

Minimal genus Seifert surfaces for alternating links

Jessica E. Banks

ABSTRACT. We give a complete proof of results announced by Hirasawa and Sakuma describing explicitly the Kakimizu complex of a non-split, prime, special alternating link.

CONTENTS

1. Introduction	445
2. The paper of Hirasawa–Sakuma	448
3. Special form	452
4. Putting surfaces into admissible special form	460
5. Distinguishing surfaces in special form	468
6. Non-special links	480
7. The realisation of $MS(L)$	485
Appendix A. Proof by picture: an example of Proposition 5.23	489
References	493

1. Introduction

The Kakimizu complex of a link records the structure of the set of minimal genus Seifert surfaces for the link. For a non-split, alternating link, the definition is analogous to that of the complex of curves of a compact, orientable surface.

Definition 1.1 ([6] p225). For a non-split, alternating link L , the *Kakimizu complex* $MS(L)$ of L is a simplicial complex, the vertices of which are the ambient isotopy classes of minimal genus Seifert surfaces for L . Distinct vertices R_0, \dots, R_n span an n -simplex exactly when they can be realised disjointly.

Definition 1.2. A metric is defined on the vertices of $MS(L)$. The distance between two vertices is the distance in the 1-skeleton of $MS(L)$ when every edge has length 1.

Received July 25, 2021.

2010 *Mathematics Subject Classification.* 57K10.

Key words and phrases. Seifert surfaces; special alternating links; Kakimizu complex.

This work formed part of my doctoral studies at the University of Oxford.

The Kakimizu complex can be defined for links in general, but not using this form of the definition. It has been calculated explicitly in certain simple cases (see, for example, [7]). However, doing so in general is complicated by the difficulty of controlling the multiple surfaces that form the vertices of a simplex. One way in which we might hope to do so is by expressing these surfaces in terms of a diagram for L . With this in mind, alternating links are a natural class to study, as applying Seifert's algorithm to any alternating diagram gives a minimal genus Seifert surface ([4] Theorem 4). In particular, the subclass of special alternating links is very well behaved.

Definition 1.3. A *special* link diagram D is one in which every Seifert circle is innermost in \mathbb{S}^2 . This determines a division of the regions of D into *black regions* (those which make up the Seifert surface given by Seifert's algorithm) and *white regions*. Thus the surface given by applying Seifert's algorithm to D is always uniquely defined.

In [5] Hirasawa and Sakuma give a complete description of $\text{MS}(L)$ for a prime link L with a reduced, special alternating diagram D via the following result. The definition of the complex $\mathcal{K}(D)$ is given in Section 2.

Theorem 1.4 ([5] Theorems 1.5, 1.1). *There is a natural isomorphism between $\text{MS}(L)$ and $\mathcal{K}(D)$. In particular, every minimal genus Seifert surface for L is given by applying Seifert's algorithm to some special alternating diagram for L .*

From this they deduce contractibility of $\text{MS}(L)$ for such links L , through the following.

Theorem 1.5 ([5] Theorem 1.6). *If L is a prime link with a special alternating diagram D then $\text{MS}(L)$ is homeomorphic to a disc, the dimension of which can be calculated from D .*

Only 'the idea of the proof' of these theorems is included in [5]. In [13] it says the following.

[Sakuma] proves the contractibility of $\text{MS}(K)$ when K is a special arborescent link. In his joint paper with Hirasawa [[5]], contractibility when K is a prime, special, alternating link is announced. Together, this partially verifies a challenging conjecture of Kakimizu's, asserting among other things that $\text{MS}(K)$ is always contractible.

The question of the contractibility of the Kakimizu complex for a general link has been answered by the following theorem of Przytycki and Schultens.

Theorem 1.6 ([11] Theorem 1.1). *The Kakimizu complex of a link is contractible.*

Here we provide a full proof of the results of [5], with a view to determining $\text{MS}(L)$ for a general alternating link L . Theorem 1.6 is needed in this proof.

1.1. Overview. Since the statement of the main result of this paper and its proof are both technical in nature, here we will take a quick tour of the main principles.

When dealing with diagrams of prime, alternating links, the Tait Flying Conjecture tells us that ‘everything is flypes’:

Theorem 1.7 ([9]). *Any two reduced, alternating diagrams of the same prime link differ by a sequence of flypes.*

We therefore need to start from the questions ‘which crossings could be used for flypes?’ and ‘where could there be crossings after doing flypes?’. Answering these questions is straight-forward, although simpler in practice than when described in the abstract.

From each reduced, special alternating diagram D for a prime link L , we can construct a minimal genus Seifert surface. The first thing that Theorem 1.4 tells us is that this process produces all possible minimal genus Seifert surfaces for L . To show this, we need to start from a Seifert surface and use it to produce instructions for a sequence of flypes to yield the corresponding diagram.

Remark 1.8. There is a subtlety to counting ‘diagrams’ when the link has symmetry. For example, the knot 7_4 has two minimal genus Seifert surfaces, given by two diagrams that differ by a flype. However, the two diagrams look the same as each other.

Having produced this list of surfaces, the next task is to eliminate duplicates (up to isotopy). There are cases where it is easy to see that performing a particular combination of flypes gives the same surface as a result. If the flype circle is in the white regions of the diagram where it passes through the flype crossing (as in Figure 3c) then the flype simply turns over part of the Seifert surface. Therefore only flype circles as in Figure 3d are of interest. In addition, the surface is unchanged up to isotopy if the flype circle encloses only a line of white bigons (as in Figure 1).




FIGURE 1

Perhaps the most intuitive interpretation of Theorem 1.4 is that there are no other, ‘hidden’ pairs of equivalent surfaces. However, the theorem also gives us information about exactly which collections of surfaces can be realised disjointly. This could be thought of as establishing which combinations of flypes

are ‘compatible’. Roughly speaking, the answer to this is that you can move one crossing from each flype-equivalence-class, so long as you move them all in the ‘same direction’. The definition of the simplicial complex $\mathcal{K}(D)$, as given in [5], encodes the precise details.

