Sewing Ribbons on Graphs in Space

Dan Archdeacon
Dept. of Math. and Stat.
University of Vermont
Burlington, VT, USA 05405
dan.archdeacon@uvm.edu

C. Paul Bonnington
Department of Mathematics
University of Auckland
Auckland, New Zealand
p.bonnington@auckland.ac.nz

R. Bruce Richter
Dept. of Combin. and Optim.
University of Waterloo
Waterloo, Ont., Canada
brichter@math.uwaterloo.ca

Jozef Širáň
Department of Mathematics
Slovak Technical University
813 68 Bratislava, Slovakia
siran@lux.svf.stuba.sk

Abstract

An *open ribbon* is a square with one side called the *seam*. A *closed ribbon* is a cylinder with one boundary component called the *seam*. We *sew* an open (resp. closed) ribbon onto a graph by identifying the seam with an open (resp. closed) walk in the graph. A *ribbon complex* is a graph with a finite number of ribbons sewn on.

We investigate when a ribbon complex embeds in 3-dimensional Euclidean space. We give several characterizations of such *spatial* complexes which lead to algorithms. We examine special cases where: 1) each edge of the graph is incident with at most three ribbons, and 2) every ribbon is closed together with a connectivity condition.

1 Introduction

The purpose of this paper is to examine when certain topological spaces called ribbon complexes are spatial, that is, when they embed in 3-dimensional Euclidean space. The problem arises from examining when an arbitrary 2-complex embeds in 3-dimensional space. Despite the topological motivation, this is a paper in discrete mathematics. One of our first results is to translate the topological problem into the existence of certain planar embeddings of a related graph. This in turn leads to several characterizations of spatial ribbon complexes. As a result, we develop algorithms to answer this question. Unfortunately, the worst-case running time of these algorithms is not bounded by a polynomial in the size of the ribbon complex. We also explore some related topics.

We form a ribbon complex by sewing ribbons onto graphs. Specifically, we identify one long side of a long, thin rectangle with an open walk in the graph. Likewise, we identify one boundary component of a cylinder with a closed walk in the graph. Sewing several ribbons on a graph forms a ribbon complex.

The principal object of our investigation is

The Ribbon Problem: *When is a ribbon complex spatial?*

This problem was posed at the *VII Vermont Summer Workshop* by Neil Robertson. He actually stated the *Closed Ribbon Problem*: the special case when each ribbon is closed. He mentioned that the problem was related to a similar one where every walk was sewn onto the boundary of a disk instead of one boundary of a cylinder. The related problem was to determine when a given 2-complex was spatial, and was motivated by the Poincare Conjecture. It is hoped that understanding spatial ribbon complexes would help in understanding spatial 2-complexes. The closed ribbons arise not by sewing a disk on a closed walk, but rather by sewing a disk with a hole in the middle. In this paper we do not focus on the arbitrary spatial complexes, but only on the ribbon problem.

It is easy to characterize when a ribbon complex embeds in 2-dimensional space. On the other hand, every ribbon complex embeds in 4-dimensional space. Which ribbon complexes embed in 3-dimensional space is the interesting problem. Thomas Tucker (personal communication) has independently and concurrently been examining the question when a given 2-complex embeds in a 3-mainfold. In particular, his work (which has not yet been prepared for publication) roughly overlaps with our material in Section 4.

The paper is organized as follows. In Section 2 we give a formal definition of ribbon complexes and show the Ribbon Problem is equivalent to the Closed Ribbon Problem. Section 3 introduces the Labeled Planarity Problem. In Section 4 we explore the relationship between this problem and the Ribbon Problem. This offers us our first characterization of spatial ribbon complexes. Since this characterization is purely combinatorial, we get an algorithm for determining when a complex is spatial. In Section 5 we introduce a partial ordering on ribbon complexes and a corresponding partial ordering on the related graphs. We use this to define minimal non-spatial ribbon complexes and give some examples. In Section 6 we give a second related signed graph and a second characterization of spatial ribbons. Section 7 explores the planarity of labeled graphs with small cut-sets. Section 8 covers the special case where each edge of the graph in a complex is incident with at most three ribbons. In Section 9 we give a polynomial-time algorithm for the planarity of some labeled graphs with an added connectivity condition. This corresponds to closed ribbons complexes. Finally, Section 10 offers some concluding remarks and directions for future research.

2 Graphs and Ribbon Complexes

We assume that the reader is familiar with the basics of graph theory. We do use the following special conventions. Our graphs are geometric objects with points for vertices and homeomorphs of line segments for edges. We allow multiple edges and loops. We also allow *one-ended edges* with only a single vertex incidence. A *half-edge* is either a one-ended edge, or one of the two ends of a two-ended edge. Two half-edges are *adjacent* if they are incident with a common vertex. They are *transverse* if they are the two ends of an edge. Define the *transverse involution* τ on half-edges by fixing a half-edge in a one-ended edge and interchanging the half-edges in a two-ended edge.

A *walk* in a graph is a sequence of half-edges such that successive half-edges alternate between being transverse and adjacent. In the first case the walk *crosses an edge.* In the second case the walk *passes through* the incident vertex (if the half-edges are different) or *backtracks* (if the half-edges are the same). The definition does not allow walks to backtrack in the middle of an edge; however, if that is desired the edge can be subdivided and the walk can backtrack at the new vertex. Our walks can start and stop at either a vertex of the graph, in the middle of an edge, or on a one-ended edge. We do not distinguish between a walk and the same walk in the reverse direction. A *closed* walk is cyclic sequence of half-edges such that each proper subsequence is a walk. We do not distinguish between the two directions on a closed walk, nor do we distinguish a starting or stopping point.

Our ribbon complexes will also have a 2-dimensional structure. Define a *closed ribbon* as a topological space homeomorphic to the compact cylinder. A closed ribbon has two boundary components, each homeomorphic to a circle. One of these is distinguished as the *seam.* An *open ribbon* is a topological space homeomorphic to a compact square. One side of this square is distinguished as the *seam.*

Let W be a closed walk in G and let R be a closed ribbon. Let S be the seam of R and let S^1 denote the unit circle. Let $f : S^1 \to W$ be a continuous function which traverses W exactly once, such that f traverses each edge of W as a homeomorphism, and let $g : S^1 \to S$ be a homeomorphism. The identification $f(x) = g(x)$, $x \in S^1$, *sews* the closed ribbon R onto the closed walk W. Similarly, we *sew an open ribbon* onto a graph by identifying (as above) the seam with an open walk. We say a ribbon crosses an edge or passes through a vertex when the corresponding walk does. A *ribbon complex* R consists of a graph with a finite number of open and closed ribbons sewn on. It is helpful to picture a closed ribbon as a cylinder with large radius and small height. Similarly, it is helpful to picture an open ribbon as a long, thin rectangle with the seam along the long side. These pictures motivate the name "ribbon". We call the underlying graph the *spine* of the ribbon complex.

A ribbon complex is *spatial* if it is homeomorphic to a subspace of 3-dimensional Euclidean space. For example, Figure 1 shows a spatial ribbon complex with a single closed ribbon wrapped four times around a single loop in the same direction.

