2009

Some results on cubic graphs

Evan Morgan
Louisiana State University and Agricultural and Mechanical College

SOME RESULTS ON CUBIC GRAPHS

A Dissertation

Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

in

The Department of Mathematics

by
Evan Morgan
B.A., Lawrence University, 2002
M.S., Louisiana State University, 2004
August 2009

Acknowledgments

This dissertation would not be possible without the boundless wisdom and patience (and of course mathematical insight) I leeched from my Ph.D. advisor, Dr. Bogdan Oporowski. The German for "Ph.D. advisor" is *Doktorvater*, which is as fitting a *Pronomen* as I can imagine.

I offer a very special thanks to my brother, Garrett, who often serves as a catalyst to my daily activities. I thank my parents, for the subtle, tactful motivation, without which this project would have been abandoned long ago. Most of all, I would like to thank Chris Oxner, for the herculean editing, and for other reasons he alone understands. Additionally, I would like to thank Melanie Win Myint, for afflatus, and for other reasons she alone understands.

This dissertation is dedicated to my family and friends, current and former, without whom I would be more lost at sea than at present. Here are a few representative members, in alphabetical order: D. Abarca, D. Aulicino, A. Baker, T. Becker, C. Chun, D. Fournet, The Hoeslys, I. Kochergin, J. Mancuso, M. Massow, C. Mathes, J. McIntosh, The Morgans, C. Oxner, M. Ross, M. Smith, A. Thomas, and M. Win Myint.

Table of Contents

Abstract

Pursuing a question of Oxley, we investigate whether the edge set of a graph admits a bipartition so that the contraction of either partite set produces a series-parallel graph. While Oxley's question in general remains unanswered, our investigations led to two graph operations (Chapters 2 and 4) which are of independent interest. We present some partial results toward Oxley's question in Chapter 3.

The central results of the dissertation involve an operation on cubic graphs called the *switch*; in the literature, a similar operation is known as the *edge slide*. In Chapter 2, the author proves that we can transform, with switches, any connected, cubic graph on n vertices into any other connected, cubic graph on n vertices. Furthermore, connectivity, up to internal 4-connectedness, can be preserved during the operations.

In 2007, Demaine, Hajiaghayi, and Mohar proved the following: for a fixed genus g and any integer $k \geq 2$, and for every graph G of Euler genus at most g, the edges of G can be partitioned into k sets such that contracting any one of the sets produces a graph of tree-width at most $O(g^2k)$. In Chapter 3 we sharpen this result, when $k = 2$, for the projective plane ($g = 1$) and the torus ($g = 2$).

During early simultaneous investigations of Jaeger's Dual-Hamiltonian conjecture and Oxley's question, we obtained a simple structure theorem on cubic, internally 4-connected graphs; that result is found in Chapter 4.

Chapter 1
Introduction

1.1 Definitions and Terminology

A *multigraph* is an ordered triple (V_M, E_M, ϕ), where V_M and E_M are disjoint sets, and ϕ is a function from E_M to the set of one- and two-element subsets of V_M. We refer to V_M and E_M as the vertex set and the edge set, respectively. We refer to the elements of V_M as *vertices*. The elements e of E_M are of three types: if $|\phi(e)| = 1$, then e is called a *loop*; if $|\phi(e)| = 2$ and $|\phi^{-1}(\phi(e))| = 1$, then e is called an *edge*; if $|\phi(e)| = 2$ and $|\phi^{-1}(\phi(e))| > 1$, then e is called a *multiedge*. If e is a loop, edge, or multiedge of a multigraph, then $|\phi^{-1}(\phi(e))|$ is the *multiplicity* of e. If G is a multigraph, then $V(G)$ refers to the vertex set of G, and $E(G)$ refers to the edge set of G. A multigraph (V_M, E_M, ϕ) is said to be *simple* if it contains no loops and ϕ is injective. If $G_1 = (V_1, E_1, \phi_1)$ and $G_2 = (V_2, E_2, \phi_2)$ are multigraphs, then a *multigraph isomorphism between G_1 and G_2* is a one-to-one correspondence $f : V_1 \longrightarrow V_2$ such that, for any vertices u and v in G_1, the following hold:

(1) There are precisely k loops at u in G_1 if and only if there are precisely k loops at $f(u)$ in G_2;

(2) $|\phi_1^{-1}(\{u, v\})| = |\phi_2^{-1}(\{f(u), f(v)\})|$.

In this case, we say that G_1 and G_2 are *isomorphic*.

In a simple multigraph, we may ignore ϕ and think of the edges as unordered pairs of vertices. In this dissertation, we deal primarily with simple multigraphs, therefore we adopt the convention that a *graph* is a simple multigraph, whose edges are unordered pairs of vertices.

A multigraph (V', E', ϕ') is a *sub-multigraph* of a multigraph (V, E, ϕ) if the following hold:

(1) $V' \subseteq V$;

(2) $E' \subseteq E$;

(3) $\phi|_{E'} = \phi'$.

If, in addition, $V' = V$, then we say that the sub-multigraph is *spanning*. If S is a subset of V, then the *subgraph induced by* S, notated $G[S]$, is the submultigraph of G whose vertex set is S, whose edge set consists of all edges whose endpoints are a subset of S, and whose ϕ function is inherited from G.

In a multigraph (V_M, E_M, ϕ), two distinct vertices u and v are *adjacent* if $\{u, v\} \in \phi(E_M)$; we say that u and v are the *endpoints* of the edges in $\phi^{-1}(\{u, v\})$. We say that two distinct edges are *adjacent* if they share at least one endpoint. We say that an edge e is *incident* to a vertex v, or equivalently, vertex v is *incident* to e, if v is an endpoint of e. The *neighborhood* of a vertex u (notated $N(u)$) is the set of all vertices, other than u, which are adjacent to u. The *degree* of a vertex u is defined to be the number of non-loop edges incident to u plus twice the number of loops incident to u. A multigraph is *cubic* if every vertex has degree three.

For notational ease, we adopt some conventions: in a graph, an edge $\{u, v\}$ will often be notated as uv, and we will often refer to edges using only their endpoints, as in "the edge uv," or "the loop at u"; in a multigraph, we will also refer to edges using only their endpoints, yet with the understanding that the reference may not be unique.

A *walk* in a multigraph consists of a sequence of vertices $(v_1, v_2, \ldots, v_k)$ and the corresponding edges with endpoints v_i, v_{i+1}, with $i \in \{1, 2, \ldots, k-1\}$; more specifically, the aforementioned walk is called a (v_1, v_k)-*path*. Note that the v_i's are not necessarily distinct. A *path* in a multigraph is a walk whose vertices are pairwise distinct. A *cycle* is a walk whose sequence $(v_1, v_2, \ldots, v_k)$ of vertices satisfies the following:

(1) $v_1 = v_k$;

(2) the vertices $v_1, v_2, \ldots, v_{k-1}$ are pairwise distinct.

If A and B are sets of vertices, then an (A, B)-*path* is any (u, v)-path such that $u \in A$ and $v \in B$; an (A, B)-*edge* is an edge with one endpoint in A and one endpoint in B. An edge e with endpoints u, v is a *chord* of a path P in G if $\{u, v\} \subseteq V(G)\backslash V(P)$ and $e \notin E(P)$. A path which has no chords is called an *induced path*. We will often refer to a path by its sequence of vertices. A *central edge* of a path $v_1, v_2, \ldots, v_k$ if one of the following holds:

(1) The edge $v_{\frac{k}{2}}v_{\frac{k}{2}+1}$, if k is even;

(2) Either edge incident to $v_{\lceil\frac{k}{2}\rceil}$, if k is odd.

Let G be a multigraph with vertex set V and edge set E. Let S be a set of vertices in G. We define a *vertex of incidence* of S to be a vertex of S which is an endpoint of an $(S, V-S)$-edge. An *edge of incidence* of S is an $(S, V-S)$-edge. A multigraph is *connected* if there is a path between any two vertices. A multigraph is *disconnected* otherwise.

A *k-separation* of G is a pair of submultigraphs $\{A, B\}$ of G such that the following hold:

(1) Each of A and B has at least k edges;

(2) $A \neq G \neq B$;

(3) $A \cup B = G$;

(4) $A \cap B \subseteq V(G)$;

(5) $|A \cap B| \geq k$.

We will often refer to a separation by the intersection of the two subgraphs, as in "the k-separation $A \cap B$." A set C of k edges is called a *k-cut* if it consists of the edges of incidence of some proper subset of vertices.

A k-separation is *verticial* if at least one partite set of the separation consists entirely of a vertex and its neighbor set. A k-separation is *nonverticial* otherwise. A k-cut is *verticial* if all k edges share a single endpoint; note that the vertex in question may be a cut-vertex with degree greater than k. A k-cut is *nonverticial* otherwise. A k-cut is *essential* if no two edges of the cut share an endpoint. A k-cut is *nonessential* otherwise. A multigraph is *connected* if it contains a (u, v)-path for every pair $\{u, v\}$ of vertices. A multigraph is *k-connected* if it is connected, contains more than k vertices, and admits no $(k-1)$-separation. A multigraph is *internally 4-connected* if it is 3-connected and each of its 3-separations is verticial. We will occasionally refer to the set of endpoints of edges in a cut C as the *endpoints of C*. A *connected component* of a multigraph is a maximal connected submultigraph, with respect to subgraph containment.

To *delete an edge e of a multigraph $G = (V, E)$*, we merely delete e from the edge set E; the resulting multigraph is notated $G \backslash \{e\}$, or occasionally $G \backslash e$. To *delete a set S of edges from G*, we merely delete S from E; the resulting multigraph is notated $G\backslash S$. To *delete a vertex v from G*, we delete v from V, and we delete from E every edge incident to v; the resulting multigraph is notated $G\backslash v$. To *delete a set S of vertices from G*, we merely delete S from V, and we delete all edges from E which are incident to some member of S; the resulting multigraph is notated $G\backslash S$.

A *cut-vertex* of G is a vertex which describes a 1-separation. A *cut-edge* is an edge whose deletion increases the number of components of G. A *block* of a multigraph G is defined to be a submultigraph B which satisfies at least one of the following:

(1) B is induced by a loop;

(2) B is induced by a cut-edge;

(3) B is maximal (with respect to containment) and 2-connected.

Suppose that G is connected. Then we can define the *block tree* T of G as follows: Let $V(T)$ consist of the blocks of G and those vertices of G which lie in more than one block; two elements u, v in $V(T)$ are adjacent if the following hold:

(1) $u \in V(G)$ and v is a block of G;

(2) u is a vertex of v.

The operation of *contraction* is slightly more complicated than the operation of deletion. Let e be an edge of a multigraph G. We denote the *contraction of e in G* as G/e, and we define it as follows:

(1) If e is a loop, then $G/e = G\backslash e$;

(2) Otherwise, if u and v are the endpoints of e, then G/e is the multigraph obtained from $G\backslash\{u, v\}$ by adding a vertex, say w, and, for each edge $f \in E(G)\backslash e$ whose set P of endpoints meets $\{u, v\}$, adding an edge with endpoints $(P\backslash\{u, v\}) \cup w$.

See Figure 1.1. If S is a set of edges of G, then to *contract the set S of edges*, we first partition S into $\{S_1, S_2, \ldots, S_k\}$, where each S_i is the edge set of a maximal connected component of $G[S]$; then, for each S_i, we delete S_i from G and identify all the endpoints of all the edges in S_i. We notate the resulting graph by G/S. If $S = \{e_1, \ldots, e_k\}$, then one may see that G/S is isomorphic to $(((\cdots(G/e_1)/e_2)/e_3)\cdots)$.

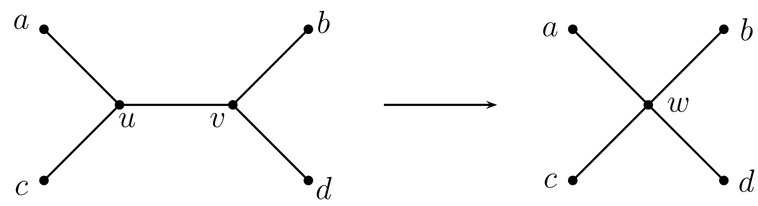

FIGURE 1.1. Contracting the edge uv.

The operation of *splitting a vertex* may be viewed as the opposite of contracting an edge; where, in the right-side of Figure 1.1, we *split* the vertex in the middle by replacing it with the scenario on the left-side of Figure 1.1. Notice that a vertex may admit distinct splittings which produce non-isomorphic multigraphs. To *suppress* a vertex of degree two or one, we contract exactly one edge incident to it. Notice that suppressing a vertex is a uniquely defined operation, up to isomorphism.

If H is a submultigraph of G, then an H-bridge of G is a submultigraph B of G such that at least one of the following holds:

(1) B consists of an edge not in H, whose endpoints lie in $V(H)$;
(2) B is a minimal submultigraph such that B contains a connected component C of $G \backslash V(H)$ and B contains all edges of incidence (in G) of C.

A vertex v of G is an *attachment of an H-bridge B* if $v \in V(B) \cap V(H)$.

One measure of the complexity of a multigraph is via the concept of tree-width, which was developed first, under a different name, by Halin [9], and later, apparently independent of Halin's paper and of each other, by the teams of Arnborg and Proskurowski [2] and Robertson and Seymour [19]. Many NP-hard problems can be solved in linear time when considering classes of multigraphs with bounded tree-width. Among the several equivalent definitions of tree-width, we present here the definition based upon the concept of a tree-decomposition, introduced by Robertson and Seymour, using the exposition from Diestel [6].

Let G be a multigraph, let T be a tree, and let $\mathcal{V} = \{V_t\}_{t \in V(T)}$ be a collection of sets of vertices of G, indexed by the vertices of T. The pair $(T, \mathcal{V})$ is called a *tree-decomposition* if the following three conditions are satisfied:

(T1) $V(G) = \bigcup_{t \in V(T)} V_t$;

(T2) For every edge $e \in E(G)$, there is a vertex $t \in V(T)$, such that $e \in G[V_t]$;

(T3) If t_1, t_2, t_3 are vertices of T such that t_2 lies on the unique (t_1, t_3)-path in T, then $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$.

