Crushtaceans: A Graph Theoretic Representation of Fully Augmented Links

Cole Twogood

June 2023

Abstract

This paper tackles a graph theoretic representation of Fully Augmented Links which we call Crushtaceans. We show that for every FAL there is a correlated crushtacean, and for every crushtacean there is a correlated FAL. We define flype equivalence on crushtaceans, which can be used to show if two FALs are isotopic.punch!

1 Introduction

The goal of our paper will be to show a powerful unique relationship between a specific class of graphs and Fully Augmented Links.

When we say FAL, we are referring to a hyperbolic Fully Augmented Link as discussed in Purcell [1]. A Hyperbolic Fully Augmented Link is a link generated by a process of augmenting a link diagram which is prime, nonsplittable, twist reduced, and at least two twist regions. The augmenting is done by putting an unknot around each twist region, which we refer to as a crossing circle. Then we "untwist" each twist region by performing Dehn twists to remove full twists, which removes crossings until there is either 0 or 1 crossing at each twist region.

In an FAL we will refer to a crossing circle as **flat** if the strands which go through it do not cross. We will refer to a crossing circle as **twisted** if the strands which go through it do cross. An FAL as a whole is flat if all the crossing circles are flat.

This paper primarily focuses on the notion of a Crushtacean, a graph designed to represent FALs. Our graph theoretic definition is inspired by paper's such as Purcell [1] and Chesebro [2]. For clarity we will explain the classic construction of the crushtacean presented in such papers, though it is critical to note that this is not the construction we assume for most of this paper.

Purcell guarantees that every FAL has a polyhedral decomposition corresponding to a circle packing. This circle packing is made of a finite collection of Euclidean circles in S^2, which meet only in points of tangency. In these circle packing there exists triangular interstices between circles, these triangular interstices are the vertices of the traditional crushtacean. We will draw an edge wherever two triangular interstices intersect. (see [1]).

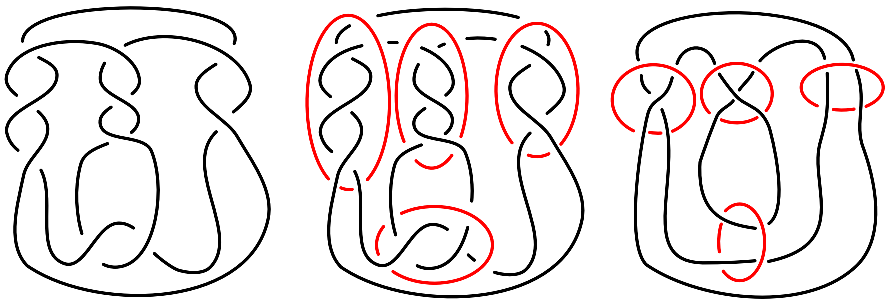

Figure 1: the basic construction of full augmentation on a link, creating an FAL.

Note that this is what Purcell refers to as the dual of the nerve, and what Chesebro et. al. refer to as a crushtacean. [1] [2]

Define Flype

It is important to define some of the terms and notation we will use in this paper. We will primarily use the symbols A, B, C to refer to links, and X, Y, Z to refer to graphs. We will use $A \sim B$ to mean that two links A and B are isotopic. We will use $X \cong Y$ to mean that two graphs are isomorphic.

2 Determining FALs under Crushtaceans

For our paper, rather than defining a crushtacean as a graph resulting from an FAL, we'll begin with a completely graph theoretic definition. Later in this paper, we will prove that this definition is equivalent to the classic definition.

First, we must recall some definitions from graph theory. A graph is **planar** if there exists a planar embedding in which no edges cross. A graph is **3-regular**, if every vertex is incident with 3 edges. Finally, a graph is **3-edge-connected** if you must remove at least 3 edges to disconnect the graph.

Definition 2.1. *A* ***Crushtacean*** *ks a if it is 3-regular, 3-edge-connected, planar graph.*

We need to establish many properties of crushtaceans before we can use them. We start with a theorem on the connectivity. Recall that the vertex-connectivity of a graph is the number of vertices that need to be removed to disconnect the graph.

Theorem 2.2. *All crushtaceans are 3-vertex-connected.*

Proof. For this proof we will refer to the vertex connectivity as c. We know from Whitney's Theorem on connected graphs that c is less than the edge-connectivity. Given the edge-connectivity of a crushtacean is 3, that means that $c \leq 3$. [3]

We know that edge-connectivity being non-zero implies that the graph is disconnected. Given the graph is connected, $c \neq 0$.

We can also show that $c \neq 1$ by contradiction. Assume there exists a vertex which would disconnect the graph if removed, we will call this vertex α. We know that removing α would split the graph into two disconnected subgraphs X_1 and X_2. Correspondingly, since crushtaceans are 3-regular, there are precisely 3 edges incident to α. By the pigeonhole principle, 2 of the edges must be incident to vertices in one subgraph, and only 1 edge is incident to the other subgraph, we will say without loss of generality that only 1 edge is incident to vertices in X_1. It is easy to see that the edge incident with α and a vertex in X_1 is a sufficient edge-cut, since it would disconnect X_1 from the rest of the crushtacean. The existence of an edge-cut of size 1 contradicts our assumption that crushtaceans are 3-edge-connected.

We can show that $c \neq 2$ by contradiction. Assume there exists two vertices which would disconnect the graph if removed, which we will call α and β. Given our graph is 2-vertex-connected, we know that neither α or β would disconnect the graph alone. We know that removing these vertices would disconnect the graph into two disconnected subgraphs X_1 and X_2. It is worth noting that X_1 and X_2 may themselves be internally disconnected, however this doesn't change the contradiction we seek to show. Since crushtaceans are 3-regular, there are 3 edges incident to each vertex. By the pigeonhole principle we can say α has 2 edges incident to vertices in one subgraph, and that only 1 edge is incident to vertices in the other subgraph, we will say without loss of generality that only 1 edge is incident to vertices in X_1.

We also know that there must be 3 edges incident to β. By the pigeonhole principle β has 2 edges incident to vertices in one subgraph, and that only 1 edge is incident to vertices in the other subgraph. If β has only 1 edge incident to a vertex in X_1 then we can construct an edge cut in which we cut the two edges of α and β which are incident to vertices in X_1, which must disconnect X_1, which would be a 2-edge-cut. If β has only 1 edge incident to a vertex in X_2 then we can construct an edge cut in which we cut the edge of β incident with a vertex in X_2, and the edge of α incident with a vertex in X_1, this also disconnects the graph, since it would create the disconnected subgraphs of $X_1 + \beta$ and $X_2 + \alpha$, thus we have shown there is a 2-edge-cut. Figure 2 gives visual to this idea. Since we have shown that $c = 2$ implies the existence of a 2-edge-cut, we have contradicted the assumption that crushtaceans are 3-edge-connected.

