OBSTRUCTIONS FOR EMBEDDING CUBIC GRAPHS ON THE SPINDLE SURFACE

DAN ARCHDEACON AND C. PAUL BONNINGTON

ABSTRACT. The *spindle surface* S is the pinched surface formed by identifying two points on the sphere. In this paper we examine cubic graphs that minimally do not embed on the spindle surface. We give the complete list of 21 cubic graphs that form the topological obstruction set in the cubic order for graphs that embed on S.

A graph G is *nearly-planar* if there exists an edge e such that $G-e$ is planar. All planar graphs are nearly-planar. A cubic obstruction for near-planarity is the same as an obstruction for embedding on the spindle surface. Hence we also give the topological obstruction set for cubic nearly-planar graphs.

1. INTRODUCTION

Kuratowski's Theorem [9] says that a graph embeds in the sphere if and only if it does not contain a subdivision of either $K_{3,3}$ or of K_5. Another way of stating Kuratowski's Theorem is that these two graphs are *obstructions* to embedding in the sphere, and any non-spherical graph must contain one of these obstructions. Kuratowski's Theorem is one of the most celebrated results in graph theory, as well as one of the most useful. Herein we study two variations on Kuratowski's theme.

A *spindle surface* is formed from the sphere by identifying two distinct points, commonly considered as the north and south poles N and S. Equivalently, a spindle surface is a quotient space formed from the torus by identifying a contractible cycle to a point. Despite its name, the spindle surface is not a surface, but rather a *pseudo-surface* with a single *pinch point* $N = S$. We assume that any graph embedded on the spindle surface has a vertex at the pinch point. We ask:

Question 1.1. *What are the obstructions to embedding on the spindle surface?*

Obstruction sets for embedding on surfaces give nice structural characterizations for these classes of graphs. Except for Kuratowski's result about planar graphs, little is known. Archdeacon [1, 2] (see also [7]) gave the complete list of 103 *topological* obstructions for embedding in the projective plane (the topological order allows taking subgraphs and supressing degree-two vertices). This is the only other surface for which a complete list is known. It follows from the work of Robertson and Seymour [15, 16] on graph minors that the set of obstructions under the *minor* order (which allows taking subgraphs and contracting edges) is finite for embedding on a particular fixed surface.

Finding obstruction sets for pseudosurfaces is more subtle. If H is a subdivision of G and G embeds in a pseudosurface, then H embeds in that pseudosurface. The

Research supported by the Marsden Fund (Grant Number UOA825) administered by the Royal Society of New Zealand.
Submitted for publication, 5 September 2001.

corresponding statement for minors H of G is not true. The minor-obstruction set for other pseudo-surfaces is not necessarily finite [17]. Knor [10] has characterized the pseudosurfaces that are minor-closed, and hence have finite obstruction sets. These resemble a pseudo-surface with a single pinch point and spheres attached in a tree-like structure. This includes our spindle surface, so it is known that the set of minor obstructions for embedding here must be finite.

Kuratowski's Theorem also characterizes those graphs that are one edge-deletion away from planarity: if a graph is non-planar, but the deletion of any edge makes the graph planar, then the graph is a subdivision of either $K_{3,3}$ or of K_5. Say a graph is *nearly-planar* (under the subgraph order) if there exists some edge whose deletion makes the graph planar. We also ask:

Question 1.2. *What are the obstructions for nearly-planar graphs?*

It is important to note that being nearly-planar is *hereditary* under the subgraph and topological orders; that is, if G is nearly-planar, then so is every topological subgraph of G. However, it is not hereditary under the minor order. An example is the graph G created by adding a single edge joining a pair of antipodal vertices of the icosahedron. This G is nearly-planar, delete this new edge. However, if we contract the new edge, then we get a graph that is not nearly-planar. The concept of obstructions makes sense only for hereditary properties. Hence it makes sense to look for obstructions for being nearly-planar when using the subgraph or topological order, but not when using the minor order.

A graph is non-nearly-planar if and only if every edge is disjoint from a Kuratowski subgraph. Call an edge of e of a non-planar graph G *essential* if it is in every Kuratowski subgraph of G, that is, if and only if $G - e$ is planar. A graph is non-nearly-planar if and only if every edge is non-essential. An edge e of G is *redundant* if every edge of $G - e$ is non-essential. Thus our second question is equivalent to finding graphs that are not nearly-planar and have no redundant edge.

Unfortunately, we can't answer either of our two main questions. However, we can answer both questions in an important special case. A *cubic* graph is one with every vertex of degree three. The *cubic order* makes smaller graphs by edge deletions followed by suppressing the resulting degree-two vertices. It is equivalent to the topological order for the class of cubic graphs, but the name emphasizes the class under consideration. In the cubic order, the analogue of Kuratowski's Theorem is that a cubic graph is planar if and only if it does not contain a $K_{3,3}$.

We answer both of the above questions for obstructions in the cubic order. It turns out that the two obstruction sets are the same, as shown by the following two lemmas.

Lemma 1.3. *A cubic graph G embeds on the spindle surface if and only if there exists an edge e such that $G - e$ is planar.*

Proof. First, suppose that G embeds on the spindle surface. Then one vertex v lies on the pinch point. Because G is cubic, either the North pole or South pole part of a neighborhood of the pinch point has only a single edge end. Deleting this edge gives a planar graph.

Conversely, suppose that there is an edge $e = uv$ with $G - e$ planar. Extend this embedding to one of G' by adding a new edge $e' = uv'$ incident with u. Identifying the two vertices v and v' in the sphere gives an embedding of G on the spindle surface. $\square$

The preceding lemma is not true for non-cubic graphs. The relation between a non-cubic graph that embeds on the spindle surface and the non-cubic analogue of near-planarity is more subtle.

We immediately get the following.

Lemma 1.4. *A cubic graph G is minimally non-spindle if and only if it is minimally non-nearly-planar.* □

We can now state our main results.

Theorem 1.5. *There are exactly 21 minimal non-nearly-planar graphs in the cubic order. They are given in Figures 3-11.*

Corollary 1.6. *There are exactly 21 minimal non-spindle graphs in the cubic order. They are given in Figures 3-11.*

The proof of Theorem 1.5 is given in Section 8. We give here a sketch of the proof of our main result, combined with an outline of this paper. When studying the embedding properties of a graph G it is helpful to study the Kuratowski subgraphs $K_{3,3}$ and K_5 of G. To this end, in Section 2 we define K-structures—disjoint subgraphs of G that contain portions of Kuratowski graphs. In Section 3 we characterize those obstructions that have rich K-structures. These correspond to obstructions with low connectivity in which one component is non-planar. Section 4 contains a technical study of obstructions that have a cyclic-cut-set of three edges. The lack of such a cut-set allows us to assume that the remaining obstructions G have a highly-cyclically-connected subgraph H. We then study the possible *bridges*, components of $G - H$. Section 5 sets up the study of these bridges, and in Section 6 we give bounds on the size of G in terms of the size of H for certain special cases. Section 7 finishes our technical analysis of graphs with a K-structure, in part involving a computer search of all small order cubic graphs.

We are ready for the proof of our main result in Section 8. The proof develops a relationship between being nearly-planar and embeddings in the projective plane. Basically, a nearly-planar obstruction that embeds in the projective plane is minimal with face-width three. A non-projective nearly-planar cubic obstruction must then contain a cubic obstruction for embedding in the projective plane. This set is known. These subgraphs are either on our list of nearly-planar obstructions, or show that the graph has a K-structure that has already been analysed.

The proof of our main result relies in part on an exhaustive computer search. In Section 9 we discuss the algorithm used in this search and give some double-checks that our work is correct. We close in Section 10 with some comments and directions for future research.

2. Bridges and K-structures

A common method of studying embeddings of a graph G is to first embed a subgraph H, then to try to extend that embedding to all of G. The following definition is useful in this context.

A (G, H)-*bridge* of a subgraph H of a graph G is the closure of a topologically-connected component of $G - H$. Each bridge consists of a connected component of the graph $G - V(H)$ together with all edges joining that component to H, or else it is a single edge. A *foot* of a bridge B is a vertex in common with B and H. A *leg* of B is an edge incident with a foot. Finally, a

The graphs we study frequently have a common structure involving portions of Kuratowski graphs. We next describe this structure.