The sets V_t are often called *bags*. Given a tree-decomposition of G, the *width* of the decomposition is one less than the size of a largest bag. The *tree-width* of G, notated as $tw(G)$, is the minimum width among all tree-decompositions of G. Note that the tree-width of a tree with at least one edge is one.

To *subdivide* an edge e of a multigraph G, we delete e and replace it with a path (or cycle, if e is a loop) with two edges. A graph H is a *subdivision* of G if H can be obtained from G by repeatedly subdividing edges. A multigraph H is a *minor* of G if we can perform edge- and vertex-deletions and edge contractions on G to obtain a multigraph isomorphic to H; we notate it $H \leq_{\mathrm{m}} G$. A class $\mathcal{G}$ of graphs is *minor-closed*, if the following holds: if $G \in \mathcal{G}$ and $H \leq_{\mathrm{m}} G$, then $H \in \mathcal{G}$. If $\mathcal{G}$ and $\mathcal{F}$ are classes of multigraphs, then $\mathcal{F}$ is a set of *forbidden minors* of $\mathcal{G}$ if the following hold:

(1) $G \not\leq_{\mathrm{m}} H$, for every pair $\{G, H\} \subseteq \mathcal{F}$;

(2) a multigraph G lies in $\mathcal{G}$ if and only if no member of $\mathcal{F}$ is a minor of G.

A graph is *series-parallel* if it contains no minor isomorphic to K_4. It is a well-known fact that K_4 is the unique forbidden minor for the class of graphs with tree-width at most two. Therefore a graph is series-parallel if and only if it has tree-width at most two.

In this dissertation, a *map* or a *mapping* is a continuous function between topological spaces. In this dissertation, multigraphs will often be viewed as topological spaces. Let G be a multigraph. For each edge e_i, with (possibly identical) endpoints x_i and y_i, let I_i be a copy of the closed unit interval, with x_i and y_i corresponding to 0 and 1, respectively. Let X be the set obtained from $V(G)$ and the I_i's by identifying $v \in V(G)$ with the endpoints of the I_i's which correspond to v. Let $v \in V(G) \subseteq X$. For each endpoint z_j of an I_i which corresponds to v, let Z_j be a half-open subinterval of I_i (using the inherited topology of I_i) which contains z. The union of these Z_j's we call an *open star*. Let $\mathfrak{B}$ consist of all sets $B \subseteq X$ such that B is an open star or B is an open subinterval of some I_i containing neither endpoint of I_i. Then $\mathfrak{B}$ is a basis for the topology on X, which describes the topology on G. From this viewpoint, we may view multigraphs as a collection of vertices and a collection of edges; the function ϕ is determined when each edge is a distinct unit interval. With this in mind, we will occasionally define multigraphs in terms of their vertices and edges.

Let S be a subset of a topological space. We denote the closure of S by $\overline{S}$. We denote the interior of S by $\mathring{S}$. We denote the boundary of S by ∂S. In this dissertation, a *disc* in a topological space is a subspace whose interior is homeomorphic to the open unit disc in $\mathbb{R}^2$ and whose closure is homeomorphic to the closed unit disc in $\mathbb{R}^2$. If $\mathcal{S}$ is a connected, compact Hausdorff topological space which is locally homeomorphic to a disc, then we call $\mathcal{S}$ a *surface*. An *embedding* of a multigraph G in a surface $\mathcal{S}$ is a one-to-one map from G to $\mathcal{S}$. If Γ is an embedding of G in some surface, we will often refer to the image of Γ as G; that is, we will speak of G as a subspace of the surface. A *curve* α in a surface $\mathcal{S}$ is a map from the unit interval to $\mathcal{S}$. A *closed curve* α in a surface $\mathcal{S}$ is a curve in $\mathcal{S}$ such that $\alpha(0) = \alpha(1)$. Given a curve α, we will often refer to the image of α as α itself; that is, we will speak of α as a subspace of the surface. A curve is *noncontractible* if it is not homotopically equivalent to a constant curve (i.e. a curve whose image is a single point). The *representativity of* G is the minimum of $|\alpha \cap G|$, over all noncontractible curves α. If G is embedded via Γ in some surface $\mathcal{S}$, then a *face* of the embedding is a connected component of $\mathcal{S}\backslash\Gamma(G)$. Note that faces are always open. An edge and a face are said to be *incident* if the edge lies on the boundary of the face. Given a face F, the *facial walk* of F is a closed walk W which satisfies the following:

(1) an edge e appears in W if and only if $e \in \partial F$;

(2) an edge e appears exactly once in W if and only if $e \in \partial F$ and $\mathring{e} \not\subseteq \mathring{\overline{F}}$; and

(3) an edge e appears exactly twice in W if and only if $e \in \partial F$ and $\mathring{e} \subseteq \mathring{\overline{F}}$.

A *triangular face* of an embedded multigraph is a face whose boundary forms a 3-cycle. A *triangulation* is an embedding, each of whose faces is a triangular face. Suppose that a multigraph G is embedded in a surface $\mathcal{S}$, and let $v \in V(G)$. Let B be an open disc in $\mathcal{S}$ such that $B \cap G$ is an open star containing v. Then the *rotation scheme at* v is the cyclic ordering, induced by $B \cap G$, of the edges incident to v, in which edges (namely, loops) may appear more than once. Suppose that another multigraph H is embedded in $\mathcal{S}$, in addition to G, and suppose that it satisfies the following:

(1) (the interior of) each face of G contains precisely one vertex of H;

(2) $|e \cap f| = |\mathring{e} \cap \mathring{f}| \leq 1$ for every $e \in E(G)$ and every $f \in E(H)$;

(3) The relation $\{(e, f) : e \in E(G); f \in E(H); |e \cap f| \neq \emptyset\}$ describes a one-to-one correspondence.

Then H is called a *surface dual* of G.

The *projective plane* is the surface obtained from the closed unit disc (with the topology inherited from the plane) by identifying antipodal points on the boundary of the disc and imbuing it with the corresponding quotient topology. For purposes of illustration, we refer to the *unit disc model of the projective plane* as a unit disc where antipodal boundary points are taken to be identical. See Figure 1.2. The torus is the surface obtained from the unit square (with the topology inherited from the plane), where, as labeled below, the boundary segments *ad* is identified with the boundary segment *bc*, and the boundary segment *ab* is identified with the boundary segment *dc*. See Figure 1.3.

1.2 Background

In 1971, Chartrand, Geller, and Hedetniemi [4] made the following:

Conjecture 1.2.1 (Chartrand, Geller, and Hedetniemi)**.** *Every planar graph admits an edge-partition into two outerplanar graphs.*

This conjecture inspired much work before it was solved in the affirmative in 2005 by Gonçalves [8].

Theorem 1.2.2 (Gonçalves)**.** *Every planar graph admits an edge-partition into two outerplanar graphs.*

A partial result toward conjecture 1.2.1, obtained in 2000 by Ding, Oporowski, Sanders, and Vertigan [7], is of relevance to us here. (This result was proved independently by K. Kedlaya [11].)

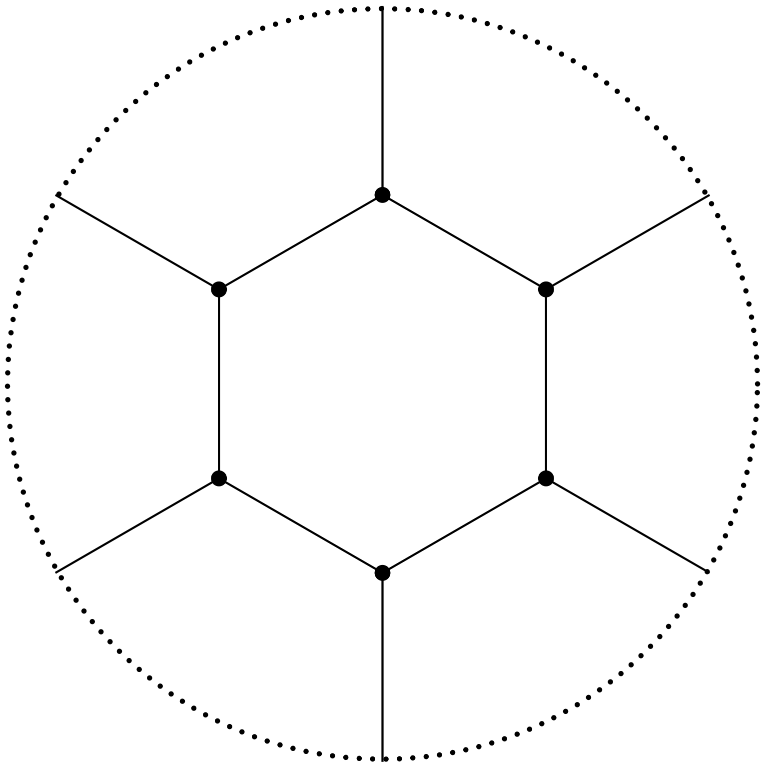

FIGURE 1.2. The complete bipartite graph $K_{3,3}$ embedded on the projective plane.

Theorem 1.2.3 (Ding, Oporowski, Sanders, and Vertigan; Kedlaya)**.** *Every planar graph has an edge partition into two series-parallel graphs.*

We may rephrase Theorems 1.2.2 and 1.2.3, respectively, as follows.

Theorem 1.2.4. *If $G = (V, E)$ is a planar graph, then E may be partitioned into two sets, E_1 and E_2, such that $G \backslash E_1$ and $G \backslash E_2$ are outerplanar.*

Theorem 1.2.5. *If $G = (V, E)$ is a planar graph, then E may be partitioned into two sets, E_1 and E_2, such that $G \backslash E_1$ and $G \backslash E_2$ are series-parallel.*

If G is a connected graph embedded on a surface, and G^* is a surface dual of G, then the reader may notice that the deletion of an edge of G corresponds to the contraction of the corresponding edge of G^*; this phenomenon is investigated more fully in Section 3.2. Deletion and contraction can, in this way, be viewed as *dual operations*. Furthermore, the reader may notice that K_4, embedded on the plane, is isomorphic to its own dual. Therefore it is an easy exercise to prove that if G is a

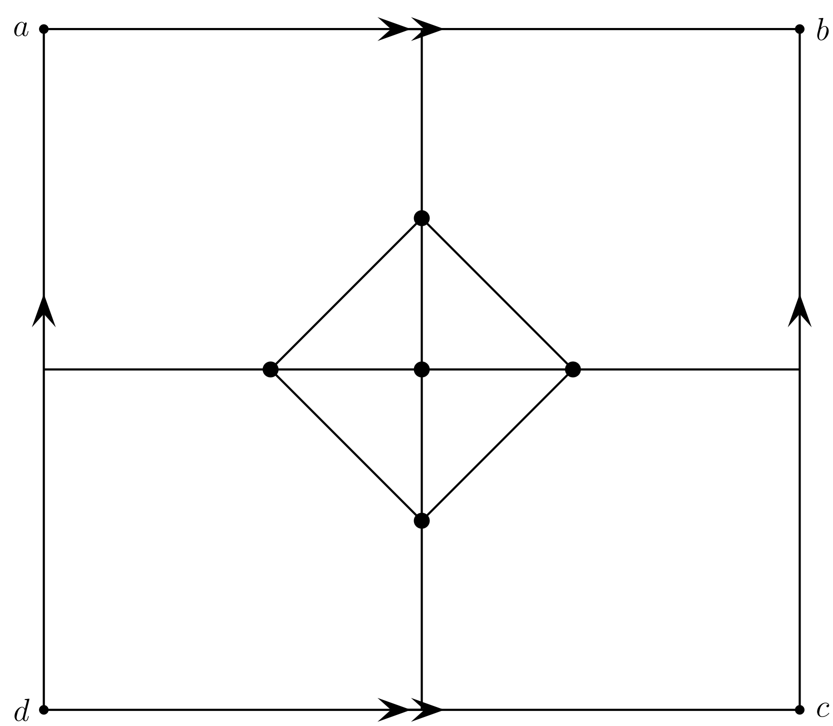

FIGURE 1.3. The complete graph K_5 embedded on the torus.

plane graph and G^* is a surface dual of G, then K_4 is a minor of G if and only if K_4 is a minor of G^*. With this in mind, we notice the following corollary of Theorem 1.2.5.

Corollary 1.2.6. *Every plane graph G admits an edge-partition $\{E_1, E_2\}$ such that G/E_1 and G/E_2 are series-parallel.*

To this end, J. Oxley [18] asked the following question.

Question 1.2.7 (Oxley)**.** *If M is a cographic matroid, can we partition the ground set of M into two sets S and T such that $M\backslash S$ and $M\backslash T$ are series-parallel?*

Oxley's question leads naturally to the following generalization of Corollary 1.2.6.

Question 1.2.8. *Does every graph G admit an edge partition $\{E_1, E_2\}$ such that G/E_1 and G/E_2 are series-parallel?*

As a first observation, we note that contracting edges in a graph will not raise its tree-width. In particular, the tree-width of a graph is equal to the maximum tree-width of its 2-connected blocks. We investigate this issue more fully in Section 3.2, and prove that it suffices to consider only 2-connected graphs.

As a second observation, we note that contracting a spanning tree in a (connected) graph results in a graph with a single vertex. Therefore, if a graph $G = (V, E)$ contains two edge-disjoint spanning trees T_1 and T_2, we can partition the edge set into $E_1 = E(T_1)$ and $E_2 = E(T_2) \cup \big(E \backslash (E(T_1) \cup E(T_2))\big)$; in this case, G/E_1 and G/E_2 are singletons. Both Nash-Williams [14] and Tutte [21] proved theorems characterizing the graphs which contain k edge-disjoint spanning trees. A nice statement and proof can be found on pages 46–48 of [6].

Theorem 1.2.9 (Nash-Williams; Tutte). *A graph contains k edge-disjoint spanning trees if and only if for every partition $(V_1, \ldots, V_l)$ of its vertex set, it has at least $k(l-1)$ distinct (V_i, V_j)-edges, where i and j are in $\{1, \ldots, l\}$.*

The following immediate corollary is relevant to our purposes.