We have proven that $c \neq 0$, 1, 2, which means that $c > 2$. Since $2 < c \leq 3$, we've shown $c = 3$. Consequently proving that every crushtacean is 3-vertex-connected. □

The property of being 3-vertex-connected allows us to invoke many theorems from other papers. 3-vertex-connected planar graphs have many special properties, such as the following.

Lemma 2.3. *All embeddings of isomorphic crushtaceans on the surface of a sphere are isotopic.*

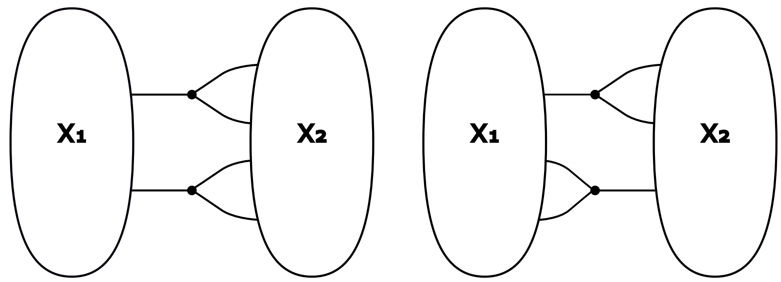

Figure 2: The two cases in which $c = 2$

Proof. This can be proven using Theorem 2.2 and a theorem proven by Whitney. Whitney's theorem implies that any two embeddings of a 3-vertex-connected graph must be isotopic on the surface of a sphere [4]. Since all crushtaceans are 3-vertex-connected this completes the proof. □

We use Lemma 2.3 to say that referring to an embedding of a crushtacean and the crushtacean itself are fundamentally the same. For this reason when we refer to crushtaceans we will always assume that it is embedded onto the surface of a sphere.

Recall that traditionally Crushtaceans can be generated using FALs, however they do not carry enough information to recreate the FAL they originated from. For this reason we must have a notion of Painting and Signing a crushtacean. A painting on a crushtacean is a selected perfect matching, such that no two painted edges are incident and that every vertex is incident with a painted edge. This is analogous to Purcell's painted nerves. [1]

To sign a crushtacean we assign a value of -1, 0, or 1 to each painted edge. The signed value of each edge represents whether or not there is a crossing at that crossing circle. A zero means that the crossing circle is flat, whereas a ± 1 means the strands cross at a crossing circle, i.e. the crossing circle is twisted. When a painted edge is signed 1 the strand which enters on the left goes on top, when the sign is -1 the strand which enters on the right goes on top. Which strand is on top remains the same no matter how you orient your crossing circle. These can be seen visually in Figure 3.

Definition 2.4. *A **Signed Painted Crushtacean** is a crushtacean in conjunction with a selected perfect matching and sign at each painted edge. As presented in prior paragraphs.*

Signed Painted Crushtaceans will be a frequent topic in our paper, and will often be shortened to SP-Crushtaceans.

We will say that two SP-Crushtaceans are isomorphic if there is an isomorphism which preserves painted edges and the signings on those edges.

Theorem 2.5. *An SP-Crushtacean uniquely determines an FAL.*

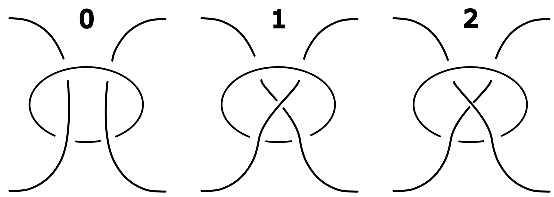

Figure 3: The appearance of different signs in $F(X)$ Change 2 to -1

Figure 4: Example of link generated from a crushtacean, as per Theorem 2.5. TODO ADD

Proof. We will assume that we have some SP-Crushtacean X. We will use only the intrinsic information of X to construct an FAL, with no ambiguity, showing the uniqueness of the link.

We do this by starting with a method similar to as described by Purcell [1]. We know that the edges of X each can be represented by a curve on the plane.

We will be constructing a link around X. We begin by constructing crossing circles around each painted edge. The disk bounded by our crossing circles will be perpendicular to the painted edge of X which determines it.

We will then contract all painted edges, by pulling the two ends of the edge together until they both represent just one point in space within the crossing disk. In doing this we have created a set of points within crossing disks and curves between each point. Each point necessarily has a corresponding signing inherited from the painted edge it came from.

At points signed with 0 we separate the curves so that they each go through the crossing circle but do not overlap. At points signed -1 we will separate the curves such that the curve on the left goes on top. At points signed 1 we will separate the curves such that the curve on the right goes on top. The edges of X now form components of a link which we will call knot circles.

It is important to see that given any crossing circle with a sign of 0 will have the strand entering on either side of the circle exit on the same side. A crossing circle with a sign of 1 or -1 will have each strand exit on the opposite side than it entered.

Once we've completed this process, we will have a completed link. Given we made no arbitrary choices, it is evident that this link is unique, and will be referred to as $F(X)$. An example of this process can be seen in Figure 5.

We now know that $F(X)$ is unique, but we still must show that $F(X)$ is a fully augmented link. In order to show this we must show that it follows the properties outlined by Purcell [1] in order to be a hyperbolic FAL. This means that we must show that $F(X)$ can be generated by fully augmenting a link diagram G which is prime, twist reduced, nonsplittable, and has at least two twist regions. We will do this by proving that G must fulfill each of Purcell's

Figure 5: An FAL constructed from a crushtacean ADD FIGURE

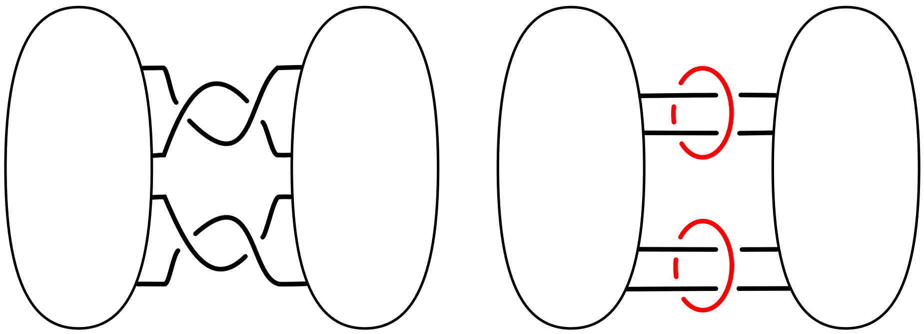

Figure 6: A non-twist-reduced generating link and the FAL it generates.

properties.

