Halin's Theorem for Cubic Graphs on an Annulus

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

Jozef Širáň
Department of Mathematics
Slovak Technical University
Radlinského 11
2813 68 Bratislava, Slovakia
siran@lux.svf.stuba.sk

June 2, 2001

Abstract

Halin's Theorem characterizes those locally-finite, infinite graphs that embed in the plane without accumulation points by giving a set of six topologically excluded subgraphs. We prove the analogous theorem for cubic graphs that embed in an annulus without accumulation points, finding the complete set of 29 excluded subgraphs.

1 Introduction

A fundamental result in graph theory is Kuratowski's Theorem [11], which says that a finite graph embeds in the plane if and only if it does not contain a subdivision of either K_5 or $K_{3,3}$. We are concerned here with embeddings of infinite graphs. Throughout this paper graphs will be connected and locally finite, hence having a countable vertex set. For the requisite background on embedding these infinite graphs we refer the reader to [4, 13]; similar work is done in [5].

Halin [9] proved that a connected, locally finite, infinite graph embeds in the plane without an accumulation point if and only if it does not contain a subdivision of one of six graphs. An equivalent form of Halin's Theorem is that exclusion of these six subgraphs characterize all graphs that embed on the 2-sphere with at most one accumulation point.

We search for analogues of Halin's Theorem for other 2-dimensional manifolds S. If an infinite graph on S has no accumulation points, then S is necessarily non-compact. A natural manifold to consider is the annulus: the 2-sphere missing two points. A characterization of graphs that embed on this manifold without accumulation points is equivalent to characterizing graphs that embed on the 2-sphere with two accumulation points, or to those that embed on the plane with at most one accumulation point.

We cannot find complete set of obstructions for infinite graphs embedding on the annulus without accumulation points. In our Main Result, we find all such graphs with maximum degree three. They are shown in Figure 1, where we attach disjoint one-way-infinite rays to the vertices of degree two.

Theorem 1.1 The Main Result *A countable locally-finite cubic graph embeds on the annulus without accumulation points if and only it does not contain a subdivision of one of the 29 graphs of Figure 1.*

L. Boza, M. Gugúndez, A. Márquez, and P. Revuelta [5] have a similar result for non-cubic graphs. However, they specify in advance which rays go to which accumulation points, which makes the problem quite different.

The annulus is not the only natural 2-manifold to consider after the plane. Another candidate is the Möbius Strip, formed by deleting a point from the real projective plane. In [3] the authors give a complete characterization of graphs that embed in the Möbius Strip without accumulation points; see also [6]. In Section 4 we discuss this result and its relation with other structural characterizations.