The proof that the definition is correct works by showing that $\mathcal{K}(D)$ is a covering space of $\text{MS}(L)$. Since Theorem 1.6 tells us that $\text{MS}(L)$ is contractible, the covering map is a homeomorphism. To establish this we show firstly that the map is surjective on the vertices, secondly that the diagrams for two adjacent vertices do not give isotopic Seifert surfaces, and thirdly that the diagrams for vertices at distance 2 from each other do not give isotopic Seifert surfaces. Proving that two surfaces cannot be isotoped to be disjoint involves first positioning one surface in a ‘minimal position’ relative to the other, and then verifying that none of the complementary regions in the link complement are ‘product regions’ (which would enable further simplification).

This paper is organised as follows. In Section 2 we recall the definition of $\mathcal{K}(D)$ from [5]. In Section 3 we define a method for describing minimal genus Seifert surfaces relative to an alternating diagram. This enables us to compare surfaces given by different diagrams. Section 4 addresses Theorem 1.1 of [5], and Section 5 extends this to a proof of [5] Theorem 1.5. In Section 6 we prove Theorem 1.2 of [5], which gives two families of examples showing that we cannot remove the adjective ‘special’ from the hypotheses of Theorem 1.4. Finally, in Section 7 we complete the verification that the Kakimizu complex is homeomorphic to a disc for prime, special alternating links.

I wish to thank Marc Lackenby for helpful conversations and encouragement. I am also grateful to Richard Webb and Jacob Rasmussen, through whom I learned of [5].

2. The paper of Hirasawa–Sakuma

There are a number of definitions we will need from [5]. We will change some of the notation and terminology. Throughout this paper we will only consider non-split, oriented links.

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 $e, e' \in E(\mathcal{G})$ 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 ρ' can be made disjoint on their interiors, the result of this process is well-defined given D .

Definition 2.1. Define $G^{\mathcal{F}}(D)$ to be the graph that results from this process.

Figure 2 shows a digram D_α of a link L_α together with the graph $G^{\mathcal{F}}(D_\alpha)$.




FIGURE 2

Definition 2.2. A flype circle ϕ is a simple closed curve in \mathbb{S}^2 that meets the link diagram D as shown in Figure 3a, 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 3b. This change is realised by an isotopy of \mathbb{S}^3 .

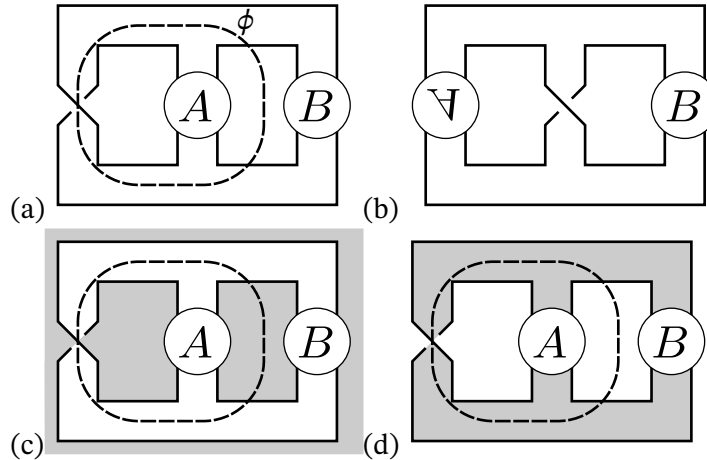


FIGURE 3

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 2.3. 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 (as in Figure 3c) 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 (as in Figure 3d).

Definition 2.4. We call the crossing through which ϕ passes the *flype crossing*, and the arc of ϕ disjoint from R the *flype arc*.

White bigons in a special alternating diagram signify plumbed on Hopf bands in the Seifert surface given by Seifert’s algorithm (see Lemma 5.16). 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 2.5. Call such flype circles *equivalent*.

Definition 2.6. Let $v, v' \in V(G^{\mathcal{F}}(D))$ such that v, v' are joined by at least two edges of $G^{\mathcal{F}}(D)$. Then the subgraph of $G^{\mathcal{F}}(D)$ consisting of v, v' and all edges joining them 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 2.7. 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 4. We similarly define the *positive* and *negative sides* of each crossing in D .




FIGURE 4

Definition 2.8. 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 $e \in E(\theta(D))$ 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, \dots, e_n . These weights are used to define a simplicial complex $\mathcal{K}(D)$. First, if $n \geq 1$, set

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

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

Definition 2.9. Let $\tilde{\theta}(D)$ be the planar graph obtained from $\theta(D)$ by cutting apart vertices to make the θ -graphs disjoint.

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 a *region of $\theta(D)$* .

Definition 2.10. 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 $V(\mathcal{K}(D))$ to $V(\mathcal{K}(D))$, given by $r^\theta(w_1, \dots, w_n) = (w'_1, \dots, w'_n)$, 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 $r^\theta(w_1, \dots, w_n)$ is defined when $w_i > 0$ for all $e_i \in \partial^- r$.

We describe the process of applying r^θ to $v \in V(\mathcal{K}(D))$ 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, \dots, v_k spans a k -simplex if there is a choice of labelling of the regions of $\theta(D)$ as r_1, \dots, r_m such that $r_j^\theta(\dots r_1^\theta(v_0) \dots)$ is defined for $1 \leq j \leq m$ and each v_i occurs as $r_j^\theta(\dots r_1^\theta(v_0) \dots)$ 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 $r_m^\theta(\dots r_1^\theta(v_0) \dots) = v_0$.

A metric is defined on $V(\mathcal{K}(D))$ in the same way as that defined on $MS(L)$. 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 $V(\mathcal{K}(D))$ to $V(\text{MS}(L))$.