Construct G by doing any nonzero number of Dehn fillings at crossing circles, provided the resulting fillings ensure at least 2 crossings at each twist region determined by a crossing circle.

First we can show that G must be twist reduced, again by assuming to the contrary. By the definition of twist reduced, being non-twist-reduced would imply that either there is a reducible crossing or there exist two twist regions in which you can flype between them. If there is a reducible crossing that would mean there is a crossing circle separating two tangles of $F(X)$, which can only be generated by a painted 1-edge-cut. If there is two flypable crossings, then as we can see in Figure 6 that this would mean that there exist two crossing circles separating two regions of $F(X)$, which can only be generated by a painted 2-edge-cut on X. The existence of either a 1-edge-cut or a 2-edge-cut contradicts the fact that crushtaceans are 3-edge-connected.

Since G is twist reduced, we know each painted edge of X represents a distinct twist region in G. We know that the smallest possible crushtacean is K_4, since no smaller graph is 3-regular. This means that each painting must have at least two painted edges. This implies by our construction that there are at least two crossing circles in $F(X)$, and correspondingly at least two twist regions in G.

We know that for each painted edge of X, there is a twist region with at least two crossings. Since this twist region bounds the strands together, connectedness on X implies non-splittability on G. Since all Crushtaceans are connected, we know that G must be nonsplittable.

Finally we can show that the link must be prime by first assuming the contrary. We assume that G is a composite link. This implies two strands of G separate two tangles and since strands of G are made of edges of X, the crushtacean X would have a 2-edge-cut, proving contradiction.

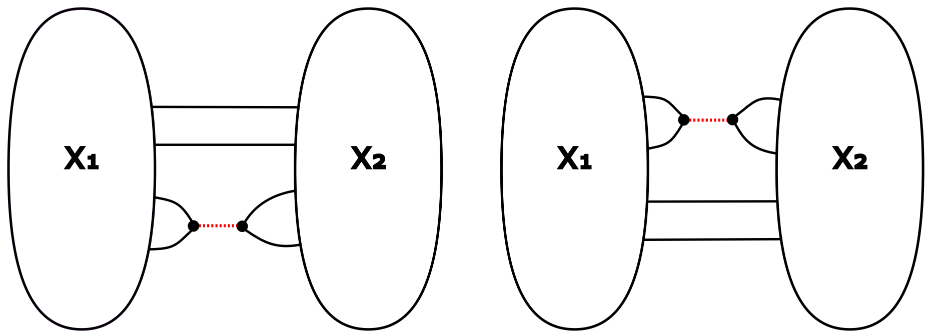

Figure 7: The generalized bridge structure required for flype modification.

Consequently, by proving that G must be nonsplittable, prime, twist reduced, and have at least two twist regions; we have fulfilled Purcell's theorem defining hyperbolic FALs, since fully augmenting G results in $F(X)$ we have shown that $F(X)$ is an FAL [1]. Thereby showing that $F(X)$ is an FAL.

Hence, an SP-Crushtacean generates a unique FAL $F(X)$. □

For the rest of our paper we will refer to the unique FAL generated by an SP-Crushtacean as $F(X)$ as we did in the proof.

We now know that for any SP-Crushtacean we can create an FAL. However we must still show that isomorphic SP-Crushtaceans always generate isotopic FALs. We can do this by first looking at the embeddings of the graph on a sphere. We are concerned primarily with embeddings on S^2 since our links are in S^3, meaning when we project our links into a diagram in S^2.

Recall that $X \cong Y$ means X and Y are isomorphic, and $A \sim B$ means A and B are isotopic.

Theorem 2.6. *Let X and Y be two SP-Crushtaceans. If $X \cong Y$ then $F(X) \sim F(Y)$.*

Proof. We show this by first defining an isotopy f of S^2 which maps X to Y, we know this exists because $X \cong Y$ implies that there exists an isotopy relating X and Y by Lemma 2.3. We can say that our FALs $F(X)$ and $F(Y)$ both exist in some neighborhood around S^2, we will refer to this neighborhood as $S^2 \times [-1, 1]$. The isotopy f can then be applied to $S^2 \times [-1, 1]$, which turns $F(X)$ into $F(Y)$. We can then extend our isotopy on $S^2 \times [-1, 1]$ to all of S^3, by making our isotopy the identity except for a regular neighborhood of $S^2 \times [-1, 1]$. Therefore showing that that if $X \cong Y$ then $F(X) \sim F(Y)$. □

We now have the tools necessary to create a Fully Augmented Link from a signed painted crushtacean.

While the tools we have are already powerful, we have no guarantee that two SP-Crushtaceans do not construct the same FAL. In fact, it is possible for two non-isomorphic graphs to construct the same FAL. However we are going to show that given a stronger equivalence relation on SP-Crushtaceans we can improve this.

In order to define this equivalence relation we will start by defining a **Flype Modification**. Flype modification is a way to take a SP-Crushtacean and create a new SP-Crushtacean. Our definition is specifically intended to mirror a flype on a link, and the precise relationship will be developed further throughout the rest of paper.

In order to flype modify an SP-Crushtacean we start by finding a non-trivial 3-edge-cut which has a single painted edge. We call a 3-edge-cut non-trivial, if not all the edges are incident with the same vertex.

We can see that the existence of this edge-cut implies the general structure shown on the left of Figure 7, since our definition of crushtacean requires a nontrivial 3-edge-cut to appear this way.

It is important to note that the any signing of the painted edge is allowed, though it changes the process. If the signing of the painted edge in the edge-cut is 0 then we change the crushtacean as is shown in figure 7. If the signing of the painted edge is 1 or -1 then we must also "flip" one of the subgraphs split by the cut. This is done by reversing the order of vertices, without loss of generality, this would be moving a to b and c to d, this is done while maintaining the structure and adjacency's of the subgraph, as in Figure 8. For an example of flype modification of this see Figure 8 or 9.

Definition 2.7. *A **Flype Modification** is a way to create a new SP-Crushtacean from an existing one. This is done by the process outlined in the above paragraphs.*

Lemma 2.8. *Let X be an SP-Crushtacean. If Y is the result of a flype modification of X then $F(X) \sim F(Y)$.*

Proof. We will need to prove this by multiple cases.

The first case we prove is where the edge which is flype modified is signed 0. In this case we know that X and Y have precisely the structure as seen in Figure 7.

We can then use the process outlined in Theorem 2.5 to create $F(X)$ and $F(Y)$. It can be trivially shown that $F(X)$ is isotopic to $F(Y)$, since we can take the crossing circle generated by the bridge structure in $F(X)$ and isotope it over the tangle and to its place in $F(Y)$. Consequently showing that for any arbitrary crushtacean, when the signing is 0, the result holds.