Corollary 1.2.10. *Every 4-connected graph contains two edge-disjoint spanning trees.*

Therefore Question 1.2.8 can be answered in the affirmative for all 4-connected graphs. The question remains: What about graphs of connectivity two and three?

E. D. Demaine, M. Hajiaghayi, and B. Mohar [5] have investigated the problem of partitioning the edge set of a graph embedded on a surface, such that contracting any partite set bounds the tree-width. They proved the following very powerful theorem.

Theorem 1.2.11 (Demaine, Hajiaghayi, and Mohar). *For a fixed genus g any integer $k \geq 2$, and for every graph G of Euler genus at most g, the edges of G can be partitioned into k sets such that contracting any one of the sets results in a graph of tree width at most $O(g^2 k)$.*

In Chapter 3, we examine a few specific surfaces—namely the plane, the projective plane, and the torus—to improve significantly the bounds obtained in [5]. We prove the following two theorems.

Theorem 1.2.12. *For any projective planar graph G, there is a bipartition $\{X, Y\}$ of $E(G)$ such that G/X and G/Y have tree-width at most three.*

Theorem 1.2.13. *If G is a toroidal graph, then there is a bipartition $\{X, Y\}$ of $E(G)$ such that $tw(G/X) \leq 3$ and $tw(G/Y) \leq 4$.*

Theorem 1.2.12 is restated and proved in Section 3.3 as Theorem 3.3.1. Theorem 1.2.13 is restated and proved in Section 3.5 as Theorem 3.5.1. Question 1.2.8 in general remains unanswered, but Theorems 1.2.12 and 1.2.13 can be considered partial results.

In 1974, F. Jaeger [10] made the following conjecture, which has become known as Jaeger's Dual-Hamiltonian Conjecture.

Conjecture 1.2.14 (Jaeger; Böhme; Oporowski)**.** *Every cubic, internally* 4*-connected graph is a union of two trees.*

Since then, both T. Böhme [3] and B. Oporowski [17] have independently made the same conjecture. Originally, we hoped to find some connection between Question 1.2.8 and Conjecture 1.2.14; none was found. Yet in our investigations, we encountered the operation we call a *switch*. Suppose that e is a non-loop edge, with endpoints u, v, of a cubic graph or cubic multigraph. If e is not a doubled edge, and if there is no loop at e, then we define a *switch* as follows, where a, c, v and u, b, d are the neighbors of u and v, respectively (note that a, b, c, d are not necessarily distinct): Let G' be the graph or multigraph obtained from G by deleting an edge cu and edge vb, and adding an edge ub and an edge cv. See Figure 1.4. Then we say G' is obtained from G via a *switch on the edge uv of the edges cu and vb*. More generally, we say that G' is obtained from G via a *switch*. If e is a doubled edge, or if there is a loop at e, then a *switch* on e is defined similarly, as illustrated in Figure 1.4.

If G is a k-connected multigraph with $k \in \{1, 2, 3\}$, and if G admits a certain switch which produces a k-connected multigraph, then that switch is called a *k-switch*. If G is internally 4-connected and admits a certain switch which produces an internally 4-connected multigraph, then that switch is called a 4*-switch*.

If we can perform a sequence of switches on a multigraph G to obtain G', then we say that G is *equivalent to G'*. If, furthermore, each switch in the sequence is a k-switch, with $k \in \{2, 3, 4\}$, then we say that G is *k-equivalent* to G'.

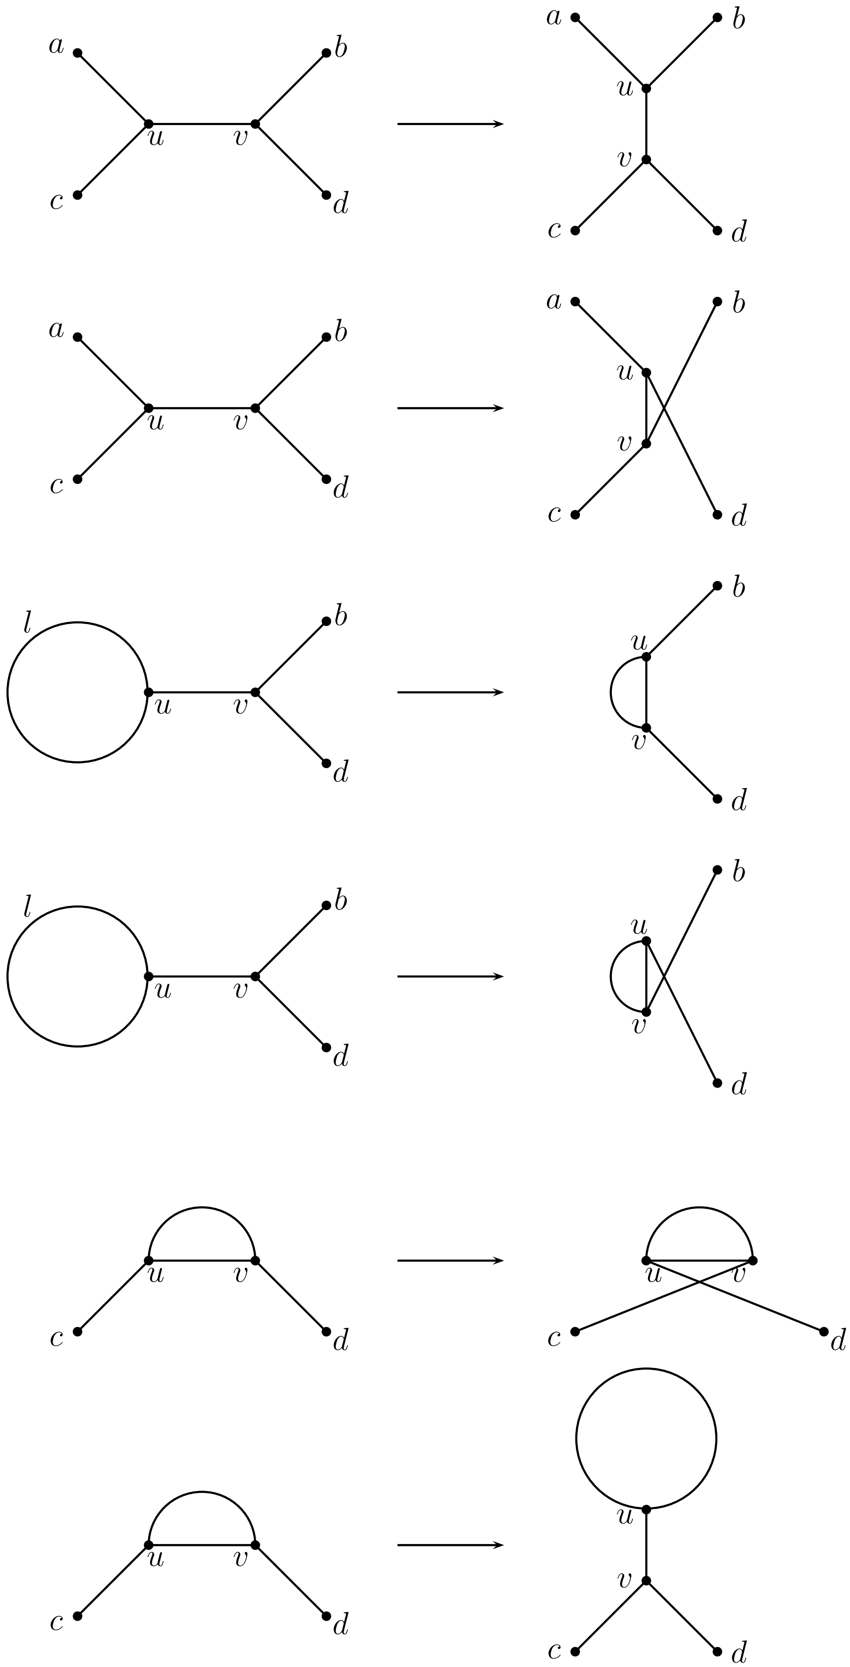

FIGURE 1.4. Top: a switch on uv, of cu and vb. Second from top: a switch on uv, of cu and vd. Third and fourth from top: two switches on uv, when there is a loop at u. Bottom and second from bottom: two switches on an edge with endpoints u, v, when that edge is doubled.

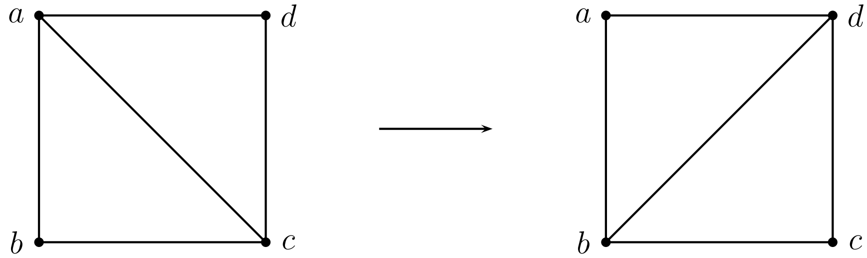

FIGURE 1.5. A diagonal flip on ac.

If u and v are adjacent vertices in G, then a *swap* on u and v is a sequence of two switches which results in a multigraph isomorphism σ, such that $\sigma(u) = v$, $\sigma(v) = u$, and $\sigma(x) = x$ when $x \notin \{u, v\}$. Therefore, if we perform a swap on u and v in $G = \{V, E\}$, then the multigraph we obtain has vertex set V, and edge set $\{\sigma(x)\sigma(y)|xy \in E\}$. If the two switches constituting a swap are 4-switches, then we call the swap a *4-swap*.

If P is a path in G, then we may also perform a switch *along* P. We will define and investigate this operation, along with the corresponding compound operation called the *path-switch*, in Section 2.3.

In the literature, the switch has been called the *edge-slide* [16], when referring to graph embedded on surfaces. If G is a cubic graph 2-cell embedded in a surface, then a surface dual G^* is a triangulation of the surface. If we perform a switch on an edge of G which respects the embedding of G, then we can consider the dual operation to the switch, which is called a *diagonal flip*. See Figure 1.5.

This operation has been studied extensively by many people, and in 1999, Negami published a vast survey of then-current results [16]. See also [13] and [15].

In Chapter 2, we demonstrate the versatility of this operation by proving the following theorems, which, unlike all prior results, are not restricted to graphs embedded on a surface.

Theorem 1.2.15. *If G and H are connected, cubic multigraphs on the same vertex set, then G and H are* 1*-equivalent.*

Theorem 1.2.16. *If G and H are cubic, internally* 4*-connected graphs on the same vertex set, then G is* 4*-equivalent to H.*

Theorem 1.2.15 and its proof appear in Section 2.5 as Corollary 2.5.2. Theorem 1.2.16 and its proof appear also in Section 2.5 as Corollary 2.5.6.

In Chapter 4, we present a structure theorem for cubic, internally 4-connected graphs, which is a strengthening of Kotzig's well-known structure theorem [12] for such graphs.

Chapter 2
Switches in Cubic Graphs

2.1 Introduction

In this chapter we investigate switches. Our ultimate goals here are Corollaries 2.5.2, 2.5.3, 2.5.4, 2.5.5, and 2.5.6. (The statements of Corollaries 2.5.2 and 2.5.6 were mentioned in Section 1.2 as Theorems 1.2.15 and 1.2.16.) Our method here is to prove that there are switches which transform any cubic, connected multigraph (on at least four vertices) into an internally 4-connected graph; we then prove that there are switches which transform that internally 4-connected graph into a circular ladder (see Figure 2.3). Furthermore, we can ensure that each switch maintains the connectivity of the multigraphs, up to internal 4-connectedness. With the ability to transform any cubic, connected multigraph into a circular ladder, our main results follow easily.

In Section 2.2, we prove a number of technical lemmas about cuts and switches that facilitate the proofs in later sections. In Section 2.3, we define and investigate the *path* 4*-switch*, which, along with the 4*-swap* from Section 2.2, is the primary tool used in the proof of Theorem 2.5.1. In Section 2.4, we find switches to transform any connected, cubic multigraph into an internally 4-connected graph. In Section 2.5, we complete the proofs of our main results.

2.2 Properties of Cuts and Switches.

In this section we prove a number of technical lemmas which are used in the proofs of subsequent lemmas and theorems. The next lemma will be used repeatedly throughout the chapter, to guide our search for switches which maintain and manipulate certain connectivity properties.

Lemma 2.2.1. *Let G be a cubic, internally* 4*-connected graph containing distinct vertices a, b, c, d, u, v and edges au, cu, uv, vb, vd. Suppose that G is isomorphic to neither K_4 nor $K_{3,3}$. If a switch on uv of cu and vb is not a* 4*-switch, then au and vb lie in some essential* 4*-cut in G.*

Proof. Let G be a cubic, internally 4-connected graph containing distinct vertices a, b, c, d, u, v and edges au, cu, uv, vb. Suppose that G is isomorphic to neither K_4 nor $K_{3,3}$. Note that G therefore has at least eight vertices. Assume that a switch on uv of cu and vb is not a 4-switch. Then the graph G' obtained from G via the aforementioned switch contains a nonverticial 3-separation $\{x, y, z\}$ and a corresponding essential 3-cut $\{e_x, e_y, e_z\}$, where vertex p is incident to e_p, for all $p \in \{x, y, z\}$. Since G' has at least eight vertices, we know that if $\{x, y, z\} \cap \{u, v\} = \emptyset$, then $\{x, y, z\}$ would be a nonverticial 3-separation of G; but this is a contradiction, since G is internally 4-connected. Therefore one of x, y, z lies in $\{u, v\}$. Without loss of generality, suppose that $x = u$. If at least one of y, z lies in $\{a, b, c, d, v\}$, say y, then by case-checking we see that $\{u, y, z\}$ is a nonverticial 3-separation of G. Hence neither y nor z lies in $\{a, b, c, d, u, v\}$, and we see that $\{u, v, y, z\}$ is a 4-separation of G. Since $\{u, y, z\}$ is a nonverticial 3-separation of G', the structure of $G'(\{a, b, c, d, u, v\})$ demonstrates that $\{au, vb, e_y, e_z\}$ is a 4-cut of G. Notice that no two of au, vb, e_y, e_z are adjacent; for otherwise we could find a nonverticial 3-cut in G. Therefore $\{au, vb, e_y, e_z\}$ is an essential 4-cut in G. $\square$

The next three lemmas express properties of cuts, which we will exploit in our search for suitable switches.