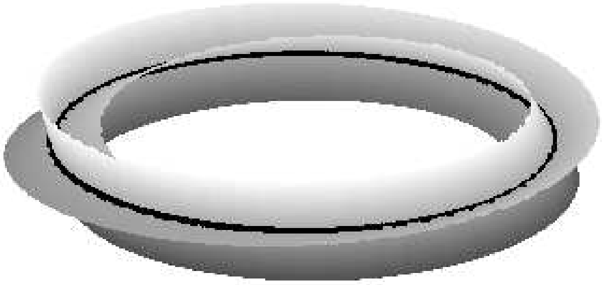

Figure 1: Wrapping four times around a loop

While some ribbon complexes are spatial, others are not. For example, let R_1 be a ribbon complex where a single closed ribbon has been wrapped twice around a loop. Then R_1 is homeomorphic to a Möbius band and is spatial. Now let R_2 be the ribbon complex created by adding a second open ribbon to R_1 which goes around the loop only once. In any spatial embedding of R_2 the two ends of this second ribbon will lie on different sides in a small neighborbood of the Möbius band. Now form R_3 by replacing the open ribbon in R_2 by a closed ribbon. It follows that R_3 is non-spatial. Figure 2 illustrates the difficulty of closing off the second ribbon.

As mentioned in the Introduction, in Section 3 we introduce labeled graphs and in Section 4 we show that a planar embedding of a labeled graph corresponds to a spatial embedding of a ribbon complex. This correspondence is not one-to-one, rather, many different ribbon complexes give the same labeled graph. However, all such complexes arise from a canonical representative by identifying vertices in the spine. We next describe this representative.

Let R be a ribbon complex on the spine H. Let $h_1, \ldots, h_m$ be the half-edges incident with a vertex v. Write $h_i \sim h_j$ whenever there is a ribbon passing through v on these half-edges. Extend $\sim$ to the *local equivalence relation* by transitive closure. The local equivalence class containing h_i is denoted $[h_i]$.

We form a new ribbon complex R' by *cleaving* the vertex v: replacing v with new vertices

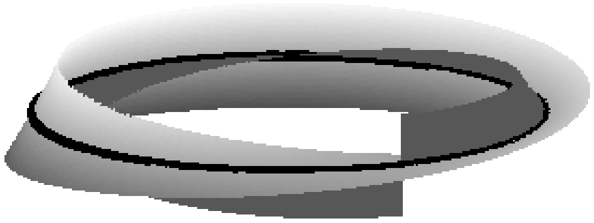

Figure 2: The second ribbon cannot be closed

corresponding to the local equivalence classes and replacing an incidence between h_i and v by an incidence between h_i and $[h_i]$. The *cleaved ribbon graph* R' corresponding to R is formed by cleaving every vertex of R. A ribbon graph is *cleaved* if it is isomorphic to its cleaved ribbon graph.

If we try to cleave a half-edge h off of a vertex v in a ribbon complex, we must also cleave off all half-edges which are paired with h as a ribbon passes through v. The cleaved ribbon complex has these forced pairings and no others.

It is not obvious that a ribbon graph R is spatial if and only if its cleaved ribbon graph R' is spatial. One direction is not too difficult: we can take a spatial embedding of R', join two vertices by a line segment in 3-space disjoint from R' except at the endpoints, and then in space contract this line segment to a point. This gives a spatial embedding of a quotient ribbon, and upon repetition, a spatial embedding of R. The other direction is less obvious, but nonetheless true (see the discussion following Theorem 4.3).

A slightly different point of view would be to extend the relation between a ribbon complex and its cleaved ribbon complex to an equivalence relation by transitive closure. In this sense we could select a representative R'' for each equivalence class which is on a spine with single vertex. We will exploit this selection later in the paper.

Recall that the *Closed Ribbon Problem* is to determine if a complex with each ribbon closed embeds in 3-space. At a first glance it may appear that allowing open ribbons gives a much more general problem. However, the following Proposition shows that the two problems are equivalent.

Proposition 2.1 *For every ribbon complex R there exists a closed ribbon complex R' such that R is spatial if and only if R' is spatial.*

Proof: We begin by subdividing every edge in the spine of R so that all ribbons begin and end at vertices of R. This clearly does not affect whether R is spatial. Next, suppose that r is a ribbon in R sewn on the open walk W. Replace r with a ribbon r' sewn on the closed walk WW^{-1}. Doing this for each open ribbon creates the desired R'.

If R' is spatial, then so is the subspace R. Conversely, if R is spatial, then we can close off the ribbon r by backing-tracking on W^{-1} along-side r. It follows that R' is spatial. ■

3 Planar Labeled Graphs

The goal of this section is to introduce a problem about certain embeddings of planar graphs. In Section 4 we will show that this problem is equivalent to the Ribbon Problem.

Let G be a graph and L be a subset of its half-edges called *labeled half-edges*. Let λ be a fixed-point-free involution on L such that if x is adjacent to y, then $\lambda(x)$ is adjacent to $\lambda(y)$. Then $G_\lambda = (G, L, \lambda)$ is called a *labeled graph*. Sometimes we will distinguish between the labeled graph G_λ and the underlying graph G, but frequently we will refer to the labeled graph G when λ is clear from context. In figures we will denote the labeled edges and the involution λ by placing a label a on one half-edge x and a second label $\bar{a}$ on $\lambda(x)$. Likewise, if v is a vertex with a labeled half-edge a, then $\bar{v}$ denotes the vertex incident with $\bar{a}$. We say that G is *totally labeled* if every half-edge has a label.

Consider a planar embedding of the graph G underlying G_λ. Suppose that $a_1, a_2, \ldots, a_k$ are the labeled half-edges incident with a vertex v and the embedding places them in that cyclic order anticlockwise around v. The involution *reverses the order of the labels* if $\bar{a}_1, \bar{a}_2, \ldots, \bar{a}_k$ appear in this cyclic order clockwise around $\bar{v}$, that is, if they appear in the reverse order with the original anticlockwise orientation. If $k \leq 2$, then the embedding always reverses the order of the labels. We

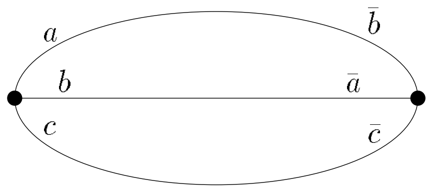

Figure 3: A small non-planar labeled graph

say that a labeled graph G_λ is *planar* if there exists a planar embedding of the underlying G which reverses the order of the labels at every vertex. Note that it might be that G_λ is non-planar as a labeled graph, even though the underlying G is planar. This is illustrated in Figure 3.

We investigate:

The Labeled Planarity Problem: *When is a labeled graph planar?*

There is a well-known algorithm for determining if an unlabeled graph is planar [7]. This algorithm runs in time bounded by a linear function in the size of the graph (see [3] or [4] for a discussion of the analysis of algorithms). However, an unlabeled graph can have many different planar embeddings. For example, consider the n-bond which has a pair of vertices and n multiple edges joining them (also called a dipole). This graph has $(n-1)!$ different planar embeddings. Checking them all to see which reverse the order of the labels would take exponential running time.

Embeddings of graphs in oriented surfaces can be described combinatorially. A *local rotation* is a cyclic permutation of the half-edges incident with a given vertex. A *rotation* is a collection of local rotations, one for each vertex. Rotations on a graph are in bijective correspondence with embeddings of that graph into oriented surfaces. For details, we refer the reader to [5] or to [2].

