proof. Suppose no such product region exists. By Proposition 1.4.3 there is a product region N between S and S'. This product region N meets one component of each of S and S', and the boundaries of these components coincide. By deleting other components if necessary, we may assume S, S' are connected.

Let $\tilde{S}$ be a lift of S to the infinite cyclic cover $\tilde{M}$. The isotopy from S to S' lifts to an isotopy from $\tilde{S}$ to a lift $\tilde{S}'$ of S'. Note that $\partial\tilde{S} = \partial\tilde{S}'$. By hypothesis, the isotopy from S' to S defined by the product region N moves the boundary of S'. As the boundaries of S and S' coincide, this isotopy therefore takes each component of $\partial S'$ once around the torus component of ∂M on which it lies. Hence the isotopy defined by N lifts to an isotopy from $\tilde{S}'$ to either $\tau(\tilde{S})$ or $\tau^{-1}(\tilde{S})$. Thus $\tilde{S}$ is isotopic to $\tau(\tilde{S})$. Again by Proposition 1.4.3 there is a product region between $\tilde{S}$ and $\tau(\tilde{S})$. This contradicts that L is not fibred. □

Remark 4.1.3. The condition that L is not fibred is not actually necessary for this result, only for our proof.

Proposition 4.1.4 ([31] Corollary 3.4)**.** *Let M_a, M_b be proper 3–submanifolds of $\tilde{M}$ such that ∂M_a and ∂M_b are unions of lifts of taut Seifert surfaces that are almost transverse with simplified intersection. If M_a can be isotoped to have interior disjoint from M_b then the interior of M_a is disjoint from M_b.*

Lemma 4.1.5 ([31] Lemma 3.5)**.** *Let R_1, R_2, R_3 be taut Seifert surfaces. Then they can be isotoped to be pairwise almost transverse and have pairwise simplified intersection.*

Proposition 4.1.6 ([31] Proposition 3.2)**.** *Following the notation of Proposition 1.3.8, if R' is almost transverse to and has simplified intersection with R then it minimises* $\max\{n : V_{R'} \cap V_n \neq \emptyset\} - \min\{n : V_{R'} \cap V_n \neq \emptyset\}$.

The following criterion allows us to test for adjacency under Definition 1.3.9.

Lemma 4.1.7. *Let R_1, R_2 be fixed, almost disjoint, taut Seifert surfaces for a link L. Then R_1, R_2 demonstrate that their isotopy classes are at most distance 1 apart in* $\mathrm{MS}(L)$ *if and only if the following holds for all pairs (x, y) of points of $R_1 \setminus R_2$.*

Choose a product neighbourhood $\mathcal{N}(R_1) = R_1 \times [-1, 1]$ in M for R_1 with $R_1 = R_1 \times \{0\}$, such that $R_1 \times \{1\}$ lies on the positive side of R_1. Let $\rho\colon \mathrm{I} \to M$ be any path with $\rho(0) = (x, 1)$ and $\rho(1) = (y, -1)$ that is disjoint from R_1 and transverse to R_2. Then the algebraic intersection number of ρ and R_2 is 1.

Proof. Suppose the condition holds for all pairs (x, y). If R_1, R_2 coincide everywhere then they are distance 0 apart. Assume otherwise, and let x_0 be a point of $R_1 \setminus R_2$. Choose a lift V_{R_2} of $M \setminus R_2$, and let $\tilde{x}_0$ be the lift of x_0 that lies in V_{R_2}. Let V_{R_1} be the lift of $M \setminus R_1$ that lies above $\tilde{x}_0$, and let $\tilde{R}_1$ be the lift of R_1 that lies between V_{R_1} and $\tau^{-1}(V_{R_1})$. Then $\tilde{x}_0$ lies on $\tilde{R}_1$, and V_{R_1} meets V_{R_2} and $\tau(V_{R_2})$. We wish to show that these are the only lifts of $M \setminus R_2$ that V_{R_1} meets. Suppose otherwise.

First suppose that V_{R_1} meets $\tau^{-1}(V_{R_2})$. Then there is a path ρ in M that runs from just above x_0 to the projection of a point in $V_{R_1} \cap \tau^{-1}(V_{R_2})$, that is disjoint from R_1 and that has algebraic intersection -1 with R_2. There is also a path ρ' from this point back to above x_0 that is disjoint from R_2. We may assume both paths are transverse to R_2. This forms a closed curve that has algebraic intersection -1 with R_2. It therefore has algebraic intersection -1 with R_1.

Consider the first point x_1 at which ρ' crosses R_1. Then $x_1 \in R_1 \setminus R_2$. If ρ' passes through R_1 in the positive direction at x_1 then the section of the path $\rho \cup \rho'$ from $(x_0, 1)$ to $(x_1, -1)$ contradicts the hypothesis as it has algebraic intersection -1 with R_2 instead of 1. Thus ρ' passes through R_1 in the negative direction at x_1. Stop the path just above x_1, at $(x_1, 1)$, and add in a path that runs to $(x_1, -1)$ in $M \setminus R_1$. This final section of path has algebraic intersection 1 with R_2 by the hypothesis. Then the complete path gives a contradiction with the hypothesis, as it has algebraic intersection 0 with R_2. Hence V_{R_1} lies entirely above $\tau^{-1}(V_{R_2})$.

Now suppose V_{R_1} meets $\tau^2(V_{R_2})$. Then there is a path ρ, disjoint from R_1, from $(x_0, 1)$ to the projection of a point in $V_{R_1} \cap \tau^2(V_{R_2})$ that has algebraic intersection 2 with R_2. Again, there is a second path ρ' from there to $(x_0, 1)$ that is disjoint from R_2. The closed curve has algebraic intersection 2 with R_2 and R_1. Consider the first point x_1 at which ρ' crosses R_1. Then $x_1 \in R_1 \setminus R_2$. If ρ' passes through R_1 in the positive direction at x_1 then the path up to this point contradicts the hypothesis, as it has intersection 2 with R_2 instead of 1. Thus it passes through R_1 in the negative direction. Stop the path just above x_1, at $(x_1, 1)$, and add in a path that runs to $(x_1, -1)$ in $M \setminus R_1$. This final section of path has algebraic intersection 1 with R_2 by the hypothesis. Thus the complete path gives a contradiction with the hypothesis, as it has intersection 3 with R_2. Hence we have the required result.

Conversely, suppose any lift of $M \setminus R_1$ meets at most two lifts of $M \setminus R_2$. Choose x, y, ρ as in the condition to be checked. Take the lift $\tilde{\rho}$ of ρ with $\tilde{\rho}(0)$ in a fixed lift V_{R_2} of $M \setminus R_2$. We may use $\tilde{\rho}$ in defining the lift V_{R_1} of $M \setminus R_1$. Therefore $\tilde{\rho}$ is contained in two lifts of $M \setminus R_2$ and the lift of R_2 between them. One of these two lifts is V_{R_2}, since this contains the lift $\tilde{\rho}(0)$ of $(x, 1)$. In addition, the lift of $(x, -1)$ lies in $\tau(V_{R_2})$. Thus the lift $\tilde{\rho}(1)$ of

$(y,-1)$ is in V_{R_2} or in $\tau(V_{R_2})$. We must show that it is in $\tau(V_{R_2})$. If not then it lies in V_{R_2}. Then the lift of $(y,1)$ lies in $\tau^{-1}(V_{R_2})$, which is a contradiction. □

Lemma 4.1.8. *Let R_a be a taut Seifert surface for L. Suppose $R_{b,0}, R'_{b,0}$ are two copies of a component of a taut Seifert surface for L that are disjoint from R_a and are not isotopic to any component of it. Then $R_{b,0}, R'_{b,0}$ are isotopic by an isotopy that does not meet R_a.*

Proof. By a small isotopy disjoint from R_a, we may ensure that $R_{b,0}, R'_{b,0}$ are transverse. Since they are isotopic, there is a product region N between them. If N meets R_a, it contains a whole component of R_a, which is then isotopic to each of the horizontal boundary components of N. This contradicts that $R_{b,0}$ and $R'_{b,0}$ are not isotopic to any component of R_a. Thus N is disjoint from R_a. If $R_{b,0} \cap R'_{b,0} \neq \emptyset$ then the isotopy defined by N reduces $|R_{b,0} \cap R'_{b,0}|$. If $R_{b,0} \cap R'_{b,0} = \emptyset$ then the isotopy makes $R_{b,0}$ and $R'_{b,0}$ coincide. □

Lemma 4.1.9. *Let R_a, R_b be adjacent vertices of* $\mathrm{MS}(L)$. *Then R_a, R_b can be isotoped so they are disjoint and realise their adjacency.*

Suppose there are components $R_{a,0}$ of R_a and $R_{b,0}$ of R_b that can be made to coincide, so there is a product region between these components. The side of R_a on which this product region lies is determined by the choice of R_a, R_b.

Proof. We regard R_b as fixed, and isotope R_a. Isotope R_a to realise the adjacency. Suppose they are not disjoint, and pick a component $R_{a,0}$ of R_a that is not disjoint from R_b. Because the surfaces realise their adjacency, $R_{a,0}$ cannot cross R_b. Therefore $R_{a,0}$ can be pushed off R_b by a small isotopy. If the two components do not coincide, there is no choice as to which direction to push $R_{a,0}$, and it is clear that the condition in Lemma 4.1.7 continues to hold.

If $R_{a,0}$ coincides with a component of R_b, it is possible to push it off in either direction, creating a product region. We will see that the choice of direction is forced upon us by wanting the condition in Lemma 4.1.7 to continue to hold.

As $R_a \neq R_b$, at least one component $R_{a,1}$ of R_a does not coincide with any component of R_b. Fix a point x_0 of $R_{a,1}$, and choose a product neighbourhood $\mathcal{N}(R_a) = R_a \times [-1,1]$ such that $R_a = R_a \times \{0\}$. For each point y of $R_{a,0}$, choose a path ρ from $(x_0,1)$ to $(y,-1)$ that is disjoint from R_a and transverse to R_b. Since ρ is contained in $M \setminus R_a$, it has algebraic intersection 0 or 1 with R_b. Suppose that this number is not both well-defined and constant on $R_{a,0}$. That is, suppose there are such points y_0, y_1 and paths ρ_0, ρ_1 such that one gives the value 0 while the other gives the value 1. Let ρ' be a path from y_0 to y_1 in $R_{a,0} \times \{-1\}$. Then $\rho' \cup \rho_0 \cup \rho_1$ forms a closed curve that has intersection 0 with R_a but intersection 1 with R_b. This is not possible. We therefore have a well-defined value for the algebraic intersection of R_b and a path as described. If this value is 1, push $R_{a,0}$ downwards, and otherwise push it upwards. Then, for $y \in R_{a,0}$, any path from $(x, 1)$ to $(y, -1)$ that is disjoint from R_a has intersection 1 with R_b.

Now pick points x, y and path ρ as in the condition in Lemma 4.1.7. Isotope ρ so that it decomposes into three paths disjoint from R_a, one from $(x, 1)$ to $(x_0, -1)$, one from $(x_0, -1)$ to $(x_0, 1)$ and one from $(x_0, 1)$ to $(y, -1)$. The outer two paths have intersection 1 with R_b while the middle one has intersection -1 with R_b. Thus ρ has intersection 1 with R_b. □

4.2 An ordering on the vertices of $\mathrm{MS}(L_i)$

In [31], a partial ordering $<_S$ is defined on the vertices of $\mathrm{MS}(L)$ for a link L, relative to a fixed vertex S. This ordering only compares adjacent vertices.

Definition 4.2.1 ([31] Section 5)**.** Let R, R', S be vertices of $\mathrm{MS}(L)$ with R, R' adjacent. Isotope the surfaces so that R, R' are almost transverse to and have simplified intersection with S, and so that R, R' are almost disjoint with simplified intersection. Set $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$. Let $\tilde{M}$ denote the infinite cyclic cover of M, and let τ be the generating covering transformation (in the positive direction). Let V_S be a lift of $M \setminus S$.