Lemma 2.2.2. *Let G be a 2-connected cubic graph, let K be an essential 2-cut of G, and let K' be an essential k-cut of G, with $k \in \{2, 3\}$. If $ab \in K$, then a and b are not endpoints of distinct edges of K'.*

Proof. Let G be a 2-connected cubic graph, let $K = \{ab, cd\}$ be an essential 2-cut of G, and let K' be an essential k-cut of G, such that $k \in \{2, 3\}$. Let ab be an edge of C. Suppose, en route to a contradiction, that ae and bf are distinct edges of K'. Let $\{A, B, C, D\}$ be a partition of the vertex set of G, where $\{A \cup B, C \cup D\}$ is the partition induced by K, and $\{A \cup C, B \cup D\}$ is the partition induced by K'. Without loss of generality, assume that $a \in A$. Since $ab \notin K'$, we know that $b \in C$. Then $e \in B$ and $f \in C$. We know that exactly one of c, d lies in $A \cup B$; suppose without loss of generality that $c \in A \cup B$.

Case 1. Suppose that $K' = \{ae, bf\}$. If $c \in A$, then B has precisely one edge of incidence. If $c \in B$, then a is a cut-vertex. In either case, the 2-connectedness of G is violated.

Case 2. Suppose that $K' = \{ae, bf, gh\}$.

Case 2a. Suppose that $c \in A$. If $d \in C$, then the member of $\{B, D\}$ containing neither of g, h has precisely one edge of incidence. If $d \in D$, then e is a cut-vertex of G. This contradicts the 2-connectedness of G.

Case 2b. Suppose that $c \in B$. If $d \in C$, then a is a cut-vertex of G. If $d \in D$, then the member of $\{A, C\}$ containing neither of g, h has precisely one edge of incidence. This contradicts the 2-connectedness of G. $\square$

Lemma 2.2.3. *Let G be a 3-connected cubic graph, let K be an essential 3-cut such that one partite set is of minimum size, and let K' be any essential k-cut of G, with $k \in \{3, 4\}$. If $ab \in K$, then a and b are not endpoints of distinct edges of K'.*

Proof. Let G be a 3-connected cubic graph, let $K = \{ab, cd, ef\}$ be an essential 3-cut such that one partite set is of minimum size, and let K' be any essential k-cut of G, where $k \in \{3, 4\}$. Suppose, en route to a contradiction, that ag and bh are edges of K'. Let $\{A, B, C, D\}$ be a partition of the vertex set of G, where $\{A \cup B, C \cup D\}$ is the partition induced by K, and $\{A \cup C, B \cup D\}$ is the partition induced by K'. Suppose without loss of generality that $A \cup B$ is the aforementioned partite set of minimum size. Furthermore, suppose without loss of generality $a \in A$. Then since $ab \notin K'$, we know that $b \in C$. Therefore $g \in B$ and $h \in D$. Furthermore, precisely one of c, d, and precisely one of e, f lies in $A \cup B$; suppose without loss of generality that c and e lie in $A \cup B$. Notice that d and f lie in $C \cup D$.

Case 1. Suppose that $K' = \{ag, bh, ij\}$. Notice that $\{i, j\} \cap \{c, d, e, f\}$ may be nonempty.

Case 1a. Suppose that c and e lie in A. Then B is nonempty and has at most two edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 1b. Suppose that $c \in A$, $e \in B$, and $d \in C$. Then one of B, D has at most two edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 1c. Suppose that $c \in A$, $e \in B$, and $d \in D$. Then $f \in D$, and we see that C has at most two edges of incidence and at most two vertices of incidence.

Case 1d. Suppose that $c \in B$, $e \in A$, $d \in C$. Then $f \in C$, and we see that D has at most two edges of incidence and at most two vertices of incidence.

Case 1e. Suppose that $c \in B$, $e \in A$, and $d \in D$. Then the member of $\{A, C\}$ which contains neither of i, j has at most three edges of incidence and precisely two vertices of incidence. This contradicts the 3-connectedness of G.

Case 1f. Suppose that c and e lie in B. Then either a is a cut-vertex of G or A has three edges of incidence and two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2. Suppose that $K' = \{ag, bh, ij, kl\}$. Notice that $\{i, j, k, l\} \cap \{c, d, e, f\}$ may be nonempty.

Case 2a. Suppose that $c \in A$, $e \in A$, and $d \in C$. Then a member of $\{B, D\}$ which contains the fewest members of $\{i, j, k, l\}$ is nonempty and has at most two edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2b. Suppose that $c \in A$, $e \in A$, and $d \in D$. Then B is nonempty and has at most two edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2c. Suppose that $c \in A$, $e \in B$, $d \in C$, and $f \in C$. Then one of B, D has precisely two edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2d. Suppose that $c \in A$, $e \in B$, $d \in C$, and $f \in D$. If one of A, B, C, D avoids $\{i, j, k, l\}$, then the vertices of incidence of that member of $\{A, B, C, D\}$ form a separation which violates the 3-connectedness of G. Otherwise, suppose that each of A, B, C, D meets $\{i, j, k, l\}$. Suppose without loss of generality that $i \in A$. We then discover an essential 3-cut $\{za, cd, ij\}$, where $z \in N(a) \backslash \{b, g\}$; one partite set of this essential 3-cut is $A - a$. This contradicts the minimality of $|A \cup B|$.

Case 2e. Suppose that $c \in A$, $e \in B$, $d \in D$. Then one of A, B is nonempty and has at most three edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2f. Suppose that $c \in B$, $e \in A$, $d \in C$, and $f \in C$. Then one of B, D has at most two edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2g. Suppose that $c \in B$, $e \in A$, $d \in C$, and $f \in D$. Then B has precisely three edges of incidence and precisely two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2h. Suppose that $c \in B$, $e \in A$, $d \in D$, and $f \in C$. If one of A, B, C, D avoids $\{i, j, k, l\}$, then the vertices of incidence of that member of $\{A, B, C, D\}$ form a separation which violates the 3-connectedness of G. Otherwise, suppose that each of A, B, C, D meets $\{i, j, k, l\}$. Suppose without loss of generality that $i \in A$. We then discover an essential 3-cut $\{za, ef, ij\}$, where $z \in N(a)\backslash\{b, g\}$; one partite set of this essential 3-cut is $A - a$. This contradicts the minimality of $|A \cup B|$.

Case 2i. Suppose that $c \in B$, $e \in A$, $d \in D$, and $f \in D$. Then one of B, C has at most two edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2j. Suppose that $c \in B$, $e \in B$, $d \in C$, and $f \in C$. Then a is a cut-vertex of G. This contradicts the 3-connectedness of G.

Case 2k. Suppose that $c \in B$, $e \in B$, $d \in C$, and $f \in D$. Then A has at most three edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2l. Suppose that $c \in B$, $e \in B$, $d \in D$, and $f \in C$. Then A has at most three edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G.

Case 2m. Suppose that $c \in B$, $e \in B$, $d \in D$, and $f \in D$. Then one of A, B has at most three edges of incidence and at most two vertices of incidence. This contradicts the 3-connectedness of G. $\square$

Lemma 2.2.4. *Let G be a cubic, internally 4-connected graph containing distinct vertices a, b, d, u, v and edges au, uv, vb, vd. If C_1 and C_2 are essential 4-cuts, each containing au, then $C_1 \cup C_2$ cannot contain both vb and vd.*

Proof. Let G be a cubic, internally 4-connected graph containing distinct vertices a, b, d, u, v and edges au, uv, vb, vd. Let C_1 and C_2 be essential 4-cuts, each containing au. Suppose, en route to a contradiction, that $C_1 \cup C_2$ contains both vb and vd. Note that neither C_1 nor C_2 contains both vb and vd. Therefore we may assume without loss of generality that $vb \in C_1$ and $vd \in C_2$.

Let $\{A, B, C, D\}$ be a partition of $V(G)$ such that $\{A \cup B, C \cup D\}$ is the bipartition induced by C_1, and $\{A \cup C, B \cup D\}$ is the bipartition induced by C_2.

Note that a and b lie in the same member of $\{A \cup B, C \cup D\}$; and in the other member lie d, u, v. Note also that a and d lie in the same member of $\{A \cup C, B \cup D\}$; and in the other member lie b, u, v. Without loss of generality, let $a \in A$. Then $b \in B$, $u \in D$, $v \in D$, and $d \in C$. No matter how the two remaining $(A \cup B, C \cup D)$-edges and the two remaining $(A \cup C, B \cup D)$-edges are arranged, one of A, B, C, D is nonempty and has at most three vertices of incidence and at most three edges of incidence. Note that none of A, B, C, D has precisely one vertex of incidence and three edges of incidence, since otherwise two of those edges would be incident to the same vertex and would lie in the same essential 4-cut. Therefore one of A, B, C, D has at least two vertices of incidence and at most three edges of incidence. This contradicts the internal 4-connectedness of G. $\square$

The next three lemmas will be used directly to find switches which increase connectivity.

Lemma 2.2.5. *Let G be a k-connected cubic graph, with $k \in \{2, 3\}$. Let G' be the graph obtained from G by performing a switch on some edge uv of G. If G' is not k-connected, then u and v are endpoints of distinct edges of an essential k-cut in G.*

Proof. Let G be a k-connected cubic graph, with $k \in \{2, 3\}$. Let G' be the graph obtained from G by performing a switch on some edge uv of G, and suppose that G' is not k-connected. Then G' has a $(k-1)$-separation. Since G' is cubic, we know that G' has an essential $(k-1)$-cut C. If at most one of u, v is an endpoint of C, then C is an essential $(k-1)$-cut of G. But since G is k-connected, we know that G has no such cut. Therefore both of u, v are endpoints of C. If $uv \notin C$, then the two edges of C incident to $\{u, v\}$ share an endpoint in G; in this case, we can find a $(k-2)$-cut in G; this contradicts the connectivity of G. Therefore $uv \in C$, and the two sets $N(u)\backslash v$ and $N(v)\backslash u$ lie in distinct partite sets induced by C. Let e and f be the two $(\{u, v\}, N(u)\backslash v)$-edges in G. We know that e, f, uv do not form a triangle, for otherwise G' would contain a doubled edge. Therefore $(C\backslash uv) \cup \{e, f\}$ is a k-cut in G. If $(C\backslash uv) \cup \{e, f\}$ were not essential, then G would not be k-connected. $\square$

Lemma 2.2.6. *Let G be a 2-connected cubic graph which contains precisely c distinct essential 2-cuts. Let e be an edge of an essential 2-cut in G. Then any switch on e is a 2-switch, and a graph obtained from G via such a 2-switch has at most $c-1$ essential 2-cuts.*

Proof. Let G be a 2-connected cubic graph which contains precisely c distinct essential 2-cuts, with $c > 0$. Let G' be the graph obtained from G via a switch on an edge e of an essential 2-cut C. Suppose, en route to a contradiction, that G' contains a cut-vertex. By Lemma 2.2.5, we know that u and v are endpoints of distinct edges ua, vb of an essential 2-cut C' of G. This contradicts Lemma 2.2.2. Therefore G' is 2-connected.

By Lemma 2.2.2, we know that every essential 3-cut in G either contains e, or shares at most one endpoint of e. Hence the essential 2-cuts of G' are precisely those essential 2-cuts of G which do not contain e. Since G contains at least one essential 2-cut which contains e (namely C), we see that G' has at most $c-1$ essential 2-cuts. □

Lemma 2.2.7. *Let G be a 3-connected cubic graph which contains more than six vertices and precisely c distinct essential 3-cuts, with $c > 0$. Then there is a 3-switch we may perform on some edge of G to obtain the graph G', such that G' contains at most $c-1$ distinct essential 3-cuts.*

Proof. Let G be a 3-connected cubic graph which contains more than six vertices and precisely c distinct essential 3-cuts, with $c > 0$. Let $C = \{ab, cd, ef\}$ be an essential 3-cut of G, one of whose partite (vertex) sets is of minimum size. Let G' be the graph obtained from G by performing a switch on ab. Suppose, en route to a contradiction, that G' contains a 2-separation. Since G' is cubic, we know that G' contains an essential 2-cut. By Lemma 2.2.5, we see that a and b are endpoints of two distinct edges of an essential 3-cut in G. This contradicts Lemma 2.2.3. Therefore G' is 3-connected.

By Lemma 2.2.3, we know that every essential 4-cut in G either contains ab, or shares at most one endpoint with ab. Hence the essential 3-cuts of G' are precisely those essential 3-cuts of G which do not contain ab. Since G contains at least one essential 3-cut which contains ab (namely C), we see that G' has at most $c-1$ essential 3-cuts. □

The next lemma proves the existence of the 4-swap. The main results of this chapter rely heavily on this operation.

Lemma 2.2.8. *Let G be a cubic, internally 4-connected graph isomorphic to neither K_4 nor $K_{3,3}$. If u and v are adjacent vertices of G, then there is a 4-swap on u and v.*

Proof. Let G be an internally 4-connected graph, and let u and v be adjacent vertices of G. Let $\{a, c, v\}$ and $\{b, d, u\}$ be the neighbor sets of u and v, respectively. Since G is triangle-free, we know that a, b, c, d are distinct vertices. Suppose, en route to a contradiction, that no switch on uv is a 4-switch. Then, in particular, the switch on uv of the edges au and vb is not a 4-switch; and the switch on uv of the edges au and vd is not a 4-switch. Then by Lemma 2.2.1, we know that au and vb lie in some essential 4-cut C_1, and au and vd lie in some essential 4-cut C_2; and by Lemma 2.2.4, we see that this is a contradiction. Hence G admits a 4-switch on uv. Without loss of generality, suppose that this 4-switch is of au and vb. We then perform the 4-switch on uv of cu to vd, which completes the 4-swap. □

2.3 The Path Switch

Let G be a cubic multigraph, and let $P = (v_0, v_2, \ldots, v_n)$ represent a chordless path in G such that $n \geq 2$. So P contains at least two edges. It will be convenient to perform switches *along* P, by which we mean the following:

(1) A *switch along P, on v_0v_1* is a switch on v_0v_1 of xv_0 and v_1v_2, where $x \in N(v_0)\backslash P$;

(2) A *switch along P, on $v_{n-1}v_n$* is a switch on $v_{n-1}v_n$ of xv_{n-1} and v_ny, where $x \in N(v_{n-1})\backslash P$ and $y \in N(v_n)\backslash P$;

(3) If $i \in \{1, \ldots, n-2\}$, then a *switch along P, on v_iv_{i+1}* is a switch on v_iv_{i+1} of xv_i and $v_{i+1}v_{i+2}$, where $x \in N(v_i)\backslash P$.