The second case is when the edge is signed either -1 or 1. This case is similar to before, with the addition of the "flip". The "flip" performed from X to Y is structurally equivalent to a flype from $F(X)$ to $F(Y)$, with the change in painted edge equivalent to the prior isotopy on the crossing circle. An example of this can be seen in Figure 9.

Consequently, we have proven both cases, showing that flype modification maintains isotopy of FALs. □

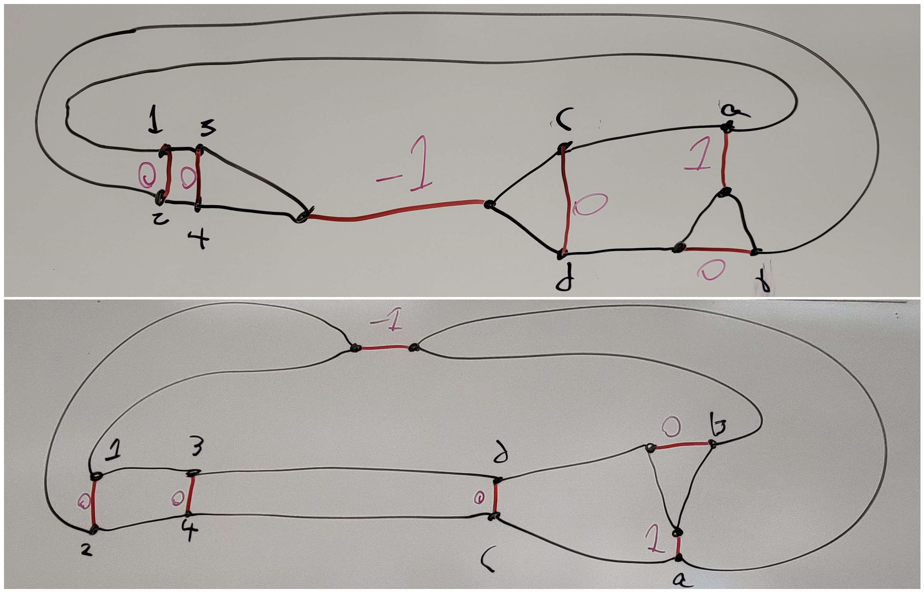

Figure 8: A SP-Crushtacean before flype modification, and the same crushtacean after flype modification. Red edges are painted edges, pink numbers are the signing for those edges, and purple labels label the name of each vertex. Rollie if you're reading this, the caption still sucks. I'll fix it once I draw it better.

Definition 2.9. *Two SP-Crushtaceans X and Y are considered* ***Flype Equivalent*** *if a (possibly empty) sequence of Flype Modifications of X generates an SP-Crushtacean isomorphic to Y.*

We will use the symbol $\overset{f}{\simeq}$ to represent flype equivalence.

Lemma 2.10. *Flype Equivalence is an equivalence relation.*

Proof. To prove flype equivalence is an equivalence relation, we will prove it is reflexive, symmetric, and transitive. Note that if we have two isomorphic SP-Crushtaceans X and Y then $X \overset{f}{\simeq} Y$, since we allowed our sequence of flype modifications to be empty.

Flype equivalence is reflexive since given a crushtacean X, we know $X \cong X$, and therefore $X \overset{f}{\simeq} X$.

Given two SP-Crushtaceans X and Y such that $X \overset{f}{\simeq} Y$. For every flype modification, there exists its inverse flype modification. We know there exists a sequence of flype modifications which creates Y from X, so there is a corresponding sequence which undoes that process. Therefore $Y \overset{f}{\simeq} X$, meaning flype equivalence is reflexive.

Given three SP-Crushtaceans X, Y, and Z such that $X \overset{f}{\simeq} Y$ and $Y \overset{f}{\simeq} Z$.

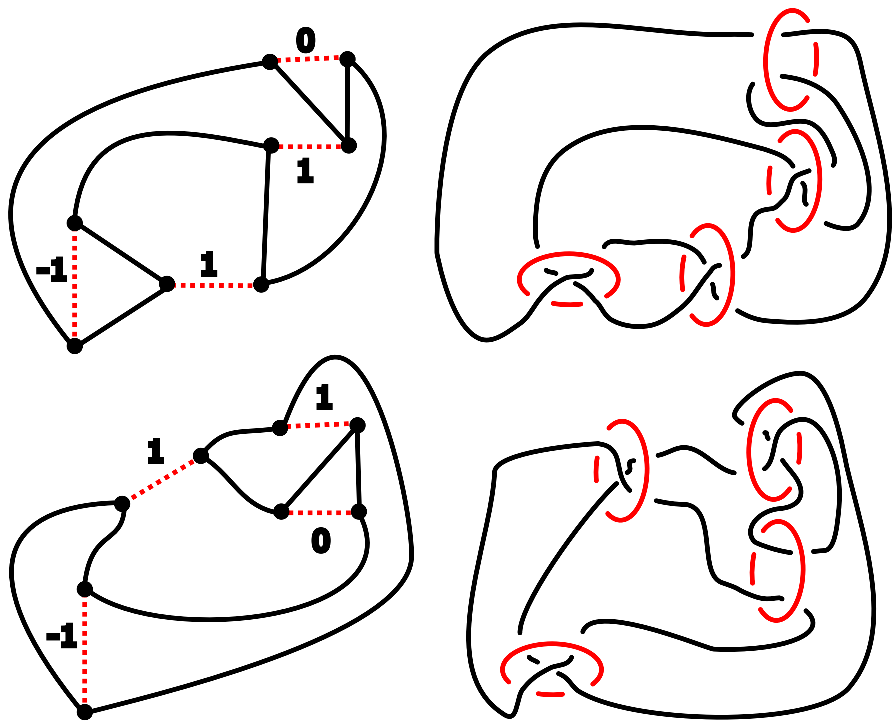

Figure 9: Visual example of two Crushtaceans related by a flype modification and the FALs corresponding. As seen in Lemma 2.8. Top FAL, left strand should be on top.

We know there must be a sequence of flype modifications that turns X into Y and a sequence that turns Y into Z, so if we do the first sequence then the second, we have a sequence that turns X into Z. Thereby showing $X \overset{f}{\simeq} Z$, and that flype equivalence is transitive.