Let V_R be the lift of $M \setminus R$ such that $V_R \cap V_S \neq \emptyset$ but $V_R \cap \tau(V_S) = \emptyset$. Finally, let $V_{R'}$ be the lift of $M \setminus R'$ such that $V_{R'} \cap V_R \neq \emptyset$ but $V_{R'} \cap \tau(V_R) = \emptyset$. See Figure 4.1.

Then $R' <_S R$ if $V_{R'} \cap V_S \neq \emptyset$.

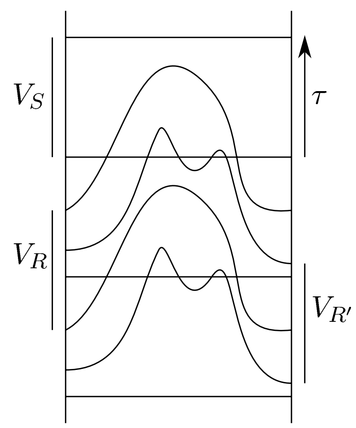

Figure 4.1

Lemma 4.2.2 ([31] Lemma 5.3)**.** *Let R, R' be adjacent vertices, and S any vertex. Then $R' <_S R$ or $R <_S R'$.*

Lemma 4.2.3 ([31] Lemma 5.4)**.** *If $\mathrm{d}_{\mathrm{MS}(L)}(R', S) < \mathrm{d}_{\mathrm{MS}(L)}(R, S)$ then $R <_S R'$.*

Lemma 4.2.4 ([31] Lemma 5.5)**.** *There are no $R_1, \ldots, R_k$, for $k \geq 2$, with $R_1 <_S R_2 <_S \cdots <_S R_k <_S R_1$.*

Now choose L, L_1, L_2 as in the statement of Theorem 1.1.10. These will remain fixed until the end of Section 4.6.

Definition 4.2.5. For $i = 1, 2$, let K_i be the link component of L_i along which the connected sum is performed. Let T_0 be a fixed copy of $\mathbb{S}^2$ that divides L into L_1 and L_2, and choose a product neighbourhood $T_0 \times [1, 2]$ such that $T_0 = T_0 \times \left\{\frac{3}{2}\right\}$ and $T_0 \times \{i\}$ lies on the same side of T_0 as L_i for $i = 1, 2$. We further require that both arcs of $L \cap (T_0 \times [1, 2])$ are of the form $\{x\} \times [1, 2]$ for some $x \in T_0$.

Next, choose a regular neighbourhood $\mathcal{N}(L)$ of L, and let $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$. In addition, let M_0 denote $M \cap (T_0 \times [1, 2])$ and for $i = 1, 2$ denote by M_i the component of $M \setminus (T_0 \times (1, 2))$ that meets L_i. See Figure 4.2a.

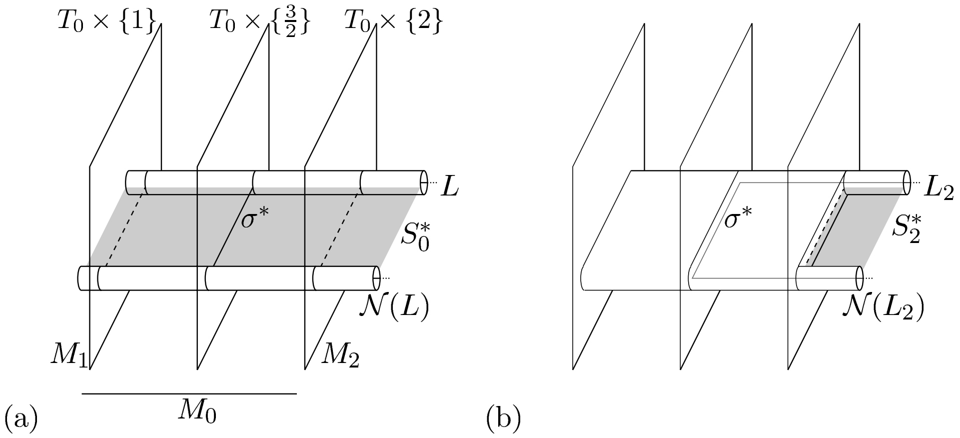

Figure 4.2

Definition 4.2.6. Choose a taut Seifert surface S_0 for L. We will use S_0 as a basepoint for $\mathrm{MS}(L)$. Isotope S_0 to have minimal intersection with T_0. Then $S_0 \cap T_0$ is a single arc σ^*. Further ensure that $S_0 \cap M_0 = \sigma^* \times [1, 2]$. Let S_0^* denote this copy of S_0, considered as a fixed surface rather than up to isotopy. Again, see Figure 4.2a.

The link made up of the part of L on the M_1 side of T_0 together with the arc σ^* is L_1, and M_1 is homeomorphic to $\mathbb{S}^3 \setminus \mathcal{N}(L_1)$. The same is true for L_2.

The sphere T_0 divides S_0^* into Seifert surfaces S_1, S_2 for L_1, L_2 respectively. Since S_0 is taut, so are S_1, S_2. Let S_1^*, S_2^* be these surfaces, again considered as fixed. Define curves $\lambda^*, \lambda_1^*, \lambda_2^*$ on $\partial M, \partial M_1, \partial M_2$ respectively, also seen as fixed, by $\lambda^* = \partial S_0^*$ and $\lambda_i^* = \partial S_i^*$ for $i = 1, 2$. By an appropriate choice of $\mathcal{N}(L_i)$ for $i = 1, 2$ we may ensure that $(\sigma^* \times \{i\}) \cap M \subset \lambda_i^*$. See Figure 4.2b.

Definition 4.2.7. Define $\leq$ on $\mathrm{V}(\mathrm{MS}(L_1))$ by $R \leq R' \Leftrightarrow (R <_{S_1} R'$ or $R = R')$. Similarly, define $\leq$ on $\mathrm{V}(\mathrm{MS}(L_2))$ by $R \leq R' \Leftrightarrow (R <_{S_2} R'$ or $R = R')$.

Corollary 4.2.8. *The pairs* $(\mathrm{MS}(L_1), \leq), (\mathrm{MS}(L_2), \leq)$ *are ordered simplicial complexes.*

Proof. Lemma 4.2.2 gives (P2)′, and Lemma 4.2.4 gives (P1). Together they give (P3). □

Definition 4.2.9. Denote by $\mathcal{Z}$ the simplicial complex with a vertex at each integer and an edge joining $n-1$ to n for each $n \in \mathbb{Z}$, so that $|\mathcal{Z}| \cong \mathbb{R}$. Order the vertices of $\mathcal{Z}$ using the usual order $\leq$ on $\mathbb{Z}$.

These three orderings allow us to apply Theorem 3.6.8. We will use this to give a triangulation of $|\mathrm{MS}(L_1)| \times |\mathrm{MS}(L_2)| \times \mathbb{R}$ that agrees with the triangulation of $\mathrm{MS}(L)$ in a natural way.

4.3 Choosing representatives of isotopy classes

Given taut Seifert surfaces R_i of L_i for $i = 1, 2$, we can isotope the surfaces so that $\partial R_i = \lambda_i^*$. Having done so, we can glue each of the two surfaces to the rectangle $\sigma^* \times [1, 2]$ to form a taut Seifert surface R of L with $\partial R = \lambda^*$.

An isotopy of any Seifert surface R for L can be split into an isotopy that fixes ∂R and an isotopy supported in a neighbourhood of ∂M. Thus an isotopy class relative to the boundary corresponds to a fixed element of the isotopy class together with a winding number for each boundary component. To measure the winding numbers, we need to decide what it means to have winding number 0. We use S_i^* as a basepoint for $\mathrm{MS}(L_i)$, setting this to have winding number 0 at each boundary component. We want to define what it means for another surface to have winding number 0. In practice we will only be concerned with the winding number at K_i, but it is convenient to choose surfaces fixed at every boundary component.

Thus our present aim is to find a fixed representative R^* for each vertex R of $\mathrm{MS}(L_i)$. We choose these such that $\partial R^* = \lambda_i^*$. We also want these representatives to interact well with regard to the ordering $\leq$.

Definition 4.3.1. Let R, R' be ∂–almost disjoint taut Seifert surfaces for L_i. Pick a component K' of L_i, and consider R, R' near K'. It may be that the components that meet K' coincide. If not, one of the two surfaces lies 'above' the other, where this is measured in the positive direction around K'. Write $R \leq_{K'} R'$ if either R' lies above R or the two coincide.

Definition 4.3.2. Define a relation $\leq_\partial$ on isotopy classes of taut Seifert surfaces relative to the boundary by $R \leq_\partial R'$ if there are representatives R_b of R and R'_a of R' such that R_b, R'_a are ∂–almost disjoint and $R_b \leq_{K'} R'_a$ for each component K' of L_i.

Lemma 4.3.3. *The relation $\leq_\partial$ is antisymmetric.*

Proof. Suppose otherwise. Choose R, R' with $R \leq_\partial R' \leq_\partial R$ and $R \neq R'$. Consider R' as fixed, and choose representatives R_a, R_b of R such that R_a shows that $R' \leq_\partial R$ and R_b shows that $R \leq_\partial R'$. Note that $\partial R' = \partial R_a = \partial R_b = \lambda_i^*$.

The surface R_a might be disconnected. However, at least one component of R_a is not isotopic to any component of R'. Remove all other components of R_a, and the corresponding ones of R_b. Also remove all components of R' that are disjoint from the new R_a.

As R_a, R_b only meet a neighbourhood of R' along their boundaries, where we know how they are positioned, we see that R_a, R_b can be put into general position by an isotopy away from a neighbourhood of R'. Choose this isotopy to minimise $|R_a \cap R_b|$.

Consider the product region N between R_a and R_b given by Corollary 4.1.2. The isotopy of R_a defined by N does not move ∂R_a, so our positioning of R_a, R_b means that N must meet R'. Let R'_0 be a component of R' that meets N. Then $R'_0 \subset N$. Since R'_0 is ∂–almost disjoint from R_a and from R_b, Proposition 1.4.2 gives that R'_0 is isotopic to R_a, which is a contradiction. □

Proposition 4.3.4. *It is possible to choose a representative R^* for each vertex R of* $\mathrm{MS}(L_i)$ *such that the following conditions hold for any pair (R_a, R_b) of adjacent vertices of* $\mathrm{MS}(L_i)$.

- $\partial R_a^* = \partial R_b^* = \lambda_i^*$.
- *There are ∂–almost disjoint copies R'_j of R_j, for $j = a, b$, that are isotopic to R_j^* via isotopies fixing $\partial R'_j$ and that demonstrate their adjacency.*
- $R_b^* \leq_\partial R_a^*$ *(equivalently, $R'_b \leq_\partial R'_a$) if and only if $R_b \leq R_a$.*

Proof. Without loss of generality, $i = 1$. Let $\tilde{M}_1$ be the infinite cyclic cover of M_1, with covering transformation τ_1. As in Proposition 1.3.8, construct a lift V_n of $M_1 \setminus S_1^*$ for $n \in \mathbb{Z}$. Let $\tilde{S}_1$ be the lift of S_1^* that lies between V_0 and V_1.

We have already chosen the representative S_1^* for S_1. Let R be a vertex of $\mathrm{MS}(L_1)$ other than S_1. Isotope R to be almost transverse to and have simplified intersection with S_1^*. Let V_R be the lift of $M_1 \setminus R$ such that $V_R \cap V_0 \neq \emptyset$ but $V_R \cap V_1 = \emptyset$, and let $\tilde{R}$ be the lift of R that lies between V_R and $\tau_1(V_R)$. From Proposition 4.1.6 we see that R minimises $\max\{n : V_R \cap V_n \neq \emptyset\} - \min\{n : V_R \cap V_n \neq \emptyset\}$. By Proposition 1.3.8 this means that $\{n : V_R \cap V_n \neq \emptyset\} = \{0, -1, \ldots, -\mathrm{d}_{\mathrm{MS}(L_1)}(S_1, R)\}$.