Let G be a graph. A *topological* $K^t_{3,3}$ is a subgraph that is a subdivision of $K_{3,3}$. A $K^{-e}_{3,3}$ is a subgraph H of G that is a subdivision of $K_{3,3} - K_2$, provided that there is a $K^t_{3,3}$, K, with $H \subset K \subset G$. Equivalently, a $K^{-e}_{3,3}$ is a subdivided K_4, provided there is a bridge with feet in the interior of two edges that are a matching on the K_4. Whether or not a graph H is a $K^{-e}_{3,3}$ depends also on the supergraph G. Similarly, a $K^{-v}_{3,3}$ is a subgraph H of G that is a subdivision of $K_{2,3}$, provided that there is a $K^t_{3,3}$, K, with $H \subset K \subset G$. Hence $K^{-e}_{3,3}$ and $K^{-v}_{3,3}$ are topologically complete bipartite graphs missing an edge and a vertex respectively, provided that the missing parts are in the whole graph G. This is similar to the idea of *k-graphs* [3], frequently used in studying embeddings.

A *K-structure* in a graph G is a pair (H_1, H_2) of disjoint subgraphs such that each H_i is either a topological $K^t_{3,3}$, $K^{-e}_{3,3}$, or $K^{-v}_{3,3}$. Observe that the K_i's completing the H_i's need not be disjoint, only the subgraphs must be disjoint.

We use the natural transitive "containment" relation that a $K^t_{3,3}$ is *richer* than $K^{-e}_{3,3}$ which is in turn *richer* than $K^{-v}_{3,3}$. In a K-structure we always list the richer of the two parts first. We extend this order to K-structures by saying that (H_1, H_2) is *richer* than (H_1', H_2') whenever H_1 is richer than H_1', or if the two are equal, then H_2 is richer than H_2' (this is the lexicographic order on pairs). Finally, we will refer to the *type* of the K-structure by using, for example, the notation $(K^{-e}_{3,3}, K^{-v}_{3,3})$-*structure*.

3. Obstructions with Rich K-structures

In this section we study nearly-planar obstructions with rich K-structures. In particular, we find all such graphs with a $(K^t_{3,3}, K^t_{3,3})$-, $(K^t_{3,3}, K^{-e}_{3,3})$-, or $(K^t_{3,3}, K^{-v}_{3,3})$-structure. These in turn correspond to all nearly-planar obstructions with either a 2-edge-cut, or a three-edge-cut with one component non-planar, as shown in Lemmas 3.3 and 3.5.

Proposition 3.1. *Let G be cubic obstruction to near-planarity with a $(K^t_{3,3}, K^t_{3,3})$-structure. Then G is the disjoint union of two $K^t_{3,3}$ as shown in Figure 3.*

Proof. If G has a $(K^t_{3,3}, K^t_{3,3})$-structure K, then every edge not in K is redundant. By minimality, $G = K$ as claimed. □

This is the only disconnected nearly-planar cubic obstruction. It is easy to show that no such obstruction has a cut edge, so henceforth we can assume that obstructions are 2-edge-connected.

Proposition 3.2. *Let G be cubic obstruction to near-planarity with a $(K^t_{3,3}, K^{-e}_{3,3})$-structure (but with no richer K-structure). Then G is one of the graphs of Figure 4.*

Proof. Let (H_1, H_2) be the $(K^t_{3,3}, K^{-e}_{3,3})$-structure. Then any edge not in H_1 is non-essential. Likewise for any edge $e \in H_1$, H_2 still completes to a $K^t_{3,3}$ in $G - e$ since H_1 is 2-edge-connected. The result now follows by an easy exhaustive search considering every way to complete H_2 to a $K^t_{3,3}$. □

As the following shows, any remaining obstructions must be at least three-edge-connected.

Lemma 3.3. *Let G be a cubic obstruction to near-planarity with edge-connectivity two. Then G has a $(K^t_{3,3}, K^{-e}_{3,3})$-structure.*

Proof. Let B be a cut set of two edges with components H_1 and H_2. Each H_i has two degree-two vertices that are the ends of edges in B. Let $\bar{H}_i$ denote the graph formed from H_i by adding an edge between these two vertices. If both H_i are planar, then deleting any edge of B yields a planar graph, and hence G is nearly-planar. Therefore, now suppose H_1 is non-planar, and hence contains a $K^t_{3,3}$. If $\bar{H}_2$ is planar, then G is nearly-planar, contrary to our assumption. So $\bar{H}_2$ is non-planar, and hence H_2 contains a $K^{-e}_{3,3}$ as desired. □

Proposition 3.4. *Let G be cubic obstruction to near-planarity with a $(K^t_{3,3}, K^{-v}_{3,3})$-structure. The G is one of the graphs of Figure 5.*

Proof. Let (H_1, H_2) be the $(K^t_{3,3}, K^{-v}_{3,3})$-structure. Then any edge not in H_1 is non-essential. Likewise for any edge $e \in H_1$, H_2 still completes to a $K^t_{3,3}$ in $G - e$ since H_1 is 2-edge-connected. The result now follows by an easy exhaustive search considering every way to complete H_2 to a $K^t_{3,3}$. The proof essentially reduces to finding all ways to insert the three degree 2 vertices corresponding to the feet of the (G, H_1)-bridge containing H_2 into the $K_{3,3}$. Each of these insertions gives one of the graphs of Figure 5 with two exceptions: *(i)* if all three vertices are in the same topological edge of H_1 then the resulting graph has a richer $(K^t_{3,3}, K^{-e}_{3,3})$-structure and is covered in the preceding proposition, and *(ii)* if the three vertices are in a perfect matching of the $K_{3,3}$, then the resulting graph contains the Petersen graph (see Figure 10) and hence is not minimally non-nearly-planar. □

As the following shows, any remaining obstructions with a cyclic three-edge cut must have both sides planar.

Lemma 3.5. *Let G be a cubic obstruction to near-planarity with a cyclic three-edge-cut having one side non-planar. Then G has a $(K^t_{3,3}, K^{-v}_{3,3})$-structure.*

Proof. Let B be a cut set of three edges with components H_1 and H_2. Each H_i has three degree-two vertices that are the ends of edges in B. Let $\bar{H}_i$ denote the graph formed from H_i by adding a vertex adjacent to each of these three degree-two vertices. If both H_i are planar, then deleting any edge of B yields a planar graph, and hence G is nearly-planar. Say H_1 is non-planar, and so contains a $K^t_{3,3}$. If $\bar{H}_2$ is planar, then G is nearly-planar contrary to assumption. So $\bar{H}_2$ is non-planar, and hence H_2 contains a $K^{-v}_{3,3}$ as desired. □

We do not yet have all of the desired obstructions with a non-trivial three-edge-cut. We examine cubic obstructions with a three-edge-cut having both sides planar in the next section.

4. Obstructions with Cyclic Edge-Connectivity Three

In this section we characterize the cubic obstructions to near-planarity with cyclic-edge-connectivity exactly three. Fix a three-edge-cut B where both components of $G-B$ are planar (the case where one component is non-planar was covered by Lemmas 3.3 and 3.5). We break into two main cases, depending upon whether one side of $G-B$ has the special property described next.

Let G be a cubic graph with three distinguished edges, which we will call the *red* edges. We say that G has the *three-red-edge* property if G is planar, but no embedding of G in the plane has two red edges on the boundary of a common face.

In Proposition 4.2 we will give the complete set of nine minimal cubic graphs with the three-red-edge property. (In our order on edge-colored graphs, deleting any red edge and supressing degree-two vertices allows either of the two new edges to become the third red edge.) First, however, we describe the relationship with cubic obstructions to planarity. If H is a cubic graph with the three-red-edge property, then let $H+v$ be the graph formed by subdividing each of the red edges and adding a new vertex v adjacent to the three degree-two vertices. Similarly, let $H+K_{2,3}$ be the graph formed by subdividing the red edges and adding three arcs making a matching between the three degree-two vertices of H and the three degree-two vertices of the $K_{2,3}$. In general, if B is any graph with all vertices of degree three except for three vertices of degree one, let $H+B$ be the cubic graph formed by subdividing the red edges of H and identifying the degree two vertices pairwise with the degree one vertices of B.