FIGURE 5

Figure 5a shows the graph $\theta(D_\alpha)$ for the diagram D_α in Figure 2, marked with the weights that come from D_α . Figure 5b shows $\tilde{\theta}(D_\alpha)$. The complex $\mathcal{K}(D_\alpha)$ is 3-dimensional, and this vertex lies in four 3-simplices, corresponding to the cycles

- $(1, 0, 2, 0, 1) \rightarrow (1, 0, 3, 0, 0) \rightarrow (0, 1, 3, 0, 0) \rightarrow (1, 0, 2, 1, 0) \rightarrow (1, 0, 2, 0, 1),$
- $(1, 0, 2, 0, 1) \rightarrow (0, 1, 2, 0, 1) \rightarrow (0, 1, 3, 0, 0) \rightarrow (1, 0, 2, 1, 0) \rightarrow (1, 0, 2, 0, 1),$
- $(1, 0, 2, 0, 1) \rightarrow (0, 1, 2, 0, 1) \rightarrow (1, 0, 1, 1, 1) \rightarrow (1, 0, 2, 1, 0) \rightarrow (1, 0, 2, 0, 1),$
- $(1, 0, 2, 0, 1) \rightarrow (0, 1, 2, 0, 1) \rightarrow (1, 0, 1, 1, 1) \rightarrow (1, 0, 1, 0, 2) \rightarrow (1, 0, 2, 0, 1).$

3. Special form

3.1. Definition. In this section we will focus on positioning a Seifert surface nicely with respect to the diagram plane, so that we can encode it on the diagram. Let L be a special alternating link with a special alternating diagram $D \subset \mathbb{S}^2 \subset \mathbb{S}^3$.

Definition 3.1. 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 [4]). Let R be the Seifert surface given by applying Seifert’s algorithm to D . Let $\mathcal{N}(L)$ be a regular neighbourhood of L in \mathbb{S}^3 and $\mathcal{N}_L(R)$ a product neighbourhood of R in $\mathbb{S}^3 \setminus \mathcal{N}(L)$. Then $\mathcal{N}(R) = \mathcal{N}(L) \cup \mathcal{N}_L(R)$ is a regular neighbourhood of R in \mathbb{S}^3 .

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 [4]). In a neighbourhood of a crossing, ∂R and $\partial R'$ are as shown in Figure 6.

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,




FIGURE 6

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. 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 3.3. 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' has minimal genus 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. 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 7. By pushing each P -arc




FIGURE 7

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 8 shows two digrams D_β and D_β^* of a link L_β that differ by a flype along the flype circle ϕ shown. Let R_β, R_β^* be the surfaces given by applying Seifert's




FIGURE 8

algorithm to D_β, D_β^* respectively. Consider the effect on R_β^* of an isotopy that takes D_β^* to D_β . By inspection, R_β^* can be put into special form with respect to D_β as shown in Figure 9a. Figure 9b shows $\partial(R_\beta^* \cap \mathbb{B}_a)$.

The special form of R_β^* 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.




FIGURE 9

Definition 3.4. Call a set of distinct flype circles with this property *coherent*. Otherwise call them *incoherent*.

Remark 3.5. Coherent flype circles must be pairwise disjoint. Figure 10 shows a set of incoherent disjoint flype circles on the diagram D_α .




FIGURE 10

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 3.6. *The arcs defined by a coherent set of flype circles form a set of P -arcs.*

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 we see that 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 11). Suppose neither is a flype crossing. They 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 . □




FIGURE 11

Definition 3.7. 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 3.8. 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 12).




FIGURE 12

Lemma 3.9. 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 13a, where A contains no flype circles. The flype ϕ changes the diagram to that in Figure 13b. This also gives R' in admissible form. By induction, the result follows. \square




FIGURE 13

3.3. Relating $\mathcal{K}(D)$ and $MS(L)$. Our aim is to relate $\mathcal{K}(D)$ to $MS(L)$. We have a map $\mathcal{A} : V(\mathcal{K}(D)) \rightarrow V(MS(L))$, as described in Section 2. To study the properties of this map, we will establish a local description at each vertex of $\mathcal{K}(D)$.

Proposition 3.10. *Let $v_0 \in V(\mathcal{K}(D))$. 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 $v_1 \in V(\mathcal{K}(D))$ with $d_{\mathcal{K}(D)}(v_0, v_1) = 1$. Then $d_{MS(L)}(R_0, \mathcal{A}(v_1)) \leq 1$.

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

Since $d_{\mathcal{K}(D)}(v_0, v_1) = 1$, 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 14). The total of these labels is 0, and the 1s and




FIGURE 14

-1 s must alternate. 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 the 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 3.9, $\mathcal{A}(v_1)$ 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 . \square

Proposition 3.11. *Let v_0 and R_0 be as above. Let R_1 be a surface given in admissible special form at D . Then there exists $v_1 \in V(\mathcal{K}(D))$ such that $d_{\mathcal{K}(D)}(v_0, v_1) \leq 1$ 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 $e \in E(\theta(D))$. 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 15 to give a new flype circle ϕ'' , we can reduce the number of flype circles running over e without changing the admissible form. Hence we may assume that at most one flype circle runs over any edge of $\theta(D)$.




FIGURE 15

Let $\Lambda \subseteq S^2$ be the positive side of the flype circles, and $\lambda \subseteq S^2$ 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 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 $d_{\mathcal{K}(D)}(v_0, v_1) = 1$. 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}_Λ with an edge from r across $e(r)$ for each r in Λ . There are only finitely many regions in Λ , so \mathcal{G}_Λ 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 3.12. *The complex $\mathcal{K}(D)$ is flag. That is, if the 1-skeleton of a simplex is contained in the complex then so is the simplex.*

Proof. Let v_0, \dots, 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, \dots, v_m given by inclusion of the sets A_0, \dots, A_m . We wish to show that v_0, \dots, v_m span a simplex in $\mathcal{K}(D)$. From the proof of Proposition 3.11 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. □

The following proposition says that $MS(L)$ is also flag. This means that each complex is defined by its 1-skeleton, so in relating them we may restrict our attention to comparing vertices and edges.

Proposition 3.13 ([12] Proposition 4.9). *Let v_0, \dots, v_n be distinct vertices of $MS(L)$. If v_i and v_j span an edge in $MS(L)$ whenever $i \neq j$, then v_0, \dots, v_n span an n -simplex in $MS(L)$.*

4. 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 ([4]), 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 16. This gives the basis for an inductive argument.