By an isotopy close to the boundary, move ∂R around $\partial \mathcal{N}(L_1)$ until $\partial \tilde{R}$ coincides with $\partial \tilde{S}_1$. Note that this does not change $\{n : V_R \cap V_n \neq \emptyset\}$. Take the resulting surface to be R^*.

Our first aim is to show that the choice of R^* is well-defined up to isotopy relative to the boundary. That is, we wish to check that the choice of winding number is independent of the choice of isotopy made when constructing R^*. Let R' be any other copy of R. Construct $(R')^*$ as described. Then $\partial \tilde{R} = \partial \tilde{R}' = \partial \tilde{S}_1$. Note that we do not know how R and R' intersect. However, R is isotopic to R' in M_1. This isotopy lifts to an isotopy in $\tilde{M}_1$ from $\tilde{R}$ to a lift of R'. This lift is $\tau_1^m(\tilde{R}')$ for some $m \in \mathbb{Z}$. The isotopy also takes V_R to $\tau_1^m(V_{R'})$.

Suppose $m \neq 0$. Without loss of generality, $m > 0$. Then the submanifold $\tau_1^m(V_{R'} \cup \tilde{R}' \cup \tau_1^{-1}(\tilde{R}'))$ of $\tilde{M}_1$ has interior disjoint from the submanifold $\tau_1^{-k}(V_0 \cup \tilde{S}_1 \cup \tau_1^{-1}(\tilde{S}_1))$, where $k = \mathrm{d}_{\mathrm{MS}(L_1)}(S_1, R)$. Thus by Proposition 4.1.4 we see that V_R is disjoint from V_{-k}. This contradicts that $\{n : V_R \cap V_n \neq \emptyset\} = \{0, -1, \ldots, -k\}$. Hence $m = 0$. We can therefore modify the isotopy from R to R' near the boundary to keep ∂R fixed throughout. Thus R^* and $(R')^*$ are isotopic relative to the boundary, as required.

It remains to show that our chosen representatives have the required properties. Let R_a, R_b be adjacent vertices of $\mathrm{MS}(L_1)$ with $R_b \leq R_a$. Position them relative to S_1^* and each other as in Lemma 4.1.5. Then by Proposition 4.1.6 we know that R_a, R_b realise their adjacency, and so are almost disjoint.

Now consider the lifts used to demonstrate that $R_b <_{S_1} R_a$, as in Definition 4.2.1. Taking V_0 as the lift of $M_1 \setminus S_1^*$ we see that V_{R_a} is the required lift of $M_1 \setminus R_a$. Let V'_{R_b} be the required lift of $M_1 \setminus R_b$. Then V'_{R_b} meets V_0 but does not meet V_1. Thus $V'_{R_b} = V_{R_b}$. This means that $\tilde{R}_b$ lies below $\tilde{R}_a$ in $\tilde{M}_1$. Now isotope R_a and R_b near the boundary to form R_a^* and R_b^*. Then it is clear that $\partial R_a^* = \partial R_b^* = \lambda_1^*$ and $R_b^* <_\partial R_a^*$. In addition, they continue to realise their adjacency and are ∂–almost disjoint. □

Remark 4.3.5. Suppose R_a and R_b are positioned as required, and that $R_b \leq R_a$. Suppose a component of R_b coincides with a component of R_a. By Lemma 4.1.9 there is a unique direction we can push the component of R_b off that of R_a so that they continue to realise their adjacency. By examining the construction above we see that this direction is downwards. To see this, note that the components must coincide when we lift them to the infinite cyclic cover, as the boundaries of the lifts coincide.

4.4 Mapping the vertices

Definition 4.4.1. Define a map $\Psi\colon \mathrm{V}(\mathrm{MS}(L_1)) \times \mathrm{V}(\mathrm{MS}(L_2)) \times \mathbb{Z} \to \mathrm{V}(\mathrm{MS}(L))$ as follows. Let $(R_1, R_2, n) \in \mathrm{V}(\mathrm{MS}(L_1)) \times \mathrm{V}(\mathrm{MS}(L_2)) \times \mathbb{Z}$. Take the copy of R_1^* in M_1, and the copy of R_2^* in M_2. Join these together by a rectangle in M_0 that winds n times around L. Here we measure the winding number of the rectangle around the arc of L that runs through M_0 and is oriented from L_1 to L_2, with respect to where λ^* lies on the boundary of the neighbourhood of this arc. Let R be the resulting surface. Note that R is a taut Seifert surface for L. If $n \neq 0$ then $\partial R \neq \lambda^*$, but $[\partial R] = [\lambda^*]$. We set $\Psi(R_1, R_2, n)$ to be the isotopy class of R.

Lemma 4.4.2. *Ψ is surjective.*

Proof. Let R be a taut Seifert surface for L. Isotope R to have minimal intersection with T_0, and so that $R \cap M_0 = \rho \times [1, 2]$ for some arc $\rho \subset T_0$. Then T_0 cuts R into taut Seifert surfaces R_i of L_i for $i = 1, 2$. Isotope R_i to R_i^*. This isotopy may move ∂R_i in M_i. However, the isotopy of M_i can be extended to an isotopy of M by also isotoping the rectangle $R \cap M_0$ in M_0 near $T_0 \times \{i\}$. After these isotopies, we can read off the winding number n of the rectangle. Then $\Psi(R_1, R_2, n) = R$. □

Lemma 4.4.3. *Ψ is injective.*

Proof. Suppose $\Psi(R_{a,1}, R_{a,2}, n_a) = \Psi(R_{b,1}, R_{b,2}, n_b)$. We first aim to show that $R_{a,1} = R_{b,1}$ and $R_{a,2} = R_{b,2}$. Let R_a, R_b be the fixed surfaces constructed by Ψ. Note that, by construction, R_a and R_b each meet T_0 in a single arc, and these arcs either coincide or are disjoint. We can ensure that the arcs are disjoint using an isotopy of R_a that breaks into an isotopy of $R_{a,1}$ and an isotopy of $R_{a,2}$. By an isotopy of R_a keeping T_0 fixed pointwise, R_a and R_b can be made almost transverse.

As the surfaces are isotopic, there is a product region N between them. If $\mathrm{int}(N)$ meets T_0 then $T_0 \cap N$ is a product disc in the sutured manifold N, and divides N into product regions N_1, N_2. The isotopy of R_a defined by N therefore breaks into an isotopy of $R_{a,1}$ and an isotopy of $R_{a,2}$. It either reduces $|R_a \cap R_b|$ or makes components of R_a and R_b coincide. Inductively we find that $R_{a,1} = R_{b,1}$ and $R_{a,2} = R_{b,2}$.

It remains to show that $\Psi(R_{a,1}, R_{a,2}, n_a) \neq \Psi(R_{a,1}, R_{a,2}, n_b)$ for $n_a > n_b$. Suppose otherwise. Without loss of generality, $n_b = 0$. Consider the fixed surfaces $R_a = \Psi(R_{a,1}, R_{a,2}, n_a)$, $R_0 = \Psi(R_{a,1}, R_{a,2}, 0)$. Push the component of the copy of $R_{a,1}$ in R_a that meets K_1 upwards off R_0, and the copy of $R_{a,2}$ downwards. Delete all other components of each surface. By the assumption that $\Psi(R_{a,1}, R_{a,2}, n_a) = \Psi(R_{a,1}, R_{a,2}, 0)$, the two remaining surfaces are isotopic, so there is a product region N between them. Note that N meets $\partial\mathcal{N}(K_1 \# K_2)$. By considering the boundary curves of the surfaces on $\partial\mathcal{N}(K_1 \# K_2)$, we can restrict the possibilities for the location of N relative to R_a, R_0. To see this, note that N only meets one side of each of the orientable surfaces R_a, R_0. Figure 4.3 shows the boundary patterns in the cases $n_a \in \{1, 2, 3, 4\}$. In general we see that N must meet the complement of one of the L_i in the complement of the surface $R_{a,i}$. Suppose this is L_1. Then $T_0 \times \{1\}$ is a product disc in N. This means $\mathbb{S}^3 \setminus \mathcal{N}(R_{a,1})$ is a product region, showing that L_1 is fibred, which is a contradiction. □

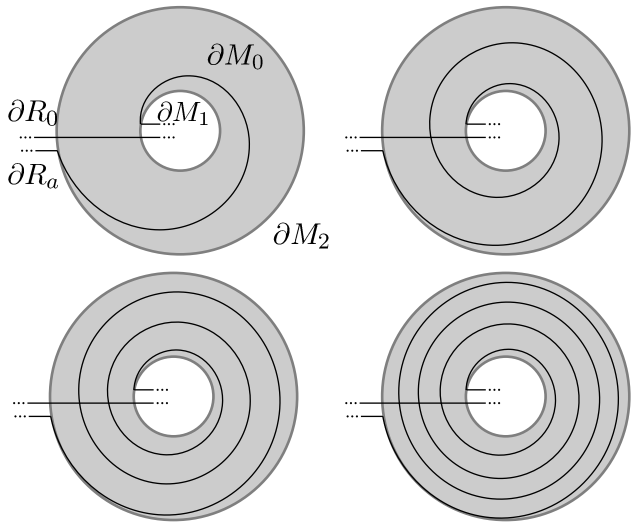

Figure 4.3

4.5 Mapping the edges

Lemma 4.5.1. *Let $R_{a,i}, R_{b,i}$ be fixed almost disjoint taut Seifert surfaces that demonstrate that their isotopy classes are adjacent in* $\mathrm{MS}(L_i)$ *for $i = 1, 2$. Suppose that there are arcs $\rho_j \subset T_0$ for $j = a, b$ such that $R_{j,i} \cap (T_0 \times \{i\}) = \rho_j \times \{i\}$ for $i = 1, 2$. Suppose further that ρ_a and ρ_b are disjoint. Let $R_j = R_{j,1} \cup R_{j,2} \cup (\rho_j \times [1, 2])$ for $j = a, b$. Then R_a, R_b demonstrate that their isotopy classes are adjacent in* $\mathrm{MS}(L)$.

Proof. The pairs $R_{a,i}, R_{b,i}$ satisfy the condition given in Lemma 4.1.7. We wish to show that R_a, R_b also satisfy this condition. First note that R_a and R_b are almost disjoint.

Choose points $x, y \in R_a \setminus R_b$, an appropriate product neighbourhood of R_a and a path ρ from $(x, 1)$ to $(y, -1)$ that is disjoint from R_a and transverse to R_b. By a small isotopy, ρ also can be made transverse to T_0. If ρ is disjoint from T_0, it has algebraic intersection 1 with R_b, as required.

Suppose otherwise. Let x_0 be the first point at which ρ meets T_0. Find an arc ρ_0 from $R_a \times \{1\}$ to $R_a \times \{-1\}$ such that ρ_0 lies entirely in T_0 and passes through x_0. We can then use ρ_0 to split ρ into three paths as in the condition in Lemma 4.1.7, each of which (up to isotopy) has strictly fewer points of intersection with T_0. The first path runs along ρ as far as x_0, and then follows ρ_0 up to $R_a \times \{-1\}$. The second is ρ_0 traversed backwards. The third runs along ρ_0 to x_0, and then runs along the rest of ρ. The first two paths can be made disjoint from T_0, and run in opposite directions. Hence removing them does not change the algebraic intersection number with R_b. In this way we can remove all points of $\rho \cap T_0$. $\square$

Corollary 4.5.2. *Let* $(R_{a,1}, R_{a,2}, n_a)$ *and* $(R_{b,1}, R_{b,2}, n_b)$ *be in* $\mathrm{V}(\mathrm{MS}(L_1)) \times \mathrm{V}(\mathrm{MS}(L_2)) \times \mathbb{Z}$ *with* $R_{b,2} \leq R_{a,2}$. *Suppose these triples are distinct and one of the following conditions holds.*

- $R_{a,1} \geq R_{b,1}$ *and* $n_a = n_b$.
- $R_{a,1} \leq R_{b,1}$ *and* $n_a = n_b - 1$.