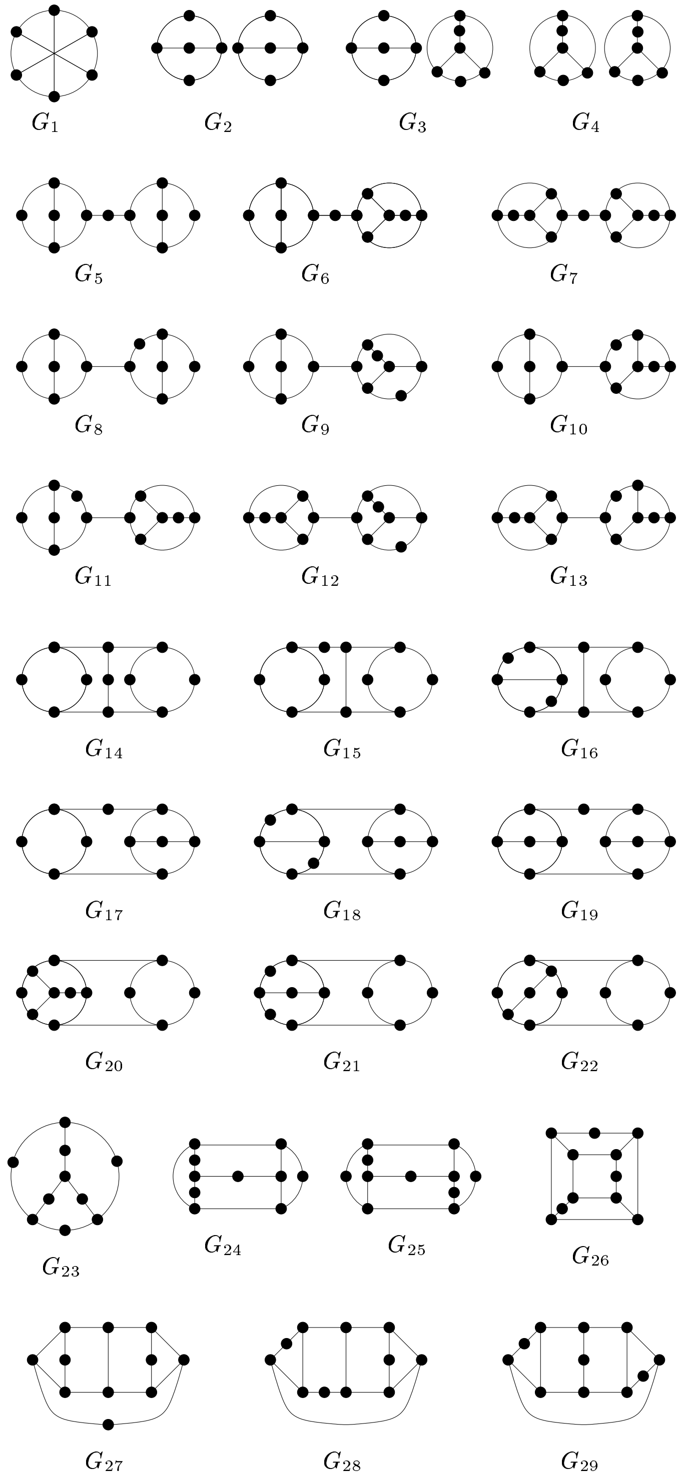

Figure 1: The 29 Obstructions (add rays to degree 2 vertices)

This paper contains four sections. After this introduction, Section 2 introduces *colored graphs*, those with a distinguished set of vertices, and relates embeddings of infinite graphs with a fixed number of accumulation points to embeddings of colored graphs with certain specified face covers. Here we reformulate the Main Result Theorem 1.1 in terms of colored graphs. In Section 3 we give the proofs of main propositions. These propositions are organized by the cyclic edge-connectivity of the graphs: in a nutshell we study (in order) graphs with edge-connectivity exactly 0,1,2, and 3. In Section 4 we combine the results of Section 3 to prove Theorem 1.1 and give some concluding remarks.

2 Colored outer-planar graphs

In this section we discuss the relation between infinite graphs and finite colored graphs. We begin with the general case, where the graphs are not necessarily cubic.

We are concerned with embedding graphs on 2-manifolds. The graphs in question will all be connected and locally finite. It follows that they have a countable vertex set. We will embed these graphs on certain manifolds. We say that a 2-manifold M has a *finite number of ends* if there is a compact subset K of M such that $M - K$ has a finite number of ends. We will only consider embeddings in a manifold with a finite number of ends. Our first task is to relate graph embeddings on these 2-manifolds to embeddings of certain subgraphs. We call a graph H *residually finite* if and only if it contains a finite subgraph K, called the *residue*, and a finite number of one-way infinite rays R, where the rays of R are pairwise disjoint and intersect H only in their endpoints.

The following is due to Bonnington and Richter [4]; see also [5].

Theorem 2.1 *Let G be a connected, locally finite graph. Let S be a surface with a finite number of ends. Then G does not embed on S without accumulation points if and only if G contains contains a residually finite subgraph H that does not embed on S without accumulation points.* ■

The existence of the residually finite subgraph H allows us to focus on only finite graphs. We record the existence of an infinite one-way ray with endpoint $v \in V(H)$ by deleting that ray and coloring the vertex *black*. More formally, a *colored graph* is a graph G together with a distinguished subset of *black* vertices. Vertices not distinguished will be called *white vertices*. The graph G^∞ is formed from a colored graph G by adding pairwise-disjoint one-way-infinite paths to each black vertex of G. It is easy to reverse the construction, taking a residually finite $H = G^\infty$ and recovering the associated finite G.

Let S be a compact surface, and let M be the non-compact surface formed by deleting a finite set of k points from S. Bonnington and Richter [4] also showed the following. (It is important to note that accumulation points in a graph embedding can only occur in the interior of faces, they cannot be points on the graph.)

Theorem 2.2 *An residually finite graph G^∞ embeds on M without accumulation points, if and only if the associated G embeds on S such that all black vertices lie on the boundary of one of k faces.* ■

We say that a colored graph G has a *k-black-cover* on a compact surface S if there is an embedding of G in S with a set of k faces such that every black vertex is incident with at least one of these faces. Using this and the previous two theorems, we can reformulate our main problem of finding the minimal infinite graphs that do not embed on the annulus without accumulation points as follows.