Proposition 3.1 *There is an (exponential-time) algorithm to determine if a labeled graph is planar.*

Proof: We consider the (exponentially many) possible rotations on the underlying graph. For each one it is easy to check if it corresponds to an embedding of the labeled graph. It is also easy to count the number of faces in that embedding and hence determine if it is planar. ∎

labeled graph	$\leftrightarrow$	cleaved ribbon complex
component	$\leftrightarrow$	vertex
vertex	$\leftrightarrow$	half-edge of the spine
half-edge	$\leftrightarrow$	occurence of half-edge in a ribbon
edge	$\leftrightarrow$	ribbon passing through a vertex
label	$\leftrightarrow$	ribbon crossing a spine edge
planar	$\leftrightarrow$	spatial

Figure 4: Relating a labeled graph and a cleaved ribbon complex

4 Relating the Two Problems

In this section we show the equivalence between the Ribbon Problem and the Labeled Planarity Problem. Specifically, we will draw a correspondence between cleaved ribbon complexes and labeled graphs. The cleaved ribbon complex will be spatial if and only if the labeled graph is planar. The relationship is summarized in Figure 4.

Let R be a ribbon complex. We describe the corresponding labeled graph G. The vertex set of G is the set of all half-edges of the spine H. For every occurence of a half-edge of H in a ribbon we introduce a half-edge in G. If two half-edges of H occur consecutively as the ribbon passes through a vertex, then we join the corresponding half-edges of G with an edge. If the ribbon backtracks at a vertex, this edge is a loop. If two half-edges of H occur consecutively as the ribbon crosses an edge, then we place labels a and $\bar{a}$ on the corresponding half-edges of G. Of course, we use a different label for each instance of a ribbon crossing an edge.

For an example of this construction, the ribbon complex in Figure 2 (in the case when both ribbons are closed) corresponds to the labeled graph in Figure 3.

There are several interesting relationships between R and G. First, if an open ribbon of R starts or stops in the middle of an edge of H, then the corresponding edge in G has one end unlabeled. If an open ribbon starts at a vertex of H, then the corresponding half-edge of G is one-ended. The relationship between the components of G and the local equivalence relation on the half-edges of H is more subtle; we explore this in the next two paragraphs.

If two vertices are joined by an edge of G, then the two half-edges on that edge correspond to adjacent half-edges of H which occur consecutively in some ribbon. It follows that each component of G consists of half-edges of H incident with a common vertex. Moreover, these half-edges lie in the same local equivalence class.

Conversely, if two half-edges h, h' of H lie in the same local equivalence class, then there exists a sequence $h = h_0, h_1, \ldots, h_m = h'$ of half-edges such that h_i and h_{i+1} occur consecutively as a ribbon passes through a vertex. This corresponds to a walk between the corresponding vertices in G.

It follows that the local equivalence classes of R correspond to the components of G.

The construction is easily reversed, except for the difficulty of distinguishing between *all* edges incident with a vertex and the local equivalence classes at that vertex. But the equivalence classes are exactly the vertices of the cleaved ribbon complex. We have shown the following.

Proposition 4.1 *Two ribbon complexes correspond to the same labeled graph if and only if they have the same cleaved ribbon complex.*

Corollary 4.2 *There is a unique cleaved ribbon complex corresponding to a given labeled graph.*

We are ready for the main theorem of this section.

Theorem 4.3 *A ribbon complex is spatial if and only if the corresponding labeled graph is planar.*

Proof: We will first show that if R is spatial, then G is planar. Consider a spatial embedding of R. We place a sphere around each vertex v of the spine H with radius ε. This ε should be chosen small enough so that R intersects the boundary of the sphere in exactly one point per edge of H and in a line segment for each ribbon passing through this vertex. It follows that the surface of the sphere contains an embedding of the components of the associated graph G corresponding to this vertex. Since the spheres can be chosen to have the same orientation and to be all exterior to each other, they define an embedding of G as a labeled graph. Thus G is planar as desired.

We will next show that if a labeled graph G is planar, then the associated ribbon complex R is spatial. The construction is the reverse of the one above. Specifically, begin with a planar embedding of G. Each vertex of R corresponds to a set of components of G. We suppose that each such set is embedded on its own oriented sphere. Moreover, we suppose that these oriented spheres are embedded in 3-space so that no one is in the interior of another.

Let c be the center of a given sphere. These centers will be the vertices of the spine H in R. For each point x on the sphere in the embedded component of G we form the line segment joining x to c. We then delete the points on the surface of the sphere not in G. Observe that each vertex of G now corresponds to line segment joining that point to the center. This line segment will be part of an edge in H. Each edge of G corresponds to a "wedge" joining the edge on the sphere to the center of the sphere. The authors like to picture a black hole at the center of the sphere, continuously sucking in the parts of the graph on the sphere.

We next join two spheres to create the edges of H as follows. Let v and $\bar{v}$ be two paired vertices of G. Suppose that v is in component B_1 and $\bar{v}$ is in component B_2. Let C_1 be a curve in the sphere corresponding to B_1 which intersects each edge incident with v at exactly one point one-third of the way along that edge. Suppose further that the interior of this curve contains v and no other vertex. Define C_2 similarly. We now add a cylinder in 3-space whose two boundary components are identified with C_1 and with C_2. The axis of symmetry of this cylinder can be chosen to identify with the vertices v and $\bar{v}$. For each pair of edges labeled a and $\bar{a}$ we can add a rectangle inside of this cylinder whose four sides are consecutively identified with a) the axis of symmetry, b) the first third of the edge with label a, c) a line segment on the boundary of the cylinder, and d) the first third of the edge labeled $\bar{a}$. Finally, we remove the cylinder which only played an auxiliary role in the description of the construction.

Because G is labeled planar, the spheres are oriented, and each is exterior to the other, it follows that the rectangles described above line up in the proper cyclic order. Doing this for each pair of labels simultaneously yields the desired spatial embedding of the ribbon complex R. ■

As an aside, the construction of Theorem 4.3 shows how to envision a spatial embedding of a cleaved complex R' from one of R. Namely, we place a small ε-sphere S around a vertex of R. Some component of $R \cap S$ can be separated from the other components by a curve in S. Form a cone from this curve to the center of S. The apex of this cone can now be pulled away from the center of the sphere, cleaving one equivalence class off of v. Repeating this process for every equivalence class creates an embedding of the cleaved complex R'. (One could also modify the proof of the 'if' part of Theorem 4.3 and construct the cleaved complex R' directly from G by embedding each component of G in a separate oriented sphere.)

Corollary 4.4 *There is an algorithm to determine if a ribbon complex is spatial.*

Proof: This follows immediately from Proposition 3.1 and Theorem 4.3. ■

5 The Partial Ordering

In this section we give a partial ordering on ribbon complexes. The ordering corresponds nicely to a partial ordering on labeled graphs. We begin by describing two reductions on ribbon complexes.

Let R be a ribbon complex and let $r = h_1, h_2, \ldots, h_m$ be an open ribbon. We *cut* r between h_i and h_{i+1} to create a new ribbon complex by deleting r and replacing it with two new ribbons $h_1, \ldots, h_i$ and $h_{i+1}, \ldots, h_m$. Similarly, we *cut* a closed ribbon between h_i and h_{i+1} by replacing it with an open ribbon $h_{i+1}, \ldots, h_m, h_1, \ldots, h_i$. It is clear that if R is spatial, then so is R'. The converse need not hold.