Then $\Psi(R_{a,1}, R_{a,2}, n_a)$ *and* $\Psi(R_{b,1}, R_{b,2}, n_b)$ *are adjacent in* $\mathrm{MS}(L)$.

Proof. Viewing $\Psi(R_{a,1}, R_{a,2}, n_a)^*$ as fixed, we build a copy of $\Psi(R_{b,1}, R_{b,2}, n_b)$ satisfying the hypotheses of Lemma 4.5.1. Without loss of generality, $n_a = 0$.

Isotope $R^*_{b,2}$ as in Proposition 4.3.4 to realise its adjacency with $R^*_{a,2}$. For any components of $R_{b,2}$ that do not coincide with those of $R^*_{a,2}$, perform a small isotopy near the boundary to move the boundary of $R_{b,2}$ to below that of $R^*_{a,2}$, making the components disjoint. Consider the pair of components that meet K_2. If these components of $R^*_{a,2}, R_{b,2}$ coincide then push that in $R_{b,2}$ downwards off $R^*_{a,2}$. By Remark 4.3.5, $R_{a,2}$ and $R_{b,2}$ still realise their adjacency.

Case 1: $R_{a,1} \geq R_{b,1}$ and $n_b = n_a = 0$. Then $R_{b,1} \leq_\partial R_{a,1}$. Perform isotopies on $R^*_{b,1}$ analogous to those performed on $R^*_{b,2}$. A flat rectangle can then be inserted connecting the boundaries of the surfaces $R_{b,1}, R_{b,2}$.

Case 2: $R_{a,1} \leq R_{b,1}$ and $n_b = n_a + 1 = 1$. Then $R_{a,1} \leq_\partial R_{b,1}$. This time isotope $R^*_{b,1}$ upwards instead of downwards, so the boundary of $R_{b,1}$ lies above that of $R^*_{a,1}$ wherever they do not coincide. Add in a rectangle that wraps nearly once around $K_1 \# K_2$, to again join up the boundaries of $R_{b,1}, R_{b,2}$.

It is clear that, in either case, $\Psi(R_{b,1}, R_{b,2}, n_b)^*$ is isotopic to the surface R_b we have constructed. We may now apply Lemma 4.5.1 to complete the proof. □

Proposition 4.5.3. *Let* $(R_{a,1}, R_{a,2}, n_a)$ *and* $(R_{b,1}, R_{b,2}, n_b)$ *be in* $\mathrm{V}(\mathrm{MS}(L_1)) \times \mathrm{V}(\mathrm{MS}(L_2)) \times \mathbb{Z}$. *Suppose* $\Psi(R_{a,1}, R_{a,2}, n_a)$ *and* $\Psi(R_{b,1}, R_{b,2}, n_b)$ *are adjacent in* $\mathrm{MS}(L)$. *Then one of the conditions in Corollary 4.5.2 holds.*

Proof. Without loss of generality, $n_a = 0$. Fix $\Psi(R_{a,1}, R_{a,2}, n_a)^*$ and isotope $\Psi(R_{b,1}, R_{b,2}, n_b)$ to be disjoint from it, realising the adjacency in $\mathrm{MS}(L)$. Since $T_0 \cap (M_0 \setminus \Psi(R_{a,1}, R_{a,2}, n_a)^*)$ is a disc, by standard methods we can also ensure that this copy of $\Psi(R_{b,1}, R_{b,2}, n_b)$ meets T_0 in a single arc. As in the proof of Lemma 4.4.2, dividing the surface along T_0 gives (fixed) Seifert surfaces $R'_{b,i}$ in $\mathrm{MS}(L_i)$ for $i = 1, 2$ and there is an integer n'_b such that $\Psi(R'_{b,1}, R'_{b,2}, n'_b)$ is isotopic to $\Psi(R_{b,1}, R_{b,2}, n_b)$. See Figure 4.5. By Lemma 4.4.3 we have in particular that $R'_{b,i}$ is isotopic to $R_{b,i}$ in M_i for $i = 1, 2$. From this we see that $R^*_{a,i}, R'_{b,i}$ demonstrate that $R_{a,i}, R_{b,i}$ are at most distance 1 apart in $\mathrm{MS}(L_i)$. We may now assume that $R_{b,2} \leq R_{a,2}$.

It now remains to verify that n_b takes the required value. Our approach is to position $\Psi(R_{b,1}, R_{b,2}, n_b)$ 'close to' $R^*_{b,i}$ for $i = 1, 2$ without affecting the relative positions of $\Psi(R_{b,1}, R_{b,2}, n_b)$ and $\Psi(R_{a,1}, R_{a,2}, n_a)^*$. Having done so, we will be able to read off the value of n_b from $\Psi(R_{b,1}, R_{b,2}, n_b)$.

For $i = 1, 2$, consider $R^*_{a,i}$ and $R^*_{b,i}$, which satisfy the conclusions of Proposition 4.3.4. As an isotopy of $R^*_{b,i}$ that fixes its boundary will not affect the winding number n_b, we may assume that no such isotopy is needed in this case. That is, we assume that $R^*_{a,i}$ and $R^*_{b,i}$ are ∂–almost disjoint and realise their adjacency. We further assume that components of $R^*_{a,i}$ and $R^*_{b,i}$ coincide whenever this is possible without moving the boundary of either surface.

Suppose that a component of $R^*_{b,i}$ and a component of $R^*_{a,i}$ bound a product region N in M_i, but that any isotopy from one to the other moves the boundary of the surface. Since the boundaries of $R^*_{a,i}$ and $R^*_{b,i}$ coincide, the component of $R^*_{b,i}$ is given by taking a parallel copy of that of $R^*_{a,i}$ and moving its boundary once around L_i, as shown in Figure 4.4. Note that N lies below $R^*_{a,i}$ if $R^*_{a,i} \leq_\partial R^*_{b,i}$ and N lies above $R^*_{a,i}$ if $R^*_{b,i} \leq_\partial R^*_{a,i}$. In

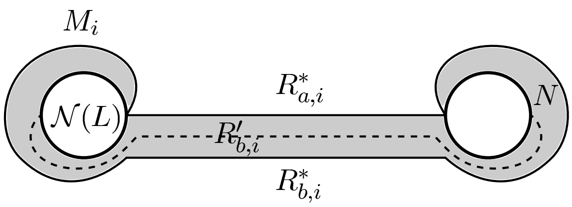

Figure 4.4

addition, a component of $R'_{b,i}$ must be contained in the product region N and be isotopic to the component of $R^*_{a,i}$. In particular, from Proposition 4.3.4 we see that if $R^*_{a,i} \leq R^*_{b,i}$ then the component of $R'_{b,i}$ lies above that of $R^*_{b,i}$ and its boundary can be seen as lying above that of $R^*_{a,i}$, whereas if $R_{b,i} \leq R_{a,i}$ then the component of $R'_{b,i}$ lies below that of $R^*_{b,i}$ and its boundary can be seen as lying below that of $R^*_{a,i}$.

For $i = 1, 2$, temporarily ignore the components of each surface that meet K_i. If a component of $R^*_{b,i}$ coincides with a component of $R^*_{a,i}$, then there is a product region between it and the corresponding component of $R'_{b,i}$. Such components will play no part in the rest of the proof, so we may assume that there are none.

For each other component of $R^*_{b,i}$, there are two possibilities. If it is isotopic to a component of $R^*_{a,i}$, we have already seen how it relates to $R'_{b,i}$. Otherwise, take a copy of this component of $R_{b,i}$ that lies parallel to $R^*_{b,i}$ and is disjoint from $R^*_{a,i}$. That is, there is a product region between $R^*_{b,i}$ and this copy of $R_{b,i}$ that is contained within a product neighbourhood of $R^*_{b,i}$. Furthermore, this product region lies above $R^*_{b,i}$ if $R^*_{a,i} \leq_\partial R^*_{b,i}$ and lies below $R^*_{b,i}$ if $R^*_{b,i} \leq_\partial R^*_{a,i}$. By Lemma 4.1.8, the corresponding component of $R'_{b,i}$ is now isotopic to that of $R_{b,i}$ by an isotopy disjoint from $R^*_{a,i}$. We may therefore assume that these components of $R_{b,i}, R'_{b,i}$ now coincide. Hence for each component of $R'_{b,i}$ away from K_i we have the following. If $R_{a,i} \leq R_{b,i}$ then the boundary of $R'_{b,i}$ lies above that of $R^*_{a,i}$. If $R_{b,i} \leq R_{a,i}$ then the boundary of $R'_{b,i}$ lies below that of $R^*_{a,i}$. Recall that we have assumed that no such components exist if $R_{a,i} = R_{b,i}$, and that $R_{b,2} \leq R_{a,2}$.

We now turn our attention to the components that meet K_i.

Case 1: This component of $R^*_{b,i}$ does not coincide with $R^*_{a,i}$ for $i = 1, 2$. Note that this means $R_{a,i} \neq R_{b,i}$, because $R^*_{a,i}$ and $R^*_{b,i}$ are not isotopic by an isotopy fixing their boundaries. In this case, the corresponding component of $R'_{b,i}$ can also be isotoped as just described. There is now an essentially unique way to join the two surfaces by a rectangle in M_0 that is disjoint from $\Psi(R_{a,1}, R_{a,2}, n_a)^*$. If $R_{b,1} \leq R_{a,1}$ then the rectangle must be flat, so $n_b = 0 = n_a$. If $R_{a,1} \leq R_{b,1}$ then the rectangle must twist nearly once around $\partial\mathcal{N}(K_1 \# K_2)$ from above $R^*_{a,1}$ to below $R^*_{a,2}$ (that is, twisting in the positive direction around $K_1 \# K_2$), so $n_b = 1 = n_a + 1$. See Figure 4.5.

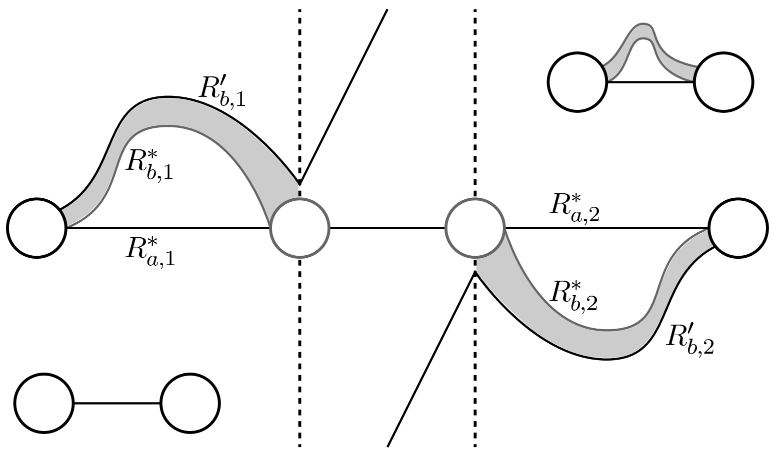

Figure 4.5

Case 2: The component of $R^*_{b,2}$ coincides with that of $R^*_{a,2}$, but those of $R^*_{a,1}, R^*_{b,1}$ do not coincide. Then $R_{a,1} \neq R_{b,1}$. Again isotope $R'_{b,1}$ as above. The component of $R'_{b,2}$ that meets K_2 is isotopic in M_2 to that of $R^*_{a,2}$, meaning there is a product region N between them in M_2. Note that the interior of N is disjoint from the three surfaces $R^*_{a,2}, R^*_{b,2}, R'_{b,2}$. Once we know which side of $R^*_{a,2}$ the product region N lies (this is well-defined, by Lemma 4.1.9), there is only one possibility for the rectangle joining the surfaces $R'_{b,1}, R'_{b,2}$, as in Case 1.