Lemma 4.1. *Let H be a minimal cubic graph with the three-red-edge property. Suppose that there exists an edge e such that $H-e$ has all three red edges on a common face. Then $H+K_{2,3}$ is non-nearly-planar. Moreover, given such an e and a bridge B with $H+B$ non-nearly-planar, then B contains a $K_{3,3}^{-v}$. If no such edge e exists, then $H+v$ is non-nearly-planar.*

Proof. We will show that in both cases the graph G augmenting H has every edge disjoint from a $K_{3,3}^t$, and so is non-nearly-planar.

First suppose that e is an edge of G that lies in the original H. If $H-e$ does not have a face with all three red edges on a common face, then it does not extend to an embedding of $(H-e)+v$. If $H-e$ does have a face with all three red edges on a common face, then it does not extend to an embedding of $(H-e)\cup K_{3,3}^{-v}$ (however, it does if B is a bridge not containing a $K_{3,3}^{-v}$). In either case, there is a $K_{3,3}^t$ disjoint from e in the appropriate extension G.

Next suppose that e is an edge of G that does not lie in the original H. Then there are two red edges r_1 and r_2 that do not correspond to that portion of $G-H$ related to e. Now $G-e$ is planar if and only if r_1 and r_2 lie on a common face. □

Before giving the complete set of minimal graphs with the three-red-property we make the following simple observation. A planar graph H with two red edges will have the *two-red-edge property* if no embedding has both edges on the same face. Subdivide the two red edges and add an edge e between the new degree-two vertices. Then H has the two-red-property if and only if $H+e$ is non-planar. It follows that the only topological minimal graph with the two-red-edge property is $K_{3,3}^{-e}$.

Proposition 4.2. *There are exactly 9 planar cubic graphs that are minimal with the three-red-edge property. They are given in Figure 1, where the red edges are denoted with solid circles.*

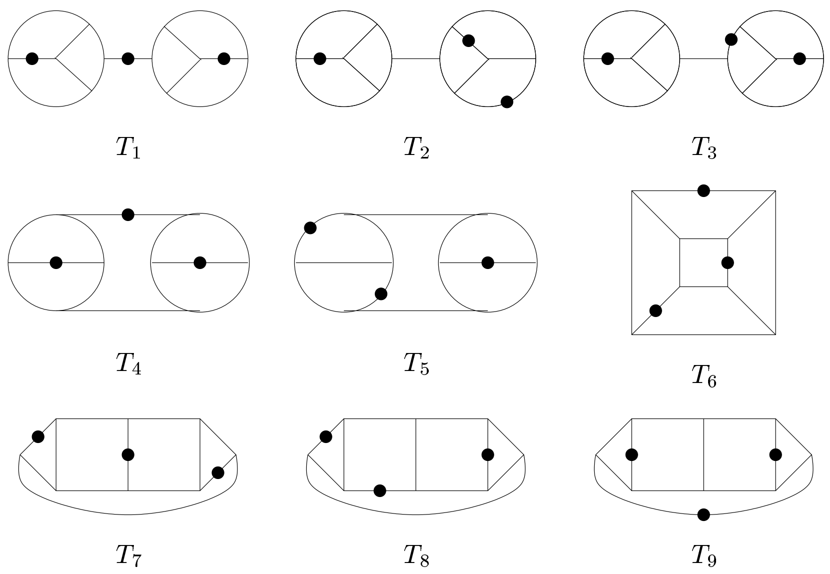

FIGURE 1. Minimal Graphs with the Three-Red-Edge Property (red edges are indicated with solid circles)

Proof. Let H be a minimal cubic planar graph with the three-red-edge property. Our analysis breaks into cases depending on the edge-connectivity κ' of H, with sub-cases depending on whether a cut set has red edges. Throughout this proof B will be a cut-set with H_1, H_2 the components of $H - B$.

Case 1: *κ'=1.* First suppose that B is a single red edge e. Observe that each H_i contains a red-edge, or else the original H was not minimal. Let H_i^+ be the graph made by supressing the degree-two vertex of H_i and coloring the resulting edge red. Then H_i has the 2-red-edge property, and by the comments preceeding this proposition is a $K_{3,3}^{-e}$. It follows that H is the graph T_1 of Figure 1.

Next, suppose that B is a single edge e that is not red. Let $\bar{H}_i$ be H_i with the degree two vertex supressed, and let H_i^+ be $\bar{H}_i$ with the resulting edge colored red, call this the *special* red edge. First, if one of the H_i^+ has two red edges on a common face and neither is special, then H has a similar embedding; this contradicts the hypothesis. If both H_i^+ have two red edges on a common face and both pairs involve the special edge, then H has an embedding with the two non-special red edges on a common face. Hence at one of the H_i^+ must be minimal with two red edges, that is, without loss of generality H_1^+ is a $K_{3,3}^{-e}$. Now, $\bar{H}_2$ must contain a $K_{3,3}^{-e}$ with one red edge, and hence H contains either T_2 or T_3 of Figure 1.

Case 2: *κ'=2.* We first observe that a 2-edge-cut B cannot contain two red edges, because they must appear together on a face.

Suppose that B has a red edge. Note by minimality, each H_i has one of the remaining red edges. Define H_i^+ as H_i with a red edge joining its two degree two vertices. If either H_i^+ has its red edges on a common face, then there is an embedding of H with this property. Hence each H_i^+ has the two-red-edge property. Again, by minimality, each H_i^+ is a $K_{3,3}^{-e}$. Hence H is the graph T_4 of Figure 1.

Next, suppose that B does not have a red edge. Define H_i^+ as before, and define $\bar{H}_i$ as H_i^+ with the new edge not colored red. If both H_1^+ and H_2^+ have two red edges on a common face, then so does some embedding of H. Hence at least one of the H_i^+ do not have two red edges on a common face. Without loss of generality this is H_1^+, and by minimality H_1^+ and $\bar{H}_2$ both have exactly two red edges. Now, $\bar{H}_2$ must have the two-red-edge property. Once more by minimality, both H_1^+ and $\bar{H}_2$ must be $K_{3,3}^{-e}$. Hence H is the graph T_5 of Figure 1.

Case 3: *$\kappa'=3$.* First note that H has at least six faces. We will show that H has exactly six faces. Call an edge e of H *reducible* if deleting e and supressing the resulting degree-two vertices results in a 3-connected graph.

Claim: *For any face f there is a reducible edge on the boundary of f.*

Proof of Claim: Thomassen ([8], page 46) has shown that any non-cubic graph contains a reducible edge. Our proof of this related result is similar in nature to his. Consider the dual statement: a triangulation contains a contractible edge. (An edge is *contractible* if and only if it lies on no non-facial triangles; after contracting an edge we remove the resulting degree-two vertices). We have to show that every vertex v is incident with a contractible edge. If v is not incident with a non-facial separating triangle, then every edge incident with v is contractible. Pick a non-facial separating triangle T incident with v enclosing the smallest area. Any edge e in the interior of T is not on a separating non-facial triangle, and so is reducible. The claim follows.

Having proved the claim, suppose that H has at least seven faces. Then there exists a face f that is not incident with a red edge. By the claim, there is an edge e in the boundary of f such that $H-e$ is three-connected. Because f was not incident with a red edge, and because each embedding of $H-e$ corresponds to an embedding of H, $H-e$ has the three-red-edge property. This contradicts minimality. We conclude that H has exactly six faces.

There are only two planar cubic three-connected graphs with six faces: the cube and a graph made from K_4 by "blowing up" two vertices into triangles. The cube has a unique way to pick three edges with no two on a common face. This gives the graph T_6 of Figure 1. The other graph has three non-isomorphic ways of selecting three red-edges with no two on a common face. These give the graphs T_7, T_8 and T_9 of Figure 1.

These three cases complete the proof of the proposition. □

Our primary application of Proposition 4.2 is the following.

Corollary 4.3. *There are exactly 4 three-edge-connected obstructions G to near-planarity with a non-trivial three-edge-cut where both sides are non-planar, but one side has the three-red-edge property. These are the graphs given in Figure 6. There are exactly 5 obstructions G to near-planarity that contain a vertex v with $G-v$ planar. These are the graphs given in Figures 4 and 11.*

Proof. In either case stated in the corollary, there is a three-edge-cut B of G so that one component H of $G - B$ has the three-red-edge property. By Lemma 4.1 either $H + K_{2,3}$ or $H + v$ is non-nearly-planar (and hence all of G) depending on whether or not there exists an edge e with $H - e$ having all three red edges on a common face. Hence G is contain in either $H + K_{2,3}$ or $H + v$ respectively (in all cases G is exactly these graphs).