FIGURE 16

Our inductive hypothesis in the proof of Proposition 4.10 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 4.1. Say that such a region is *cuttable*.

Remark 4.2. We will call a link diagram *prime* if the link is prime. That is, we do not require a ‘prime link diagram’ to be reduced.

Theorem 4.3 ([8] Theorem 1). *Let D be a reduced, alternating link diagram of a link L . Then L is not prime if and only if this is visible in D . That is, if L is not prime then there is a simple closed curve ρ in \mathbb{S}^2 that is disjoint from the crossings of D and meets the edges of D exactly twice transversely such that there is at least one crossing on each side of ρ .*

Lemma 4.4. *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 17a, where neither A nor B is trivial or consists of a single line of crossings surrounding 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 17b. 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




FIGURE 17

r such that ρ_r is innermost. Then no white region r' lies entirely inside ρ_r . This contradicts the choice of ρ_r . \square

4.1. Reducing a diagram. Once we have cut the link and the surface along a crossing, the resulting diagram may not be reduced; there could be one or two lines of crossings that can be removed by a sequence of type I Reidemeister moves. In the proof of Proposition 4.10 we will need to know how the set of P -arcs changes as we remove these crossings. In this section we consider the specific situations that will arise because of this.

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 18a. Let p be the point of D^c corresponding to c in D . One way to remove c is to pull the undercrossing




FIGURE 18

arc upwards into \mathbb{B}_a and then untwist. This leaves $R \cap \mathbb{S}^2$ as shown in Figure 18b, 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 19. The P -arcs in D^c , without ρ , give a special form for R .




FIGURE 19

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 20.




FIGURE 20

Definition 4.5. 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 4.6. 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.

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

Lemma 4.7. 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 minimal genus). 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 3.9, the surface given by the P -arcs on D_0^c is R , and in particular is minimal genus.

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 20), in \mathbb{B}_a the P -arcs on D^c near p are connected as shown in Figure 21. Since there is no P -arc across c in D , σ_1 cannot consist




FIGURE 21

entirely of P -arcs. 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 we have Figure 22. There is a P -arc across c_1 , so a neighbourhood of c_1 is as shown




FIGURE 22

in Figure 23a. 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 23b. Inductively we see that the arc σ_2 will return to ∂r after every crossing. Thus D^c has the structure shown in Figure 24. Because D^c is reduced and alternating, this contradicts that it is prime. \square




FIGURE 23




FIGURE 24

Lemma 4.8. *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 4.7 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 25. In the




FIGURE 25

first case, the P -arcs around c' are not connected as in Figure 20. Therefore the situation in Figure 25b 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 4.7, we may assume that there are no flype circles needed and every crossing in D^c has a P -arc across it. Further following the proof of Lemma 4.7 shows that D has the structure shown in Figure 26, contradicting that it is prime.




FIGURE 26

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 4.7 gives that there is a P -arc across c in D . □

Lemma 4.9. *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 20) the arcs of $R \cap \mathbb{B}_a$ on D must connect as in Figure 27a, where σ does not consist entirely of P -arcs. This contradicts that there is a P -arc across c' in D^c (see Figure 27b). □




FIGURE 27

4.2. Adjacent surfaces can be put into admissible special form.

Proposition 4.10. *Let D be a prime, reduced, special alternating diagram of a link L . Let R be the (minimal genus) Seifert surface for L given by applying Seifert's algorithm to D . Let R' be a minimal genus 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 4.4, 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 16 to give a new diagram D_c and a new minimal genus 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 28. Again, the special form of R' at D is admissible.




FIGURE 28

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 in Figure 29. As D is reduced and D_c is prime, either A or B consists of a single line of crossings. Thus cutting along c leaves one or two lines of crossings joined by black bigons in a single white region of what is otherwise a reduced diagram (for example as shown in Figure 30).

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 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 4.7 or Lemma 4.8 hold in relation to D_1 . One of these lemmas followed by repeated




FIGURE 29




FIGURE 30

application of Lemma 4.9 gives that the special form of R'_c at D_c has a P -arc across every crossing in the one or two lines. 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 crossings of D_1 not in D_0 lies between c' and ρ . This then means that in fact c and all crossings in the line(s) of 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. \square

Corollary 4.11 ([5] Theorem 1.1). *Let L be a prime, special alternating link with a reduced, special alternating digram D . Let R be a minimal genus 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 4.10, Lemma 3.9 and the fact that $MS(L)$ is connected. \square

Corollary 4.12. *Every minimal genus Seifert surface for L is connected.*

5. Distinguishing surfaces in special form

Our current aim is to prove Theorem 1.4. To do so, we wish to show that the map $\mathcal{A} : V(\mathcal{K}(D)) \rightarrow V(MS(L))$ extends to an isomorphism of the simplicial complexes. By Proposition 3.10 we know that \mathcal{A} can be extended to a map on the edges of the complexes. Since $MS(L)$ is flag, this gives a simplicial map $\mathcal{A} : \mathcal{K}(D) \rightarrow MS(L)$. We will show that \mathcal{A} is a local isomorphism at each point, and so is a covering map. As $MS(L)$ is simply connected, \mathcal{A} is therefore an isomorphism.

Because both $MS(L)$ 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 $MS(L)$. From Proposition 4.10, together with Proposition 3.11, 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 1.4.

5.1. Product regions and product discs.

Definition 5.1. Let M be a connected 3-manifold, and let S, S' be surfaces properly embedded in M in general position. Say S and S' *bound a product region* if the following holds. There is a compact surface T , a finite collection $\rho \subseteq \partial T$ of arcs and simple closed curves and a map of $N = (T \times I) / \sim$ into M that is an embedding on the interior of N and has the following properties.

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

Here \sim collapses $\{x\} \times I$ to a point for each $x \in \rho$.

Proposition 5.2 ([12] Proposition 4.8). *Let M be a Haken manifold with incompressible boundary. Suppose S, S' are incompressible, ∂ -incompressible surfaces properly embedded in M in general position. Suppose further either that $S \cap S' \neq \emptyset$ but S can be isotoped to be disjoint from S' , or that S, S' are isotopic. Then S, S' bound a product region.*

Remark 5.3. If $M = \mathbb{S}^3 \setminus \mathcal{N}(L)$ for a non-split link L other than the unknot then M is Haken and ∂ -irreducible. Furthermore, if S, S' are minimal genus Seifert surfaces for L then they are properly embedded, incompressible and ∂ -incompressible.