If $R_{a,2} = R_{b,2}$ then we may further assume that $R_{b,1} \leq R_{a,1}$, and in particular that the boundary of $R'_{b,1}$ lies below that of $R^*_{a,1}$. If N lies below $R^*_{a,2}$ then the rectangle is flat, so $n_b = 0 = n_a$. If N lies above $R^*_{a,2}$ then the rectangle runs in the negative direction around $K_1 \# K_2$, from below $R^*_{a,1}$ to above $R^*_{a,2}$. Hence $n_b = -1 = n_a - 1$. Note that this satisfies the condition in Corollary 4.5.2 with the two surfaces switched.

Suppose instead that $R_{a,2} \neq R_{b,2}$. If N lies below $R^*_{a,2}$ then there are two possible situations. One is that $R_{b,1} \leq R_{a,1}$ and $n_b = n_a$. The other is that $R_{a,1} \leq R_{b,1}$ and $n_b = n_a + 1$. It therefore remains to rule out the possibility that N lies above $R^*_{a,2}$. For this we must return to the definitions of $R^*_{a,2}, R^*_{b,2}$. Originally we chose these to be ∂–almost disjoint, with $R^*_{b,2} \leq_\partial R^*_{a,2}$. As they now coincide, the components that meet K_2 were isotopic by an isotopy keeping their boundaries fixed. Corollary 4.1.2 therefore shows there was a product region N' between these components that lay below $R^*_{a,2}$ and above $R^*_{b,2}$. However, Lemma 4.1.9 says that the side of $R^*_{a,2}$ on which N lies is determined by the choice of surfaces $R_{a,2}, R_{b,2}$. Therefore N cannot lie above $R^*_{a,2}$.

Case 3: The component of $R_{b,1}$ coincides with that of $R^*_{a,1}$, but those of $R^*_{a,2}, R_{b,2}$ do not coincide. This is similar to case 2.

Case 4: Both pairs of components coincide. This case again uses the same ideas as cases 1 and 2. The only situation that is very different is when $R_{a,1} = R_{b,1}$ and $R_{a_2} = R_{b,2}$, when we must show that $n_a \neq n_b$. However this is true since $\Psi(R_{a,1}, R_{a,2}, n_a) \neq \Psi(R_{b,1}, R_{b,2}, n_b)$. □

Figure 4.6 is a schematic picture of the local structure of MS(L). Figure 4.6a focuses on the edges radiating from a single vertex. Figure 4.6b shows one of the smaller cubes in Figure 4.6a, with all edges included.

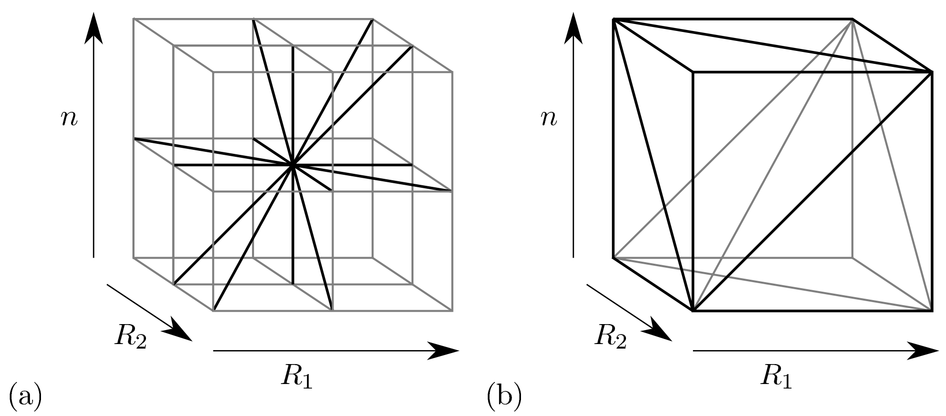

Figure 4.6

4.6 Completing and interpreting the proof

Theorem 1.1.10. *Let L_1, L_2 be non-split, non-fibred links in $\mathbb{S}^3$, and let $L = L_1 \# L_2$. Then $|\operatorname{MS}(L)|$ is homeomorphic to $|\operatorname{MS}(L_1)| \times |\operatorname{MS}(L_2)| \times \mathbb{R}$.*

Proof. Define an ordering $\leq_1$ on $\operatorname{V}(\operatorname{MS}(L_1)) \times \operatorname{V}(\mathcal{Z})$ by

$$(R_b, n_b) \leq_1 (R_a, n_a) \iff (R_b \leq R_a \text{ and } n_b \leq n_a)\,.$$

Let $\mathcal{X}$ be the complex $(\operatorname{MS}(L_1) \times \mathcal{Z}, \leq_1)$. Applying Theorem 3.6.8 to the ordered simplicial complexes $(\operatorname{MS}(L_1), \leq)$, $(\mathcal{Z}, \leq)$ and $(\mathcal{X}, \leq_1)$ gives that

$$|\mathcal{X}| \cong |\operatorname{MS}(L_1)| \times |\mathcal{Z}|.$$

Now define a second ordering $\leq_2$ on $\operatorname{V}(\mathcal{X}) = \operatorname{V}(\operatorname{MS}(L_1)) \times \operatorname{V}(\mathcal{Z})$ by

$$(R_b, n_b) \leq_2 (R_a, n_a) \iff \big((R_b \leq R_a \text{ and } n_b = n_a) \text{ or } (R_a \leq R_b \text{ and } n_a = n_b - 1)\big).$$

Note that this corresponds to the conditions in Corollary 4.5.2. It can be checked that $\leq_2$ has properties (P1), (P2)$'$, (P3).

Define $\leq_3$ on $\operatorname{V}(\operatorname{MS}(L_1) \times \mathcal{Z}) \times \operatorname{V}(\operatorname{MS}(L_2))$ by

$$\big((R_{a,1}, n_a), R_{a,2}\big) \leq_3 \big((R_{b,1}, n_b), R_{b,2}\big) \iff \big((R_{a,1}, n_a) \leq_2 (R_{b,1}, n_b) \text{ and } R_{a,2} \leq R_{b,2}\big).$$

Let $\mathcal{Y}$ be the complex $(\mathcal{X} \times \operatorname{MS}(L_2), \leq_2)$. Applying Theorem 3.6.8 to $(\mathcal{X}, \leq_2)$, $(\operatorname{MS}(L_2), \leq)$ and $(\mathcal{Y}, \leq_3)$ gives that

$$|\mathcal{Y}| \cong |\mathcal{X}| \times |\operatorname{MS}(L_2)|.$$

Thus

$$|\mathcal{Y}| \cong |\operatorname{MS}(L_1)| \times |\operatorname{MS}(L_2)| \times \mathbb{R}.$$

By Lemmas 4.4.2 and 4.4.3, the map $\psi\colon \operatorname{V}(\operatorname{MS}(L_1)) \times \operatorname{V}(\operatorname{MS}(L_2)) \times \operatorname{V}(\mathcal{Z}) \to \operatorname{V}(\operatorname{MS}(L))$ defined in Definition 4.4.1 is a bijection. Recall that

$$\operatorname{V}(\mathcal{Y}) = \operatorname{V}(\operatorname{MS}(L_1)) \times \operatorname{V}(\operatorname{MS}(L_2)) \times \operatorname{V}(\mathcal{Z}).$$

From Corollary 4.5.2 and Proposition 4.5.3 we see that ψ extends to an isomorphism between the 1–skeleta of the complexes $\operatorname{MS}(L)$ and $\mathcal{Y}$. It remains only to note that both of these complexes are flag. For $\operatorname{MS}(L)$ this is the case by definition. For $\mathcal{Y}$ it follows from the fact that the three complexes $\operatorname{MS}(L_1)$, $\operatorname{MS}(L_2)$ and $\mathcal{Z}$ are flag. □

By examining the proof of Theorem 3.6.8 in [10] we can give the following geometric description of the extension of Ψ to $\mathcal{Y} = (\operatorname{MS}(L_1) \times \mathcal{Z}) \times \operatorname{MS}(L_2)$.

Remark 4.6.1. Let $x \in |\operatorname{MS}(L_1)| \times |\operatorname{MS}(L_2)| \times \mathbb{R}$. Without loss of generality, $\pi_{\mathcal{Z}}(x) \in [0,1)$. Let $\pi_{\operatorname{MS}(L_1)}(x) = a_0 A_0 + \cdots + a_m A_m$ where $A_i \in \operatorname{V}(\operatorname{MS}(L_1))$ and $a_i > 0$ for $0 \le i \le m$, with $\sum_{i=0}^{m} a_i = 1$ and $A_0 \le A_1 \le \cdots \le A_m$. Similarly let $\pi_{\operatorname{MS}(L_2)}(x) = b_0 B_0 + \cdots + b_n B_n$.

Consider the surfaces $A_0, \ldots, A_m$. As in Lemma 4.1.5, they can be positioned in M_1 so they are pairwise almost disjoint with simplified intersection. By Proposition 4.1.6, they then realise their adjacencies. As in the proof of Lemma 4.1.9, the surfaces can be made disjoint while still realising their adjacencies, and it can be shown that the boundaries of the surfaces occur in order around ∂M_1. For $i = 0, \ldots, m$, thicken the Seifert surface A_i to a product region $A_i \times [0, a_i]$, and view this as a 'continuum of surfaces'.

Do the same for the Seifert surfaces $B_0, \ldots, B_n$ in M_2. Glue the thickened surfaces to give thickened Seifert surfaces for L in M. In doing so, instead of aligning $A_0 \times \{0\}$ with $B_0 \times \{0\}$, introduce a shift of length $\pi_{\mathcal{Z}}(x)$. This creates a finite set of vertices of $\operatorname{MS}(L)$, each with a weight given by its thickness. Applying Lemma 4.5.1 shows that these vertices span a simplex.

4.7 Incompressible surfaces

In addition to $\operatorname{MS}(L)$, Kakimizu defined a larger complex $\operatorname{IS}(L)$, which records all incompressible Seifert surfaces for L rather than just taut ones.

Definition 4.7.1 (see [19], [31])**.** Let L be a link, and let $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$. Define $\operatorname{IS}(L)$ of L to be the following flag simplicial complex. Its vertices are ambient isotopy classes of incompressible Seifert surfaces for L. Two vertices span an edge if they have representatives R, R' such that a lift of $M \setminus R'$ to the infinite cyclic cover of M intersects exactly two lifts of $M \setminus R$.

Note that $\operatorname{MS}(L)$ is a subcomplex of $\operatorname{IS}(L)$. Proposition 1.3.8 holds for $\operatorname{IS}(L)$ as well as for $\operatorname{MS}(L)$, and so do Propositions 4.1.4 and 4.1.6, and Lemma 4.1.5. The same is true of Lemmas 4.2.2, 4.2.3 and 4.2.4.

Let R be an incompressible Seifert surface for L. Isotope R to have minimal intersection with T_0. Then $R \cap T_0$ is a single arc. Splitting R along T_0 gives incompressible Seifert surfaces R_1, R_2 for L_1, L_2 respectively.

Now consider the converse situation. That is, take incompressible Seifert surfaces R_1, R_2 for L_1, L_2 respectively, and join them along an arc in T_0 to form a Seifert surface R for L.

Lemma 4.7.2. *R is incompressible.*

Proof. Suppose otherwise. Choose a compressing disc S for R that minimises its intersection with T_0 over all compressing discs for R. Then $S \cap T_0$ does not include any simple closed curves. In addition, S is not disjoint from T_0, as otherwise it would be a compressing disc for either R_1 or R_2. Let ρ be an arc of $S \cap T_0$ that is outermost in the disc $T_0 \cap (M_0 \setminus R)$. Then ρ cuts off a subdisc S_T of $T_0 \cap (M_0 \setminus R)$ that is disjoint on its interior from S. It also divides S into two discs S_1 and S_2. Since $|S \cap T_0|$ cannot be reduced, each of the discs $S_1 \cup S_T$ and $S_2 \cup S_T$ is a compressing disc for R. Furthermore, each can be isotoped to have a smaller intersection with T_0 than S does, which is a contradiction. □

Now replacing taut Seifert surfaces with incompressible Seifert surfaces in the proof of Theorem 1.1.10 gives the following.