The proof proceeds by examining the nine graphs of Figure 1 and looking for the desired edge e. Exactly four of the graphs have such an edge: T_4, T_5, T_7, T_8. In each case $T_i + K_{2,3}$ is an obstruction for near-planarity. These are respectively the four graphs 16.6, 16.2, 16.5, and 16.3 of Figure 6. Two more graphs T_6 and T_7 do not have such an edge. Here $T_v + v$ is an obstruction for near-planarity. These are respectively the graphs 12.3 and 12.1 of Figure 11. (These graphs will arise again in a later part of the argument.) Finally, three of the graphs T_1, T_2, T_3 have $T_i + v$ an obstruction to near-planarity with edge-connectivity two. These are respectively graphs 14.1, 14.9, and 14.8 of Figure 4.

A nice feature of this result is that the nine minimal graphs with the three-red-edge property correspond exactly with nine obstructions to near-planarity. □

We next turn to nearly-planar obstructions with an edge-cut B of size three, but where neither side of $G - B$ has the three-red-edge property. We establish some terminology.

Let e_a, e_b, e_c be the three cut edges, let H_1, H_2 the two components of $G - B$, and let a_i, b_i, c_i be the *colored* topological edges of H_i, that is, those containing the ends of e_a, e_b, e_c respectively. We'll say two colored edges are a *facial pair* if there exists an embedding of H_i with these two edges on a common face. Because B does not have the three-red-property, each H_i has a facial pair. We study possible facial pairs more closely using the following definition.

Consider a graph H with two distinguished *red* edges and one distinguished *blue* edge. We also allow one edge to be colored red twice, in other words, the two red edges are not necessarily distinct. Then H has the *red-and-blue-edge property* if no embedding of H has a red edge and the blue edge on a common face. The two red edges are allowed to lie on a common face, provided that that face does not contain the blue edge. We characterize minimal graphs with the red-and-blue-edge property in the following lemma.

Lemma 4.4. *A graph that is minimal with the red-and-blue-edge property is one of the four graphs of Figure 2. The red edges are those with one (or two) filled-in circles; the blue edge is marked with an open square.*

Proof. Fix one red edge a_1 and let c_1 be the blue edge. Because a_1 and c_1 are not a facial pair, H must contain a $K_{3,3}^{-e}$. If the two red edges coincide, then we have the graph A in the Figure 2. The second red edge b_1 must lie in one of the two faces on either side of a_1. There are exactly three ways to add a chord that does not re-embed on a face with c_1 and that not contain the graph A. These are the graphs B, C, and D of the figure. □

We return to our analysis of obstructions to near-planarity.

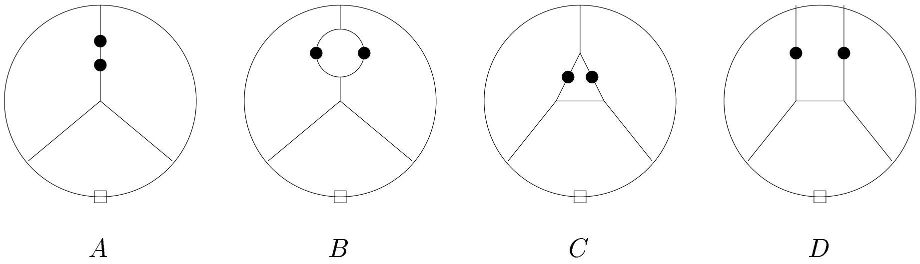

FIGURE 2. Minimal Graphs with the Red-and-Blue-Edge Property

Proposition 4.5. *Let G be a cubic obstruction to near-planarity with cyclic-edge-connectivity three. Suppose that no cyclic cut has a non-planar component or a component with the three-red-edge property. Then G is the graph of Figure 7.*

Proof. Let $\{e_a, e_b, e_c\}$ be a three-edge-cut with components H_1, H_2; label the ends of each e_a in H_i as before. By assumption H_1 has a facial pair, say it is (a_1, b_1). Note that (a_2, b_2) is not a facial pair in H_2, or else $G - e_c$ is planar. If we color a_1, b_1 red and c_1 blue, then H_1 has the red-and-blue-edge property. Hence H_1 contains one of $\{A, B, C, D\}$ in Figure 2.

Also color a_2, b_2 red and c_2 blue in H_2. Now, H_2 doesn't have the red-and-blue-edge property, but instead has the property that the two red edges do not appear on a common face. There are two minimal graphs with this property. The first we'll call X, which is a $K_{3,3}^{-e}$ where the blue edge and one red edge are the same. The second, Y, is also a $K_{3,3}^{-e}$ with the red edges a matching, but now the blue edge is distinct.

The key observation is that any one of A, B, C, D together with any one of X, Y and the bond joining them has no essential edges. Because G is minimal, it must be one of these 8 possible graphs. We examine each in turn.

First suppose that G is A together with either X or Y. This A has another three-edge-cut: use the two segments at the ends of the red edge and e_c. This cut has H_2 non-planar, contrary to our assumption. Likewise, B with either X or Y has another 3-edge-cut using the two edges that disconnect the red edges from the rest of B. Again this has a non-planar component, contrary to our assumption. Either C with X or D with X has a non-planar H_1 using a 3-edge-cut involving the ends of the common red/blue edge in X. Next, C together with Y has no other 3-edge-cuts. However, it has a redundant edge whose deletion yields graph 14.4 of Figure 5. This leaves only D together with Y, which is the desired graph of Figure 7. □

By the results of Sections 3 and 4, we can assume that the remaining cubic obstructions for near-planarity are cyclically 4-edge-connected.

5. Overlap Graphs

In this section we examine bridges of a graph with respect to a cycle and their relation with planarity. The basic approach is related to the overlap graph introduced by Tutte [18], we include it here for completeness.

Throughout this section G will be a fixed graph with a fixed cycle C. We assume that the cycle C has some distinguished degree-two vertices that divide C into paths called the *sides of* C.

A *θ-graph* is a subgraph homeomorphic to $K_{2,3}$. Each foot of a (G,C)-bridge lies in a unique side of C. The pair (G,C) is *θ-less* if there does not exist a θ-graph of G disjoint from $C-S$, where S is a side of C.

There are two parts to this section. In the first part we consider only embeddings of G where all of the (G,C)-bridges have to go on the same side of C (a *disk embedding*). In the second part we allow (G,C)-bridges to embed on either side of C (a *planar embedding*).

A *disk embedding* of the pair (G,C) is an embedding of G in a disk such that C is on the boundary of the disk. The pair (G,C) is *disk-critical* if G has no disk embedding, but $G-e$ has a disk embedding for every edge $e \in G-C$. A (G,C) bridge is *solid* if it has feet in at least two sides of C. Two bridges B_1 and B_2 are *skew* if there exists distinct feet u_i, v_i of B_i that appear in cyclic order u_1, u_2, v_1, v_2 along C.

Lemma 5.1. *Let (G,C) be a disk-critical θ-less pair with every bridge solid. Then there are exactly two different (G,C)-bridges, they are skew, and both bridges are a single solid edge. In particular, there are exactly four vertices in G of degree three.*

Proof. We begin with the case that there is a single (G,C)-bridge B, so that $C \cup B$ is non-planar. If B has at most three feet, then there is a θ-graph disjoint from C in contradiction of the hypothesis. Hence B has at least four feet; say these are v_1, v_2, v_3, v_4 appearing in that order along C.

Chose P to be a shortest v_1v_3-path in B. Let $B_1, \ldots, B_k$ be the $(G, C \cup P)$-bridges. The half of $C - \{v_1, v_3\}$ that contains v_2 will be called P_2 and the half containing v_4 will be P_4. Each B_i must contain a foot in either P_2 or P_4, or else we either have parallel edges or we contradict the hypothesis that (G,C) is θ-less.

Suppose that some B_i has feet in both P_2 and P_4. Then there is a second path P' in B_i joining without loss of generality v_2 and v_4. Hence the original bridge B contains an H-tree on 5 vertices, where the two vertices adjacent to one of the cubic vertices in the H are skew with the two vertices adjacent to the other cubic vertex of H. The two vertices adjacent to one of the cubic vertices are not feet on a common side S of C, or else there exists a θ-graph disjoint from $C-S$. Hence if we delete the edge e joining these two degree-two vertices, then the pair $(G-e, C)$ satisfies the other hypotheses of this lemma. This contradicts minimality.