The second reduction on a ribbon complex R consists of *deleting a short ribbon*, removing a ribbon $r = h_1$ consisting of a single half-edge. Observe that R is spatial if and only if R' is spatial.

We use these two reductions to put a partial ordering on all ribbon complexes. Specifically, we write $R > R'$ if we can get from R to R' by a sequence of cutting ribbons and deleting short ribbons. This relation is clearly a partial ordering. Call R' a *minor* of R. The reader should not confuse this minor ordering with the graph minor ordering on the spines.

We examine the corresponding partial ordering on labeled graphs. First, suppose that R' is formed from R by cutting a ribbon crossing an edge. Let G' and G be the corresponding labeled graphs. The ribbon crossing the edge in R corresponds to a pair of labels $a, \bar{a}$ in G. After cutting, the ribbon no longer crosses the edge and so these two labels are missing in G'. In other words, G' is formed from G by deleting a pair of labels $a, \bar{a}$.

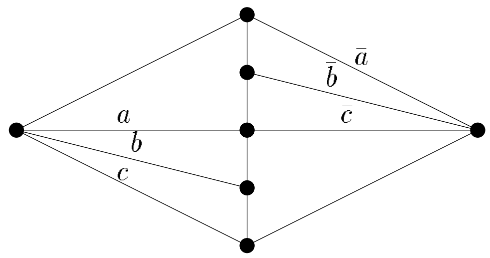

Figure 5: A sporadic non-planar labeled graph

Second, suppose that R' is formed from R by cutting a ribbon passing through a vertex. This corresponds to an edge in G. After cutting, the ribbon no longer passes through the vertex and so that edge is missing from G'. However, the half-edges and any labels on them still remain. In other words, G' is formed from G by replacing a (labeled) 2-ended edge with two (labeled) one-ended edges.

Third, suppose that R' is formed from R by deleting a short ribbon. This short ribbon corresponds to an unlabeled one-ended edge in G. Then G' is formed by deleting this unlabeled edge.

These operations on labeled graphs, erasing a label, replacing an edge with two half-edges, and deleting an unlabeled one-ended edge, create a partial ordering on the set of all labeled graphs. The ordering is the same as that on the corresponding ribbon complexes. Note that if $G' < G$ and G is labeled planar, then G' is also. We call G' a *minor* of G. The reader should not confuse this with the graph minor operation on the underlying graphs. The graph minor operation as defined here is not useful in labeled graphs since it is unclear what to do with the labels.

A ribbon complex R is *minimal non-spatial* if it is non-spatial, but any proper minor is spatial. Analogously, a labeled planar graph G is *minimal non-planar* if it is non-planar, but any proper minor is non-planar.

The Minimal Problem: *Find the minimal non-spatial ribbon complexes, or equivalently, find the minimal non-planar labeled graphs.*

For example, Figure 3 is a minimal non-planar labeled graph. Minimal non-planar graphs can be surprising, as illustrated by Figure 5. Our partial ordering is not a well-quasi-ordering. In

labeled graph	$\leftrightarrow$	triple graph
claw	$\leftrightarrow$	vertex
theta graph	$\leftrightarrow$	θ-edge
planar theta	$\leftrightarrow$	positive sign
nonplanar theta	$\leftrightarrow$	negative sign
reversing local rotation	$\leftrightarrow$	local switch
labeled claws	$\leftrightarrow$	λ-edges
linked components	$\leftrightarrow$	κ-edges
planar	$\leftrightarrow$	twist-free balance

Figure 6: Relating a labeled graph and a triple graph

Section 8 we will construct infinitely many non-comparable minor-minimal ribbon complexes.

6 Theta Graphs and the Triple Graph

In this section we give a second combinatorial characterization of when a ribbon complex is spatial. In particular, for a labeled graph G we construct a signed graph $T^{\pm}(G)$. This signed graph has a special type of balance if and only if G is labeled planar. This characterization will be useful in subsequent sections, particularly in trying to determine if certain classes of problems can be solved in polynomial time. The relationship is summarized in Figure 6.

A *claw* in a graph G is three half-edges incident with a common root v. A *theta graph* is a pair of distinct vertices joined by three internally-disjoint paths. Each theta graph contains two claws. Form the *triple graph* $T = T(G)$ whose vertices correspond to the claws of G and whose edges join the two claws in a theta subgraph. Two claws of G may lie in many different theta subgraphs, but are joined by just a single edge in T.

A *signature* on a graph places a *plus* or *minus* sign on each edge. A graph with a signature is called a *signed graph.* A *balanced cycle* is one with an even number of negative edges. Cycles with an odd number of negative edges are *unbalanced.* A signed graph is *balanced* if it contains no unbalanced cycle. A *balance* is a signature that yields a balanced graph; a balance that assigns a plus sign to every edge is *all-positive.*

A *local switch* on a signature toggles the sign of each edge incident with a fixed vertex. Local switches do not change the balance of any cycles. Two signed graphs are *equivalent* if there is a sequence of local switches changing one signature into the other. It is easy to see that each balanced graph is equivalent with a graph that admits an all-positive balance.

Theta graphs have exactly two embeddings in an oriented plane which differ only in the fixed orientation. Each embedding is described in terms of the cyclic permutations of the half-edges in a claw. These cyclic permutations are called *claw rotations*. If we toggle the claw rotation at exactly one claw of a planar theta, then we get a non-planar embedding.

We now use embeddings of theta graphs in G to put a signature on the triple graph $T(G)$; in this process some of the edges can be doubled. The resulting *signed triple graph* is denoted $T^{\pm}(G)$. As before, claws in G correspond to vertices in $T^{\pm}(G)$. Arbitrarily fix a rotation on each claw. If there is a theta graph whose two fixed claw rotations give a planar embedding, then join the corresponding vertices in $T^{\pm}(G)$ with a positive edge. If the claw rotations give a non-planar embedding, then join the corresponding vertices with a negative edge. The signed graph $T^{\pm}(G)$ may have multiple edges, but only if they have different signs.

The graph $T^{\pm}(G)$ depends on the fixed claw rotations of G. However, toggling a particular claw rotation in G simultaneously toggles the sign on all edges incident with the corresponding vertex in $T^{\pm}(G)$. Hence this corresponds to a local switch and so the two signatures on $T^{\pm}(G)$ are equivalent.

The following is the main result of [1].

Theorem 6.1 *An unlabeled graph G is planar if and only if the triple graph $T^{\pm}(G)$ is balanced.*

One direction is not too hard. A planar embedding induces a rotation on each claw and these claw rotations balance $T^{\pm}(G)$; in fact, they induce an all-positive balance on $T^{\pm}(G)$. The converse follows by relating $T^{\pm}(G)$ with $T^{\pm}(H)$ for subgraphs H of G, proving that $T^{\pm}(K_{3,3})$ and $T^{\pm}(K_5)$ are unbalanced, and citing Kuratowski's Theorem [8] (see [1] for details).

We would like to point out that it is *not* the case that every all-positive balance of $T^{\pm}(G)$ corresponds to a planar embedding of G. This can not be so, as the n-bond has only $(n-1)!$ planar embeddings but there are $2^{n(n-1)(n-2)/6}$ ways to obtain an all-positive balance on the corresponding triple graph $T^{\pm}$ (which is a matching with $\binom{n}{3}$ edges). It follows that many all-positive balances are not induced by rotations, i.e., cyclic permutations of half-edges at each vertex. This feature will be explored later in this section.