Theorem 4.7.3. *Let L_1, L_2 be non-split, non-fibred, links in $\mathbb{S}^3$, and let $L = L_1 \# L_2$. Then $|\operatorname{IS}(L)|$ is homeomorphic to $|\operatorname{IS}(L_1)| \times |\operatorname{IS}(L_2)| \times \mathbb{R}$.*

Chapter 5

Locally infinite Kakimizu complexes

5.1 A knot with locally infinite MS

Theorem 5.1.1. *Let L_κ be the twisted Whitehead double of the trefoil shown in Figure 5.1. Then* $\mathrm{MS}(L_\kappa)$ *is not locally finite.*

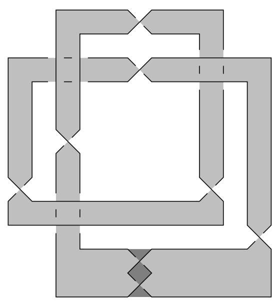

Figure 5.1

Proof. Let R be the genus 1 Seifert surface for the knot L_κ shown in Figure 5.1 (note that every Whitehead double has such a Seifert surface). We construct an infinite family of genus 1 Seifert surfaces for L_κ that are disjoint from R.

Let $M = \mathbb{S}^3 \setminus \mathcal{N}(L_\kappa)$. Let T be the torus that bounds the trefoil knot companion of L_κ, such that L_κ lies in the solid torus bounded by T. In addition, let M_1 be the part of M outside of T as drawn in Figure 5.2 (that is, the side away from the knot), and M_0 the part on the inside. Let μ be a meridian of $T \subset \mathbb{S}^3$. There is a Möbius band properly embedded in M_1, the boundary of which is a longitude λ of the solid torus bounded by T. Then λ and μ are as shown in Figure 5.2. Let S_1 be the annulus properly embedded in M_1 that is contained in the boundary of a regular neighbourhood of this Möbius band in M_1. Then ∂S_1 is two copies of λ, with opposite orientations. Let S_T be one of the two annuli into which T is divided by ∂S_1.

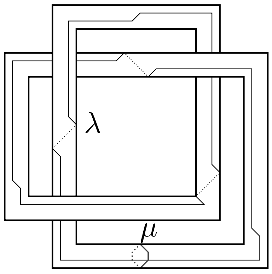

Figure 5.2

R is a plumbing of two annuli S_0 and S_0' in M_0, where S_0 is parallel to S_T in $\mathbb{S}^3 \setminus \text{int}(M_1)$. Isotope R in M so that $R \cap T = S_T$, keeping ∂R fixed. Let R_0 be the Seifert surface for L_κ given by removing S_T from R and replacing it with S_1. Then $|R_0 \cap T| = 2$. In addition, R_0 can be made disjoint from R.

Express a regular neighbourhood $\mathcal{N}(T)$ of T as $\mathbb{S}^1 \times \mathrm{I} \times \mathbb{S}^1$, where $\mathbb{S}^1 \times \{\frac{1}{2}\} \times \{1\} = \mu$ and $\{1\} \times \{\frac{1}{2}\} \times \mathbb{S}^1 = \lambda$, and let S be the annulus $\mathbb{S}^1 \times \mathrm{I} \times \{1\}$. Let $\psi \colon S \to S$ be a Dehn twist. Define $\Psi \colon \mathbb{S}^3 \setminus \mathcal{N}(L_\kappa) \to \mathbb{S}^3 \setminus \mathcal{N}(L_\kappa)$ by

$$\Psi(x) = \begin{cases} (\psi(y), z) & \text{if } x = (y, z) \in S \times \mathbb{S}^1 = \mathcal{N}(T) \\ x & \text{else.} \end{cases}$$

For $n \in \mathbb{Z}$ let $R_n = \Psi^n(R_0)$. Then, for each n, R_n is a taut Seifert surface for L_κ that can be made disjoint from R. It remains to show that $R_n \neq R$ and $R_n \neq R_m$ for $m \neq n$ when viewed as vertices of $\mathrm{MS}(L_\kappa)$.

Fix $n \in \mathbb{Z}$. To show that $R_n \neq R$ we will show that R_n cannot be made disjoint from T. In this case we may assume $n = 0$. First note that M is ∂–irreducible, R_0 and T are incompressible, and T is obviously ∂–incompressible. R_0 is also ∂–incompressible as it is orientable, incompressible and not ∂–parallel and ∂M is a torus. $M \setminus \mathcal{N}(R_0 \cup T)$ has three components. One of these is $M_0 \setminus \mathcal{N}(R_0)$. This is not a product manifold between R_0 and T since R_0 meets L_κ in M_0 whereas T does not. The other two components lie in M_1. One is homeomorphic as a sutured manifold to that shown in Figure 5.3, and the other is homeomorphic to its complement. Neither of these is a product manifold. By Proposition 1.4.3, R_n cannot be isotoped to be disjoint from T.

Now fix $m \in \mathbb{Z}$. Again we may assume $n = 0$. Let R_0' be a copy of R_0, isotoped to be disjoint from R_0 (except along its boundary). Then $R_m' = \Psi^m(R_0')$ is isotopic to R_m. Figure 5.4 shows a cross-section of $\mathcal{N}(T)$ in the case $m = 2$, where L_κ lies on the inside of T as shown. The components of $M \setminus (R_0 \cup R_m')$ are of five types, as marked. Outside $\mathcal{N}(T)$, those marked $M_{0,b}$ and $M_{1,b}$ are each part of the parallel region between R_0 and R_0'. It is therefore clear that neither of $M_{0,b}, M_{1,b}$ is a product region as they each have disconnected

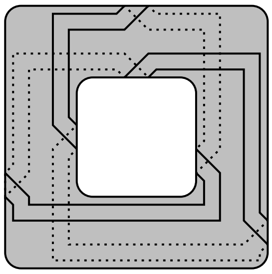

Figure 5.3

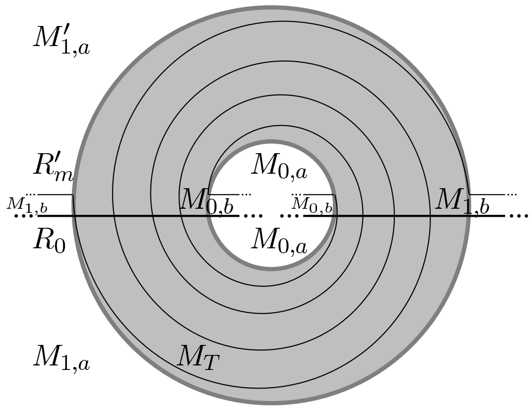

Figure 5.4

intersection with R_0. For the same reason, the components of the same type as M_T are not product regions, and neither is $M_{0,a}$. The manifolds $M_{1,a}$ and $M'_{1,a}$ are sutured manifolds and are the same as the components of $M \setminus (R_0 \cup T)$ in M_1. Hence, again by applying Proposition 1.4.3, we see that $R_0 \neq R_m$.

Thus $\mathrm{MS}(L_\kappa)$ is locally infinite at R. □

Remark 5.1.2. In [18], Kakimizu constructs incompressible Seifert surfaces for a Whitehead double of a knot K using two copies of a Seifert surface for K. Although expressed differently, the above construction is very similar to that used by Kakimizu, with the two Seifert surfaces replaced by the annulus S_1.

5.2 A restriction on links with locally infinite MS

In this section we prove Theorem 1.1.13. Our proof relies heavily on the work of Wilson in [37], which uses normal surfaces.

Definition 5.2.1. Let S be a triangulated surface. A simple closed curve ρ in S is a *normal curve* if it is transverse to the 1–skeleton of the triangulation and, for every triangle T in the triangulation, $\rho \cap T$ consists of arcs whose endpoints lie on different edges of T.

Definition 5.2.2. Let M be a triangulated 3–manifold. A surface S properly embedded in M is a *normal surface* (or is *in normal form*) if it is transverse to the 2–skeleton of the triangulation and, for each tetrahedron V in the triangulation, each connected component of $S \cap V$ is a disc whose boundary meets ∂V in one of the two patterns shown in Figure 5.5.

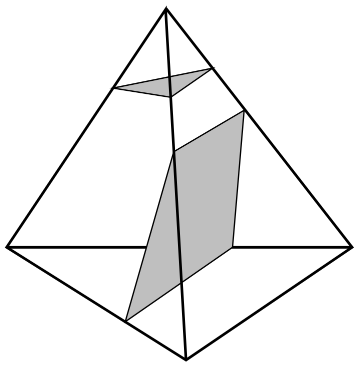

Figure 5.5

Each such disc may be one of 7 types (4 types of triangle, 3 types of square), where these are defined by which edges of the 1–skeleton of V they meet. In this way, the surface S defines a vector in $\mathbb{Z}^{7n}_{\geq 0}$, where n is the number of tetrahedra in the triangulation of M. Conversely, given a vector v in $\mathbb{Z}^{7n}_{\geq 0}$, there is at most one way to join together the set of discs defined by v to form a normal surface properly embedded in M, up to isotopy.

Definition 5.2.3. Let S and S' be normal surfaces in M with corresponding vectors v and v' respectively in $\mathbb{Z}^{7n}_{\geq 0}$. If there exists a properly embedded normal surface corresponding to the vector $v + v'$ then this surface is the *Haken sum* of S and S', denoted $S + S'$.

The surface $S + S'$ can be constructed from S and S' by a 'cut-and-paste' operation in M as follows. We may assume that S and S' are in general position. Let ρ be a curve of $S \cap S'$. Change the surface $S \cup S'$ in $\mathcal{N}(\rho)$ as follows. Delete a neighbourhood of ρ in each of S and S'. Replace these two annuli with two new annuli joining the newly created boundary components so that the resulting surface is embedded in $\mathcal{N}(\rho)$. There are two possible positions for these two new annuli, up to isotopy, corresponding to the choices of pairings of the new boundary components. The choice between these two positions is made using the positions of S and S' in the triangulation of M. Repeating this process for each curve of $S \cap S'$ gives the embedded surface $S + S'$.

Definition 5.2.4. Given normal surfaces S, S_1, S_2 with $S = S_1 + S_2$, say that S_1 and S_2 are in *reduced form* if they have been isotoped to minimise $|S_1 \cap S_2|$ while maintaining the equation $S = S_1 + S_2$.

Definition 5.2.5. Let S be a normal surface. Its *weight* is the number of times it meets the 1–skeleton of the triangulation. Call S *minimal* if it has minimal weight among normal surfaces isotopic to S by an isotopy fixing ∂S.

Lemma 5.2.6 ([16] Lemma 2.1)**.** *Let S be a minimal normal surface in a closed, irreducible 3–manifold M. If S is incompressible and $S = S_1 + S_2$ is in reduced form then each component of $S_1 \setminus (S_1 \cap S_2)$ is incompressible and is not a disc.*

Definition 5.2.7. A compact surface S embedded in $\mathbb{S}^3$ with no closed components is a *spanning surface* for an unoriented link L if $\partial S = L$.

In [37], Wilson states the following.

Theorem 5.2.8 ([37] Main Theorem 1.1)**.** *Let K be a non-trivial knot, and let $M = \mathbb{S}^3 \setminus \mathcal{N}(K)$. Then there is a finite set $\{R_1, \ldots, R_m\}$ of incompressible Seifert surfaces for K and a finite set $\{S_1, \ldots, S_n\}$ of closed surfaces in M that are not boundary parallel such that any incompressible Seifert surface R is isotopic to a Haken sum $R = R_i + a_1 S_1 + \cdots + a_n S_n$, where $a_1, \ldots, a_n$ are non-negative integers.*

The surfaces $R_1, \ldots, R_m$ that arise from Wilson's proof are spanning surfaces for K. However, he does not consider the orientability of these surfaces, which is necessary to conclude, as he does, that they are in fact Seifert surfaces. In fact it is possible to require these surfaces to be orientable, as shown by the following two lemmas.