From the last three paragraphs we conclude that there are at least two (G,C)-bridges B_1, B_2, and that $C \cup B_i$ is planar for each i. Because G is non-planar, there must exist a single pair of skew bridges B_1, B_2. We will show that these must each have feet in different sides of C, that is, that they are solid.

Let u_1, v_1 and u_2, v_2 be feet of B_1 and B_2 respectively such that they occur in cyclic order $u_1u_2v_1v_2$ around C. These four vertices divide C into four paths which we will denote (u_1,u_2), (u_2,v_1), (v_1,v_2), and (v_2,u_1). By way of contradiction, suppose that u_1 and v_1 lie in the same side S of C. By supposition, B_1 has a foot w_1 not on S. If w_1 lies in (v_1,v_2), then we can delete the leg of B_1 incident with v_1 and still have skew bridges. If w_1 lies on (v_2,u_1), then we can similarly delete the leg incident with u_1.

The conclusion of the lemma has been shown. $\square$

We next turn our attention to embeddings where bridges can go on either side of C, that is, planar instead of disk embeddings.

The *overlap graph* $L = L(G,C)$ has vertices the (G,C)-bridges, with edges joining skew bridges. The original concept is due to Tutte [18], but in the more general catagory of non-cubic graphs. Tutte showed the following.

Theorem 5.2. *A graph G is planar if and only if for every cycle C, the overlap graph $L(G,C)$ is bipartite.* □

Lemma 5.3. *Let G be a non-planar graph containing a fixed cycle C. Suppose that there are at least two (G,C)-bridges, and that for any leg e of any bridge B, $G-e$ is planar. Then*

(1) *there is an odd number of bridges $B_1, \ldots, B_k$,*
(2) *each B_i is skew to B_{i-1} and to B_{i+1} but to no other bridge (subscripts are read cyclically), and*
(3) *each B_i has exactly two feet.*

Proof. We first observe that $C \cup B$ is planar for every bridge B, since this graph is contained in $G-e$ where e is a leg of another bridge. We form the overlap graph $L = L(G,C)$. If L is bipartite, then as in Tutte's Theorem, G is planar (we use the bipartition to say whether bridges go inside or outside of C). Hence L contains an odd cycle. Let O be the shortest odd cycle of L. Observe that O is chordless, as any chord makes a shorter odd cycle. Also note that if there were a bridge B that is not a vertex of O, then for any leg e of B, we have the contradiction that $G-e$ is still non-planar. We conclude that the overlap graph is exactly a simple odd cycle, which gives Conclusions 1 and 2 of this lemma.

Next, suppose that there is a bridge with $|B_i| \geq 3$, where $|B_i|$ denotes the number of feet (or equivalently legs) of B_i. Let S_{i+1} denote the set of legs of B_i such that $B_i - e$ is still skew to B_{i+i} for any $e \in S_{i+1}$. Define S_{i-1} similarly. It is easy to see that $|S_{i+1}|$ and $|S_{i-1}|$ are both of cardinality at least $|B_i|-1$. Since $|B_i|$ is at least three, there is a leg $e \in S_{i-1} \cap S_{i+1}$. Deleting this leg still leaves a non-planar graph, a contradiction. Hence we have shown Conclusion 3 of this lemma. □

We use this lemma only for the following corollary, whose proof now follows immediately.

Corollary 5.4. *Under the hypotheses of the preceding lemma, there is a $K_{3,3}^{-v}$-graph contained in $G-B$ for every bridge B of a fixed cycle C.* □

6. Extending Planar Embeddings

In this section we use the results of the previous section to restrict (under certain conditions) the way a planar embedding of a subgraph can extend to a planar embedding of the whole graph. These results will be used in the next section to bound the size of a nearly-planar graph.

Let G be a graph with a subgraph H. Suppose that H is a three-connected planar graph and that G is non-planar. The pair (G,H) is *critical* if for every edge $e \in G-H$, $G-e$ is planar (there is no such restriction on planarity for $e \in H$). Recall that a pair is *θ-less* if no edge $e \in H$ has a θ-graph disjoint from $H-e$. The pair is *minimal* if every subgraph K of G that is homeomorphic to H has

$|V(G) \cup K| \geq |V(G) \cup H|$, that is, if we have the topological copy of H in G that has the fewest possible number of vertices in G.

In any embedding of a three-connected graph H there is at most one face containing two topological edges of H. Recall that a (G, H)-bridge is *solid* if it has feet in at least two topological edges of H. One of our goals is to restrict where a (G, H)-bridge B can go when extending the unique planar embedding of H by ensuring that B is solid. The following lemma achieves this.

Lemma 6.1. *Suppose that (G, H) is a critical minimal θ-less pair. Then every (G, H)-bridge is solid.*

Proof. Let B be a bridge with at least three feet. Then B is solid unless all of its feet lie on a common topological edge e of H. But in this case, there is a θ-graph disjoint from $H - e$, contradicting θ-less.

Let B be a bridge with exactly two feet. By minimality, B consists of a single edge e. If B is not solid, there are either parallel edges in G, contradicting criticalness, or we can replace a portion of e with B, contradicting minimality. □

We turn our attention to ensuring that a pair is θ-less.

Lemma 6.2. *Suppose that (G, H) is a minimal critical pair that has a θ-graph T disjoint from $H - e$ for some edge e (i.e., (G, H) is not θ-less). Then there is a single bridge B, the one containing T. Moreover,*

$$|V(G)| \leq |V(H)| + 8.$$

Proof. We first examine the case that T is a $K_{3,3}^{-v}$-graph. Consider H together with the $(H - e)$-bridge B containing the $K_{3,3}^{-v}$ graph. Then $H \cup B$ is non-planar; by minimality it is all of G. But adding G to H increases the number of vertices of H by at most 8: the five topological vertices of the $K_{3,3}^{-v}$ and the three feet of B.

We next show that a θ-graph T disjoint from $H - e$ implies that there is a $K_{3,3}^{-v}$-graph disjoint from $H - e$, so that the arguments of the previous paragraph apply. This supposed T contains three cycles. At least one of the cycles C must have more than one bridge, or else T is a $K_{3,3}^{-v}$-graph as desired. By criticality, each leg e of a (G, C)-bridge has $G - e$ planar. Hence (G, C) satisifies the hypotheses of Lemma 5.3. By Corollary 5.4 there is a $K_{3,3}^{-v}$-graph disjoint from $H - B$, where B is the (G, C) bridge containing $H - e$. The conclusion follows. □

We begin our examination of θ-less pairs with the following lemma.

Lemma 6.3. *Suppose that (G, H) is a critical minimal θ-less pair. Then either*

(1) *there is exactly one bridge B consisting of a single edge whose ends do not lie on a common face of H, or*
(2) *there are exactly two bridges which both embed in a common face f of H but which are skew.*

In either case, we have $|V(G)| \leq |V(H)| + 4$.

Proof. If there is a (G, H)-bridges B with feet in topological edges e_1, e_2 of H that do not lie on a common face of the planar embedding of H, then $H \cup B$ is non-planar. By criticality, B is a single edge, $G = H \cup B$, and so Conclusion 1 holds. Note that $|V(G)| \leq |V(H)| + 2$ as desired.

Next suppose that for each bridge B the feet of B are lie in a common face $F = F(B)$. By Lemma 6.1 each bridge is solid, so the choice of F is unique. Partition the (G, H)-bridges according to their corresponding face. For each face, consider the graph G_F consisting of the boundary of G and all bridges corresponding to F. Then G is planar if and only if every G_F has a disk embedding. It follows that at least one $G_{F'}$ does not have a disk embedding. By Lemma 5.1 there are exactly two bridges corresponding to F' and they are skew solid edges. By criticality there are no bridges corresponding to faces other than F'. Hence Conclusion 2 follows. □

7. Obstructions with Poor K-Structures

In this section we complete our analysis of obstructions to near-planarity that have a K-structure. There are three main remaining cases: $(K_{3,3}^{-e}, K_{3,3}^{-e})$-, $(K_{3,3}^{-e}, K_{3,3}^{-v})$-, and $(K_{3,3}^{-v}, K_{3,3}^{-v})$-structures. We examine them in Propositions 7.1, 7.2, and 7.3 respectively.