Problem 2.3 *Find all colored graphs that minimally do not embed on the sphere with a 2-black-cover.*

We are primarily interested in graphs with all vertices of degree 2 or 3. By repeatedly *smoothing* degree-two vertices (replacing a path on two edges whose midpoint is a degree-two vertex by a single edge) we reduce down to the case that G is cubic (ignoring the trivial case that the graph has no degree three vertices). Our infinite rays will attach only to degree-two vertices. Instead of distinguishing a black vertex of G^∞, we will distinguish the *red* topological edge of G containing that vertex (we switch colors from black to red to emphasize we're now working with edges instead of vertices). Hence we're working in the category of cubic graphs with coloured vertices, with a distinguished subset of *red* edges.

If we recolor a red edge of G white, then the resulting H is defined to be *smaller* than G. If we delete a red edge e of G, then we can color either one of the two resulting topological edges containing an end of e red, and consider the resulting H as *smaller* than G. The subgraph order, together with smoothing and our two extra operations, form a partial order on all colored cubic graphs. It is this order that we'll consider for the remainder of this paper when using the words "smaller", "contains", or "minimal". The following lemma shows that this order agrees with the topological order on residually finite graphs.

Lemma 2.4 *An edge-colored cubic graph H is smaller than G if and only if the associated infinite graph H^∞ is a topological subgraph of G^∞.* ∎

Infinite graphs G^∞ that embed in the plane without accumulation points correspond to reduced graphs G that embed in the sphere with a 1-red-cover. We'll call these latter *red-outer-planar* graphs. Infinite graphs that embed in the annulus without accumulation points correspond to reduced graphs that embed in the plane with a 2-red-cover.

Two special graphs will be important to us. Let $K_{3,3}^{-e}$ denote $K_{3,3} - K_2$, where the two degree-two vertices are smoothed and the resulting edges colored red. Similarly, form $K_{3,3}^{-v}$; it is a graph with two degree-three vertices and three red edges joining them. The following is immediate by Halin's Theorem [9].

Theorem 2.5 *A edge-colored cubic planar graph G is red-outer-planar if and only if it does not contain $K_{3,3}^{-e}$ or $K_{3,3}^{-v}$.* ∎

As mentioned in the introduction, we cannot find the minimal colored graphs without a 2-black cover. However, we do solve the following problem, which by Lemma 2.4 is equivalent to our main result.

Problem 2.6 *Find the minimal edge-colored* cubic *graphs that do not have a 2-red-cover.*

3 The Main Propositions and Their Proofs

In this section we solve Problem 2.6, although the actual proof is postponed until Section 4. This proof proceeds by a sequence of four propositions organized by the edge-connectivity of G. The first proposition covers non-planar and disconnected graphs.

Proposition 3.1 *Let G be a edge-colored cubic graph that is minimal without a 2-red-cover. If G is non-planar, then G is the graph G_1 of Figure 1. If G is disconnected, then G is one of the graphs G_2, G_3, G_4 of Figure 1.*

Proof: By a result of Erdös (see [8]) a locally-finite graph on a countable vertex set embeds in the plane if and only if it does not contain a $K_{3,3}$ or K_5 subgraph. A non-planar cubic graph must contain $K_{3,3}$, which is G_1 of Figure 1.

If G is disconnected, then by minimality no component can be red-outer-planar. Moreover, if there are two non-red-outer-planar components, then G cannot have a 2-red-cover. The two minimal non-red-outer-planar graphs are $K_{3,3}^{-v}$ and $K_{3,3}^{-e}$. The three ways of combining these two graphs give G_2, G_3, and G_4. ■

We next turn our attention to connected graphs with a cut-edge. We use the following notation. Let e be a cut edge. The components of $G-e$ will be H_1 and H_2. If we smooth the degree two vertex in H_i and color the resulting edge red, then the graph is denoted H_i^r.

Proposition 3.2 *Let G be an edge-colored cubic graph that is minimal without a 2-red-cover. If G is 1-edge-connected, then G is one of the nine graphs G_5-G_{13} of Figure 1.*

Proof: By minimality, each H_i^r has a 2-red-cover. If say H_1^r is red-outer-planar, then we can combine an embedding of H_1^r with all red edges on a common face with an embedding of H_2^r having a 2-red-cover to get an embedding of G with a 2-red-cover. Hence each H_i^r is non-red-outer-planar, and contains either a $K_{3,3}^{-v}$ or a $K_{3,3}^{-e}$.

We break into two cases, depending on whether the cut edge is red or not.