We extend the definition of the signed triple graph from unlabeled to labeled graphs. Let G_λ be a labeled graph on the underlying graph G. The signed triple graph $T = T^{\pm}(G_\lambda)$ again has vertex set the claws of G. As before, we fix a claw rotation on each claw. Edges come in three flavors. A *θ-edge* is defined as before: it joins two claws if they lie in a common theta subgraph with sign determined by whether the two claw rotations give a planar embedding of the theta subgraph. A *λ-edge* joins a claw with ordered labels abc to one with ordered labels $\bar{a}\bar{b}\bar{c}$. The edge is signed negative if the claw rotation permutes the labels in the orders above, signed positive if exactly one of the orders is reversed, and signed negative if both orders are reversed. The third flavor of edge comes from cut vertices. Let h_1, h_2, g_1, g_2 be half-edges incident with a cut-vertex v in such a way that h_1, h_2 are in the same block and g_1, g_2 are in a different common block. If the fixed claw rotations of the claws $\{h_1, h_2, g_1\}$ and $\{h_1, h_2, g_2\}$ are $(h_1 h_2 g_1)$, $(h_1 h_2 g_2)$, we join the two claws with a positive *κ-edge*; we do the same if both claws have the reverse orientation fixed, and we join the two claws with a negative κ-edge in the remaining two cases. Observe that the signed triple graph $T^{\pm}(G)$ of the unlabeled G is a subgraph of our new graph $T^{\pm}(G_\lambda)$ induced by θ-edges.

If the graph G_λ is labeled planar, then in a corresponding plane embedding the local rotations at vertices induce claw rotations on all claws and these claw rotations sign each θ-edge positive. Since the embedding reverses the labels at every pair of vertices, it also signs each λ-edge positive. No half-edges h_1, h_2 in one block and g_1, g_2 in another appear in cyclic order h_1, g_1, h_2, g_2 around a cut-vertex, so each κ-edge is positive. Hence a labeled planar graph G_λ has its signed triple graph $T^{\pm}(G_\lambda)$ balanced; in fact, the claw rotations induced by a labeled planar embedding of G_λ endow the graph $T^{\pm}(G_\lambda)$ with an all-positive balance.

Unfortunately, the converse is not true. Figure 7 gives an example of a labeled graph G_λ where $T^{\pm}(G_\lambda)$ is balanced, but no *all-positive* balance corresponds to a labeled planar embedding of G_λ. The signed triple graph here consists of a 2-bond on claws $\{abc, \bar{a}\bar{b}\bar{c}\}$ and a hexagon with θ- and λ-edges alternating. The difficulty comes from the fact that no all-positive balance corresponds to local rotations on the two vertices in Fig. 7.

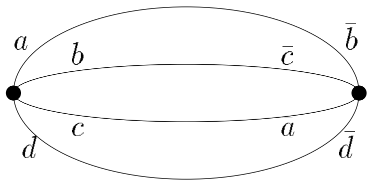

Figure 7: Every balance has a twist

Our next goal is to determine which all-positive balances correspond to labeled planar embeddings. This is done in terms of the claw rotations which define the signed triple graph. Assigning one of the two possible claw rotations to each claw of a graph gives a *system of claw rotations.* A *twist* is a set of three claws four adjacent half-edges x, a, b, c which have the claw rotations xab, xbc, and xca. A system of claw rotations in a graph is *twist-free* if no subset of claws form a twist.

Recall that a rotation on a graph is collection of local rotations, each local rotation being a cyclic permutation of the half-edges incident with a vertex. These local rotations induce claw rotations on each claw. We will say that a system of claw rotations in a graph is *rotational* if all the claw rotations are induced by a rotation on the graph.

Lemma 6.2 *A system of claw rotations in a graph is rotational if and only if it is twist-free.*

Proof: The rotation on the graph clearly induces a twist-free system of claw rotations. For the other direction, suppose that we are given a twist-free system of claw rotations on a graph. Let v be an arbitrary vertex of the graph. We will show that there is a cyclic permutation σ on the set $1, \ldots, n$ of half-edges incident with v such that the claw rotations at v are induced by σ.

A *tournament* is a digraph resulting from assigning a direction on each edge of a complete graph. If a tournament contains a directed cycle, then it contains a directed triangle. Also, any tournament which contains no directed cycle corresponds to a linear order on the vertices.

Construct a tournament T as follows. The vertex set of T is the half-edges $2, \ldots, n$. Direct an edge (i, j) in T whenever the corresponding claw has rotation $(1, i, j)$. The tournament has no directed triangle since the system of claw rotations in our graph is assumed to be twist-free. By the above comments, the tournament has no directed cycles and corresponds to a linear ordering $i_2, \ldots, i_n$ on the set of half-edges. Define the local rotation on v by $\sigma = (1, i_2, \ldots, i_n)$.

We need to show that the local rotation σ is consistent with the given system of claw rotations in our graph. This is true for each each claw with edge 1 by definition. Suppose by way of contradiction that a claw with rotation (i, j, k) appears in σ as $(1, \ldots, i, \ldots, k, \ldots, j, \ldots)$. This implies the claw rotations $(k, 1, i)$ and $(k, j, 1)$. Together with the hypothesised (k, i, j) these give a twist, contradicting the hypothesis that the system of claw rotations is twist-free. ■

As a corollary we obtain the following characterization of planarity for labeled graphs.

Proposition 6.3 *A graph G_λ is labeled planar if and only if the signed triple graph $T^{\pm}(G_\lambda)$ has a twist-free all-positive balance.*

Proof: Suppose that the graph G_λ is labeled planar. Then the corresponding set of local rotations induce a twist-free all-positive balance on $T^{\pm}(G_\lambda)$.

Conversely, suppose that we have a twist-free all-positive balance on $T^{\pm}(G_\lambda)$. By Lemma 6.2 this all-positive balance is consistent with a rotation which determines an embedding of the underlying graph of G_λ. If the embedding is planar, then we are done. Any non-planar embedding contains either a non-planar theta subgraph or a subgraph homeomorphic to one vertex with two loops, where each loop is in different block and is embedded as a non-contractible curve. The first case contradicts the all-positive balance on the θ-edges in $T^{\pm}(G_\lambda)$; the second case contradicts the all-positive balance on the κ-edges. ■

Observe that Proposition 6.3 gives a second algorithm for determining if a ribbon complex is spatial. Namely, construct the associated labeled graph G_λ and form its signed triple graph $T^{\pm}(G_\lambda)$; let the triple graph have a total of m components. If a component is unbalanced then the ribbon complex is not spatial. If all components are balanced, look for an *all-positive* balance that is twist-free. If no such all-positive balance exists then the ribbon complex is not spatial; if a twist-free all-positive balance is found then a spatial embedding of the ribbon complex is obtained.

All but one of these steps can be done in polynomial time. The trouble arises in that the number m of components of the signed triple graph may be linear in the size of the ribbon complex (i.e., linear in the size of G_λ), and one may need to generate all the 2^m ways to all-positively balance $T^{\pm}(G_\lambda)$ to check for twist occurence.