Proposition 7.1. *There does not exist a cubic obstruction to near-planarity that is cyclically-4-edge-connected and that has a $(K_{3,3}^{-e}, K_{3,3}^{-e})$-structure (but no richer K-structure).*

Proof. We will actually show that any such G has at most 24 vertices. We then conclude the non-existence of G by an exhaustive search.

We start with the hypothesised K-structure: the two subgraphs H_1 and H_2. By assumption each H_i completes to a $K_{3,3}^t$. Hence there two arcs for H_1 to H_2 needed to complete H_1 to $K_{3,3}^t$. If the two arcs intersect in $G - H_2$, then G has a $(K_{3,3}^t, K_{3,3}^{-e})$-structure contrary to our assumption. Similarly select another two arcs (not necessarily distinct) to complete H_2. By 4-edge-connectivity there are at least four distinct arcs from H_1 to H_2. If any of the previously selected arcs coincide, select enough distinct arcs from H_1 to H_2 to ensure cyclically 4-edge-connectivity. Let K be the resulting graph. Select K as the graph between H and G with a minimal number of edges in $K_{3,3}^t \subset K$.

If every edge of K is essential in K, then $G = K$ and $|V(G)| \leq 16$ as desired. There are at most two non-essential edges. Let e_1 be one such edge, and if necessary let e_2 be the other such edge. Set $H = K - e_1$. Then H is a three-connected planar graph with $|V(H)| = 14$. Note that every edge in $G - H$ except for e_1 is essential, and every edge of H except for possible e_2 is also essential. Let $\bar{H}$ graph with $H \subset \bar{H} \subset G$ such that $(H, \bar{H})$ is a critical pair. Renaming H if necessary, we can assume that this pair is minimal.

We first consider the case that $(\bar{H}, H)$ is not θ-less. Let T be a θ-graph disjoint from $H - e$ (where possible $e = e_2$). By Lemma 6.2, $|V(\bar{H})| \leq |V(H)| + 8$. If $e \neq e_2$, then every edge in $\bar{H}$ is essential. Hence $G = \bar{H} \cup e_1$ and $|V(G)| \leq 14 + 8 + 2 = 24$ as desired. If $e = e_2$, then we consider the $(H_1 \cup H_2)$-bridge B containing T. By the arguments in Lemma 6.2 T is a $K_{3,3}^{-v}$-graph. Two of the paths P_1, P_2 completing this $K_{3,3}^{-v}$ to a $K_{3,3}^t$ must connect to the same H_i, say H_1. But now $(T \cup P_1 \cup P_2, H_2)$ is a $(K_{3,3}^{-e}, K_{3,3}^{-e})$-structure with a K having only one essential edge, contradicting our choice of K.

Next consider the case that $(\bar{H}, H)$ is θ-less. Then by Lemma 6.3, $|V(\bar{H})| \leq |V(H)| + 4 = 18$. If e_1 were the only essential edge, then $|V(G)| \leq 18 + 2 = 20$ as desired. However, if H contained a second essential edge e_2, then it is no longer essential in $\bar{H}$ unless at least two feet of $(\bar{H}, H)$-bridges lie in e_2. But in this case we have a different homeomorph K in $\bar{H}$. This has fewer essential edges, so it again contradicts our choice of K.

We have shown that any such G as hypothesised in the proposition has at most 24 vertices. The proof is completed by a computer search on all cubic graphs with at most 24 vertices to see if any are obstructions to near-planarity. The details and double-checks of the algorithm are presented in Section 9. □

Proposition 7.2. *Let G be a cyclically 4-edge-connected cubic obstruction to near-planarity with a $(K_{3,3}^{-e}, K_{3,3}^{-v})$-structure (but no richer K-structure). Then G is the graph of Figure 8.*

Proof. The proof mimics that of the last proposition. We will actually show that any such G has at most 22 vertices. We then conclude that G is the graph named by an exhaustive search.

We begin again with the hypothesised K-structure: the two subgraphs H_1 and H_2. As before, we select the subgraph K of G that is cyclically-4-edge-connected, contains $H_1 \cup H_2$, and such that the number of edges that are non-essential in K is minimized.

If every edge of K is essential in K, then $G = K$ and $|V(G)| \leq 14$ as desired (observe that this is where the graph of Figure 8 arises). There are at most two non-essential edges. Let e_1 be one such edge, and if necessary let e_2 be the other such edge. Set $H = K - e_1$. Then H is a three-connected planar graph with $|V(H)| = 14$. As before, every edge in $G - H$ except for e_1 is essential, and every edge of H except for possible e_2 is also essential. Let $\bar{H}$ graph with $H \subset \bar{H} \subset G$ such that $(H, \bar{H})$ is a critical pair. Renaming H if necessary, we can assume that this pair is minimal.

We first consider the case that $(\bar{H}, H)$ is not θ-less. Let T be a θ-graph disjoint from $H - e$ (where possible $e = e_2$). By Lemma 6.2 $|V(\bar{H})| \leq |V(H)| + 8$. If $e \neq e_2$, then every edge in $\bar{H}$ is essential. Hence $G = \bar{H} \cup e_1$ and $|V(G)| \leq 12 + 8 + 2 = 22$ as desired. If $e = e_2$ we consider the $(H_1 \cup H_2)$-bridge B containing T. As before, by the arguments in Lemma 6.2 T is a $K_{3,3}^{-v}$-graph. Two of the paths P_1, P_2 completing this $K_{3,3}^{-v}$ to a $K_{3,3}^{t}$ must connect to the same H_i, say H_1. But now $(T \cup P_1 \cup P_2, H_2)$ is a (possibly richer) K-structure with a K having only one essential edge, a contradiction.

Next consider the case that $(\bar{H}, H)$ is θ-less. Then by Lemma 6.3, $|V(\bar{H})| \leq |V(H)| + 4 = 16$. If e_1 were the only essential edge, then $|V(G)| \leq 16 + 2 = 18$ as desired. However, if H contained a second essential edge e_2, then it is no longer essential in $\bar{H}$ unless at least two feet of $(\bar{H}, H)$-bridges lie in e_2. But in this case we have a homeomorph of K in $\bar{H}$. Since this has fewer essential edges, it again contradicts our choice of K.

We have shown that any such G as hypothesised in the proposition has at most 22 vertices. The proof is completed by a computer search on all cubic graphs with at most 22 vertices to see if any are obstructions to near-planarity. The details and double-checks of the algorithm are presented in Section 9. □

Proposition 7.3. *Let G be a cyclically 4-edge-connected cubic obstruction to near-planarity with a $(K_{3,3}^{-v}, K_{3,3}^{-v})$-structure (but no richer K-structure). Then G is one of the graphs of Figure 9.*

Proof. The proof mimics the proofs of the last two propositions. We will actually show that any such G has at most 20 vertices. We then conclude that G is one of the graphs claimed by an exhaustive search of all graphs on 20 vertices.

We again begin with the hypothesised K-structure: the two subgraphs H_1 and H_2 and a subgraph K of G with the aforementioned properties.

If every edge of K is essential in K, then $G = K$ and $|V(G)| \leq 12$ as desired. There are at most three non-essential edges. Let e_1 be one such edge, and if necessary let e_2, e_3 be the other such edges. Three of the edges joining H_1 to H_2 must form a bijection between the topological edge of H_1 with those of H_2. Label the topological edges A_i, B_i, C_i such that each X_i is in H_i (where $X \in \{A, B, C\}$) and each X_1 is bijectively matched with X_2.

First, suppose that an edge of K connects say A_1 to B_2. Then the only remaining essential edge is one joining C_1 to C_2. The analysis proceeds as in the previous propositions, yielding $|V(G)| \leq 12+8 = 20$. (This is where the first graph of Figure 9 arises.)

Second, suppose that all remaining arcs from H_1 to H_2 join X_1 to X_2. As soon as we get corresponding arcs between two corresponding edges of H_i, then both are non-essential. Hence there are at most three additionial such arcs, and $|V(G)| \leq 16$. (This is where the second graph of Figure 9 arises.)

As before, we have bounded the number of vertices in G. The proof is completed by an exhaustive search. □

We have finished the three propositions covering the three remaining K-structures, and hence have found all obstructions to near-planarity under the cubic order that contain a K-structure.

8. Projective Planarity and the Main Proof

In this section we give the proof of our Main Theorem 1.5. The proof uses an interesting connection between nearly-planar graphs and graphs that embed in the projective plane. We need a preliminary definition.