See Figure 2.1. Notice that when we perform a switch along P, we can consider P to have "shrunk" by one edge. In this section—specifically Lemma 2.3.3 and Corollary 2.3.4—we establish the existence of a sequence of $(n-1)$ separate 4-switches along the ever-shrinking path P, which "shrinks" P down to a single edge. We call that sequence of 4-switches a *path 4-switch* on P. We

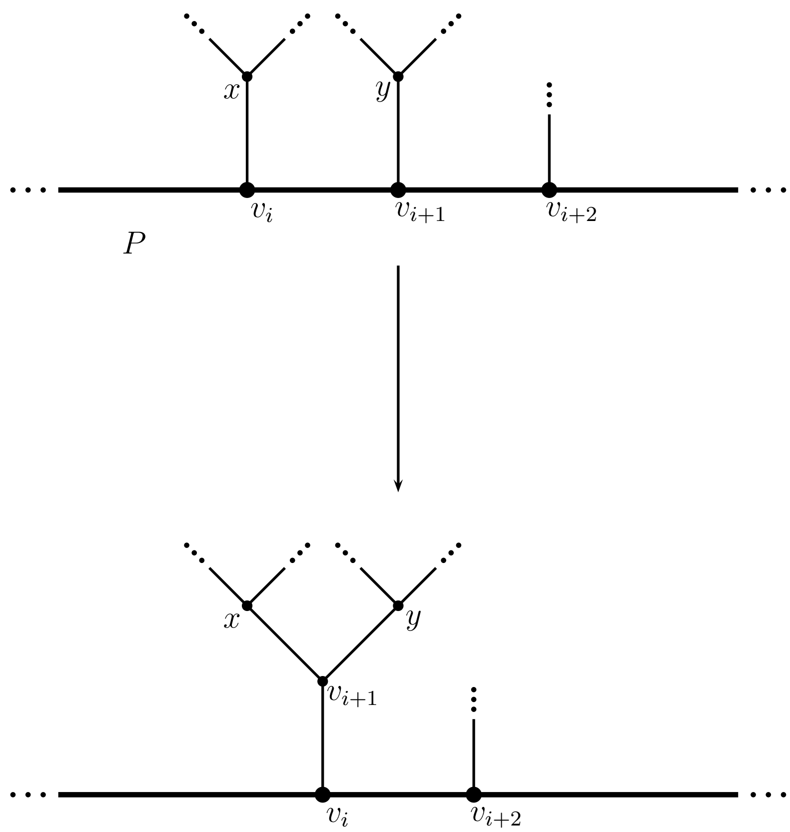

FIGURE 2.1. A switch along P, on $v_i v_{i+1}$.

first define *path switch* precisely. Let $G_1 = G$, and let $P_1 = P$. A sequence $(s_1, s_2, \ldots, s_{n-1})$ of switches is a *path switch on* P if the following three conditions are satisfied:

(1) s_i is a switch in G_i along P_i for all i;

(2) G_{i+1} is the multigraph obtained from G_i by performing s_i, for all i;

(3) For all i, if s_i is a switch on an edge whose endpoints are u_j and u_{j+1}, and P_i is a path described by $(u_1, u_2, \ldots, u_k)$, then P_{i+1} is the path described by $(u_1, \ldots, u_j, u_{j+2}, \ldots, u_k)$ in G_{i+1}.

Furthermore, we say that the path switch *respects* E' if, when P is described by $(u_1, \ldots, u_{n-1})$, the following hold:

(1) E' is a set of edges in $E(G)\backslash P$, each of which is incident to an endpoint of P;

(2) Every G_i, with $i \in \{1, \ldots, n\}$, contains a sets of edges E'_i and a one-to-one correspondence $\zeta : E' \longrightarrow E'_i$, such that $\phi_{G_i}(\zeta(e)) = \phi_{G_1}(e)$ for every $e \in E'$.

If all of the switches in a path switch are 4-switches, then we call that path switch a *path 4-switch*. The next lemma establishes the existence of the path switch.

Lemma 2.3.1. *Let G be a connected, cubic multigraph, and let P be a path described by $(u_1, \ldots, u_k)$ in G. If E' is a set of at most three distinct non-loop edges not in P, each of which is incident to an endpoint of P, then there is a path switch on P respecting E'.*

Proof. Let G be a connected multigraph, and let P be a path described by $(u_1, \ldots, u_k)$ in G. If P contains only one edge, then the conclusion follows immediately. We proceed by induction. Suppose that P contains precisely k edges, with $k \geq 2$. Let E' be a set of at most three distinct non-loop edges not in P, each of which is incident to an endpoint of P. Since G is cubic, we know that there is some edge e_3 not in $P \cup E'$ which is incident to an endpoint u_i of P. Let $e_1 \in P$ be the edge incident to u_i, and let $e_2 \in P - e_1$ be the edge adjacent to e_1. Let G' be the multigraph obtained from G by performing a switch on e_1, of e_2 and e_3. By the induction principle, the conclusion follows. □

The following technical lemma describes a primary mechanism behind the existence of the path 4-switch. It is, therefore, a primary facilitator of the main results of this chapter.

Lemma 2.3.2. *Let G be a cubic, internally 4-connected graph not isomorphic to $K_{3,3}$, and let P be a chordless path in G. If P contains at least two edges, and some switch along P on $uv \in E(P)$ is not a 4-switch, then any switch along P on any edge of P which is adjacent to uv is a 4-switch.*

Proof. Let G be a cubic and internally 4-connected graph not isomorphic to $K_{3,3}$, and let P be a chordless path in G. Suppose that uv and vw are distinct, incident edges of P, and suppose furthermore that some switch s_1 along P on uv is not a 4-switch. Finally suppose, en route to a contradiction, that a switch s_2 along P on vw is not a 4-switch.

We claim that there exist distinct vertices x, y, z incident to u, v, w, respectively, which do not lie on P. Since P has no chords, we know that each of u, v, w has a neighbor which does not lie on P. Since G is triangle-free, we know that u and v do not share a neighbor; similarly, v and w do not share a neighbor. Suppose that the neighbor sets of u and w are equal. We see that their neighbor set is a nonverticial 3-separation of G, unless the single neighbor of v not contained in P shares their neighbor set as well; but in that case G is isomorphic to $K_{3,3}$. Either case leads to a contradiction. Therefore the neighbor sets of u and w are not equal. Hence we can find distinct vertices x, z which do not lie on P, such that x is a neighbor of u, and z is a neighbor of w. Let y be the unique neighbor of v not contained in P.

By Lemma 2.2.1, we know that ux and vy lie in some essential 4-cut C_1, and we know that vy and wz lie in some essential 4-cut C_2. If $wz \in C_1$, then we find an essential 3-cut in G, namely $\{mu, wn, e\}$, where $m \in N(u)\backslash\{x, v\}$, $n \in N(w)\backslash\{z, v\}$, and e is the unique member of $C_1\backslash\{ux, vy, wz\}$. See Figure 2.2. This contradicts the internal 4-connectedness of G.

Therefore $wz \notin C_1$. By a similar argument, we may assume that $ux \notin C_2$. Since G is internally 4-connected, we know also that no two edges of C_1 are incident, for we could then find a nonverticial 3-cut in G. Similarly, we know that no two edges of C_2 are incident.

We partition the vertices of G into four components, A, B, C, D, where C_1 consists of the $(A \cup B, C \cup D)$-edges, and C_2 consists of the $(A \cup C, B \cup D)$-edges. Without loss of generality, suppose that $u \in A$ and $x \in C$. Since $v \in C_1 \cap C_2$, we know that either $v \in A$ and $y \in D$, or $v \in B$ and $y \in C$. Since no two edges of C_1 are incident, we know that either $v \in A$ or $v \in B$. If $v \in B$, then $uv \in C_2$, which is not true. Therefore $v \in A$ and $y \in D$. Since $wz \in C_2$, we know that

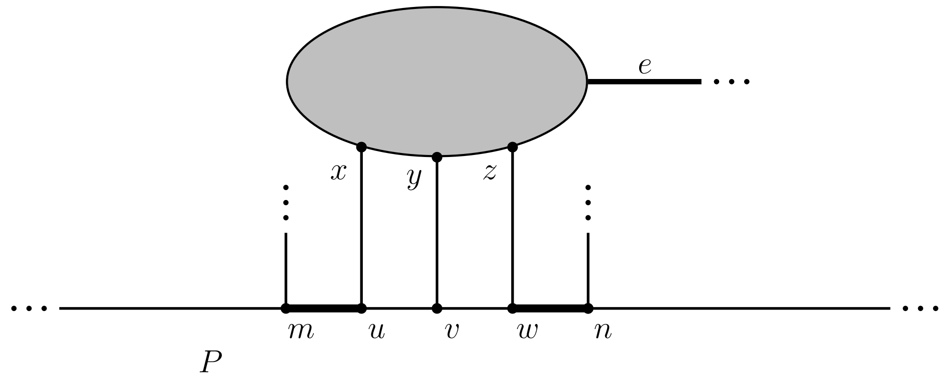

FIGURE 2.2. When ux, vy, and wz lie in C_1, we find a nonverticial 3-cut $\{mu, wn, e\}$.

$vw \notin C_2$. And since $wz \notin C_1$, we know that $w \in A$ and $z \in B$. There are at most four edges in $(C_1 \cup C_2)\backslash\{ux, vy, wz\}$; regardless of where the endpoints of these edges lie, one of B, C, D is nonempty and has at most three edges of attachment, no three of which share an endpoint. This contradicts the internal 4-connectedness of G. $\square$

The next two results are book-keepers which streamline the proof of Theorem 2.5.1.

Lemma 2.3.3. *Let G be a cubic, internally 4-connected graph, and let P be a chordless path $(v_1, v_2, \ldots, v_k)$ in G. If $x \in N(v_1)\backslash P$ and $y \in N(v_k)\backslash P$, such that $x \neq y$, then there is a path 4-switch on P respecting xv_1 and v_ky.*

Proof. Let G be a cubic, internally 4-connected graph, and let P be a chordless path $(v_1, \ldots, v_k)$ in G. Let $x \in N(v_1)\backslash P$ and $y \in N(v_k)\backslash P$. If P contains exactly one edge, then the conclusion follows. So suppose that $k \geq 2$. We proceed by induction.

Case 1. Suppose $k = 3$. If v_1v_2 admits a 4-switch of wv_1 and v_2v_3, where w is the unique member of $N(v_1)\backslash\{x, v_2\}$, then we may perform that switch to obtain the final 4-switch of the path 4-switch on P. If v_1v_2 does not admit such a 4-switch, then by Lemma 2.3.2, we know that there is a 4-switch on v_2v_3 of wv_2 and yv_3, where w is the unique member of $N(v_2)\backslash P$; this switch is the final 4-switch of the path 4-switch on P.

Case 2. Suppose $k > 3$. Let v_iv_{i+1} be a central edge of P. If v_iv_{i+1} admits a 4-switch along P, of wv_i and $v_{i+1}v_{i+2}$, where $w \neq x$, then let $P' = (v_1, \ldots, v_i, v_{i+2}, \ldots, v_k)$. If v_iv_{i+1} does not admit such a 4-switch, then by Lemma 2.3.2, there is a 4-switch along P, on $v_{i+1}v_{i+2}$, of wv_{i+1} and $v_{i+2}z$, where z is the unique member of $(N(v_{i+2})\backslash\{v_{i+2}\}) \cap (P \cup \{y\})$; let $P' = (v_1, \ldots, v_i, v_{i+2}, \ldots, v_k)$. $\square$

The following is an immediate corollary of Lemma 2.3.3.

Corollary 2.3.4. *If G is a cubic, internally 4-connected graph, and P is a chordless (u,v)-path in G, then there is a path 4-switch on P.*

2.4 Introductory Results

In this section we prove that there are switches which increase, monotonically, the connectivity of a multigraph, up to internal 4-connectedness.

Lemma 2.4.1. *Any connected, cubic multigraph G with at least four vertices is equivalent to some connected, cubic graph.*

Proof. Let $G = \{V, E, \phi\}$ be a connected, cubic multigraph. If G is simple, the conclusion follows. Otherwise, we must find switches to perform which eliminate the loops and multiple edges. Let l be a loop in G, and let e be the edge adjacent to l. Then e has multiplicity one and is not a loop. Let u and v be the vertices such that $\phi(e) = \{u, v\}$ and $\phi(l) = \{u\}$. Since G has at least four vertices, we know that there is no loop at v. Let G_1 be the multigraph obtained from G by performing a switch on e.

Then u and v, in G_1, will be joined by an edge of multiplicity two. Notice that the only switch which increases the number of loops in a multigraph is a switch on an edge of multiplicity greater than one. Therefore, G_1 contains one fewer loops than G. We can continue inductively to obtain a multigraph G_k which contains no loops.

Now we must eliminate multiple edges. Notice that in a cubic multigraph with more than two vertices, every vertex is incident to an edge with multiplicity one. Let u, v, w be vertices of G_k. If u, v, w lie in a 3-cycle, and u and v are joined by an edge of multiplicity two, then after a switch on the edge incident to w and u, a pair of multiple edges will remain whose set of endpoints is either $\{w, v\}$ or $\{u, v\}$. But if u, v, w do not lie in a 3-cycle, then any switch on the edge incident to w and u will result in a multigraph which has no multiple edges between any pair of w, u, v; and furthermore, after the switch, w would not be incident to any multiple edges, even if it was so before the switch. We see then that a switch on wu, when u, v, w do not lie in a 3-cycle, reduces the number of collections of multiple edges in G_k by at least one. Therefore, it suffices for the induction to show that we may perform switches on G_k to obtain a multigraph G_k' and three vertices a, b, c of G_k' which satisfy the following:

(1) a is adjacent to b;

(2) b and c are adjacent via an edge of multiplicity two;

(3) a, b, c do not lie in a 3-cycle;

(4) G_k' is loopless and contains precisely as many multiple edges as G_k.