Case 1: *The cut edge is red.* There are three ways to join $K_{3,3}^{-v}$ and $K_{3,3}^{-e}$ by first removing the red color from an edge of each, and adding a red cut-edge between those two edges. These give the three graphs G_5, G_6, G_7. The resulting graphs have no 2-red-cover, and so constitute all of G.

Case 2: *The cut edge is not red.* If both H_i are red-outer-planar, then G has a 2-red-cover. Hence without loss of generality we can assume that H_2 is non-red-outer-planar. It must contain either a $K_{3,3}^{-v}$ or a $K_{3,3}^{-e}$. Form G by removing the red color from one edge in H_1, then adding a non-red edge from that edge in H_1 to any edge in H_2. There are exactly six ways to join

H_2 to H_1^r in this manner. These give graphs G_8-G_{13}. No such graph has a 2-red-cover, and these are minimal with this property. ■

By the preceding proposition we can assume that our graphs G are 2-edge-connected. Suppose that B is a 2-edge cut of G with components H_1 and H_2 of $G - B$. Let u_i, v_i be the ends of B in H_i. We consider the graph $H_i \cup \{u_i v_i\}$. Color the new edge $u_i v_i$ *blue* to readily distinguish it from the edges of H_i. In the literature this blue edge is sometimes called a *virtual* edge.

There are several possibilities for a red color on the virtual edge. We consider $H_i \cup \{u_i v_i\}$ as a subgraph of G. The red and blue colors along the topological edge $u_i v_i$ correspond to a red edge of G and a 2-edge-cut of G. We distinguish the order in which these are encountered along $u_i v_i$ by ordering the two colors along this edge. In particular, if neither edge of B is red, then $u_i v_i$ has color blue but not red on $u_i v_i$. If say the edge incident with u_i is red, then $u_i v_i$ is colored both red and blue with the red color on the end incident with u_i. By minimality, there is at most one red edge in B, so that we'll never assign two red colors to one topological edge. Denote this coloring on $H_i \cup \{u_i v_i\}$ by H_i^b. In our figures, the red edges are formed by suppressing a degree-two vertex and a square indicates the blue edge.

Let H_1^r, H_2^r result from breaking G along a 2-edge-cut as described above.

We need to relate red covers of H_i^b to face covers of the original G. There are several possibilities for a red cover with two faces: neither face is incident with B, exactly one face is so incident, or both faces are so incident. Faces of G incident with an edge in B correspond to faces incident with the blue edge in the corresponding embedding of H_i^b. This motivates the following definition. A graph with a set of red edges and one blue edge has a *blue cover* if the two faces incident with the blue edge together are incident with every red edge.

The following lemma characterizes with certain combinations of red-outer-planarity and blue covers. It refers to the graphs of Figure 2.

Lemma 3.3 *Let G be an edge-colored 2-edge-connected cubic graph that is minimal without 2-red-cover. Let H^b be one of the H_i^b described above.*

i) If H^b is red-outer-planar, then the face covering the red edges is not incident with the blue edge.

ii) H^b contains either R_1 or R_2. If H^b has a blue-cover, then it contains R_1. If H^b is 3-connected and has no blue-cover, then it contains R_2.

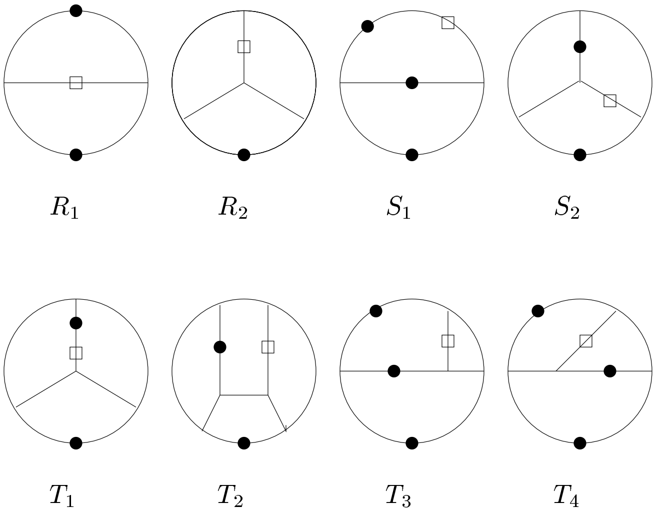

Figure 2: The Graphs of Lemma 3.3

iii) If H^b is not red-outer-planar and has a blue cover, then H^b contains either S_1 or S_2.

iv) Suppose that H^b is minimal with the property that it is not red-outer-planar and has no blue cover. If H^b is 3-edge-connected, then H^b is either $T_1, T_2, T_3,$ or T_4.