Using this, we can show that two disjoint Seifert surfaces are not isotopic, or that two non-disjoint surfaces cannot be made disjoint, by showing that they do not bound a product region. In other words, we need to check whether the pieces of the link complement given by cutting along the Seifert surfaces have a particular form, which we can describe using the language of sutured manifolds.

Definition 5.4. 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.

Definition 5.5. A *product sutured manifold* is a sutured manifold (M, s) that is homeomorphic to $S_+(M) \times [-1, 1]$ with $s = \partial S_+(M) \times \{0\}$.

Remark 5.6. A product region 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' .

One method for establishing whether a sutured manifold is a product sutured manifold is to relate the sutured manifold to a single Seifert surface for a link.

Definition 5.7. For a sutured manifold (M, s) embedded in \mathbb{S}^3 , the *complementary sutured manifold* (M', s') is defined by $M' = \mathbb{S}^3 \setminus \text{int}_{\mathbb{S}^3}(M)$ 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 .

Remark 5.8. 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.

It will therefore be of interest to us to be able to tell whether a given link is fibred.

Definition 5.9. Given a special alternating link diagram D , let $G(D)$ be the planar graph with a vertex in each white region of D and an edge through each crossing.

Theorem 5.10 (see [10] Theorem 3(5) and [1] Theorem 5.1). *Let D be a special alternating diagram of a prime link L . Then L is fibred if and only if the graph $G(D)$ can be reduced to a single vertex by a sequence of the following moves.*

- Delete a loop (that is, an edge with both endpoints at the same vertex).
- Contract an edge, one endpoint of which is at a vertex with valence 2.

Informally, $G(D)$ satisfies this condition exactly when it is a ‘tree of loops’.

Corollary 5.11. *If D is a reduced, special alternating diagram of a prime link L , other than the unknot, where L is fibred then D is as shown in Figure 31.*




FIGURE 31

Proof. As L is fibred, the description in Theorem 5.10 applies to $G(D)$. Since D is reduced, the graph cannot contain a loop.

Suppose that $G(D)$ contains a cut vertex. Consider a simple closed curve ρ in the diagram plane that passes through the cut vertex and is otherwise disjoint from $G(D)$. Then at least one vertex lies on each side of ρ . By reconstructing D from $G(D)$, we can see that the curve ρ shows that L is not prime. This is a contradiction, so $G(D)$ cannot contain a cut vertex.

The graph $G(D)$ must therefore form a single simple closed curve divided into at least two edges. This shows that D is as required. \square

To reach the point where the sutured manifold of interest can be viewed as the complementary sutured manifold to a Seifert surface, and to simplify the Seifert surface where we can, we will cut along product discs.

Definition 5.12. 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.

Definition 5.13. 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 $s' = (s \cap M') \cup (\rho \times \{\pm 1\})$.

Remark 5.14. M' is a product sutured manifold if and only if M is.

Remark 5.15. Suppose 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' .

Lemma 5.16. Let D be a special alternating diagram, and let r be a white bigon region of D . Then r defines a product disc in the complementary sutured manifold to the Seifert surface R given by applying Seifert's algorithm to D . The effect of the product disc decomposition is to remove the region r , replacing the two crossings of r with a single crossing. Such a change on D has the effect of deplumbing a Hopf band from R (see Figure 32).

Proposition 5.17 ([7] Proposition 1.4). Let R be a connected minimal genus 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 $\text{MS}(L)$ is isomorphic to the link of R' in $\text{MS}(L')$. Furthermore, the isomorphism is induced by the procedure described in Remark 5.15.




FIGURE 32

We will use Lemma 5.20 in Section 6.

Definition 5.18. Say a product disc is *inessential* if it is separating and the disc decomposition along it creates a component that is a 3–ball with a single suture. Otherwise, it is *essential*.

Remark 5.19. If a product disc T in a sutured manifold M is inessential then $\partial T \cap S^+(M)$ is inessential in $S^+(M)$, and the same is true for $S^-(M)$.

Lemma 5.20. 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 in S_2 . \square

5.2. Adjacent surfaces are distinct.

Definition 5.21. 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 5.22. Let D be a reduced, special alternating diagram for a prime link L . Let $v, v' \in V(\mathcal{K}(D))$ such that $d_{\mathcal{K}(D)}(v, v') = 1$ 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 $d_{\text{MS}(L)}(R, R') = 1$.

Proof. Since $d_{\mathcal{K}(D)}(v, v') = 1$, the proof of Proposition 3.10 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 3.11). 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 5.2.

We first 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 33, where the black region r' shown is an a -region, then M^+




FIGURE 33

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 .

Considering again Figure 33, the arcs shown in bold in Figure 34 form part of the 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 , given the positioning of the P -arcs. Hence the arc ρ_1 defines a product disc in M^+ contained in \mathbb{B}_a , and the product disc in M^+ defined by ρ_2 is contained in \mathbb{B}_b . Decomposing M^+ along these product discs, as well as the analogous discs from the other flype circles,




FIGURE 34

cuts it into pieces that each lie either on the positive side or on the negative side of the flype circles. Each piece on the negative side of the flype circles is a product region, since we have seen that R and R' are ‘parallel’ there with M^+ lying between them. Thus we only need to consider the pieces of M^+ left by removing these product regions. We will see that at least one of the remaining pieces is not a product sutured manifold.

Let Φ be the union of the flype circles. The remaining pieces of M^+ correspond to the components of $S^2 \setminus \Phi$ 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 35 along




FIGURE 35

each 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_Λ constructed in this way is of a fibred link with fibre given by the surface R_Λ . Again by symmetry, we need only consider D_Λ .

Consider $G(D)$ and $G(D_\Lambda)$ (as in Definition 5.9), and a fixed flype circle ϕ_Λ that bounds Λ . As can be seen from in Figure 36, $G(D_\Lambda)$ is obtained from $G(D)$ by collapsing everything on the negative side of the flype circles. There may be other flype circles in A that are collapsed to give A' .