A *projective-planar map* is a fixed embedding of a graph in the projective plane. The *face-width* of a projective-planar map G is the minimum $|C \cup G|$ taken over all non-contractible cycles C.

Lemma 8.1. *A cubic projective-planar graph is nearly-planar if and only if it has face-width at most two.*

Proof. Fiedler, Huneke, Richter, and Robertson [5] have shown the remarkable result that the orientable genus of a graph embedded in the projective-plane with face-width $w \geq 3$ is $\lfloor w/2 \rfloor$. Note that the face-width can vary between different projective-plane embeddings, but not by more than one if one embedding has face-width at least three. The result now follows. □

It follows from this lemma that the obstructions to near-planarity that embed in the projective plane do so with face-width three, and are minimal with that property. It is known [19] that there are exactly six minor-minimal graphs for the property of having face-width at least three in the projective plane. These are six of the seven members of the Petersen Family (the remaining member of the Petersen Family is non-projective-planar). Any topological obstruction to this property in the cubic order must be a splitting of one of these graphs. It is straightforward to check their possible splittings and conclude the following.

Proposition 8.2. *Let G be cubic obstruction to near-planarity that embeds in the projective plane. Then G is one of the graphs of Figure 10.*

Proof. The proof follows from the preceeding comment. However, this has also been noted in the literature as we will now describe. Any graph with crossing number one is nearly-planar. Hence any obstruction has crossing number at least 2. Any such graph that embeds in the projective plane must do so with face-width at least three. McQuillan and Richter ([12], see also [14]) have shown that there are exactly two minimal cubic graphs that embed in the projective plane with face-width at least three. These are the two graphs claimed. □

We are now ready for our main result.

Proof of the Main Theorem 1.5: Let G be an obstruction to near-planarity under the cubic order.

If G embeds in the projective plane, then by Proposition 8.2 this is one of the graphs of Figure 10.

If G does not embed in the projective plane, then it must contain a cubic obstruction for non-projective-planarity. Thanks to Glover and Huneke [6] this set is known: it contains exactly six graphs. (See [1, 2] for the non-cubic case, from which this result also follows.)

We examine these six graphs in turn. The first graph is the disjoint union of two $K_{3,3}^t$ as shown in Figure 3, and so is already on our list of near-planar obstructions. This is the graph F_{42} in the notation of [7]. The next two graphs are shown in Figure 11, and so are already on our list of near-planar obstructions. They are F_{12} and F_{13} in the notation of [7]. The fourth graph (F_{14} in [7]) is minimal non-projective-planar, but is not minimal non-nearly-planar: it has an edge whose deletion gives a homeomorph of the Petersen graph. The remaining two graphs (F_{11} and G_1 in [7]) have a $(K_{3,3}^{-e}, K_{3,3}^{-e})$- and a $(K_{3,3}^{-v}, K_{3,3}^{-v})$-structure respectively. Hence any non-projective obstruction to near-planarity is either on our list or has a K-structure.

In Sections 3 and 4 we found all cubic obstructions to near-planarity that have cyclic-edge-connectivity at most three or that contain one of three specific K-structures. In Section 7 we found all cubic obstructions to near-planarity that are cyclically-4-edge-connected and have one of the remaining three K-structures. This completes the proof that our list of obstructions is complete. □

In Figures 3 through 11 we give the complete set of 21 obstructions to near-planarity under the cubic order. By Proposition 1.4 this is also the complete set of 21 obstructions to embedding on the spindle surface under the cubic order.

9. The Algorithm and our Double Checks

The main result of this paper relies on an exhaustive computer search of cubic graphs of small order. We give here several double-checks that this search was done successfully.

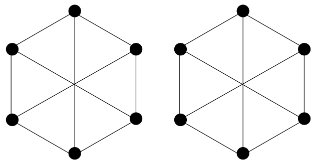

$2K_{3,3}$

FIGURE 3. Obstruction with a $(K^t_{3,3}, K^t_{3,3})$-Structure

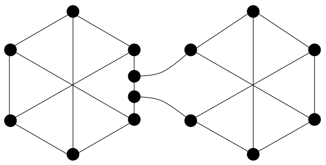

14.1

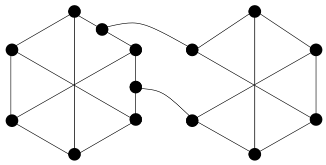

14.8

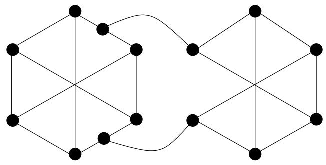

14.9

FIGURE 4. Obstructions with a $(K^t_{3,3}, K^{-e}_{3,3})$-Structure

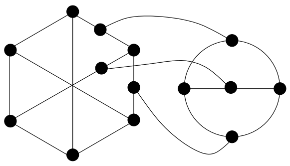

14.2

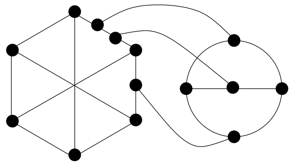

14.5

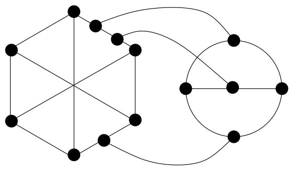

14.4

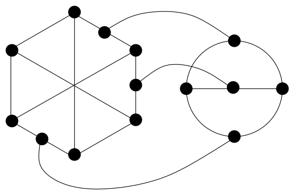

14.6

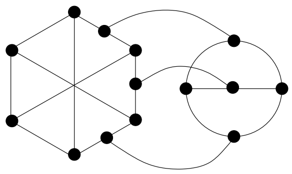

14.7

FIGURE 5. Obstructions with a $(K^t_{3,3}, K^{-v}_{3,3})$-Structure

We begin by describing the algorithm used for determining which small cubic graphs are obstructions for near-planarity. The first program used is the well-tested planarity algorithm for graphs included in the LEDA package [11]. Using this program, it is easy to make a program to determine if a graph is nearly-planar: we merely test each subgraph $G - e$ to determine if one is planar. Similarly, it is easy to make a program to test for an obstruction for near planarity: merely test that a given graph G is not nearly-planar, but that all subgraphs $G - e$ are.

To check all small cubic graphs we generated all cubic graphs using Brendan McKay's Nauty and Gtools packages [13]. We generated all cubic graphs on up to

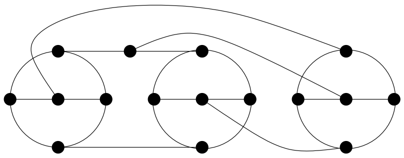

16.6

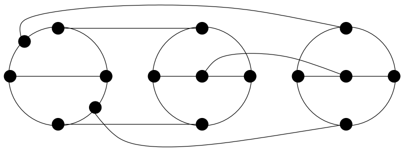

16.2

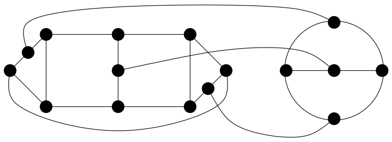

16.5

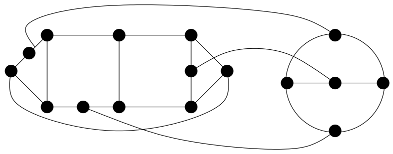

16.3

FIGURE 6. Obstructions from the Three-Red-Edge Property

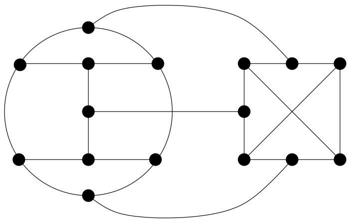

16.4

FIGURE 7. Obstruction from the Red-and-Blue-Edge Property

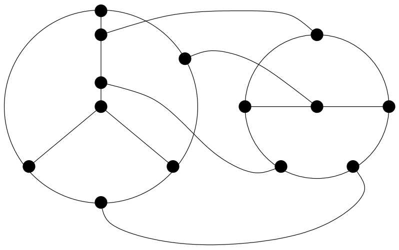

14.10

FIGURE 8. Obstruction with a $(K_{3,3}^{-e}, K_{3,3}^{-v})$-Structure

24 vertices and checked each to see if it was an obstruction to near-planarity. The peculiar numbering of the graphs reflects the number of vertices and the order in which they are generated by Gtools.