We know that u and v are joined by an edge of multiplicity two. If u and v do not lie in a 3-cycle, then the conclusion follows. Suppose, then, that u, v, w lie in some 3-cycle of G_k. Let x be the unique vertex adjacent to w and not in $\{u, v\}$. Then the edge xw has multiplicity one. And since $N(u) = \{v, w\}$ and $N(v) = \{u, w\}$, we know that $N(x) \cap \{u, v\} = \emptyset$. Therefore the edge xw does not lie in a 3-cycle. And since G_k contains no loops, we know there is no loop at x. We may then perform a switch on xw of the edges yx and wv, where $y \in N(x) - w$, to obtain a multigraph G_k' for which the following hold:

(1) The vertices w and u are adjacent;

(2) The vertices u and v are joined by an edge of multiplicity two;

(3) $\{u, w, v\}$ does not induce a 3-cycle;

(4) G_k' is loopless and contains precisely as many multiple edges as G_k.

The conclusion therefore follows. □

Lemma 2.4.2. *Any connected, cubic graph G is equivalent to some 2-connected cubic graph.*

Proof. Let G be a connected, cubic graph. We perform induction on the number of blocks of G. If G is a block, then the conclusion follows. For the induction step, assume that G has $k+1$ blocks, and let B be a block which is also a leaf of the resulting block-tree. Since G is cubic, there is an edge uv which separates B from $G - B$. Suppose, without loss of generality, that u is contained in the block corresponding to a leaf of the tree. Let G' be the graph obtained from G

by performing a switch on uv. Then we obtain the block tree of G' from the block tree of G by deleting uv and combining the vertices in the block which contains u to the block which contains v. Then G' has k blocks. By the induction assumption, G is equivalent to some 2-connected cubic graph. This proves the first statement of the Lemma. □

Lemma 2.4.3. *Any* 2*-connected, cubic graph* G *is* 2*-equivalent to a* 3*-connected graph.*

Proof. Let G be a 2-connected, cubic graph. We proceed by induction on the number of distinct, essential 2-cuts of G. If G contains no essential 2-cuts, then since G is cubic, we know that G is 3-connected. Suppose now that G has k distinct essential 2-cuts. Then by Lemma 2.2.6, there is a 2-switch on G we may perform to obtain a graph G' which contains at most $k-1$ distinct essential 2-cuts. By the induction assumption, we see that G is 2-equivalent to some graph G'' which contains no essential 2-cuts. Since G'' is cubic, we know that G'' is 3-connected. □

Lemma 2.4.4. *Any* 3*-connected, cubic graph* G *is* 3*-equivalent to an internally* 4*-connected graph.*

Proof. Let G be a 3-connected, cubic graph. We proceed by induction on the number of distinct, essential 3-cuts of G. If G contains no essential 3-cuts, then since G is cubic and 3-connected, we know that G is internally 4-connected. Suppose now that G has k distinct essential 3-cuts. By Lemma 2.2.7, we know that there is a 3-switch we may perform on G to obtain a graph G' which contains at most $k-1$ essential 3-cuts. By the induction assumption, we see that G is 3-equivalent to some graph G'' which contains no essential 3-cuts. Since G'' is cubic, we know that G'' is internally 4-connected. □

2.5 Main Results

We are now ready to prove that any cubic, internally 4-connected graph is 4-equivalent to the circular ladder.

Theorem 2.5.1. *Every cubic, internally* 4*-connected graph* G *with vertex set* $\{v_1, v_2, \ldots, v_{2n}\}$*, with* $n \geq 4$*, is* 4*-equivalent to the circular ladder,* L_n*, as specified in Figure 2.3.*

Proof. Let G be a cubic, internally 4-connected graph with at least four vertices. Let $\{v_1, v_2, \ldots, v_{2n}\}$ be the vertex set of G. Let P_1 be a chordless (v_1, v_2)-path. Let G_1 be the graph obtained from

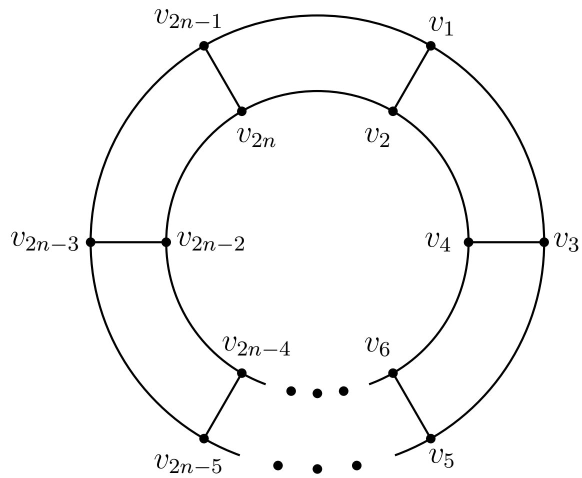

FIGURE 2.3. The circular ladder L_n.

G by performing the path 4-switch from Corollary 2.3.4 on P_1. Let $H_1 = \{v_1, v_2\}$. Then we see that $G_1(H_1)$ is one "rung" of the ladder L_n. Notice that H_1 has four distinct vertices of incidence (with respect to G_1). We proceed by induction. Suppose that we have G_i and H_i, and H_i has four distinct vertices of incidence (with respect to G_i), and G_i is internally 4-connected.

Case 1. Suppose that $i \leq n-3$. If every (v_{2i+1}, v_{2i+2})-path meets H_i, then $G_i \backslash H_i$ is disconnected, and some proper subset of the vertices of incidence of H_i forms a nonverticial separation of G_i. This is a contradiction. Hence there is a (v_{2i+1}, v_{2i+2})-path which avoids H_i; let P_2 be a shortest such path. Notice that P_2 is chordless. Let G_i' be the graph obtained from G_i by performing the following procedure: if P_2 contains more than one edge, perform the path 4-switch along P_2 given by Corollary 2.3.4. Since P_2 avoids H_i in G_i, we may consider H_i as a subgraph of G_i', since it was not modified by the path 4-switch along P_2. By a slightly extended version of Menger's Theorem ([6], p. 62), we know that if there do not exist two vertex-disjoint $(\{v_{2i-1}, v_{2i}\}, \{v_{2i+1}, v_{2i+2}\})$-paths which avoid $H_i \backslash \{v_{2i-1}, v_{2i}\}$, then a single vertex z separates $\{v_{2i-1}, v_{2i}\}$ from $\{v_{2i+1}, v_{2i+2}\}$ in $G_i' \backslash (H_i \backslash \{v_{2i-1}, v_{2i}\})$; in this case, we see that $\{z, v_{2i-1}, v_{2i}\}$ forms a nonverticial 3-separation in G_i'; and this, we see, is a contradiction. Hence there are two vertex-disjoint $(\{v_{2i-1}, v_{2i}\}, \{v_{2i+1}, v_{2i+2}\})$-paths P_3, P_4 which avoid $H_i \backslash \{v_{2i-1}, v_{2i}\}$. Let P_5 be the (v_{2i-1}, v_{2i})-path formed by first tracing P_3, then tracing the edge $v_{2i+1}v_{2i+2}$, then tracing P_4 backwards. Let G_i'' be the graph formed from G_i' by performing 4-swaps (obtained from Lemma 2.2.8) on interior pairs of adjacent vertices of P_5, so that v_{2i+1} is brought into the neighborhood of v_{2i-1}, and v_{2i+2} is brought into the neighborhood of v_{2i}. Let P_6 be the (v_{2n+1}, v_{2n+2})-subpath of this reordered P_5. Let G_{i+1} be the graph formed from G_i'' by performing the path 4-switch on P_6 respecting $v_{2n-1}v_{2n+1}$ and $v_{2n+2}v_{2n}$ obtained from Lemma 2.3.3. Let $H_{i+1} = H_i \cup \{v_{2i+1}, v_{2i+2}\}$. Notice that G_{i+1} and H_{i+1} are internally 4-connected, and H_{i+1} has four distinct vertices of attachment to G_{i+1}.

Case 2. Suppose that $i = n-2$. Then H_i is a 4-cycle, and there are precisely four $(H, \{v_1, v_2, v_{2n-4}, v_{2n-5}\})$-edges. Using Lemma 2.2.8, we may perform 4-swaps on pairs of vertices of H_i, to obtain a graph G_i' in which v_{2n-2} is adjacent to v_{2n-4}. Similarly, we may perform 4-swaps on pairs of vertices of $H_i \backslash v_{2n-2}$, if necessary, to obtain a graph G_i'' in which v_{2n-3} is adjacent to v_{2n-5}. If v_{2n-2} and v_{2n-3} are adjacent, then we let $H_{i+1} = H_i \cup \{v_{2n-2}, v_{2n-3}\}$ and $G_{i+1} = G_i''$. If v_{2n-2} and v_{2n-3} are not adjacent, then H_i is a 4-cycle with cyclic ordering $(v_{2n}, v_{2n-3}, v_{2n-1}, v_{2n-2})$; in this case, we let G_i''' be the graph obtained from G_i'' by performing a 4-switch on $v_{2n-1}v_{2n-3}$ of $v_{2n-1}v_{2n-2}$ and $v_{2n}v_{2n-3}$; we then let $H_{i+1} = H_i \cup \{v_{2n-2}, v_{2n-3}\}$ and $G_{i+1} = G_i'''$.

Case 3. Suppose that $i = n-1$. Then v_{2n} and v_{2n-1} are adjacent. If v_1 and v_{2n-3} are adjacent to v_{2n-1}, then G_i equals L_n. If v_1 and v_{2n-3} are adjacent to v_{2n}, then we let G_n be the graph obtained from G_i by first performing the 4-switch on $v_{2n}v_{2n-1}$ of v_1v_{2n} and v_2v_{2n-1}, and then performing the 4-switch on $v_{2n}v_{2n-1}$ of $v_{2n-1}v_{2n-2}$ and $v_{2n}v_{2n-3}$. Then G_n equals L_n. If v_1 and v_{2n-2} are adjacent to v_{2n}, then we let G_n be the graph obtained from G_i by performing the 4-switch on $v_{2n-1}v_{2n}$ of v_1v_{2n} and v_2v_{2n-1}. Then G_n equals L_n. Finally, if v_1 and v_{2n-2} are adjacent to v_{2n-1}, then we let G_n be the graph obtained from G_i by performing the 4-switch on $v_{2n-1}v_{2n}$ of $v_{2n-1}v_{2n-2}$ and $v_{2n}v_{2n-3}$. Then G_n equals L_n. □

The following corollaries are the primary goals of this chapter.

Corollary 2.5.2. *If G and H are connected, cubic multigraphs on the same vertex set, then G is* 1-*equivalent to H.*

Proof. Let G and H be connected, cubic multigraphs on some vertex set V. Using Lemmas 2.4.1, 2.4.2, 2.4.3, 2.4.4, and Theorem 2.5.1, we obtain sequences of switches $\{s_1, \ldots, s_m\}$ and $\{t_1, \ldots, t_n\}$ which transform G and H, respectively, into the circular ladder. To obtain H from

G, we perform the switches $s_1, \ldots, s_m$ and then, with each switch reversed, the switches $t_n, \ldots, t_1$. $\square$

Corollary 2.5.3. *If G and H are connected, cubic graphs on the same vertex set, then G is 1-equivalent to H.*

Proof. Let G and H be cubic graphs on some vertex set V. Using Lemmas 2.4.2, 2.4.3, 2.4.4, and Theorem 2.5.1, we obtain sequences of switches $\{s_1, \ldots, s_m\}$ and $\{t_1, \ldots, t_n\}$ which transform G and H, respectively, into the circular ladder. To obtain H from G, we perform the switches $s_1, \ldots, s_m$ and then, with each switch reversed, the switches $t_n, \ldots, t_1$. $\square$

Corollary 2.5.4. *If G and H are 2-connected, cubic, graphs on the same vertex set, then G is 2-equivalent to H.*

Proof. Let G and H be 2-connected, cubic graphs on some vertex set V. Using Lemmas 2.4.3, 2.4.4, and Theorem 2.5.1, we obtain sequences of 2-switches $\{s_1, \ldots, s_m\}$ and $\{t_1, \ldots, t_n\}$ which transform G and H, respectively, into the circular ladder. To obtain H from G, we perform the 2-switches $s_1, \ldots, s_m$ and then, with each switch reversed, the 2-switches $t_n, \ldots, t_1$. $\square$

Corollary 2.5.5. *If G and H are 3-connected, cubic graphs on the same vertex set, then G is 3-equivalent to H.*

Proof. Let G and H be 3-connected, cubic graphs on some vertex set V. Using Lemmas 2.4.4, and Theorem 2.5.1, we obtain sequences of 3-switches $\{s_1, \ldots, s_m\}$ and $\{t_1, \ldots, t_n\}$ which transform G and H, respectively, into the circular ladder. To obtain H from G, we perform the 3-switches $s_1, \ldots, s_m$ and then, with each switch reversed, the 3-switches $t_n, \ldots, t_1$. $\square$

Corollary 2.5.6. *If G and H are internally 4-connected, cubic graphs on the same vertex set, then G is 4-equivalent to H.*

Proof. Let G and H be internally 4-connected, cubic graphs on some vertex set V. Using Theorem 2.5.1, we obtain sequences of 4-switches $\{s_1, \ldots, s_m\}$ and $\{t_1, \ldots, t_n\}$ which transform G and H, respectively, into the circular ladder. To obtain H from G, we perform the 4-switches $s_1, \ldots, s_m$ and then, with each switch reversed, the 4-switches $t_n, \ldots, t_1$. $\square$

Chapter 3
Bounding Tree-Width under Contraction

3.1 Introduction

In Sections 2 and 3 of this chapter, we will work intimately with the *unit disc model* of the projective plane; for notational ease, we refer to the unit disc as U. When we do so, we refer specifically to the unit disc as a subspace of the plane $\mathbb{R}^2$, in which ∂U is the unit circle. From U, of course, we obtain the projective plane by identifying antipodal boundary points. Yet for certain topological arguments, we will need to speak of the boundary of U, as a subspace of $\mathbb{R}^2$, in relation to a multigraph embedded in the unit disc model of the projective plane.