As D is prime and reduced, the same is true of D_Λ . Suppose D_Λ is fibred. Then it must be as in Figure 31. Since the flype ϕ_Λ is essential, at least one vertex of $G(D_\Lambda)$ in A' comes from collapsing a flype circle in A . By ‘uncollapsing’ all such vertices in $G(D_\Lambda)$ we see that $G(D)$ decomposes as shown in Figure 37. We




FIGURE 36

can then combine two flype circles as in Figure 15, contradicting our choice of flype circles. Thus D_Λ is not fibred. \square




FIGURE 37

Proposition 5.23. *Let D be a reduced, special alternating diagram for a prime link L , and let $v_0 \in V(\mathcal{K}(D))$ be given by D . Let $v_{-1}, v_1 \in V(\mathcal{K}(D))$ such that $d_{\mathcal{K}(D)}(v_{-1}, v_0) = d_{\mathcal{K}(D)}(v_0, v_1) = 1$ and $d_{\mathcal{K}(D)}(v_{-1}, v_1) = 2$. Let $R_i = \mathcal{A}(v_i)$. Then $d_{MS(L)}(R_{-1}, R_1) = 2$.*

Remark 5.24. 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 Appendix A.

Proof. Let $\mathcal{N}(L)$ denote the regular neighbourhood of L . Following the proof of Proposition 3.10, 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. 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 3.11, 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 $d_{\mathcal{K}(D)}(v_{-1}, v'_0) \leq 1$ and $d_{\mathcal{K}(D)}(v'_0, v_1) \leq 1$. Thus $d_{\mathcal{K}(D)}(v_{-1}, v'_0) = d_{\mathcal{K}(D)}(v'_0, v_1) =$

1 since $d_{\mathcal{X}(D)}(v_{-1}, v_1) = 2$. 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. In addition, we can isotope all flype circles to be mutually disjoint. Also choose the number of flype circles to be minimal. Denote the set of all flype circles by Φ .

We wish to apply Proposition 5.2 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 42 gives an example of the type of picture we will build up.

First note that $\partial R_{-1} = \partial R_1$. Let $\bar{\rho}$ 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, $\bar{\rho}$ 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 $\partial \bar{\rho}$ denote the (as yet undefined) set of endpoints of these arcs.

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 38a. All P -arcs in λ can be copied to D_λ . If D_λ together with the P -arcs is not connected,




FIGURE 38

choose a simple closed curve ψ around each component, separating it from the other components. We can choose these 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 38b. 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 now arrange the P -arcs of each surface. Let c be a crossing in D . Our positioning of the flype circles ensures that c lies in the interior of either λ_{-1} or λ_1 , and so it has at least one P -arc across it on the positive side.

Definition 5.25. 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 39, 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




FIGURE 39

same is true of the analogous curve C_1 in $\partial(R_1 \cap \mathbb{B}_a)$.

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 $\hat{\cdot}: A \rightarrow A$ such that if $a_1 \in A$ then \hat{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, \hat{a}_1 and a_2, \hat{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

40). We therefore create a subset a of A with a new involution where this does




FIGURE 40

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 \hat{a}_1 with \hat{a}_2 . Repeat this as many times as possible. Note that it is possible, as in Figure 40, to reach a pair where $\hat{a}_1 = a_2$. In such a case the two points are simply deleted from A . Since A is finite, this process will terminate. Notice that the final result will not depend on the order in which pairs of points are chosen. Let a be the resulting set of points, and $\tau : a \rightarrow a$ the resulting involution. The following three properties of A are also true in a . If $a_1 \in a$ then a_1 and \bar{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 \bar{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, \bar{a}_1 and a_2, \bar{a}_2 do not interleave on C_{-1} or on C_1 . In addition, a has the property that the inside P -arcs at the two crossings adjacent to a point of a have different colours.

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

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 41a. We can then connect the P -arcs along the overcrossing arc at c' without them crossing (see Figure 41a). Suppose instead that 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 41b. Note that, since there are crossings where the blue P -arc




FIGURE 41

is inside, and crossings where the red P -arc is inside, we know that $a \neq \emptyset$. In this way we can connect all the P -arcs to form two sets of simple closed curves, as required. Figure 42 is an example of the result of this process.




FIGURE 42

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 $\bar{\rho}$ connects points $a_1, a_2 \in a$ if and only if $\bar{a}_1 = a_2$. Recall that a_1 and \bar{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, \bar{a}_1 and a_3, \bar{a}_3 do not interleave along these curves. The same holds in \mathbb{B}_b . Thus there is no obstruction to our choice of $\bar{\rho}$. This completes the construction of $R_{-1} \cup R_1$.

It now remains to check that R_{-1} and R_1 do not bound a product region. Then, since R_{-1} and R_1 are not disjoint, it follows from Proposition 5.2 that $d_{\text{MS}(L)}(R_{-1}, R_1) > 1$. Let M be the complement in $\mathbb{S}^3 \setminus \mathcal{N}(L)$ of a regular neighbourhood of $R_{-1} \cup R_1$.

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 $\bar{\rho}$. 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 $\partial\bar{\rho}$. 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 $\partial\bar{\rho}$ 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 43. Suppose M_L is a product region between R_{-1} and R_1 . Then a closed




FIGURE 43

regular neighbourhood of $\bar{\rho} \cup \partial\mathcal{N}(L)$ in ∂M_L is of the form $\partial S_L \times I$ for some surface S_L , and in particular is a collection of annuli. Since any component of $\partial\mathcal{N}(L) \cap \partial M_L$ is an annulus, and each endpoint of an arc of $\bar{\rho}$ 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 $R_{-1} \cap R_1$.

First 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 5.22, 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 .