Consequently showing that flype equivalence is an equivalence relation. □

Theorem 2.11. *Let X and Y be two SP-Crushtaceans. If $X \overset{f}{\simeq} Y$ then $F(X) \sim F(Y)$.*

Proof. Given any two arbitrary crushtaceans $X \overset{f}{\simeq} Y$, there exists a sequence of flype modifications which creates Y from X. We will refer to this sequence of crushtaceans created by flype modifications as $X_1, X_2, ...X_{n-1}$ where n is the number of modifications needed to show equivalence. We complete this sequence by saying that it starts with X and ends with Y, such that the sequence becomes $X, X_1, X_2, ...X_{n-1}, Y$. We can then create a sequence $F(X), F(X_1), F(X_2), ...F(X_{n-1}), F(Y)$.

It is easy to see that any two neighbors in this sequence are isotopic by Lemma 2.8, and by the transitive property of isotopy we can see that all members of this sequence are isotopic. Therefore $F(X) \sim F(Y)$. Thus any FAL generated by flype equivalent crushtaceans is isotopic, and it is unique by Theorem 2.5. □

3 Determining Crushtaceans under FALs

We've already seen that flype equivalence on an SP-Crushtacean determines isotopy on an FAL. In order to make this relationship stronger we must show that the opposite direction expresses similar properties, while some of these proofs may seem trivial based on prior work they are not. The definition of crushtacean we are using is slightly different than that of previous work, and therefore these concepts must be re-proven.

When we talk about FALs we must have some notion of what a crossing circle is and what a knot circle is. This is why we will be looking at what we call Painted FALs or PFALs for short. A PFAL is an FAL whose crossing circles are specified. Any component which is not a crossing circle is called a knot circle. This definition is important because it gives us the tools required to reconstruct a link which generated the PFAL.

All crossings in a PFAL happen within a crossing circle. No two crossing circles ever cross each other. These crossing circles are the same as those created by the process of fully augmented, and can be used to reconstruct generating links.

Definition 3.1. *A **Painted FAL** is an FAL in conjunction with a decision about which components are crossing circles.*

We say that A_p and B_p are equivalent PFALs if and only if there is an isotopy from A to B preserving chosen crossing circles. We will denote PFAL equivalence as $A_p \overset{p}{\sim} B_p$. Note that Purcell's construction yields a PFAL since we know which components are crossing circles.

Conjecture 3.2. *Given any FAL A. If A_{p1} and A_{p2} are any two paintings of A then $A_{p1} \overset{p}{\sim} A_{p2}$.*

SOMEHOW INCLUDE EXAMPLE FIGURE

Conjecture 3.2 is important, since if it is true then many of our theorems can be trivially extended from statements about PFALs, to statements about any FALs. We will refer to Conjecture 3.2 as the One Painting Conjecture in this paper. More details addressing the truth of Conjecture 3.2 can be found in Section 5.

Theorem 3.3. *A PFAL uniquely determines a flype equivalent class of crushtaceans.*

Proof. We determine the crushtacean by a construction precisely inverted from Theorem 2.5.

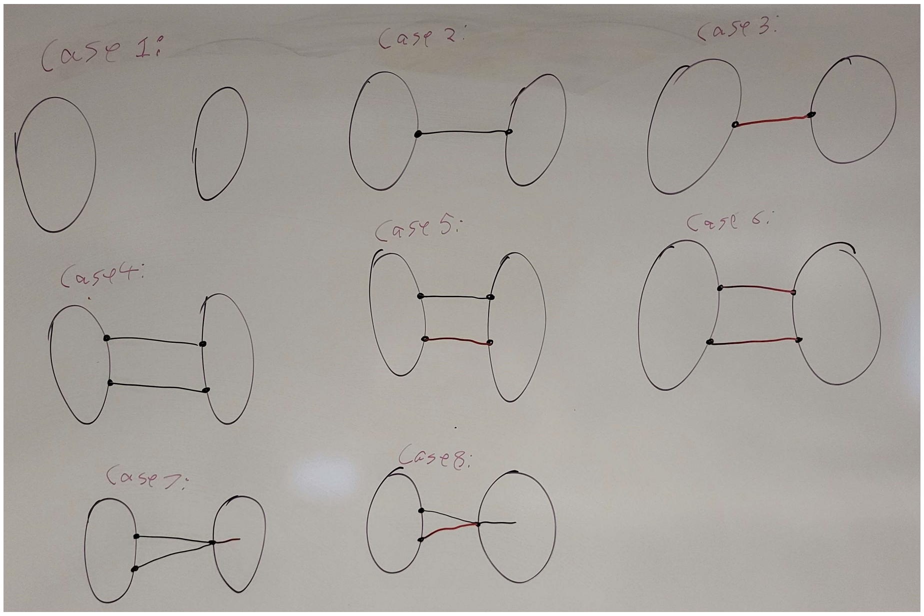

Figure 10: The cases for proof of Theorem 3.3

In that light, we begin by looking at the strands bounded by crossing circles. We will for this construction view all components of the PFAL as curves. We begin by bringing together the curves which pass through any crossing circle so that they meet at a point within the crossing circle. Mark this point with a signing corresponding to which way the strands were positioned as we did earlier.

Extend these points into painted edges perpendicular to the crossing circle, before removing the crossing circles. Making sure there is a vertex at the end of each painted edge. Thereby creating a graph with paintings and signings.

Notice that while a graph $F^{-1}(A)$ exists we do not know that it is a crushtacean. We will show that $F^{-1}(A)$ fulfills Definition 2.1 because it fulfills each necessary property.

It is easy to see that $F^{-1}(A)$ is 3-regular, since each vertex is incident with 1 painted edge, and two edges created by knot curves. It is trivial to see that $F^{-1}(A)$ is planar.

A much more difficult property to show is that $F^{-1}(A)$ is 3-edge-connected. We do this by contradiction.

However, even assuming to the contrary is difficult and this proof will have to be done in several cases. A visual of all the cases is given in figure 10.

The first case is a very simple one where the Crushtacean is 0-edge-connected.

case 1: $F^{-1}(A)$ is 0-edge-connected. This is equivalent to $F^{-1}(A)$ is disconnected, and it can be seen that this implies that A is splittable. This creates contradiction because A is hyperbolic, which implies it is nonsplittable.