Lemma 5.2.9. *Let $A \subseteq \mathbb{Z}^n_{\geq 0}$. Then there is a finite set $B \subseteq A$ such that any vector in A can be obtained from some vector in B by increasing some coordinates.*

Proof. We proceed by induction on the dimension n.

First suppose each vector has only one coordinate (that is, suppose $n = 1$). In this case, take B to be the single smallest element.

Now suppose we can find a suitable subset B of A whenever we are working in a space of dimension at most k, and let $A \subseteq \mathbb{Z}^{k+1}_{\geq 0}$. Choose an element $a = (a_1, \ldots, a_{k+1})$ of A. For $1 \leq i \leq k+1$ and $0 \leq j < a_i$ let $A_{i,j} = \{(v_1, \ldots, v_{k+1}) \in A : v_i = j\}$. Each set $A_{i,j}$ has dimension at most k and so by the inductive hypothesis has a finite subset $B_{i,j}$ with the given property. Let $B = \{a\} \cup \bigcup B_{i,j}$. Then B is finite.

Let $v = (v_1, \ldots, v_{k+1}) \in A$. If $v_i \geq a_i$ whenever $1 \leq i \leq k+1$ then v can be obtained from a by increasing some coordinates. Else $v_l < a_l$ for some $l \leq k+1$. Then $v \in A_{l,v_l}$, so there exists $b \in B_{l,v_l}$ such that v can be obtained from b by increasing some coordinates other than coordinate l. Therefore B has the required property, and the induction is complete. □

Lemma 5.2.10. *Let the surfaces R_i and S_j be as in Theorem 5.2.8, except that the R_i are not necessarily orientable. Then there exists a finite set $\{R'_1, \ldots, R'_{m'}\}$ of incompressible Seifert surfaces for K such that any incompressible Seifert surface for K can be expressed as $R'_i + a_1S_1 + \cdots + a_nS_n$, where $a_1, \ldots, a_n$ are non-negative integers.*

Proof. For each incompressible Seifert surface R for K, choose an expression $R = R_i + a_1S_1 + \cdots + a_nS_n$. Let v_R be the vector $(i, a_1, \ldots, a_n)$. Let $A \subseteq \mathbb{Z}^{n+1}_{\geq 0}$ be the set of all such vectors. For $1 \leq i \leq n$ let A_i be the set of elements of A with i as the first coordinate. By Lemma 5.2.9 there is a finite subset B_i of A_i such that any element in A_i can be obtained from some element of B_i by increasing some coordinates other than the first. Let $B = \bigcup B_i$. Then B is finite. Take $R'_1, \ldots, R'_{m'}$ to be be the surfaces defined by the elements of B. □

It is also worth noting the nature of the isotopy referred to in Theorem 5.2.8. In his proof, Wilson isotopes the chosen Seifert surface R into normal form based on the following lemma.

Lemma 5.2.11 ([37] Lemma 3.3)**.** *Let K be a knot, let $M = \mathbb{S}^3 \setminus \mathcal{N}(K)$ and let R be an incompressible Seifert surface for K in M. Suppose that M is triangulated, and ∂R meets each triangle of the triangulation in at most one normal arc. Then R can be put into normal form by an isotopy fixing ∂R.*

The proof of this lemma gives the stronger conclusion that the isotopy puts the surface into minimal normal form. This is important because minimality is a key hypothesis of [16] Theorem 2.2, which is used in the proof of Theorem 5.2.8.

Aside from these points, Wilson's proof is actually stronger than the statement of Theorem 5.2.8 suggests. In particular, by following the proof with M as the complement of a taut Seifert surface for a link, it gives the following.

Theorem 5.2.12. *Let L be an oriented link such that every taut Seifert surface for L is connected. Let R be a taut Seifert surface for L, let $M = \mathbb{S}^3 \setminus \mathcal{N}(R)$, and fix a set $\rho_1, \ldots, \rho_k$ of core curves of the annuli $\partial M \cap \partial\mathcal{N}(L)$, one for each link component. There is a triangulation of M such that every Seifert surface R' for L disjoint from R can be put into normal form with $\partial R' = \bigcup_{i=1}^{k} \rho_i$.*

Furthermore, there is a finite set $\{R_1, \ldots, R_m\}$ of surfaces in M with non-empty boundary contained in $\bigcup_{i=1}^{k} \rho_i$, and a finite set $\{S_1, \ldots, S_n\}$ of closed surfaces in M, such that all these surfaces are incompressible and in normal form, and the following holds. Any taut Seifert surface R' for L in M with $\partial R' = \bigcup_{i=1}^{k} \rho_i$ and in minimal normal form can be expressed as $a_1R_1 + \cdots + a_mR_m + b_1S_1 + \cdots + b_nS_n$ for some $a_i, b_i \in \mathbb{Z}_{\geq 0}$.

If L has more than one component, it is possible that, for a given $j \leq m$, ∂R_j is a strict subset of $\bigcup_{i=1}^{k} \rho_i$. However, only finitely many combinations of $R_1, \ldots, R_m$ will yield the correct boundary. Hence we may assume that $\partial R_j = \bigcup_{i=1}^{k} \rho_i$. Then $\sum_{i=1}^{m} a_i = 1$. In addition, by Lemma 5.2.10 we may assume that R_j is orientable.

We will also need the following proposition in the proof of Theorem 1.1.13.

Proposition 5.2.13 ([6] 15.26)**.** *Let K be a knot, and let $M = \mathbb{S}^3 \setminus \mathcal{N}(K)$. Suppose there is an annulus S properly embedded in M that is not ∂–parallel. If neither component of ∂S bounds a disc in ∂M then K is a torus knot, a cable knot, or a connected sum.*

Theorem 1.1.13. *Let L be an oriented link such that every taut Seifert surface for L is connected. If* $\mathrm{MS}(L)$ *is locally infinite then L is a satellite of either a torus knot, a cable knot or a connected sum, with winding number 0.*

Proof. Let R be a taut Seifert surface for L such that $\mathrm{MS}(L)$ is locally infinite at R. That is, there are infinitely many taut Seifert surfaces for L that can be made disjoint from R. Let $M = \mathbb{S}^3 \setminus \mathcal{N}(R)$, and fix a set $\rho_1, \ldots, \rho_k$ of core curves of the annuli $\partial M \cap \partial \mathcal{N}(L)$, one for each link component. Then Theorem 5.2.12 applies. In addition, it is clear that none of the R_i is a disc and that, since R is connected, M is irreducible.

By discarding surfaces if necessary, we may ensure that, for any $j \leq n$, the sets $\{R_1, \ldots, R_m\}$ and $\{S_1, \ldots, S_n\} \setminus \{S_j\}$ do not satisfy the conclusions of Theorem 5.2.12. We may also assume that S_1 has minimal genus among the S_i. Let R' be a taut Seifert surface in minimal normal form such that $R' = R_1 + b_1 S_1 + \cdots + b_n S_n$ with $b_1 > 0$, and set $T = S_1$. Let $R^- = R_1 + (b_1 - 1)S_1 + b_2 S_2 + \cdots + b_n S_n$, so that $R' = R^- + T$, and isotope R^- and T into reduced form. Since the isotopy keeps $\partial R'$ fixed and T is closed, this will leave ∂R^- unchanged. Then, by Lemma 5.2.6, no curve of $R^- \cap T$ bounds a disc in either R^- or T. Note that although Lemma 5.2.6 is proved in [16] only for closed surfaces, the same proof works in this case because T is closed.

Suppose that T is a 2–sphere. Then, after the isotopy, it must be disjoint from R^-. This contradicts that R' is connected. Since there are infinitely many taut Seifert surfaces in minimal normal form in M, it follows that T is a torus.

Let M_0 be the component of $M \setminus \mathcal{N}(T)$ containing ∂M, and M_1 the other component. The orientation that R' inherits from L induces an orientation on each component of $R' \cap M_0$ and hence on each curve of $R^- \cap T$. Let ρ be a curve on T that meets each curve of $R^- \cap T$ once. Because T is disjoint from R, the algebraic intersection $\rho \cdot R$ of ρ and R is 0. As $[R'] = [R]$ in $H_2(\mathbb{S}^3 \setminus \mathcal{N}(L), \partial \mathcal{N}(L))$, this gives that $\rho \cdot R' = 0$, and so $\rho \cdot (R^- \cap T) = 0$ on T. Therefore half the curves of $R^- \cap T$ are oriented in one direction, and half are oriented in the other direction. In particular, $|R^- \cap T|$ is even. Find adjacent curves with opposite orientations, and surger R^- along the subannulus of T between them. Repeating this to remove all curves of $R^- \cap T$ gives a new Seifert surface R'' for L, together with a closed, possibly disconnected, surface S''. Note that $R'' \subset M_0$ and S'' is orientable. As R' is taut, $\chi(R') \geq \chi(R'') = \chi(R^-) - \chi(S'') = \chi(R') - \chi(T) - \chi(S'')$, so $\chi(S'') \geq 0$. The components of $(R^- \cup T) \setminus (R^- \cap T)$ from which S'' is constructed each have boundary, and none of them is a disc. Therefore each of these components is an annulus, and in particular this includes every component of $R^- \cap M_1$.

Let S be one such annulus in M_1, and suppose it is parallel to a subannulus S_T of T. If there are other curves of $R^- \cap T$ in S_T, they must also bound annuli parallel to T. Hence we may assume $R^- \cap \mathrm{int}(S_T) = \emptyset$. At each of the two boundary curves of S_T, the cut-and-paste operation that creates R' from R^- and T might go one of two ways (see Figure 5.6). If

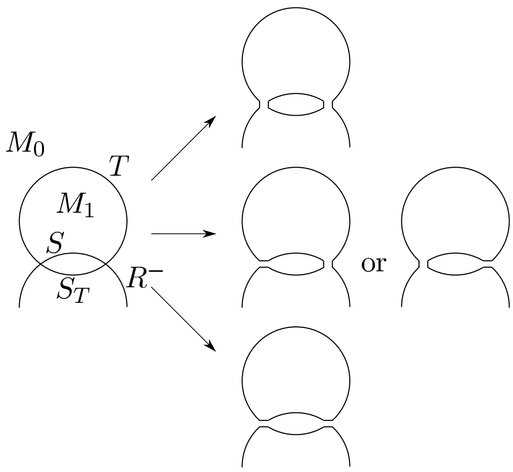

Figure 5.6

both join together S and S_T then this creates a torus component of R', contradicting that R' is connected. If both go the other way, we see that an isotopy of R^- and T could reduce $R^- \cap T$ without changing R', contradicting the choice of R^- and T. If only one joins the two annuli, an isotopy along the product region reduces the weight of R', again giving a contradiction.

Thus S is not ∂–parallel in M_1. Note that the part of $\mathbb{S}^3 \setminus \mathcal{N}(T)$ containing L is a solid torus V. Let K be the core curve of V. Since $R \subset V$ and T is incompressible, L is a satellite of K with winding number 0. Because S is not parallel to T, the knot K satisfies the hypotheses of Proposition 5.2.13. $\square$

Remark 5.2.14. In [14], Gustafson gives an example of a hyperbolic knot K with incompressible Seifert surfaces of arbitrarily high genus. By inspection, each of these is disjoint from a single taut Seifert surface for K. Therefore $\mathrm{IS}(K)$ is locally infinite, contrasting with Theorem 1.1.13.

5.3 Other examples

Theorem 1.1.13 raises the question of whether each of the three possibilities given can occur. Theorem 5.1.1 shows that K can be a torus knot, and clearly the same construction will work for a cable knot. We now give an example where K is a connected sum of two trefoil knots.