From the construction of A , it is clear that there is no point $a_1 \in A$ such that x and x' interleave with a_1 and \hat{a}_1 on C_{-1} or C_1 . It follows that there is no point

$a_2 \in a$ such that x and x' interleave with a_2 and $\overline{a_2}$ on C_{-1} or C_1 . Therefore 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 35. 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 5.22, 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 P -arcs in $\lambda_{-1} \cap \lambda_1$. Then r lies between a red P -arc and a blue P -arc parallel to it. 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 , that is contained in D and the P -arcs, from r to a flype circle ϕ on the boundary of λ . 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 $\partial \bar{\rho} = a \cup b$. From this we see that this is a flype circle of R_{-1} . It also tells us that λ is coloured blue, which means there is a path σ , contained 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 $\bar{\rho}$ 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. \square

6. Non-special links

Theorem 1.2 of [5] gives two families of links that illustrate the importance of assuming, in Theorem 1.4, 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 6.1. *If L_{γ_n} is the link shown in Figure 44 then $MS(L_{\gamma_n})$ contains an $(n - 1)$ -simplex in which exactly one vertex is given by applying Seifert's algorithm to an alternating diagram.*




FIGURE 44

Proof. There is only one non-trivial flype possible on the diagram given, and it is not an essential flype. Hence only one minimal genus Seifert surface, R , for L_{γ_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'_{γ_n} shown in Figure 45. By Theorem 1.4 the link of the vertex R' in $MS(L'_{\gamma_n})$ is an $(n - 2)$ -simplex, and so the same is true of R in $MS(L_{\gamma_n})$ by Proposition 5.17. \square




FIGURE 45

Proposition 6.2. *If L_{δ_n} is the link shown in Figure 46 then $MS(L_{\delta_n})$ is as shown in Figure 47, 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 minimal genus 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 48. In




FIGURE 46




FIGURE 47

Figure 46 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 46 (see Figure 49).




FIGURE 48




FIGURE 49

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 50. Thus we see that R_1 is disjoint from R_0 . We




FIGURE 50

want to show that $R_0 \neq R_1$, and that no other minimal genus Seifert surface is disjoint from R_0 . To do this, we show that the link of R_0 in $MS(L_{\delta_n})$ consists of only one vertex, and this vertex is R_1 .

The white region marked * in Figure 50 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 1.4 the link of $R_0^{(n)}$ is a single vertex. Hence, by repeated application of Proposition 5.17, the link of R_0 is also a single vertex, as required. By following the constructions in Theorem 1.4 and Proposition 5.17 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 51a and b respectively. Also as before, after n product disc decompositions we may apply Theorem 1.4, 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 vertex. Each white bigon 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 52, where the horizontal twisted band lies above the rest of the surface. Thus we are interested in a plumbing of two surfaces, S_a and S_b , shown in Figure 53. 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 minimal genus Seifert surfaces for L_a




FIGURE 51




FIGURE 52




FIGURE 53

and $L_b = \partial S_b$ respectively, and also that neither link is fibred. This is enough to show there is at least one minimal genus 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.

To do this, we consider the arcs ρ_a and ρ_b shown in Figure 53. 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 [7] 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 54 shows the result of the first decomposition. This leaves




FIGURE 54

two copies of the 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 5.20, there is no product disc T with $\rho_b \subset \partial T$. \square

Remark 6.3. These two propositions together complete the proof of [5] Theorem 1.2.

7. The realisation of $MS(L)$

Let L again be a prime link with a reduced, special alternating diagram D . In this section we turn our attention to Theorem 1.6 of [5], which asserts that $MS(L)$ is homeomorphic to a ball in \mathbb{R}^n for some n . We will show, firstly, that

each θ -graph in $\theta(D)$ gives us a simplicial complex that is a subdivision of a single simplex and, secondly, that combining these gives a product simplicial complex. The results we will use refer to simplicial complexes with a partial ordering on the vertices. Thus we will need to give such an ordering 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 $\theta(D)$.

Imagine cutting up an n -simplex using cuts parallel to the faces, such that each edge is cut into m equal pieces. It is straight-forward to identify the locations of the vertices. The following definition sets out which higher-dimensional simplices need to be included.

Definition 7.1 ([2]). Fix $m, n \in \mathbb{N}$, and let v_0, \dots, v_n be the (ordered) vertices of an n -simplex $[v_0, \dots, v_n]$.

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

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

Remark 7.2. In [2] Section 3 it is proved that $\text{Esd}_m([v_0, \dots, v_n])$ is indeed a subdivision of $[v_0, \dots, v_n]$.

Proposition 7.3 ([5] 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.*

That is, the more places there are that we could put a crossing, the higher the dimension of the simplex, and the more crossings we have to distribute between those places, the finer the subdivision.

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 $\hat{\theta}(D)$ is $\theta(D)$). Label the edges of $\theta(D)$ as e_0, \dots, e_n where $e_0 < e_1 < \dots < e_n$. 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, \dots, v_n be the (ordered) vertices of an n -simplex $[v_0, \dots, v_n]$. Given a vertex $u = (w_0, \dots, w_n)$ 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 $\text{Esd}_m([v_0, \dots, v_n])$. Thus there is a map B from the vertices of $\mathcal{K}(D)$ to the vertices of $\text{Esd}_m([v_0, \dots, v_n])$. 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, \dots, u_n be the vertices of an n -simplex in $\mathcal{K}(D)$ with $u_0 < u_1 < \dots < u_n$. Then there is a permutation $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ 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, \dots, 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 $\text{Esd}_m([v_0, \dots, v_n])$.

Conversely, choose an n -simplex of $\text{Esd}_m([v_0, \dots, v_n])$ 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 $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ 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, \dots, 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. \square

Definition 7.4 ([3] 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 7.5. 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 $(u \leq v \text{ or } v \leq u)$.

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

Definition 7.6 ([3] 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 $V(\mathcal{X}_1) \times V(\mathcal{X}_2)$. Vertices $(u_0, v_0), \dots, (u_n, v_n)$ span an n -simplex if the following hold.

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

Remark 7.7. The projection maps on the vertices extend to simplicial maps of the complexes.

Theorem 7.8 ([3] Chapter II Lemma 8.9). *The map $|\mathcal{X}_1 \times \mathcal{X}_2| \rightarrow |\mathcal{X}_1| \times |\mathcal{X}_2|$ induced by projection is a homeomorphism.*

Remark 7.9. With the ordering on the vertices defined above, $\mathcal{K}(D)$ is an ordered simplicial complex.