7 Connectivity of a Labeled Graph

In this section we will examine how vertex and edge cuts of a labeled graph affect its planarity. Our goal is to reduce (in some cases) to 3-connected labeled graphs and n-bonds. This makes our problems easier to solve, since a 3-connected graph has a unique embedding in the unoriented plane and we can analyze the embeddings of n-bonds. We can easily determine labeled planarity for 3-connected graphs, reduce 2-connected graphs to the 3-connected case, and reduce cut-edges. The difficulty comes in determining labeled planarity for graphs with cut vertices. Our approach is inspired by Tutte's theory of 3-connected components [11] and uses our theory of balance in triple graphs.

We begin with an examination of the edge-connectivity of G. Suppose that a labeled graph G has a cut edge $e = uv$ separating it into two components H_u, H_v. Let h_u, h_v be the ends of e incident with u, v respectively. Form G' by *breaking the 1-edge cut*: replacing e by two half-edges h_u and h_v which receive the same labels as in G.

Lemma 7.1 *A labeled graph G is planar if and only if G' formed by breaking a 1-edge cut is planar.*

Proof: The lemma follows immediately from Proposition 6.3 since $T^{\pm}(G)$ is isomorphic to $T^{\pm}(G')$. ■

Next, suppose that a 2-edge-connected labeled graph G has a 2-edge cut $e_1 = uv$ and $e_2 = wx$, where possibly $u = w$ and $v = x$. Let H_u and H_v be the components of $G - \{e_1, e_2\}$ where $u, w \in H_u$ and $v, x \in H_v$. Form G' by *breaking the 2-edge cut*: replacing the half-edges h_u, h_v in e_1 and h_w, h_x in e_2 with half-edges forming new 2-ended edges $h_u h_w$ and $h_v h_x$. These four half-edges of G' receive the same labels as their counterparts in G.

Lemma 7.2 *The labeled G is planar if and only if G' formed by breaking a 2-edge cut is planar.*

Proof: The lemma follows immediately from Proposition 6.3 since $T^{\pm}(G)$ is isomorphic to $T^{\pm}(G')$. ■

Now we briefly discuss 2-cuts. By Tutte's theory [11] of 3-connected components, if G is not 3-connected, a cycle, or a bond, one may find a decomposition $G = G_1 \cup_{\{u,v\}} G_2$ such that the graph $G_2' = G_2 + uv$ obtained from G_2 by adding a new *virtual* edges uv is either 3-connected, a simple cycle, or a bond. Cleaving off G_2' from G along the cut $\{u, v\}$ then yields a smaller graph $G_1' = (G_1 \cup \{uv\})$ (again with a virtual edge). Iterating this process one might hope that the labelled planarity problem could be reduced to (disconnected) graphs whose each component is either 3-connected, or a simple cycle, or a bond; in what follows such graphs will be called *nearly 3-connected.*

However, the problem with applying this idea to *labelled* graphs is that when we cleave a vertex v into two vertices v_1 and v_2, it is unclear what to do with the labels. In particular, we must also cleave the vertex $\bar{v}$ so that the labels still preserve adjacency. While we know what labeled edges must be incident with $\bar{v}_1$ and $\bar{v}_2$, we do not know which unlabeled edges must be so incident. In Sections 8 and 9 we examine some special cases where we can reduce across an arbitrary vertex 2-cut as indicated above. For those sections, note that [6] gives an algorithm for finding the 3-connected components of a graph in linear time.

We conclude this section with an observation about how much freedom we have to pick rotations in nearly 3-connected labeled planar graphs. Let $C_1, \ldots, C_m$ be the n-bonds and planar 3-connected components. A vertex $v \in C_i$ is *rotation fixing* if either a) C_i is 3-connected and v is incident with at least three labeled half-edges, or b) C_i is an n-bond in which every half-edge is labeled. Observe that if v is rotation fixing, then specifying the rotation on the labeled half-edges incident with v in a planar embedding of C_i determines the rotation on all other vertices in that component.

We can use the labels to pass these fixed rotations to other components. Specifically, suppose that v and $\bar{v}$ are both rotation fixing vertices in separate components C_1 and C_2. Then putting a rotation on any rotation fixing vertex in C_1 also determines the rotations on all vertices in C_2.

The *fixing graph* of a labeled G has vertex set the components of G and an edge joining C_1 to C_2 for every pair $v \in C_1$, $\bar{v} \in C_2$ of rotation fixing vertices. We call two vertices *rotationally linked* if they lie in components C_i, C_j which are joined by a path in the fixing graph. The following fact is easily seen to be true.

Lemma 7.3 *Let G be a nearly 3-connected labeled graph and let v_1 and v_2 be two vertices which* *are rotationally linked. If ϕ_1 and ϕ_2 are two labeled planar embeddings of G with the same local rotation on v_1, then they have the same local rotation on v_2.*

8 Solving the Cubic Problem

In this section we investigate the Cubic Ribbon Problem, the special case where at most three ribbons are sewn onto an edge. We give a reduction process which allows us to characterize the minor-minimal non-spatial ribbons. We also give a polynomial-time algorithm for determining if a ribbon is spatial. The proof of the characterization and presentation of the algorithm are similar and are given simultaneously.

We work in the corresponding class of labeled graphs. Since there are at most three ribbons on an edge, the corresponding labeled graph is of maximum degree three. We will reduce to the case where the graphs are nearly-3-connected.

We first check that the underlying graph is planar. If not, then it contains a homeomorph of $K_{3,3}$ and we are done.

There are several easy reductions which do not change the planarity of a labeled graph. First, if any vertices are incident with two or fewer labeled half-edges, then we can erase these labels. Second, if there is a vertex of degree one, we can delete it and its incident half-edge. Third, if a component contans a single vertex, we can delete that component and erase any incident labels. Fourth, if there is a vertex of degree two incident with two unlabeled half-edges, then we can replace the two incident half-edges with a single edge labeled in the obvious way.

If any of the easy reductions apply, then the labeled graph is not minimal non-planar. For the labeled planarity testing we apply all reductions as long as possible. The resulting graph is regular of degree three and every vertex is incident with either none or three labeled half-edges.

Suppose now that G_λ has a cut edge. By Lemma 7.1 there is a related graph G'_λ formed by breaking this edge. The graph G_λ cannot be minimal, as a reduction applies in G'_λ which corresponds to a labeled minor of G_λ. For the planarity testing we make this reduction and proceed. By repeating this process, we can assume that G_λ is 2-edge-connected.

Suppose that G_λ has an edge-cut of size two. By Lemma 7.2 we can break along these edges and create a related graph G'_λ. For the characterization of minimal graphs we keep track of these new edges. As in Tutte 2-cuts, we call them *virtual edges*. We repeat this process (and earlier reductions) until we obtain a nearly 3-edge-connected graph. Since the graph is cubic, it is also nearly 3-connected. In particular, every vertex is rotation fixing.

We next form the fixing graph F as in Section 7. Observe that G_λ is labeled planar if and only if the subgraphs corresponding to each component of F is planar. If G_λ is minimal, then F is connected. For the planarity algorithm we check the components of F one at a time.

We put a signature on F as follows. Each component of G_λ embeds in the nonoriented plane in a unique way. Arbitrarily fix an orientation on each component of this embedding. An edge of F occurs whenever one component has a vertex v and the other $\bar{v}$. Label this edge plus if the fixed orientations disagree and minus otherwise. The signature of F depends on the fixed orientation, but reversing this orientation is just a local switch on the signed graph. This signed graph is related to the signed triple graph of Section 6. Here each component of the subgraph of $T^\pm$ induced by θ-edges becomes a vertex of F. The λ-edges of $T^\pm$ become edges of F. There are no κ-edges in a graph of maximum degree three.