Our goals in this chapter are to prove Theorem 1.2.12 (which is restated in Section 3 as Theorem 3.3.1) and Theorem 1.2.13 (which is restated in Section 5 as Theorem 3.5.1). In Section 2, we prove a variety of technical lemmas used in the proof of Theorem 3.3.1. Our method overall, in Sections 2 and 3, is roughly as follows:

(1) Reduce the problem to the case of cubic, 2-connected graphs;

(2) Look at a surface dual G^* of an arbitrary cubic, 2-connected projective plane graph G;

(3) Find a disc in the projective plane which contains all the vertices of G^* and which induces a connected (spanning) subgraph;

(4) Decompose G^* into nested subgraphs called *distance layers*;

(5) Obtain a bipartition of $E(G^*)$ by grouping edges in alternating distance layers;

(6) Prove that the corresponding bipartition $\{X, Y\}$ of $E(G)$ satisfies the theorem: namely, that G/X and G/Y have tree-width at most three.

In Section 4, we prove some technical lemmas used in the proof of Theorem 3.5.1. Our method, overall, in Sections 4 and 5, is roughly as follows:

(1) Reduce the problem to cubic, 2-connected toroidal graphs;

(2) Prove that certain 4-connected plane triangulations admit edge partitions into two outerplanar graphs;

(3) Use (2) to prove that all planar graphs admit a special edge partition into two series-parallel graphs;

(4) Find a suitable set of pairwise non-adjacent edges in our toroidal graph whose deletion produces a planar graph;

(5) Use the partition from (3) on the planar graph from (4) to produce an edge partition $\{X, Y\}$ of our toroidal graph G;

(6) Prove that $\{X, Y\}$ satisfies the theorem: namely, that $tw(G/X) \leq 3$ and $tw(G/Y) \leq 4$.

3.2 The Case of the Projective Plane – Introductory Results

Arnborg, Corneil, and Proskurowski [1] proved the following forbidden-minor characterization of graphs with tree-width at most three.

Theorem 3.2.1 (Arnborg, Corneil, and Proskurowski)**.** *A graph G has tree-width at most three if and only if none of K_5, M_6, M_8, M_{10} is a minor of G; where M_6 is the octahedron, M_8 is the Moebius ladder on eight vertices (also called the Wagner graph), and M_{10} is the pentagonal prism, as depicted in Figure 3.1.*

The next lemma is a basic fact about embeddings in the projective plane. It will be used in the proof of a subsequent lemma.

Lemma 3.2.2. *If G is a 2-connected, non-planar multigraph embedded in the projective plane, then every edge of G lies on the boundary of two distinct faces of G.*

Proof. Let G be a 2-connected, non-planar multigraph embedded in the projective plane. Suppose, en route to a contradiction, that some edge e of G does not lie on the boundary of two distinct faces. Since G is 2-connected, we know that e is not a loop. Let F be the unique face on whose boundary e lies. Then there is a simple closed curve α such that the following hold:

(1) $G \cap \alpha$ is a single point in the interior of e;

(2) $\alpha \backslash G \subsetneq F$.

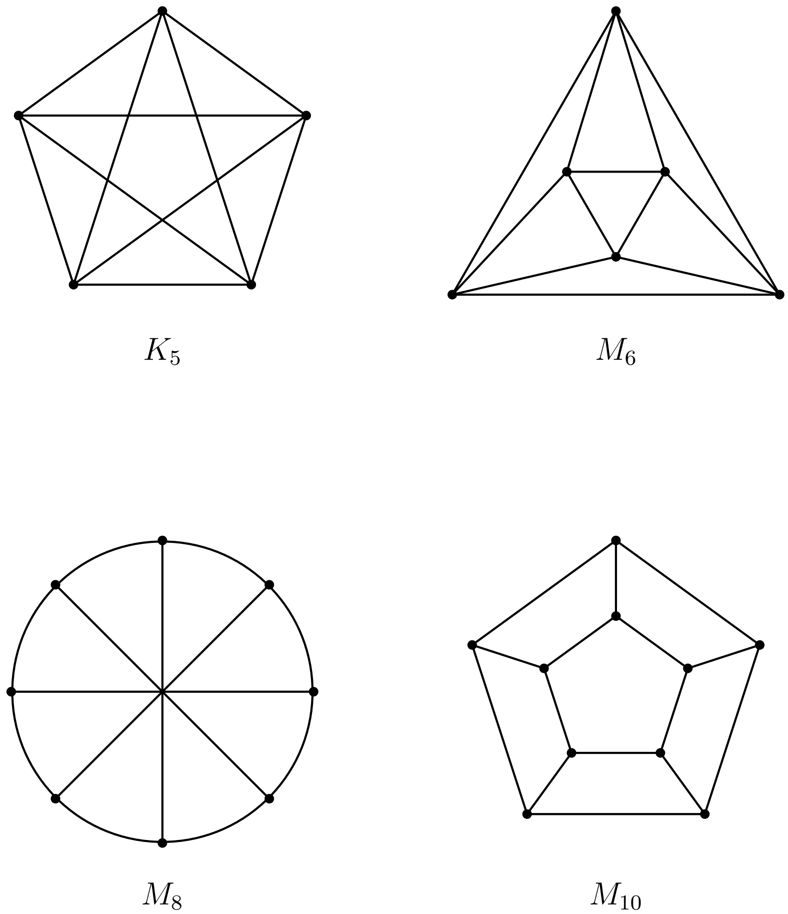

FIGURE 3.1. The forbidden minors for graphs with tree-width at most three.

If α were contractible, then one endpoint of e would necessarily be a cut-vertex; since G is 2-connected, it follows that α is non-contractible. Therefore we can find a disc D in the projective plane which bounds $G \backslash e$. Then we can map D to the plane, to obtain a planar embedding of $G \backslash e$ in which both endpoints of e lie on the boundary of the infinite face. We can then embed e in the infinite face, thus obtaining a planar embedding of G. This is a contradiction. □

The next four lemmas express basic facts about tree-width, and they show that it suffices to prove Theorems 3.3.1 and 3.5.1 for cubic, 2-connected graphs.

Lemma 3.2.3. *If G is a multigraph and H is a 2-connected minor of G, then H is a minor of some 2-connected block of G.*

Proof. Let G be a multigraph, and let H be a 2-connected minor of G. Then we can delete a set D of vertices and edges from G, and then contract a set C of edges of $G \backslash D$ to obtain a multigraph isomorphic to H.

Let S consist of all the edges ab of G, such that none of a, b, ab lies in $C \cup D$. If S lies within a single block of G, then the result follows. Suppose then, en route to a contradiction, that there are two edges ab, cd which lie in distinct blocks of G, such that none of a, b, c, d, ab, cd lies in $C \cup D$. Let P be the unique path in the block-tree of G between the components containing ab and cd. Since H is connected, we know that $G \backslash D$ is connected. Therefore D does not contain any cut-vertex of G which corresponds to a vertex of P. Therefore every path in $G \backslash D$ from an endpoint of e to an endpoint of f contains a vertex which is a cut-vertex of G; let Q be some such path, containing a cut-vertex v of G. The vertex of $(G \backslash D)/C$ which corresponds to v yields a 1-separation of $(G \backslash D)/C$. This contradicts the 2-connectedness of $(G \backslash D)/C$.$\square$

Lemma 3.2.4. *Let G be a connected multigraph, let $G_1, G_2, \ldots, G_k$ be the blocks of G, and let $m \geq 1$ and $n \geq 1$. Then if each $E(G_i)$, for $i \in \{1, \ldots, k\}$, admits a bipartition $\{X_i, Y_i\}$ such that $tw(G_i/X_i) \leq m$ and $tw(G_i/Y_i) \leq n$, then $E(G)$ admits a bipartition $\{X, Y\}$ such that $tw(G/X) \leq m$ and $tw(G/Y) \leq n$.*

Proof. Let G be a multigraph, and let $G_1, G_2, \ldots, G_k$ be the blocks of G, and let $m \geq 1$ and $n \geq 1$. If $k = 1$, then the conclusion holds. We proceed by induction. Let G_k be a leaf on the block tree of G. Suppose that $k \geq 2$, and that $E(G_1 \cup G_2 \cup \cdots \cup G_{k-1})$ and $E(G_k)$ admit bipartitions $\{X, Y\}, \{X', Y'\}$, respectively, such that the following hold:

(1) $tw\big((G_1 \cup G_2 \cup \cdots \cup G_{k-1})/X\big) \leq m$;
(2) $tw\big((G_1 \cup G_2 \cup \cdots \cup G_{k-1})/Y\big) \leq n$;
(3) $tw(G_k/X') \leq m$;
(4) $tw(G_k/Y') \leq n$.

Let $X'' = X \cup X'$ and $Y'' = Y \cup Y'$. Let $\big(T^X, \{V_t\}_{t \in V(T^X)}\big)$, $\big(T^Y, \{V_t\}_{t \in V(T^Y)}\big)$, $\big(T^{X'}, \{V_t\}_{t \in V(T^{X'})}\big)$, $\big(T^{Y'}, \{V_t\}_{t \in V(T^{Y'})}\big)$ be tree-decompositions of $\big(G_1 \cup G_2 \cup \cdots \cup G_{k-1}\big)/X$, $\big(G_1 \cup G_2 \cup \cdots \cup G_{k-1}\big)/Y$, $G_k/X', G_k/Y'$, respectively, of minimum width.

Let c be the unique vertex which constitutes $(G_1 \cup G_2 \cup \cdots \cup G_{k-1}) \cap G_k$. Let c_X and c_Y be the vertices of $\big((G_1 \cup G_2 \cup \cdots \cup G_{k-1})/X\big)$ and $\big((G_1 \cup G_2 \cup \cdots \cup G_{k-1})/Y\big)$, respectively, to which c is contracted. Let $c_{X'}$ and $c_{Y'}$ be the vertices of G_k/X' and G_k/Y', respectively, to which c is contracted. Let $t_X, t_Y, t_{X'}, t_{Y'}$ be vertices of $T^X, T^Y, T^{X'}, T^{Y'}$, respectively, such that $c_X \in V_{t_X}, c_Y \in V_{t_Y}, c_{X'} \in V_{t_{X'}}$, and $c_{Y'} \in V_{t_{Y'}}$.

Let $T^{X''}$ be the tree with vertex set $V(T^X) \cup V(T^{X'})$ and edge set $E(T^X) \cup E(T^{X'}) \cup t_X t_{X'}$. Let $T^{Y''}$ be the tree with vertex set $V(T^Y) \cup V(T^{Y'})$ and edge set $E(T^Y) \cup E(T^{Y'}) \cup t_Y t_{Y'}$. Then $\big(T^{X''}, \{V_t\}_{t \in V(T^{X''})}\big)$ and $\big(T^{Y''}, \{V_t\}_{t \in V(T^{Y''})}\big)$ satisfy (T1) and (T2). For (T3), we see that for any (t_1, t_3)-path P in $T^{X''}$ or $T^{Y''}$, the set $V_{t_1} \cap V_{t_3}$ is non-empty only if $\{t_1, t_3\}$, and thus P, is contained in one of $T^X, T^{X'}, T^Y, T^{Y'}$. Then since $\big(T^X, \{V_t\}_{t \in V(T^X)}\big)$, $\big(T^Y, \{V_t\}_{t \in V(T^Y)}\big)$, $\big(T^{X'}, \{V_t\}_{t \in V(T^{X'})}\big)$, $\big(T^{Y'}, \{V_t\}_{t \in V(T^{Y'})}\big)$ all satisfy (T3), we see that (T3) holds for $\big(T^{X''}, \{V_t\}_{t \in V(T^{X''})}\big)$ and $\big(T^{Y''}, \{V_t\}_{t \in V(T^{Y''})}\big)$. Therefore $\big(T^{X''}, \{V_t\}_{t \in V(T^{X''})}\big)$ and $\big(T^{Y''}, \{V_t\}_{t \in V(T^{Y''})}\big)$ are tree-decompositions. Since the bags of $\big(T^{X''}, \{V_t\}_{t \in V(T^{X''})}\big)$ and $\big(T^{Y''}, \{V_t\}_{t \in V(T^{Y''})}\big)$ are merely bags of $\big(T^X, \{V_t\}_{t \in V(T^X)}\big)$, $\big(T^Y, \{V_t\}_{t \in V(T^Y)}\big)$, $\big(T^{X'}, \{V_t\}_{t \in V(T^{X'})}\big)$, $\big(T^{Y'}, \{V_t\}_{t \in V(T^{Y'})}\big)$, we know that $tw\big((G_1 \cup \cdots \cup G_k)/(X'')\big) \leq m$ and $tw\big((G_1 \cup \cdots \cup G_k)/(Y \cup Y'')\big) \leq n$. $\square$

In the next lemma, we prove that "tree-width at most k" is a *minor-closed property*. That is, we prove that if the tree-width of G is at most k, then the tree-width of all minors of G is at most k.

Lemma 3.2.5. *Let G be a multigraph, let X be a set of edges in G, and let Y be a set of vertices in G. Then $tw\big((G\backslash X)\backslash Y\big) \leq tw(G)$ and $tw(G/X) \leq tw(G)$. If G' is a subdivision of G, then $tw(G') \leq \max\{tw(G), 2\}$. If G' is obtained from G by adding leaves, then $tw(G') \leq \max\{tw(G), 1\}$.*

Proof. Let G be a multigraph, let X be a set of edges in G, and let Y be a set of vertices in G. Let $\big(T, \{V_t\}_{t \in V(T)}\big)$ be a tree-decomposition of G of width $tw(G)$. Clearly $\big(T, \{V_t \backslash Y\}_{t \in V(T)}\big)$ is a tree-decomposition of $(G\backslash X)\backslash Y$. Therefore $tw\big((G\backslash X)\backslash Y\big) \leq tw(G)$.