In Sections 3 and 4 we were careful to complete the analysis without using a computer search. The graphs that we found with these structures agreed with the

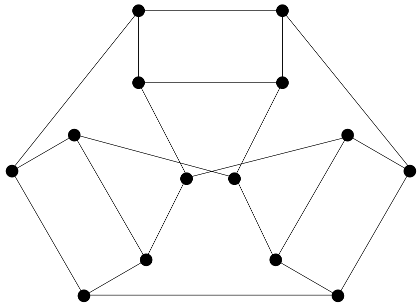

14.3

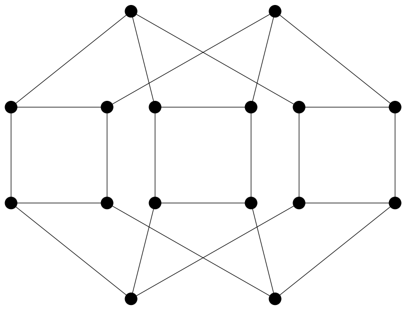

16.1

FIGURE 9. Obstructions with a $(K_{3,3}^{-v}, K_{3,3}^{-v})$-Structure

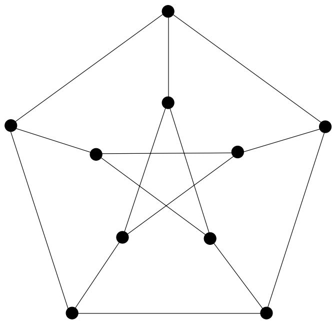

10.1 (Petersen)

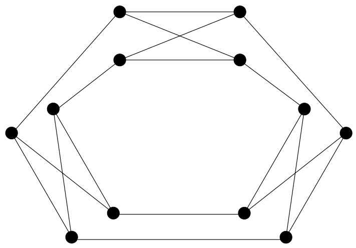

12.2

FIGURE 10. Obstructions that Embed in the Projective-Plane

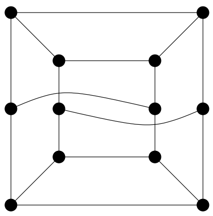

12.1

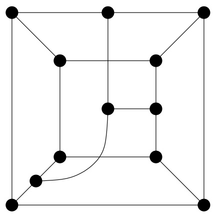

12.3

FIGURE 11. Obstructions that do not Embed in the Projective-Plane

graphs found by our computer search. This is valuable evidence that the search was correctly programmed.

There are several classes of nearly-planar obstructions that are easy to find independently, for example those with small edge connectivity. We had also found ten or so examples of obstructions to near-planarity by hand. In each case our hand calculations agreed with the computer search.

Finally, we checked each of the 21 computer-generated graphs by hand to verify that each was an obstruction to near-planarity.

Perhaps the most compelling evidence of the correctness of our results is the fact that there were no obstructions found for graphs of order 18 through 24. It is difficult to believe that there may be a larger example, given that there are none of these orders.

10. Conclusion

Most of the arguments given in this paper were geared towards finding the cubic obstruction set for near-planarity. Some of the arguments hold in greater generality. We mention a few cases.

The techniques of Proposition 4.2 generalize nicely. Say that a cubic planar G has the k-red-edge property if it has k red edges and no planar embedding of G has any two red edges on a common face. It's possible to show that minimal graphs with the k-red-edge property are either built along 1- or 2-edge cuts from examples with the k'-red-edge property, $k' < k$, or have exactly $2k$ faces. In particular, the following holds.

Proposition 10.1. *Let G be a minimal cubic planar graph with the k-red-edge property, and let m be the number of cut-edges of G. Then $|V(G)| = 4(k-1) + 2m$.* ☐

The material of Section 5 on overlap graphs can be generalized to the non-cubic case fairly easily. We leave the details for another day.

The material in Section 6 on extending embeddings could be used for other embedding problems. We worked hard to guarantee that our embeddings were three-connected. Thus, when we were dealing with solid bridges (those with feet in two different topological edges), there was at most one face where that bridge could lie. The same idea would work for any H embedded in any surface with face-width three.

To a large extent the techniques used in this paper are common for such structure theorems. If you are only interested in proving that a set is finite, then you can often apply cruder bounds. If you don't use a computer search the casework analysis in Section 7 is daunting. A combination of a careful analysis of bounds and a computer search can be useful.

A natural question is to determine the obstruction set for embedding (possibly non-cubic) graphs in the spindle surface. The authors feel that finding this set may be within reach using the techniques developed herein. We note the work of [4] who under a finer partial order have found some of these obstructions. However, it is known that their set is not complete.

Finally, we have not discussed the algorithmic implications of our work. It is easy to create a quadratic-time algorithm to embed a cubic graph in the spindle surface: merely delete each edge in turn and run a linear-time planarity algorithm on the resulting graph. Is it possible to adapt these techniques to give a linear-time algorithm to embed a cubic graph in the spindle surface?

Acknowledgement: The first author would like to thank the University of Auckland, which hosted him during the time this research was conducted. Both authors are thankful for support from the Marsden Fund (Grant Number UOA825) administered by the Royal Society of New Zealand.

References

[1] D. Archdeacon, A Kuratowski Theorem for the projective plane, *Ph.D. Thesis*, The Ohio State University (1980).
[2] D. Archdeacon, A Kuratowski Theorem for the projective plane, *J. Graph Theory* **5** (1981) 243–246.
[3] D. Archdeacon and J.P. Huneke, A Kuratowski theorem for nonorientable surfaces, *J. Combin. Theory Ser. B* **46-2** (1989) 173–231.
[4] R. Bodendiek and K. Wagner, On the minimal basis of the spindle-surface $\mathcal{S}_1$, *Combinatorics and graph theory*, Banach Center Publ. **25**, Warsaw (1987) 19–31.
[5] J. Fiedler, J.P. Huneke, R.B. Richter, and N. Robertson, Computing the orientable genus of projective planar graphs, *J. Graph Theory* **20-3** (1995) 279–308.
[6] H. Glover and J.P. Huneke, Cubic irreducible graphs for the projective plane, *Discrete Math.* **13-4** (1975) 341–355.
[7] H. Glover, J.P. Huneke, and C.S. Wang, 103 Graphs that are irreducible for the projective plane, *J. Combin. Th. Ser. B* **27-3** (1979) 332–370.
[8] Jonathan Gross and Thomas Tucker, Topological Graph Theory, Wiley, New York (1987)
[9] K. Kuratowski, Sur le probléme des courbes gauches en topologie, *Fund. Math.* **5** (1930) 271–283.
[10] M. Knor, Characterization of minor-closed pseudosurfaces, *Ars Combin.* **43** (1996) 246–256.
[11] LEDA: A library of efficient datatypes and algorithms, Max Planck Institute for Computer Science. `http://www.mpi-sb.mpg.de/LEDA/`
[12] D. McQuillan and R.B. Richter, On 3-regular graphs having crossing number at least 2, *J. Graph Theory* **18-8** (1994) 831–839.
[13] B. McKay, A NAUTY user's guide, Technical report TR-CS-90-02, Computer Science Department, Australian National University. `http://cs.anu.edu.au/people/bdm/nauty/`
[14] R.B. Richter, Cubic graphs with crossing number two, *J. Graph Theory* **12-3** (1988) 363–374.
[15] N. Robertson and P.D. Seymour, Graph minors VIII: A Kuratowski theorem for general surfaces, *J. Combin. Th. Ser. B* **48-2** (1990) 255–288.
[16] N. Robertson and P.D. Seymour, Graph minors XX: Wagner's conjecture, manuscript.
[17] J. Širáň and P. Gvozdjak, Kuratowski-type theorems do not extend to pseudosurfaces, *J. Combin. Th. Ser. B* **54-2** (1992) 209–212.
[18] W.T. Tutte, Graph Theory, *Encyclopedia of Mathematics and its Applications* **21**, Cambridge University Press, Cambridge (2001).
[19] R. Vitray, The 2- and 3-representative graphs on the projective plane, *J. Combin. Th. Ser. B* **54-1** (1992) 1–12.
[20] K. Wagner, Über eine eigenschaft der ebenen complexe, *Math. Ann.* **114** (1937) 570-590.

Department of Mathematics and Statistics, University of Vermont, Burlington, VT, USA 05405
E-mail address: `dan.archdeacon@uvm.edu`

Department of Mathematics, University of Auckland, Auckland, New Zealand
E-mail address: `p.bonnington@auckland.ac.nz`