The graph G_λ is labeled planar if and only if the associated F is balanced. For the algorithm this can be easily checked and we are done with its description. We have shown the following.

Theorem 8.1 *There is a polynomial-time algorithm to determine planarity for the class of labeled graphs with maximum degree three.*

Corollary 8.2 *If a ribbon complex has at most three ribbons on each edge of the spine, then we can determine if it is spatial in polynomial time.*

For the characterization of minimal graphs, note that F is unbalanced if and only if it contains an unbalanced cycle. It follows that a minimal non-planar G_λ has F a single unbalanced cycle. We further examine the structure of such minimal G_λ.

A vertex of F is incident with exactly two edges, so each component of G_λ has exactly two labeled vertices. Since G_λ is nearly 3-connected, these two vertices have three disjoint paths joining them.

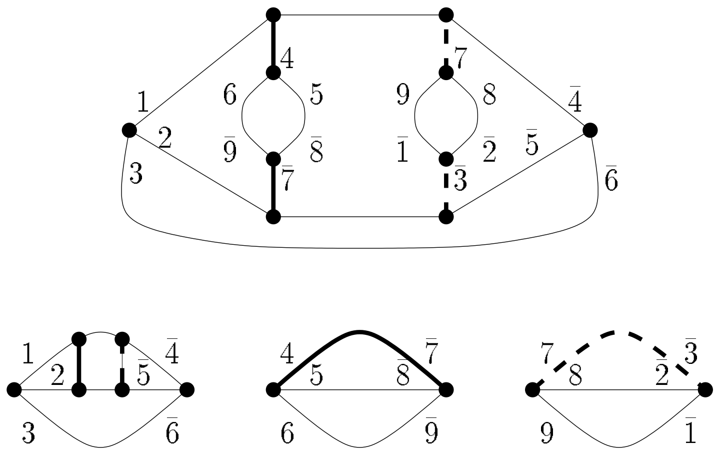

Figure 8: A minimal cubic non-planar labeled graph

As G_λ is cubic, the labeled claws lie in a common theta subgraph. Moreover, a minimal such G_λ consists only of this theta graph together with the virtual edges arising from 2-edge-cuts; in such a case we will briefly refer to a *labeled theta with virtual edges*. For ease of statement we also say that a sequence of components is *unbalanced* if the corresponding cycle in F is unbalanced.

Theorem 8.3 *A cubic labeled G_λ is minimal non-labeled-planar if and only the nearly 3-connected graph formed by repeated 2-edge-cuts has an unbalanced sequence of components and each component is a labeled theta with virtual edges.*

Figure 8 gives an example of a minimal non-planar labeled graph. The top half of this figure is the graph. The bottom half is the result after 2-edge-cuts. The (dotted) bold line in the first component of the bottom is a virtual edge which is paired with the (dotted) bold line in the second (respectively third) component.

We close this section by observing that, based on the example in Fig. 8 it is easy to construct infinitely many minor-minimal cubic graphs no one of which is a minor of another.

9 Solving a Totally Labeled Case

In this section we show that determining the planarity of a totally labeled graph with a certain connectivity condition can be done in polynomial time. The connectivity condition is that G does not contain a cut-vertex incident with 3 or more blocks. Translating to ribbon complexes, our results show that determining the spatialness of a large class of closed ribbon complexes can be done in polynomial time. Unfortunately, this class does not correspond to those arising from open ribbons via Proposition 2.1. We cannot find a satisfactory geometric characterization of these ribbon complexes.

We begin with some reductions similar to those in Section 8. Let v be a vertex of a labeled graph G_λ and let $V_1 \cup V_2$ be a partition of all half-edges incident with v. By *splitting* the vertex v with respect to $V_1 \cup V_2$ we mean forming a new graph from G_λ by replacing v with a new edge v_1v_2 where the new vertex v_i is incident to precisely the half-edges in V_i, $i = 1, 2$.

First suppose that G has a cut-vertex v with exactly two blocks B_1 and B_2. Let $V_1 \cup V_2$ be the partition of the half-edges incident with v induced by the blocks and let $\bar{V}_1 \cup \bar{V}_2$ be the corresponding partition at the vertex $\bar{v}$. Form a new graph G^* by splitting v with respect to $V_1 \cup V_2$, and simultaneously split the vertex $\bar{v}$ with respect to $\bar{V}_1 \cup \bar{V}_2$. Label the edges of G^* in the natural way.

Lemma 9.1 *G is labeled planar if and only if G^* is labeled planar.*

Proof: Suppose that G is labeled planar. Because there are only two blocks, half-edges in V_1 occur consecutively in the local rotation at v. Since G is labeled planar, so must the half-edges $\bar{V}_1$ at $\bar{v}$. Hence we can split these in the embedding and get a labeled embedding of G^*.

Conversely, suppose that G^* is labeled planar. Then we can contract the two edges and construct a labeled planar embedding of G. ■

Next suppose that G is 2-connected and 3-edge-connected but not 3-connected. Then G has a 2-cut uv such that $G = G_1 \cup_{\{u,v\}} G_2$ where each G_i contains at least two edges. Set $G_i' = G_i \cup \{uv\}$, that is, G_i' is obtained from G_i by adding an extra *virtual* edge joining the vertices u and v. As mentioned earlier, it follows from Tutte's decomposition theory [11] that the 2-cut can be chosen in such a way that (say) G_2' is either 3-connected, or a simple cycle, or a bond. Moreover, taking the 3-edge-connectivity assumption of G into account we see that G_2' cannot be a simple cycle; so we may assume that G_2' also is 3-edge-connected. For $i = 1, 2$ let U_i and V_i be the half-edges incident with u and v in G_i, respectively; it follows that the sets V_2 and U_2 each have cardinality at least two. Invoking 3-edge-connectivity of G again, one of V_1 or U_1 has cardinality at least two, say V_1 does. Form G^{**} from G by splitting the vertex v with respect to the partition $V_1 \cup V_2$. Simultaneously split the corresponding labeled half-edges at $\bar{v}$ and put labels on the two new edges in a natural way.

Lemma 9.2 *G is labeled planar if and only if G^{**} is labeled planar.*

Proof: Suppose that G is labeled planar. Using the fact that G_2' is either 3-connected or a bond, u and v occur together on a unique face of G_2. This face must contain all of G_1. It follows that the half-edges in V_2 occur consecutively in the local rotation at v. Since G is labeled planar, so must the half-edges $\bar{V}_2$ at $\bar{v}$. Hence we can split these in the embedding and get a labeled embedding of G^{**}.

Conversely, suppose that G^{**} is labeled planar. Then we can contract the two edges and construct a labeled planar embedding of G. ■

By repeated application of the above lemmas together with the 1- and 2-edge-cut breaking Lemmas 7.1 and 7.2 we can assume that G is nearly 3-connected, that is, each component of G is either 3-connected, or a simple cycle, or a bond (possibly, all with labeled one-ended edges). For the ease of explanation assume that there are no labeled one-ended edges; the reader will quickly realize that their presence does not invalidate the arguments that follow.