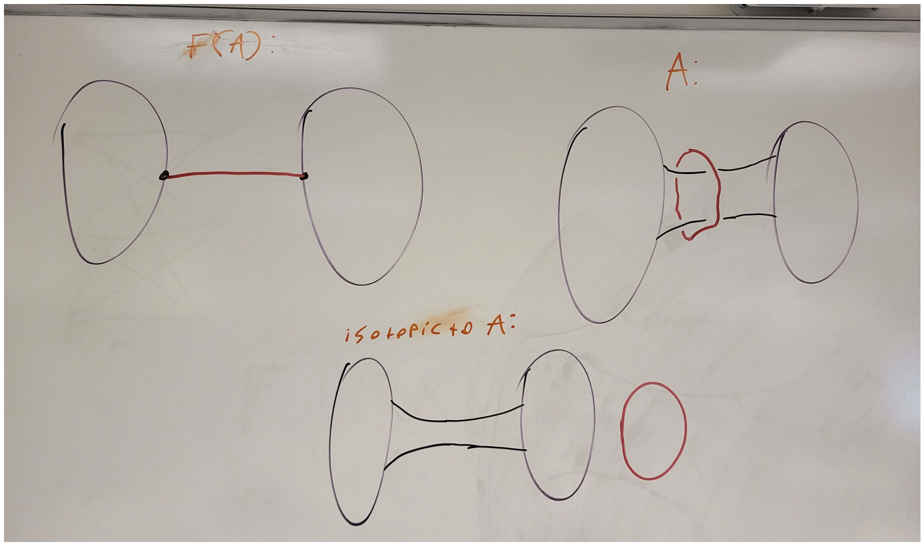

Figure 11: Case 3 of Theorem 3.3 and the associated splittable link A

In cases 2 and 3 there is a 1-edge-cut.

case 2: $F^{-1}(A)$ has an unpainted 1-edge-cut. Given the unpainted edge represents a part of a knot circle, it is part of a closed loop in the plane. This closed loop intersects the boundary of the tangle(which must be another closed loop) in 1 point on the plane. We know that every closed loop must intersect another loop on the plane in an even number of points, thus proving contradiction.

case 3: $F^{-1}(A)$ has a painted 1-edge-cut. We know that A must be non-splittable, since it is hyperbolic. However the 1 edge-cut $F^{-1}(A)$, would imply A has a crossing circle which may be split from the diagram. This can be seen in Figure 11.

For cases 4, 5, and 6, we look at cases where a 2-edge-cut exists such that the 2 edges are not incident with each other.

case 4: $F^{-1}(A)$ has a completely unpainted 2-edge-cut, in which both edges are not incident. We can see that $F^{-1}(A)$ could only be generated by a composite link since there is a line in A which only intersects the link in 2 places. We know A must be prime, since A is hyperbolic, thereby proving A is both composite and prime.

case 5: $F^{-1}(A)$ has a 2-edge-cut in which one edge is painted, and both edges are not incident. We can see that this shows that a 3-edge-cut exists where no edges are painted. Notice that this is similar to case 2, where a closed loop in the plane intersects the boundary of the tangle(which must be another closed loop) in 3 points, when the number of intersections should be even.

case 6: $F^{-1}(A)$ has a 2-edge-cut in which both edges are painted, and both edges are not incident. The two crossing circles corresponding to the edges of the 2-cut must be generated from two twist regions which can be flyped between,

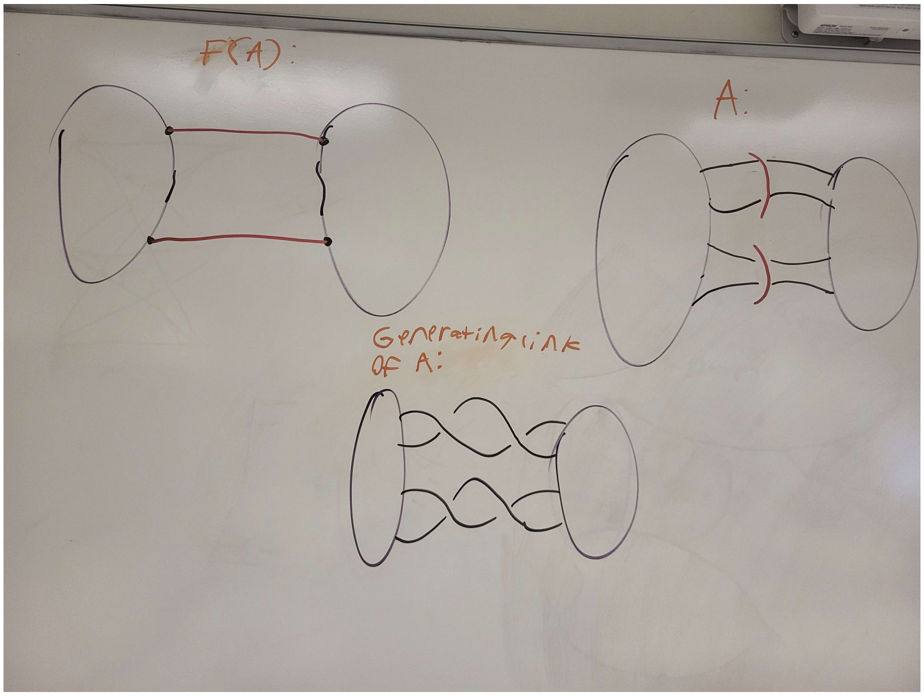

Figure 12: Case 6 of Theorem 3.3 and the associated non-twist-reduced generating link

meaning that the link that generates A is not twist reduced, this contradicts the assumption that it must be twist reduced. (This can be seen in Figure 12).

Finally in cases 7 and 8, we look at cases where there is a 2-edge-cut and the edges are incident with one another.

case 7: $F^{-1}(A)$ has a 2-edge-cut in which one edge is painted, and the edges are incident. As can be seen in Figure 10 this means that the single vertex must be connected to a 1-edge-cut which is not painted, reducing to case 2.

case 8: $F^{-1}(A)$ has a 2-edge-cut in which no edges are painted, and the edges are incident. As can be seen in Figure 10 this means that the single vertex must be connected to a 1-edge-cut which is painted, reducing to case 3.

Our method of generating a graph can never create a graph with a 2-edge-cut in which both edges are incident to 1 vertex and they are painted, therefore this case does not exist.

Thus, we have proven that the graph must be 3-edge-connected. In combination with the other properties, this graph must be a crushtacean, according to definition 2.1.