Proof: We divide the proof into the four parts given in the statement of the lemma.

Part i) Suppose that H_1^b has a face f_1 incident with all red edges and the blue edge. Let e be any edge of H_1 and find a 2-red-cover of $G - e$. At least one of these faces f_2 is incident with a red vertex in H_i. We can combine this embedding with the supposed embedding of H_i^b so that the two faces f_1 and f_2 merge to one. This gives a 2-red-cover of G, a contradiction.

Part ii) Recolor the blue edge of H^b red and call the result H^r. By Part (i), H^r is non-red-outer-planar, so it contains either $K_{3,3}^{-e}$ or $K_{3,3}^{-v}$. If the blue edge is on a topological edge of $K_{3,3}^{-v}$, then we get R_1. If it is on a red topological edge of $K_{3,3}^{-e}$, then we get R_2. If it is on a non-red edge of $K_{3,3}^{-e}$, then the graph properly contains R_1. Finally, if it is in a bridge of $K_{3,3}^{-v}$ or $K_{3,3}^{-e}$, then that bridge has at least two feet. Any way of selecting these two feet give either a R_1 or R_2 subgraph.

If H^b has a blue cover, then it cannot contain R_2. The second statement now follows from the first.

If H^b is 3-connected and has no blue cover, then let H^b_+ denote the (non-cubic) graph formed by adding in edges from a fixed degree-two vertex in the blue edge to a degree-two vertex in each red edge. Because H^b has no blue cover, this graph is non-planar, and hence contains a $K_{3,3}$. It follows that H^b contains R_2 as desired.

Part iii) Because H^b is non-red-outer-planar, it contains either $K_{3,3}^{-v}$ or $K_{3,3}^{-e}$. If these do not contain the blue edge, then H^b cannot have a blue cover contrary to assumption. There are two possibilities for where this blue edge can be and have a blue cover, giving S_1 and S_2 respectively.

Part iv) By Part (ii), H^b contains R_1. If the blue edge e_b is also red, then H^b is T_1 as desired. So henceforth we assume that e_b is not red. We will first establish the following claim.

Claim: If $H^b - e_b$ is red-outer-planar, then H_b contains T_1. If the red-outer-planar embedding of $H^b - e_b$ extends to the unique embedding of H_b, then either the outer face still covers all red edges, or it is divided into two faces which form a blue cover. Both cases contradict the hypotheses. It follows that the embedding of $H^b - e_b$ does not extend, and hence e_b is in a non-trivial 3-edge-cut B. Label the edges of B, their incident faces, and the components of $H^b - B$ as shown in Figure 3.

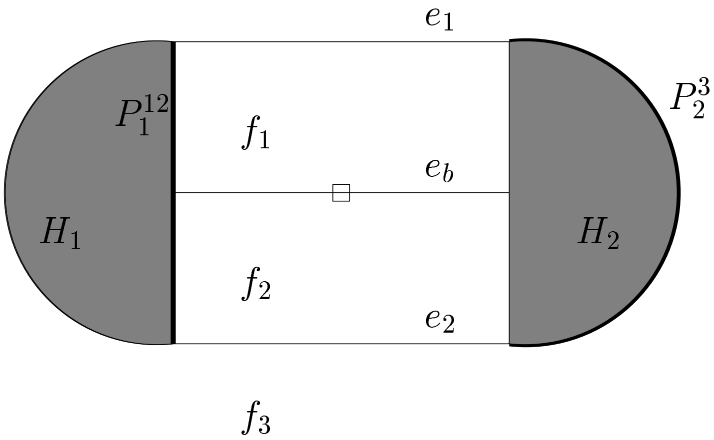

Figure 3: The Graph in the Claim

Let P_i^{12} denote the edges in H_i incident with $f_1 \cup f_2$, and let P_i^3 denote those edges incident with f_3. The paths P_1^{12} and P_2^3 are shown as bold edges in Figure 3. Let $C_i = P_i^{12} \cup P_i^3$.

First, we note that each H_i must contain a red edge. If not, then we can replace H_i by a single vertex incident with three edges, where these three edges inherit the colors on the edges of B. Any red-outer-planar embedding of this new graph can be easily modified to a red-outer-planar embedding of H^b. The same holds for blue covers of these two graphs. The new graph is a strict subgraph of H^b, contradicting minimality.