Cycles present no obstruction to labeled planarity, so we assume there are none. In the remaining two cases, each vertex is rotation fixing. Form the fixed graph F as in Section 7. Because every vertex is rotation fixing, G is planar if and only if each subgraph corresponding to a component of F is planar. Hence we can check these subgraphs one at a time, or equivalently, proceed under the assumption that F is connected.

If G contains one component which is 3-connected, then there are only two possible rotations on that component. Since F is connected, either one of these two rotations completely determine the rotations on every other vertex. It follows that we can check the labeled planarity of F quickly.

Suppose instead that G contains only n-bonds. If there is a single component, then it is labeled planar if and only if it corresponds to a set of ribbons sewn on a single loop in such a way that each ribbon is wrapped the same number of times around the loop. This can be quickly checked. If there is more than one n-bond, then without loss of generality the first contains vertices v_1 and $\bar{v}_2$, the second v_2 and $\bar{v}_3$, and so forth until the last contains v_k and $\bar{v}_1$. We now glue the first two of these n-bonds together and eliminate the vertices v_2 and $\bar{v}_2$. The resulting graph is labeled planar if and only if G is. It follows from induction that we can quickly check the labeled planarity of G.

We have shown the following:

Theorem 9.3 *There is a polynomial-time algorithm to determine labeled planarity for the class of totally labelled graphs without cut-vertices incident with 3 blocks.*

We note the following two special cases of Theorem 9.3.

Corollary 9.4 *There is a polynomial-time algorithm to determine if a nearly 2-connected totally labeled graph is planar.*

Corollary 9.5 *There is a polynomial-time algorithm to determine if a totally labeled graph of maximum degree 5 is planar.*

Proof: Any cut-vertex in a graph of maximum degree 5 incident with 3 or more blocks must contain a cut edge. ■

Theorem 9.3 is tantalizing close to completely solving the totally labeled case. To illustrate the difficulty of the remaining case we offer the following example. Let G be a totally labeled graph with two vertices $v, \bar{v}$ each incident with 3 loops. Suppose that v is incident with loops ab, cd, ef and $\bar{v}$ is incident with loops $\bar{b}\bar{c}, \bar{d}\bar{e}, \bar{f}\bar{a}$. Then we cannot split v or $\bar{v}$ so that each embedding of G corresponds to an embedding of G^*.

10 Conclusion

There are many interesting questions concerning ribbon complexes. In this section we explore just a few.

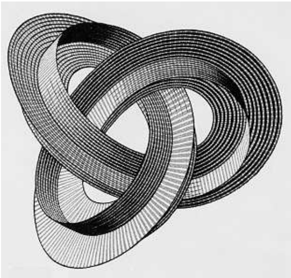

Figure 9: A knotty ribbon

We first consider the problem of determining all possible spatial embeddings of a given ribbon complex, where two embeddings are *equivalent* if there is an ambient isotopy of 3-dimensional space which carries one to the other. This is related to knot theory, or more specifically to the theory of embedding graphs in 3-space [10], because an ambient isotopy cannot change the *knot type* of the spine. Note also that an ambient isotopy cannot change the oriented planar embedding of the corresponding labeled planar graph.

Define a *Dehn twist* of a ribbon complex embedded in 3-space as follows. First cut the ribbon complex along a plane perpendicular to an edge of the spine. Twist one side of this plane n revolutions, or equivalently $2\pi n$ radians, and re-identify the ribbons. The result does not change the type of the ribbon complex, but does change the spatial embedding.

Theorem 10.1 *Two embeddings of a ribbon complex into 3-dimensional space are equivalent up to Dehn twists if and only if 1) the spines have the same knot type, and 2) the oriented embeddings of the corresponding labeled planar graphs are the same.*

Proof: Together with the comments above, the proof follows the proof of Theorem 4.3. ■

The authors were delighted to discover that M.C. Escher had anticipated these embeddings in his picture shown in Figure 9. This figure shows a single closed ribbon wrapped four times around a trefoil knot.

It would be interesting to find a geometric characterization of cubic ribbons. To some extent this problem is answered by Theorem 8.3. The geometric role of the different components in the labeled graph just corresponds to subdividing the underlying spine. However, the geometric role of 2-edge-cuts and virtual edges is less clear. Similarly, it would be interesting to find a geometric characterization of the ribbon complexes corresponding to the labeled graphs of Theorem 9.3.

Left open is the computational complexity of the general ribbon problem. Specifically, consider the following decision problem:

The Ribbon Problem —

Instance: *a ribbon complex R.*

Question: *is R spatial?*

The ribbon problem is in NP by Corollary 4.4.

Conjecture 10.2 *The Ribbon Problem is NP-complete.*

A ribbon complex can be described as a set of words on the half-edges in the spine. Our original goal was to find combinatorial conditions on these words that would ensure that the ribbon complex was spatial. This is in the spirit of Gauss' Word Problem [9]. We were unable to find such a characterization of spatial ribbons. One consideration is a variation on a Dehn twist where an edge incident with n ribbons is cut and rotated $2\pi/n$ radians before being re-identified. This does not change the spatialness of a ribbon complex, but seems to greatly change the words describing that complex.

In closing, we mention the possibility of generalizing the ribbon problem to higher dimensions. Suppose that we sew small portions of half-3-space on a 2-dimensional simplicial complex. When does the resulting structure embed in 4-dimensional space?

Acknowledgement: The authors worked in many subsets hosted by one or another of the four home institutions. Each author gratefully thanks the other hosts. The third author gratefully acknowledges the support of NSERC. The fourth author acknowledges partial support by the VEGA grant No. 1/3034/96 and the GAČR grant No. 201/98/1451. The proof of Lemma 6.2 is due to Andrej Zlatoš working with the Combinatorics Group at Comenius University, Bratislava. We also thank Gelasio Salazar, Jay Adamsson, Steve Fisk, and Tom Tucker for interesting discussions on this material. Finally, we thank M.C. Escher for his wonderful picture which we have included as Figure 9.

References

[1] D. Archdeacon and J. Širáň, Characterizing planar graphs using theta graphs, *J. Graph Theory* **27**(1998)17-20.

[2] C.P. Bonnington and C.H.C. Little, “Foundations of Topological Graph Theory”, Springer, New York (1995).

[3] T.H. Cormen, C.E. Leiserson, and R.L. Rivest, “Introduction to Algorithms”, MIT Press, Cambridge, Mass (1989)

[4] M.R. Garey and D.S. Johnson, “Computers and Intractability: A Guide to the Theory of NP-Completeness”, W.H. Freeman (1979).

[5] J. Gross and T. Tucker, “Topological Graph Theory”, Wiley, New York (1987).

[6] J.E. Hopcroft and R.W. Tarjan, Dividing a graph into triconnected components, *SIAM J. Comput.* **2**(1973)135-158.

[7] J.E. Hopcroft and R.E. Tarjan, Efficient planarity testing, *Jour. ACM* **21**(1974)549-568

[8] K. Kuratowski, Sur le problème des courbes gauches en topologie, *Fund. Math.* **15**(1930)271-283.

[9] L. Lovász and M. Marx, A forbidden substructure characterization of Gauss codes, *Acta Sci. Math.* **38**(1976)115-119.

[10] N. Robertson, P.D. Seymour, and R. Thomas, Sachs' linkless embedding conjecture, *J. Combin. Theory Ser. B* **64**(1995)185-227.

[11] W.T. Tutte, “Connectivity in Graphs”, Oxford University Press, London (1966).