The painting of the crushtacean is determined by which edges were constructed from crossing circles. The signing is determined by the crossings at those same crossing circles. □

Figure 13: An augmented flype on an FAL. TODO CREATE EXAMPLE

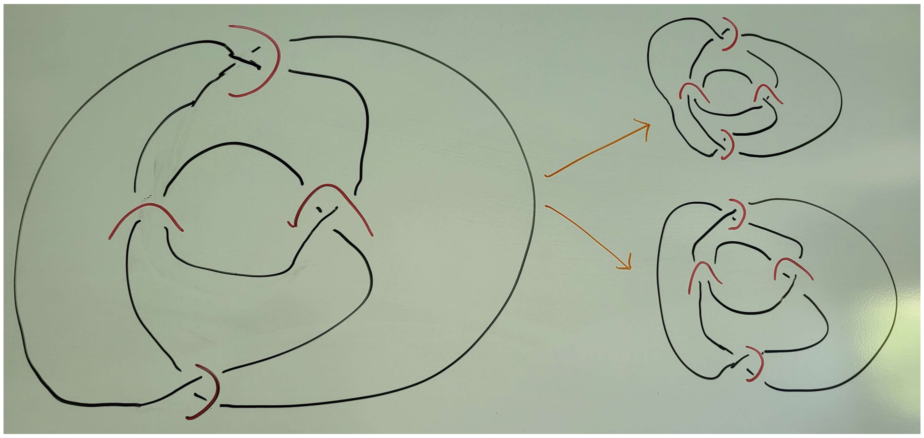

Figure 14: I will change it to an FAL and two alternating generating links.

Now that we know we can construct SP-Crushtaceans from PFALs, it's important to show that equivalent PFALs always construct flype equivalent Crushtaceans. We begin by constructing some of the tools for our proof.

The most important tool is the notion of an augmented flype. An augmented flype is a move on an FAL which is related to a flype on a link diagram. The move is done with respect to a crossing circle. If the crossing circle is flat then we shift the circle along and over a tangle and onto another valid location. If the crossing circle is twisted, we do the same thing in conjunction with a flype to move the crossing to the same place as the circle. An example of this can be seen in Figure 13.

Definition 3.4. *An **augmented flype** is a move on an FAL as described in the above paragraph.*

We seek to show that if we have two equivalent PFAL diagrams A and B, that they are related by a sequence of augmented flypes.

Theorem 3.5. *Given a PFAL A we can construct two alternating links A_{g1} and A_{g2} such that (a) A is a full augmentation of both A_{g1} and A_{g2} (b) A_{g1} and A_{g2} have either 2 or 3 crossings in each twist region (c) A_{g1} and A_{g2} are twist reduced.*

Proof. We can prove this by construction. We will use Dehn fillings on A in order to recreate the generating links, this is as presented by Purcell in [1]. We can choose the direction of twist for the first crossing circle arbitrarily, one of these rotations will be used to generate A_{g1} and the other to generate A_{g2}.

We do the rest of the Dehn fillings taking care to retain an alternating diagram at each stage.

Figure 15: TODO general flype picture

The links A_{g1} and A_{g2} are both twist reduced since they are valid generating links of A. ☐

Theorem 3.6. *Given two PFAL diagrams A and B, such that $A \overset{p}{\sim} B$. There exists an isotopy $f : A \to B$, which is a composite of augmented flypes.*

Proof. We can see this by creating the alternating generating links for each PFAL A and B, which we will call A_{g1}, A_{g2}, B_{g1}, and B_{g2}. It must be that either $A_{g1} \sim B_{g1}$ or $A_{g1} \sim B_{g2}$IS THIS TRUE?, without loss of generality we will say $A_{g1} \sim B_{g1}$ since the labels are arbitrary. There is much question about the validity of this proof, I'll finish it later

By Tait's flyping theorem, there exists an isotopy $g : A_{g1} \to B_{g1}$ which is a composite of flypes. We can then construct an isotopy $f : A \to B$, by performing analogous augmented flypes to the ones in g.

Thereby guaranteeing that any isotopy from A to B is a composite of augmented flypes. ☐

Lemma 3.7. *Given a PFAL A. If B is constructed by a single augmented flype on A then $F^{-1}(B)$ can be constructed by a flype modification on $F^{-1}(A)$.*

Proof. An augmented flype moves a crossing circle (and possibly a crossing) from one pair of separating strands to another pair of separating strands. The start and end of the augmented flype when made into a crushtacean become the beginning and end of the definition of flype modification on a crushtacean. The general structure of this can be seen in Figure 15. Examples for this concept are given in Figures 16 and 9. ☐

Theorem 3.8. *Given two PFALs A and B. If $A \sim B$ then $F^{-1}(A) \overset{f}{\simeq} F^{-1}(B)$.*

Proof. Flype equivalence is defined as repeated flype modification. Any two isotopic PFAL diagrams are related by a sequence of augmented flypes by Lemma 3.6. Lemma 3.7 implies their crushtaceans must be related by a sequence of flype modifications. Therefore any two crushtaceans of isotopic PFALs are flype equivalent crushtaceans. ☐

4 Putting it Together

We can now see that a SP-Crushtacean under flype equivalence and Fully Augmented links have a close relationship which implies many things.

Theorem 4.1. *Given an FAL B, then $B \sim F(F^{-1}(B))$.*

Given a crushtacean X, then $X \cong F^{-1}(F(X))$.

Put simply, F and F^{-1} are inverse operations, as their names imply.

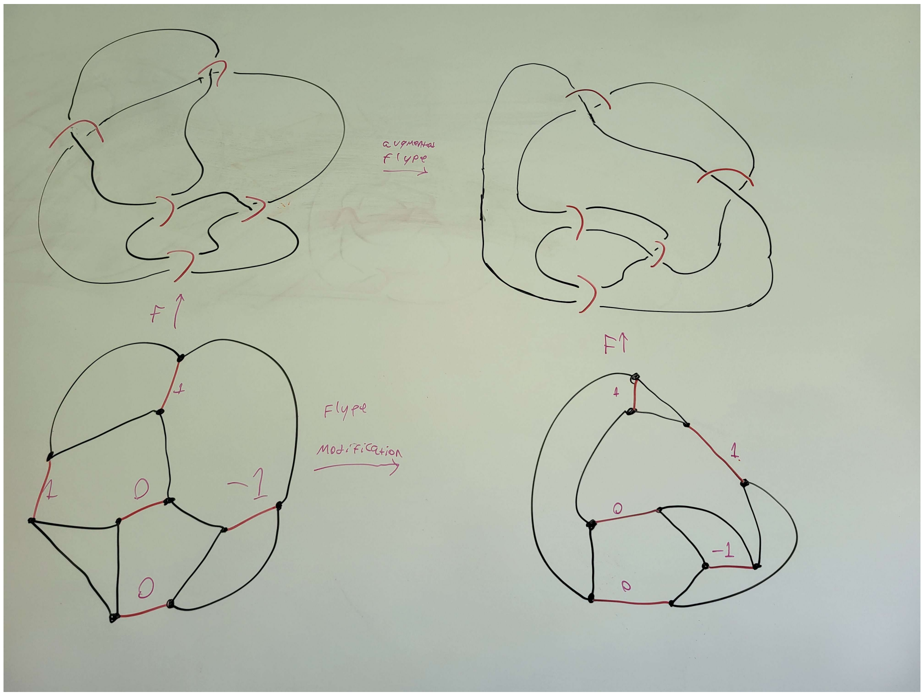

Figure 16: An example of Lemma 3.7.