Second, we note that each red edge in H_i is contained C_i. If not, then we replace the other H_j with a single vertex, where either one of the edges e_1 or e_2 are colored red. This new graph is not red-outer-planar (no face can cover both red edges) nor has a blue cover (the two faces are not incident with the red edge in $H_i - C_i$). Again the new graph is a strict subgraph of H^b, contradicting minimality.

Now, observe that if all the red edges of $H_1 \cup H_2$ are contained in $P_1^{12} \cup P_2^{12}$, then H^b has a blue cover. If all of the red edges are in $P_1^3 \cup P_2^3$, then H^b is red-outer-planar. Hence, without loss of generality, there is a red edge in P_1^{12} and one in P_2^3. The resulting graph contains T_1 as desired. The claim is demonstrated.

The proof of Part *iv)* is now easy. Any graph not containing T_1 must have the blue edge disjoint from either a $K_{3,3}^{-e}$ or from a $K_{3,3}^{-v}$. Moreover, the blue edge forms the whole of the only bridge, because adding any such edge cannot have a blue cover. If the blue edge is disjoint from a $K_{3,3}^{-e}$, then the only choice for a 3-connected graph that does not contain T_1 gives T_2. If the blue edge is disjoint from a $K_{3,3}^{-v}$, there are two possible ways to add in an edge to give a 3-connected graphs. These give graphs T_3 and T_4. ∎

We are now ready to characterize the desired graphs with edge-connectivity exactly two.

Proposition 3.4 *Let G be an edge-colored cubic graph that is minimal without a 2-red-cover. If G has edge-connectivity exactly two, then G is one of the graphs G_{14}-G_{22} of Figure 1.*

Proof: We break the proof into three main cases. The first is when G has two different 2-edge-cuts that share an edge. The second is when G has two 2-edge-cuts that do not share an edge. The third is when G has a unique 2-edge-cut.

Case 1: *There are two 2-edge cuts with an edge in common.* Let $B_1 = \{e_2, e_3\}$ and $B_2 = \{e_1, e_3\}$. Note that $B_3 = \{e_1, e_2\}$ is also a 2-edge cut. Hence $G - \{e_1, e_2, e_3\}$ has three components H_1, H_2, H_3. Label these so that B_i separates H_i from the remaining two H_j's.

At least one of the three H_i^b does not have a blue cover, or else the two faces on $\{e_1, e_2, e_3\}$ are a 2-red-cover of G. Say H_3^b does not have a blue cover. To cover the red vertices in H_3 requires at least one face not incident with B_3. To cover the red vertices in $H_2 \cup B_2$ requires either one face not incident with B_2, or two faces.

Now, delete a single non-red edge from H_1. The resulting graph has a 2-red-cover. It must involve one face incident with H_1 in order to cover the red edge in H_1, this face may also be incident with $e_1 \cup e_2 \cup e_3$. There are two possibilities for the other face. If it is incident with these three edges, then we contradict that H_3^b does not have a blue cover. If it is not incident with one of these three edges, then we contradict Lemma 3.3(i) applied to H_2^b.

Case 2: *There are two disjoint 2-edge-cuts.* Let $G - B_1$ have components H_1 and $\bar{H}_1$, $G - B_2$ have components H_3 and $\bar{H}_3$, and label the components so that $H_2 = \bar{H}_1 \cap \bar{H}_3$ is non-empty. So B_1 is a bond joining H_1 to H_2, and B_2 is a bond joining H_2 to H_3.

Suppose that both H_1^b and H_3^b are red-outer-planar. Then we can cover all of the red edges in $G - H_2$ with these two faces. Hence H_2 has a red edge. By Lemma 3.3(ii) both H_i^b contain either R_1 or R_2 of Figure 2. This combination using both $H_i^b = R_1$ together with a single red edge in H_2 give the graph G_{14} of Figure 1. Using R_2 for one of the H_i^b gives an edge whose deletion contains G_{17}. Using R_2 for both H_i^b's gives an edge whose deletion contains G_{19}.

Without loss of generality we can suppose that H_1^b is non-red-outer-planar. If H_1^b has no blue cover, then consider a 2-red-cover of $G - e$ for some edge e in H_2. If both of these faces are incident with B_1, then we contradict that H_1^b has no blue cover. If only one of these faces are incident with B_1, then we contradict Lemma 3.3(ii) applied to H_3^b.