Theorem 7.10 ([5] 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 7.3.

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 $\mathbb{S}^2 \setminus \theta_0$ 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 7.8, $|\mathcal{K}(\theta_1) \times \mathcal{K}(\theta_2)| \cong |\mathcal{K}(\theta_1)| \times |\mathcal{K}(\theta_2)|$. 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 $[w_0, \dots, w_n]$ 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 $[w_0, \dots, w_n]$ 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 $[(u_0, v_0), \dots, (u_n, v_n)]$ is a simplex in $\mathcal{K}(\theta_1) \times \mathcal{K}(\theta_2)$.

Conversely, consider an n -simplex $[(u_0, v_0), \dots, (u_n, v_n)]$ 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 $(u_0, v_0) < (u_1, v_1) < \dots < (u_n, v_n)$. 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 $[u_0, \dots, u_n]$ 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 $[(u_0, v_0), \dots, (u_n, v_n)]$ is an n -simplex in $\mathcal{K}(\theta_0)$. \square

Appendix A. Proof by picture: an example of Proposition 5.23

Let L_ε be the link with minimal genus Seifert surfaces R_{-1} and R_1 given in admissible special form in Figure 55a and b respectively. We show the steps of the proof of Proposition 5.23 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 55c.
- 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 56a and b.
- iii. Colour the components of $\mathbb{S}^2 \setminus \Psi$. Define a and b , and the corresponding involutions. See Figure 56c (the pairs of points in a are marked with stars).
- iv. Position the P -arcs and connect them. Figure 57 shows the resulting picture in \mathbb{B}_a .




FIGURE 55




FIGURE 56




FIGURE 57

References

- [1] BANKS, JESSICA E. Homogeneous links, Seifert surfaces, digraphs and the reduced Alexander polynomial. *Geom. Dedicata* **166** (2013), 67–98. MR3101161, Zbl 1312.57003, arXiv:1101.1412, doi:10.1007/s10711-012-9786-1. 469
- [2] EDELSBRUNNER, HERBERT; GRAYSON, DANIEL R. Edgewise subdivision of a simplex. *Proceedings of the Fifteenth Annual Symposium on Computational Geometry* (Miami Beach, FL, (1999)), 24–30. *ACM, New York*, 1999. MR1802191, Zbl 0968.51016, doi:10.1007/s004540010063. 486
- [3] EILENBERG SAMUEL; STEENROD, NORMAN. Foundations of algebraic topology. *Princeton University Press, Princeton, New Jersey*, (1952). xv+328 pp. MR0050886, Zbl 0047.41402, doi:10.1515/9781400877492. 487, 488
- [4] GABAI, DAVID. Genera of the alternating links. *Duke Math. J.*, **53** (1986) no. 3, 677–681. MR860665, Zbl 0631.57004, doi:10.1215/S0012-7094-86-05336-6. 446, 452, 460
- [5] HIRASAWA, MIKAMI; SAKUMA, MAKOTO. Minimal genus Seifert surfaces for alternating links. *KNOTS '96* (Tokyo), 383–394. *World Sci. Publ., River Edge, NJ*, 1997. MR1664976, Zbl 0972.57009. 446, 448, 467, 480, 485, 486, 488
- [6] KAKIMIZU, OSAMU. Finding disjoint incompressible spanning surfaces for a link. *Hiroshima Math. J.* **22** (1992), no. 2, 225–236. MR1177053, Zbl 0774.57006, doi:10.32917/hmj/1206392900. 445
- [7] KAKIMIZU, OSAMU. Classification of the incompressible spanning surfaces for prime knots of 10 or less crossings. *Hiroshima Math. J.* **35** (2005), no. 1, 47–92. MR2131376, Zbl 1083.57009, doi:10.32917/hmj/1150922486. 446, 470, 485
- [8] MENASCO, WILLIAM W. Closed incompressible surfaces in alternating knot and link complements. *Topology* **23** (1984), no. 1, 37–44. MR0721450, Zbl 0525.57003, doi:10.1016/0040-9383(84)90023-5. 460
- [9] MENASCO, WILLIAM W.; THISTLETHWAITE, MORWEN. The classification of alternating links. *Annals of Mathematics* **138** (1993), no. 1, 113–171. MR1230928, Zbl 0809.57002, doi:10.2307/2946636. 447
- [10] MURASUGI, KUNIO; STOIMENOW, ALEXANDER. The Alexander polynomial of planar even valence graphs. *Adv. in Appl. Math.* **31** (2003), no. 2, 440–462. MR2001624, Zbl 1029.05039, doi:10.1016/S0196-8858(03)00021-6. 469
- [11] PRZYTYCKI, PIOTR; SCHULTENS, JENNIFER. Contractibility of the Kakimizu complex and symmetric Seifert surfaces. *Trans. Amer. Math. Soc.* **364** (2012), no. 3, 1489–1508. MR2869183, Zbl 1239.57025, arXiv:1004.4168, doi:10.1090/S0002-9947-2011-05465-6. 446
- [12] SAKUMA, MAKOTO. Minimal genus Seifert surfaces for special arborescent links. *Osaka J. Math.* **31** (1994), no. 4, 861–905. MR1315011, Zbl 0871.57010. 460, 468
- [13] SAKUMA, MAKOTO; SHACKLETON, KENNETH J. On the distance between two Seifert surfaces of a knot. *Osaka J. Math.* **46** (2009), no. 1, 203–221. MR2531146, Zbl 1177.57006, arXiv:math/0701468. 446

(Jessica E. Banks) DEPARTMENT OF MATHEMATICAL SCIENCES, THE UNIVERSITY OF LIVERPOOL, MATHEMATICAL SCIENCES BUILDING, LIVERPOOL L69 7ZL, ENGLAND
Jessica.Banks@Liverpool.ac.uk
Jessica.Banks@LMH.oxon.org

This paper is available via <http://nyjm.albany.edu/j/2022/28-18.html>.