Proof. trivial, since the constructions in Theorem 3.3 and 2.5 are opposite one another, and each construction is invariant. □

Theorem 4.2. *If $X \overset{f}{\simeq} Y$ then $F(X) \sim F(Y)$. If $A \overset{p}{\sim} B$ then $F^{-1}(A) \overset{f}{\simeq} F^{-1}(B)$.*

In other words, flype equivalence and isotopy are precisely related notions of equivalence.

Proof. This is simply a combination of theorem 3.8 and theorem 2.11. □

5 On the One Painting Conjecture

While as a whole we are not able to prove Conjecture 3.2, we are able to prove it in some cases. The most general condition under which Conjecture 3.2 is true, is when the FAL has only a single reflection plane.

Theorem 5.1. *Given any FAL A, which has only a single reflection plane. If A_{p1} and A_{p2} are any two paintings of A then $A_{p1} \sim A_{p2}$. If A has single ref plane... Yeah, that's in proof? unique choice of crossing circles.*

Proof. This comes from the idea that we can create a single valid painting based on the information from the reflection plane.

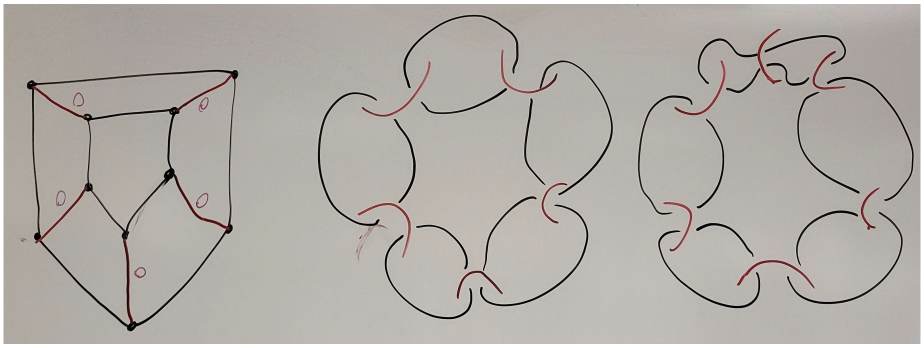

Figure 17: left: 5-prism graph with spokes painted, center: 5-chain link, right: 5-Chain link with extra crossing circle

meridional

Crossing circles are strands whose tubular neighborhoods intersect the reflection plane in two meridional curves. Knot circles are strands whose tubular neighborhoods intersect the plane in generalized longitudinal curves.

Therefore, when there is only a single reflection plane, there is only one valid painting. Therefore two paintings of A are isotopic trivially. □

There are only 3 classes of flat links which have more than one reflection plane [5]. The n-chain links, which is the link constructed from an n-prism graph where the "spokes" are painted with a signing of 0 as in Figure 17. The n-chain links with an extra perpendicular crossing circle as pictured on the right in Figure 17. The Borromean rings. It's important to note that the first two are infinite classes, where as the Borromean rings are only one link. We can now prove these three cases individually

Theorem 5.2. *Given a flat FAL A. If A_{p1} and A_{p2} are any two paintings of A then $A_{p1} \sim A_{p2}$*

Proof. This can be proven by cases. These cases are when there is a single reflection plane or one of the three cases highlighted earlier which have more than one projection plane.

The case in which there is one reflection plane is trivially completed by citing Theorem 5.1.

When considering the unpainted n-chain links, there are 2 ways to paint the link. This is because we must declare every other component a crossing circle, and the choice of which alternating sequence is arbitrary. These two paintings are equivalent since you can isotope from one to the other by a simple rotation of the link.

When considering the n-chain links with an extra perpendicular crossing circle, there are 2 ways to paint the link. The argument for why is identical to the n-chain-link, since the perpendicular crossing circle does not change it.

Finally, the borromean rings have 3 ways they can be painted. This is because we must declare that 2 of the circles are crossing circles, and there are 3 different choices, which all create valid PFALs. All of these paintings are equivalent because there exists an isotopy which brings any component □

Further proof of more cases may be possible by reducing the problem to that of Theorem 5.1 and simple cases, the way Theorem 5.2 does. It is important to attempt further proof of the one painting conjecture, since it makes the theorems in this paper more general and applicable.

6 Possible Applications

There are many possible uses to this relationship between flype equivalence and isotopy. We will mention a few possibilities here.

1. It is possible to create an algorithm to generate SP-Crushtaceans, and to test for flype equivalence. Such an algorithm could then be used to enumerate all FALs. It may be also possible to generate specifically some notion of a "prime" FAL such as belted-sum-prime, given a detection algorithm for them on the crushtacean. This is all much easier with this tool given it is much easier to algorithmically reference graphs compared to links.

2. We could determine which of the crushtaceans share a complement. If we take the absolute value of the signs of a crushtacean we can create a crushtacean which we will refer to as the absolute signing of a crushtacean. It is clear that if two crushtaceans have flypes equivalent absolute signings, then they share a complement. This is because the absolute signing of a crushtacean may be made by Dehn twists on the crushtacean, which maintains complements. Based on this we propose the following conjecture:

   Conjecture 6.1. *Given two FALs A and B, where $S^3 \backslash A$ refers to the complement of A. $S^3 \backslash A$ is homeomorphic to $S^3 \backslash B$ iff the absolute signing of $F^{-1}(A)$ is flype equivalent to the absolute signing of $F^{-1}(B)$.*

3. We could also use this to seek finding some useful classes of FALs and determining facts about them. We can do this by looking at various classes and types of crushtaceans.

4. This may have close relationships with symmetry groups of FALs. It could be shown easily that automorphisms of the crushtaceans are symmetries of an FAL. However it may also be possible to define symmetries of the flype modifications to extend this further.

5. If we find a way to create a Crushtacean from a smaller crushtacean, which is often possible with graph classes, then we could use such a construction to make an inductive argument on FALs.

6. Finally, it may be possible to generalize this result further to be able to construct a relationship between any prime nonsplittable link and crush-taceans by reconstruction of the generating link of the FAL.

References

[1] J. S. Purcell, "An introduction to fully augmented links," 2010.

[2] J. Chesebro Eric; Deblois and H. Wilton, "Some virtually special hyperbolic 3-manifold groups," 2010.

[3] H. Whitney, "Congruent graphs and the connectivity of graphs.," 1932.

[4] H. Whitney, "2-isomorphic graphs.," 1932.

[5] C. Millichap and R. Trapp, "Flat fully augmented links are determined by their complements," 2023.