We conclude that H_1^b is non-red-outer-planar and has a blue cover. By Lemma 3.3(iii) H_b^1 contains either S_1 or S_2. By emma 3.3(ii) H_b^3 contains either R_1 or R_2. Also, H_2 contains at least one edge. We check the four ways to combine these graphs. Using S_1 and S_2 with R_1 give G_{15} and G_{16} of Figure 1 respectively. Using S_1 with R_2, we can delete the edge in H_2 and get G_{17}. Using S_2 with R_2, we can delete the edge in H_2 and get G_{18}.

Case 3: *There is a unique 2-edge-cut.* We note that both H_i^b are 3-edge-connected. As before, at least one of the H_i^b do not have a blue cover, assume that H_1^b does not.

We first consider the possibility that H_1^b is red-outer-planar. Then H_2^b cannot be red-outer-planar. By Lemma 3.3(iii) and (iv) H_2^b either contains S_1, S_2, or it contains T_1, T_2, T_3, T_4. Each of T_2, T_3, T_4 contain S_1 or S_2. We conclude that H_2^b contains one of S_1, S_2, T_1. By Lemma 3.3(ii) H_1^b contains R_2 because it is 3-connected. Combining these graphs in the three possible ways give G_{17}, G_{18}, and G_{19} of Figure 1 respectively.

We next consider the possibility that H_1^b is not red-outer-planar. Then by Lemma 3.3(iv) H_1^b contains one of T_1-T_4. By Part (i) of that lemma, H_2^b contains either R_1 or R_2. There are 8 possible ways to combine these parts. Combining T_1 with R_1 again gives G_{17}. Combining T_2-T_4 with R_1 gives G_{17}, G_{18}, and G_{19} of Figure 1 respectively. Combining T_1 with R_2 again gives G_{19}. Combining T_2 or T_4 with R_2 both give graphs that properly contain G_{18}. Finally, combining T_3 with R_2 gives a graph that properly contains G_{17}.

This completes the cases and the proof of the proposition. ∎

By Proposition 3.4 we can assume that G is 3-edge-connected. Before covering this remaining case in Proposition 3.6 we need one more lemma.

Lemma 3.5 *Let R^* a graph with the property that no two vertices cover all of its edges. Moreover, suppose that R^* is minimal with this property under edge deletions and deleting isolated vertices. Then R^* is either K_4, $K_3 \cup K_2$, $K_2 \cup K_2 \cup K_2$, or C_5.*

Proof: Cattell and Dinneen [7] found the minor-minimal graphs without a k-vertex cover for $k \leq 5$. Their set for $k = 2$ is the four graphs above. Because the maximum degree of these graphs is at most 3, they are also minimal without a 2-vertex-cover under our coarser graph order. ∎

Proposition 3.6 *Let G be a planar edge-colored cubic graph that is minimal without a red-cover. If G is 3-edge-connected, then G is one of the graphs G_{23}-G_{29} of Figure 1.*

Proof: We first establish that every face f of G is incident with a red edge. For suppose not. Then there is an edge e incident with f such that $G - e$ is still 3-edge-connected (this is related to ear-decompositions, see [10]).

Deleting this edge gives a 2-red-cover. This now makes a 2-red-cover of G that doesn't use the face f.

Consider the graph R^* whose vertices are the faces of G, with two vertices joined by an edge if and only if the faces share a common red edge of G. By the preceding paragraph there are no isolated vertices in R^*. Because G has no 2-red-cover, no two vertices of R^* cover all edges of R^*. By minimality, any edge-deletion from R^* does have this property. By Lemma 3.5 R^* is one of four graphs. We consider the graphs in turn.

If $R^* = K_4$, then G has exactly four faces and hence is K_4. All edges must be red, giving G_{23}.

If $R^* = K_3 \cup K_2$, then G has exactly five faces and six vertices. There is a unique such graph, the 3-prism $C_3 \times K_2$. It is easy to show that up to isomorphism there is exactly one way to pick R^* in this graph, giving G_{24}.

If $R^* = C_5$, then again G is the 3-prism, and the choice of R^* is unique up to isomorphism. This gives G_{25}.

If $R^* = K_2 \cup K_2 \cup K_2$, then G has exactly eight vertices. There are two graphs to consider. The first is the 3-cube. There is a unique way to pick a matching in the dual of the 3-cube, up to isomorphism, giving G_{26}. The second graph is the one underlying G_{27}, G_{28}, and G_{29}. There are three non-isomorphic ways to pick a matching in the dual $K_5 - K_2$. These give G_{27}, G_{28}, and G_{29} respectively. The details are left to the reader. ■

4 The Proof of the Main Theorem and Concluding Remarks

We begin this section by combining the propositions of the previous section to prove our main result, Theorem 1.1.