Lemma 5.3.1. *Let L_ν be the knot shown in Figure 5.7. Then* $\mathrm{MS}(L_\nu)$ *is locally infinite.*

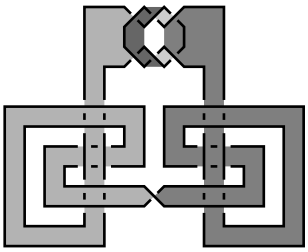

Figure 5.7

Proof. The proof is essentially that of Proposition 5.1.1, where R_ν is the genus 1 surface shown in Figure 5.7 and S_1 is the annulus dividing the two trefoil components of the connected sum, as shown in Figure 5.8. In this case, the two components of $M \setminus \mathcal{N}(R_0 \cup T)$

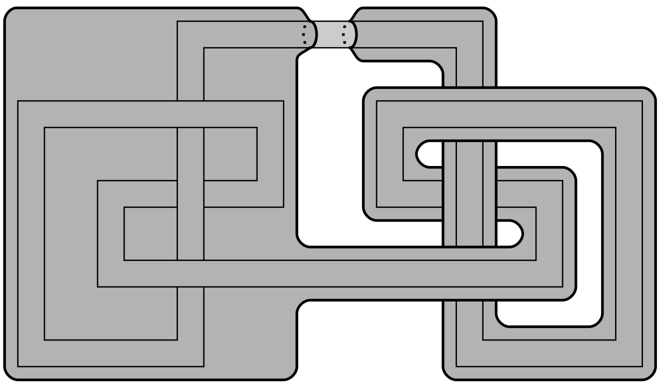

Figure 5.8

in M_1 are each a sutured manifold homeomorphic to the complement of a trefoil in $\mathbb{S}^3$, with two sutures that are meridians of the trefoil. This is not a product sutured manifold. □

We now give a final example to show that in Theorem 1.1.13 it is necessary to include the assumption that every taut Seifert surface for L is connected.

Proposition 5.3.2. *Let* L_ξ *be the three component link shown in Figure 5.9. Then* $\mathrm{MS}(L_\xi)$ *is not locally finite.*

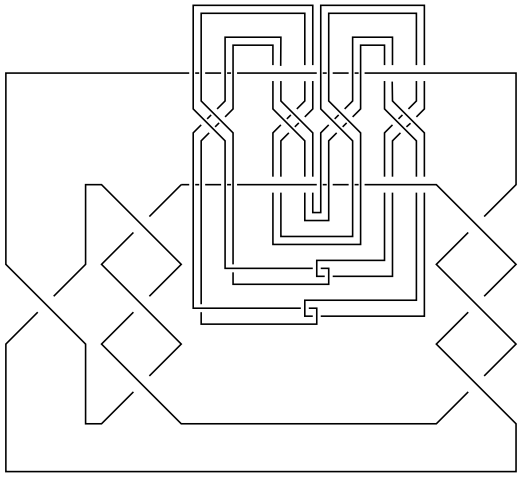

Figure 5.9

Proof. The construction in this case is again very similar to that in Proposition 5.1.1, and we will use the same notation. Let $M = \mathbb{S}^3 \setminus \mathcal{N}(L_\xi)$. Each of the three components $L_{\xi,i}$ of L_ξ has a genus 1 Seifert surface $R_{\xi,i}$ given by applying Seifert's algorithm to the diagram shown. These together form a Seifert surface for L_ξ in the obvious way. Call this surface R_ξ. Let T be the torus shown in Figure 5.10 separating $L_{\xi,0}$ from $L_{\xi,1} \cup L_{\xi,2}$. Set the

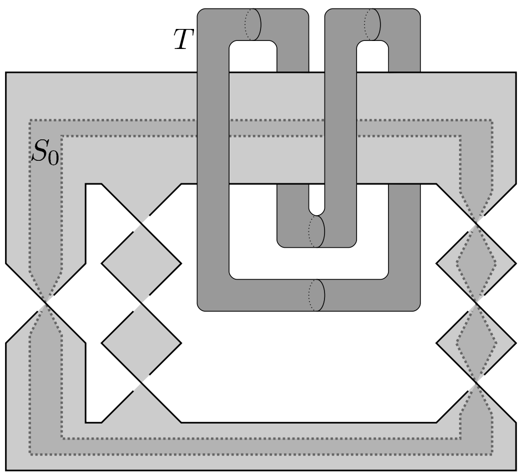

Figure 5.10

component of $M \setminus \mathcal{N}(T)$ containing $L_{\xi,0}$ as M_0, and the component containing $L_{\xi,1}$ and $L_{\xi,2}$ as M_1. Take S_1 to be the annulus properly embedded in M_1 that runs between $L_{\xi,1}$ and $L_{\xi,2}$. Then the subannulus S_T of T is parallel in M_0 to the annulus S_0 on $R_{\xi,0}$ shown in Figure 5.10.

The same proof as above gives an infinite family of distinct Seifert surfaces R_n for L_ξ disjoint from R_ξ. The only difference this far is in considering the components of $M \setminus \mathcal{N}(R_0 \cup T)$ in M_1. In this case these are not product sutured manifolds as they have disconnected intersection with R_0.

It remains to show that R_ξ is taut. For $n \in \mathbb{Z}$, this will show that R_n is also taut since $\chi(R_n) = \chi(R_\xi)$. It is sufficient to show that L_ξ has a taut Seifert surface with three connected components. We can see that this is the case using some of the methods from the proof of Theorem 1.1.13. Let R' be a taut Seifert surface for L_ξ. Isotope R' to have minimal intersection with the torus T. Then $R' \cap T$ consists of parallel essential curves on T, and no component of $R' \setminus T$ is a disc. As R_ξ is disjoint from T and $[R'] = [R_\xi]$, half of the curves in $R' \cap T$ are oriented in one direction and the other half are oriented in the other direction. We may therefore surger R' along subannuli of T to give a new Seifert surface R'' for L_ξ that is disjoint from T, together with a closed, possibly disconnected, possibly empty, surface S''. As in the proof of Theorem 1.1.13, $\chi(S'') \geq 0$. We therefore see that $\chi(S'') = 0$. Hence R'' is also taut. This process can be repeated using the torus $S_T \cup S_1$, to give a taut Seifert surface R''' for L_ξ with three connected components. $\square$

References

[1] J. W. Alexander. Topological invariants of knots and links. *Trans. Amer. Math. Soc.*, 30(2):275–306, 1928.

[2] Jessica E. Banks. Homogeneous links, Seifert surfaces, digraphs and the reduced Alexander polynomial, 2011, arXiv:1101.1412 [math.GT].

[3] Jessica E. Banks. The Kakimizu complex of a connected sum of links, 2011, arXiv:1109.0965 [math.GT].

[4] Jessica E. Banks. Minimal genus Seifert surfaces for alternating links, 2011, arXiv:1106.3180 [math.GT].

[5] Jessica E. Banks. On links with locally infinite Kakimizu complexes. *Algebr. Geom. Topol.*, 11(3):1445–1454, 2011.

[6] Gerhard Burde and Heiner Zieschang. *Knots*, volume 5 of *de Gruyter Studies in Mathematics.* Walter de Gruyter & Co., Berlin, 1985.

[7] P. R. Cromwell. Homogeneous links. *J. London Math. Soc. (2)*, 39(3):535–552, 1989.

[8] Richard Crowell. Genus of alternating link types. *Ann. of Math. (2)*, 69:258–275, 1959.

[9] Herbert Edelsbrunner and Daniel R. Grayson. Edgewise subdivision of a simplex. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry (Miami Beach, FL, 1999)*, pages 24–30 (electronic), New York, 1999. ACM.

[10] Samuel Eilenberg and Norman Steenrod. *Foundations of algebraic topology.* Princeton University Press, Princeton, New Jersey, 1952.

[11] Julian R. Eisner. Knots with infinitely many minimal spanning surfaces. *Trans. Amer. Math. Soc.*, 229:329–349, 1977. Also has an addendum that contains the figures.

[12] David Gabai. The Murasugi sum is a natural geometric operation. In *Low-dimensional topology (San Francisco, Calif., 1981)*, volume 20 of *Contemp. Math.*, pages 131–143. Amer. Math. Soc., Providence, RI, 1983.

[13] David Gabai. Genera of the alternating links. *Duke Math. J.*, 53(3):677–681, 1986.

[14] Richard F. Gustafson. A simple genus one knot with incompressible spanning surfaces of arbitrarily high genus. *Pacific J. Math.*, 96(1):81–98, 1981.

[15] Mikami Hirasawa and Makoto Sakuma. Minimal genus Seifert surfaces for alternating links. In *KNOTS '96 (Tokyo)*, pages 383–394. World Sci. Publ., River Edge, NJ, 1997.

[16] William Jaco and Ulrich Oertel. An algorithm to decide if a 3-manifold is a Haken manifold. *Topology*, 23(2):195–209, 1984.

[17] Andras Juhasz. Knot Floer homology and Seifert surfaces. *Algebr. Geom. Topol.*, 8(1):603–608, 2008.

[18] Osamu Kakimizu. Doubled knots with infinitely many incompressible spanning surfaces. *Bull. London Math. Soc.*, 23(3):300–302, 1991.

[19] Osamu Kakimizu. Finding disjoint incompressible spanning surfaces for a link. *Hiroshima Math. J.*, 22(2):225–236, 1992.

[20] Osamu Kakimizu. Classification of the incompressible spanning surfaces for prime knots of 10 or less crossings. *Hiroshima Math. J.*, 35(1):47–92, 2005.

[21] Tsuyoshi Kobayashi. Uniqueness of minimal genus Seifert surfaces for links. *Topology Appl.*, 33(3):265–279, 1989.

[22] Marc Lackenby. The volume of hyperbolic alternating link complements. *Proc. London Math. Soc. (3)*, 88(1):204–224, 2004. With an appendix by Ian Agol and Dylan Thurston.

[23] W. B. Raymond Lickorish. *An introduction to knot theory*, volume 175 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.

[24] W. Menasco. Closed incompressible surfaces in alternating knot and link complements. *Topology*, 23(1):37–44, 1984.

[25] William Menasco and Morwen Thistlethwaite. The classification of alternating links. *Ann. of Math. (2)*, 138(1):113–171, 1993.

[26] William W. Menasco and Morwen B. Thistlethwaite. The Tait flyping conjecture. *Bull. Amer. Math. Soc. (N.S.)*, 25(2):403–412, 1991.

[27] K. Murasugi and A. Stoimenow. The Alexander polynomial of planar even valence graphs. *Adv. in Appl. Math.*, 31(2):440–462, 2003.

[28] Kunio Murasugi. On the genus of the alternating knot. I, II. *J. Math. Soc. Japan*, 10:94–105, 235–248, 1958.

[29] Kunio Murasugi. On alternating knots. *Osaka Math. J.*, 12:277–303, 1960.

[30] Roberto Carlos Pelayo. *Diameter Bounds on the Complex of Minimal Genus Seifert Surfaces for Hyperbolic Knots*. PhD thesis, California Institute of Technology, 2007.

[31] Piotr Przytycki and Jennifer Schultens. Contractibility of the Kakimizu complex and symmetric Seifert surfaces. *Trans. Am. Math. Soc.*, 364(3):1489–1508, 2012.

[32] Robert Riley. A finiteness theorem for alternating links. *J. London Math. Soc. (2)*, 5:263–266, 1972.

[33] Dale Rolfsen. *Knots and links*, volume 7 of *Mathematics Lecture Series*. Publish or Perish Inc., Houston, TX, 1990. Corrected reprint of the 1976 original.

[34] Makoto Sakuma. Minimal genus Seifert surfaces for special arborescent links. *Osaka J. Math.*, 31(4):861–905, 1994.

[35] Martin Scharlemann and Abigail Thompson. Finding disjoint Seifert surfaces. *Bull. London Math. Soc.*, 20(1):61–64, 1988.

[36] Friedhelm Waldhausen. On irreducible 3-manifolds which are sufficiently large. *Ann. of Math. (2)*, 87:56–88, 1968.

[37] Robin T. Wilson. Knots with infinitely many incompressible Seifert surfaces. *J. Knot Theory Ramifications*, 17(5):537–551, 2008.