Proof: (of Theorem 1.1.) By the collective results of Section 2, it suffices to find the minimal edge-colored cubic graphs without a 2-red-cover (Problem 2.6). Proposition 3.1 finds all such graphs that are either non-planar or disconnected. Proposition 3.2 finds all such connected planar graphs with a cut-edge. Proposition 3.4 finds all such graphs with edge-connectivity exactly two. Finally, Proposition 3.6 finds these graphs with edge-connectivity exactly three.

Combining the lists of the four propositions gives exactly the graphs of Figure 1. ■

There is an interesting connection between Halin's Theorem and Kuratowski's Theorem. Namely, consider the four possible graphs formed from $K_{3,3}$ and from K_5 by deleting either an edge (or a vertex), then coloring all incident (or adjacent) vertices black. The four graphs formed in this manner are exactly the four planar graphs of Halin's Theorem.

For the Möbius band, a colored graph G embeds with a 1-black-cover if and only if the corresponding G^+ embeds in the projective plane. Using this relation and the known minimal graphs that do not embed in the projective plane, Archdeacon et al [3] find the minimal colored graphs that embed in the projective plane with a 1-black-cover. This corresponds to Halin's Theorem for graphs that embed on the Möbius band without an accumulation point.

The *spindle surface* S is formed from the sphere by identifying two distinct points, commonly known as the *north* and *south pole* as a common *pinch point*. Given a colored graph G, let G^+ denote the graph with one additional vertex v_+ adjacent to every colored vertex of G. If G has a 2-black-cover, then G^+ embeds in the spindle surface. The converse is not necessarily true, as it is possible that G^+ embeds in the spindle surface with V^+ not on the pinch point. This embedding does not necessarily correspond to an embedding of G with a 2-black cover.

It is tempting, nevertheless, to try to relate the minimal non-spindle graphs to the minimal non-2-black-cover graphs. We do not know the exact relation.

A graph is *outer-cylindrical* if it embeds on the plane so that every vertex is on the boundary of one of two faces. The vertices are not colored, so that such an embedding is equivalent to a 2-black-cover where all vertices are colored black. The set of minimal non-outer-cylindrical graphs is known[2]. These graphs contain subgraphs that are minimal with no 2-black-cover formed by coloring only a subset of the vertices. However, not all graphs without a 2-black-cover arise in this way, and the relation between the two sets is again unclear.

We close by asking to reader to extend these techniques, or invent new ones, to completely characterize the (non-cubic) colored graphs without a 2-black-cover.

Acknowledgement: The first and third author would like to thank the University of Auckland, which hosted them during the time this research was conducted. All three 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, *J. Graph Theory* **5**(1981)243–246.

[2] D. Archdeacon, P. Bonnington, N. Dean, N. Hartsfield, and K. Scott, Obstruction sets for outer-cylindrical graphs, *J. of Graph Theory* (to appear).

[3] D. Archdeacon, P. Bonnington, M. Debowsky, and M. Prestidge, Halin's Theorem for the Möbius strip, manuscript (2001).

[4] C.P. Bonnington and R.B. Richter, Graphs embedded in the plane with finitely many accumulation points, manuscript (2000).

[5] L. Boza, M. Gugúndez, A. Márquez, and P. Revuelta, Embedding graphs without accumulation points in generalized cylinders, manuscript (2001).

[6] J. Cáceres, Diversos tipos de planaridad de grafos, *Tesis doctoral*, Universidad de Almería (1996).

[7] K. Cattell and M.J. Dinneen, A characterization of graphs with vertex cover up to five, *Orders, Algorithms and Applications, ORDAL'94* (V. Bouchitte and M. Morvan eds.), in *Lecture Notes on Computer Science* **831**, Springer-Verlag (1994)86-99.

[8] G.A. Dirac and S. Schuster, A theorem of Kuratowski, *In dag. Math.* **16** (1954)343-348.

[9] R. Halin, Zur häufungspunktfreien Darstellung abzählbarer Graphen in der Ebene, *Arch. Math. (Basel)* **17** (1966)239–243.

[10] A. Kelmans, On convex embeddings of planar 3-connected graphs, *J. of Graph Theory* **33** (2000)120–124.

[11] K. Kuratowski, Sur le probléme des courbes gauches en topologie, *Fund. Math.* **5**(1930)271-283.

[12] B. Mohar and C. Thomassen, Graphs on Surfaces, Johns Hopkins University Press (2001).

[13] C. Thomassen, Infinite graphs. In *Selected topics in graph theory 2*, Academic Press (1983)129-160.