We now prove the contraction result. Let $e \in X$, let P consist of the endpoints of e, and let v be the vertex of G/e corresponding to the contraction of e. By induction, it suffices to show that $tw(G/e) \leq tw(G)$. For each $t \in V(T)$, if $V_t \cap P = \emptyset$, let $V'_t = V_t$; otherwise, let $V'_t = (V_t \backslash P) \cup v$. We prove now that $\big(T, \{V'_t\}_{t \in V(T)}\big)$ is a tree-decomposition. Clearly $\big(T, \{V'_t\}_{t \in V(T)}\big)$ satisfies (T1) and (T2). For (T3), suppose that t_1, t_2, t_3 are vertices of T such that t_2 lies on the unique (t_1, t_3)-path in T. We know that $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$, since $\big(T, \{V_t\}_{t \in V(T)}\big)$ is a tree-decomposition of G. Clearly $(V_{t_1} \cap V_{t_3})\backslash P \subseteq V_{t_2}\backslash P$. This implies that $(V'_{t_1} \cap V'_{t_3})\backslash v \subseteq V'_{t_2}\backslash v$. If $V_{t_1} \cap V_{t_2} \cap P$ is nonempty, then $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$, and (T3) holds. Suppose, then, that $V_{t_1} \cap V_{t_2} \cap P = \emptyset$. If either of $V_{t_1} \cap P$ or $V_{t_2} \cap P$ is empty, then (T3) holds. Suppose, then, that $V_{t_1} \cap P \neq \emptyset$ and $V_{t_2} \cap P \neq \emptyset$. Let V_{t_4} be the bag of $\big(T, \{V_t\}_{t \in V(T)}\big)$ such that $e \in G[V_{t_4}]$. Then $P \subseteq V_{t_4}$. Since T is a tree, we know that t_2 is contained in either the unique (t_1, t_4)-path, or the unique (t_3, t_4)-path; in either case, the verity of (T3) in $\big(T, \{V_t\}_{t \in V(T)}\big)$ ensures that $V_{t_3} \cap P$ is nonempty. Therefore $v \in V'_{t_3}$, and we see that $V'_{t_1} \cap V'_{t_3}) \subseteq V'_{t_2}$. Hence (T3) holds in $\big(T, \{V'_t\}_{t \in V(T)}\big)$, which is therefore a tree-decomposition of G/e. Since $|V'_t| \leq |V_t|$ for each $t \in V(T)$, we see that $tw(G/e) \leq tw(G)$.

For the second part of the lemma, it suffices, by induction, to show the following: if G' is a multigraph obtained from G by subdividing one edge, then $tw(G') \leq \max\{tw(G), 2\}$. Therefore, let e be an edge of G, let P be the set of endpoints of e, and let z be the new vertex created in the subdivision. Let $\big(T, \{V_t\}_{t \in V(T)}\big)$ be a tree-decomposition of G with width $tw(G)$. Let t be a vertex of T such that $P \subseteq V_t$. Let T' be a tree obtained from T by adding a vertex t' and an edge tt'. Let $V_{t'} = P \cup z$. Then $\big(T', \{V_t\}_{t \in V(T')}\big)$ satisfies (T1) and (T2). Let t_1, t_2, t_3 be vertices of T' such that t_2 lies on the unique (t_1, t_3)-path W of T'. If $V_{t_1} \neq V_{t'} \neq V_{t_3}$, then $V_{t_2} \neq V_{t'}$, and we know that $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$, in which case (T3) holds. Suppose, then, that $V_{t_1} = V_{t'}$. If $z \in V_{t_3}$, then $t_1 = t_2 = t_3$, and (T3) holds. Suppose, then, that $z \notin V_{t_3}$. Then t_2 lies on the unique (t, t_3)-path in T', and $V_{t_1} \cap V_{t_3} \subseteq V_t \cap V_{t_3}$. Then we know that $V_t \cap V_{t_3} \subseteq V_{t_2}$. Therefore $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$, and thus (T3) holds. Hence $\big(T', \{V_t\}_{t \in V(T')}\big)$ is a tree-decomposition of G' with width $\max\{tw(G), 2\}$.

For the third part of the lemma, let G' be a graph obtained from G by adding leaves. Let $\big(T, \{V_t\}_{t \in V(T)}\big)$ be a tree-decomposition of G of width $tw(G)$. Let $S = V(G')\backslash V(G)$. For each $v \in S$, let t_v be a vertex in T such that V_{t_v} contains the unique neighbor of v in G', and let e_v be the unique edge of G' with v as an endpoint. Then $\Big(T \cup S \cup \{vt_v\}_{v \in S}, \{V_t\}_{t \in V(T)} \cup \{v, e_v\}_{v \in S}\Big)$ is a tree-decomposition of G', with width $\max\{tw(G), 1\}$. $\square$

Lemma 3.2.6. *Let G and G' be graphs, such that G' is cubic and can be obtained from G by suppressing vertices of degree two and one, and by repeatedly splitting vertices of degree greater than three. If $E(G')$ admits a partition $\{X', Y'\}$ such that $tw(G'/X) \leq m$ and $tw(G'/Y) \leq n$,* *where $m \geq 2$ and $n \geq 2$, then $E(G)$ admits a partition $\{X, Y\}$ such that $tw(G/X) \leq m$ and $tw(G/Y) \leq n$.*

Proof. Let G be a graph, let G^- be a graph obtained from G by suppressing vertices of degree two and one, and let G' be a cubic graph obtained from G^- by repeatedly splitting vertices of degree greater than three. Let S be a subset of $E(G)$ such that $G/S = G^-$, and let T be a subset of $E(G')$ such that $G'/T = G^-$.

Suppose that $E(G')$ admits a partition $\{X', Y'\}$ such that $tw(G'/X) \leq m$ and $tw(G'/Y) \leq n$, where $m \geq 3$ and $n \geq 3$. Then we can obtain a bipartition $\{X^-, Y^-\}$ of $E(G^-)$ which corresponds to $\{X', Y'\}$ by contracting T in G'. Since G^-/X^- is isomorphic to $G'/(T \cup X')$, and since $G'/(T \cup X')$ is a minor of G'/X', then by Lemma 3.2.5, we know that $tw(G^-/X^-) \leq m$. Since G^-/Y^- is isomorphic to $G'/(T \cup Y')$, and since $G'/(T \cup Y')$ is a minor of G'/X', then by Lemma 3.2.5, we know that $tw(G^-/Y^-) \leq n$.

For each edge e of G^-, there is a maximal path or cycle P_e in G, such that all internal vertices of P_e have degree two, and the contraction of all but one edge of P_e results in the edge e in G^-. Let X be the set of all edges $e \in E(G)$ such that $e \in P_f$ for some edge $f \in X^-$. Let $Y = E(G)\backslash X$. Then G/X is a multigraph obtained from G^-/X^- by adding leaves and subdividing edges. And G/Y is a multigraph obtained from G^-/Y^- by subdividing edges. Then $tw(G/X) \leq tw(G^-/X^-) \leq m$ and $tw(G/Y) \leq tw(G^-/Y^-) \leq n$, by Lemma 3.2.5. □

To clarify some proofs which follow, we use an alternate notion of projective planarity. A *plane with a crosscap* is a plane with a specified point P, called its *crosscap*. We say that a graph G can be embedded in the plane with a crosscap if we can map G to the plane such that the following are satisfied:

(PC1) The image of $V(G)$ is one-to-one;

(PC2) No vertex of G is mapped to P;

(PC3) If x and y are distinct points in $G\backslash V(G)$ whose images are equal, then x and y are mapped to P;

(PC4) If the images of distinct edges e and f intersect (necessarily at P), then there is a closed curve in the plane with a crosscap which separates P from the image of $V(G)$ and which alternately meets the images of e and f, each twice.

The edges which meet P are called *cap-edges*. We will prove that if a graph G is mapped via Γ to the plane with a crosscap and satisfies (PC1), (PC2), (PC3), and the following condition (PC4′), then the mapping may be altered to produce an embedding in the plane with a crosscap.

(PC4′) If x is an interior point of an edge e and $\Gamma(x) = P$, then there is an open subsegment s of e containing x such that $\Gamma(s)$ is a straight line segment.

Lemma 3.2.7. *If a graph G is mapped to the plane with a crosscap and satisfies (PC1), (PC2), (PC3), and (PC4′), then the mapping may be altered to produce an embedding in the plane with a crosscap.*

Proof. Let G be a graph and Γ be a mapping of G to the plane with a crosscap. Suppose that Γ satisfies (PC1), (PC2), (PC3), and (PC4′).

Case 1. Suppose that the restriction of Γ to any edge is one-to-one. Let e and f be distinct edges such that $\Gamma(e)$ and $\Gamma(f)$ contain P. Let s_e and s_f be open subsegments of e and f, respectively, given by (PC4′), and let s'_e and s'_f be closed subsegments of s_e and s_f, respectively, which contain P as an interior point. Let $\mathfrak{C}$ be a collection of open discs in the plane such that the following hold:

(1) $\bigcup_{B\in\mathfrak{C}} B \supsetneq \Gamma(s'_e \cup s'_f)$;
(2) $\left(\bigcup_{B\in\mathfrak{C}} B\right) \cap \Gamma(s'_e \cup s'_f) \subseteq \Gamma(s_e \cup s_f)$;
(3) The closure of each $B \in \mathfrak{C}$ is disjoint from $\Gamma(V(G))$.

Since $s'_e \cup s'_f$ is compact in the plane, we may suppose that $\mathfrak{C}$ is finite. Then the closure of $\bigcup_{B\in\mathfrak{C}} B$ contains a closed disc D such that the closed curve bounding D separates P from $\Gamma(V(G))$. Since $D \cap e$ and $D \cap f$ are straight line segments, we see that the closed curve bounding D alternately meets e and f, each twice. Therefore Γ satisfies (PC4) and is an embedding of G in the plane with a crosscap.

Case 2. Let e be an edge of G, let $C \subsetneq e$ consist of all points x such that $\Gamma(x) = P$, and suppose that $|C| \geq 1$. We will show that we can modify the mapping so that the restriction to e is one-to-one (thus invoking Case 1). The compactness of $\Gamma(e)$ shows that C has only finitely many connected components. In each connected component C', we can pick a point c' around which to modify Γ so that the following hold:

(1) The mapping is one-to-one on C';

(2) The image of C' is equal to $\Gamma(C')$;

(3) c' is the only element of C' mapped to P.

Thus we obtain a new map Γ' which satisfies the following:

(1) $\Gamma'(G) = \Gamma(G)$;

(2) $\Gamma'|_{G-C} = \Gamma|_{G-C}$;

(3) $\Gamma'^{-1}\{P\}$ is finite.

We proceed by induction. Let $\{c_1, c_2, \ldots, c_k\} = \Gamma'^{-1}\{P\}$, and suppose that the open subsegment of e between c_i and c_{i+1} contains no other c_j, for $1 \leq i \leq k-1$; this ordering can be obtained by traversing e from endpoint to endpoint. Let L be the image under Γ' of the closed subsegment of e between c_1 and c_2. Then L is a loop containing P.

Case 2a. Suppose that the open disc bounded by L is disjoint from $\Gamma'(G)$. Let a be a point of e such that a lies in the open subsegment of e that contains no c_i and lies between c_1 and an endpoint of e. Let b be a point of e which lies in the open subsegment of e between c_1 and c_2. Then there is a topological $\big(\Gamma'(a), \Gamma'(b)\big)$-path in the plane whose image intersects $\Gamma'(G)$ only at $\Gamma'(a)$ and $\Gamma'(b)$. Let Γ'' be the map from G to the plane with a crosscap which maps the open subsegment between a and b to the image of the topological $\big(\Gamma'(a), \Gamma'(b)\big)$-path, and which is identical to Γ' everywhere else. Then $|\Gamma''^{-1}\{P\}| < |\Gamma'^{-1}\{P\}|$, and by induction we see that we may modify the mapping so that the restriction to e is one-to-one.

Case 2b. Suppose that the open disc D bounded by L contains a point of $\Gamma'(G)$. We know that there is an open disc B about P whose intersection with $\Gamma'(G)$ consists solely of straight line segments, all meeting at P. Let $M = \Gamma'(G) \cap (D \cup B)$. With this in mind, we can view $\Gamma'(G)$ as having a "twist" at P. We can modify Γ' by "untwisting" so that D avoids G. See Figure 3.2.

The resulting map Γ'' is identical to Γ' outside of $\Gamma'^{-1}(M)$ and, importantly, the open disc bounded by L (which Γ'' inherited unchanged from Γ') is disjoint from $\Gamma''(G)$. The result follows from Case 2a. □

We now prove the desired equivalence.

Lemma 3.2.8. *A multigraph is projective planar if and only if it admits an embedding in the plane with a crosscap.*

Proof. For the left-to-right implication, suppose that G is a projective planar graph, embedded via Γ in the projective plane following the unit disc model, centered at the origin, labeled with polar coordinates. We may suppose that $\Gamma(G)$ avoids the origin, since it is easy to find a homeomorphism of U which fixes ∂U and ensures that no point of G is mapped to the origin. Furthermore, via small perturbations of Γ, we may suppose the following:

(1) No vertices of G are mapped to ∂U;

(2) For each point $x \in G \backslash V(G)$ such that $\Gamma(x) \in \partial U$, there is an open disc about $\Gamma(x)$ whose intersection with $\Gamma(G)$ consists of a straight line segment perpendicular to ∂U.

We map $\Gamma(G)$ from the unit disc model of the projective plane to the plane with a crosscap P, where P is located at the origin, via the following stereographic projection:

$$(r, \theta) \longmapsto \left(\frac{1-r}{r}, \theta\right). \tag{3.2.1}$$

Clearly, the projection restricted to $\Gamma(V(G))$ is one-to-one. Furthermore, no element of $\Gamma(V(G))$ is mapped to P, since $\Gamma(V(G))$ avoids ∂U. And we can see that (PC3) holds, since the following hold:

(1) The projection restricted to $\mathring{U}$ is one-to-one;

(2) All of ∂U is mapped to P;

(3) $\Gamma(V(G))$ avoids ∂U.

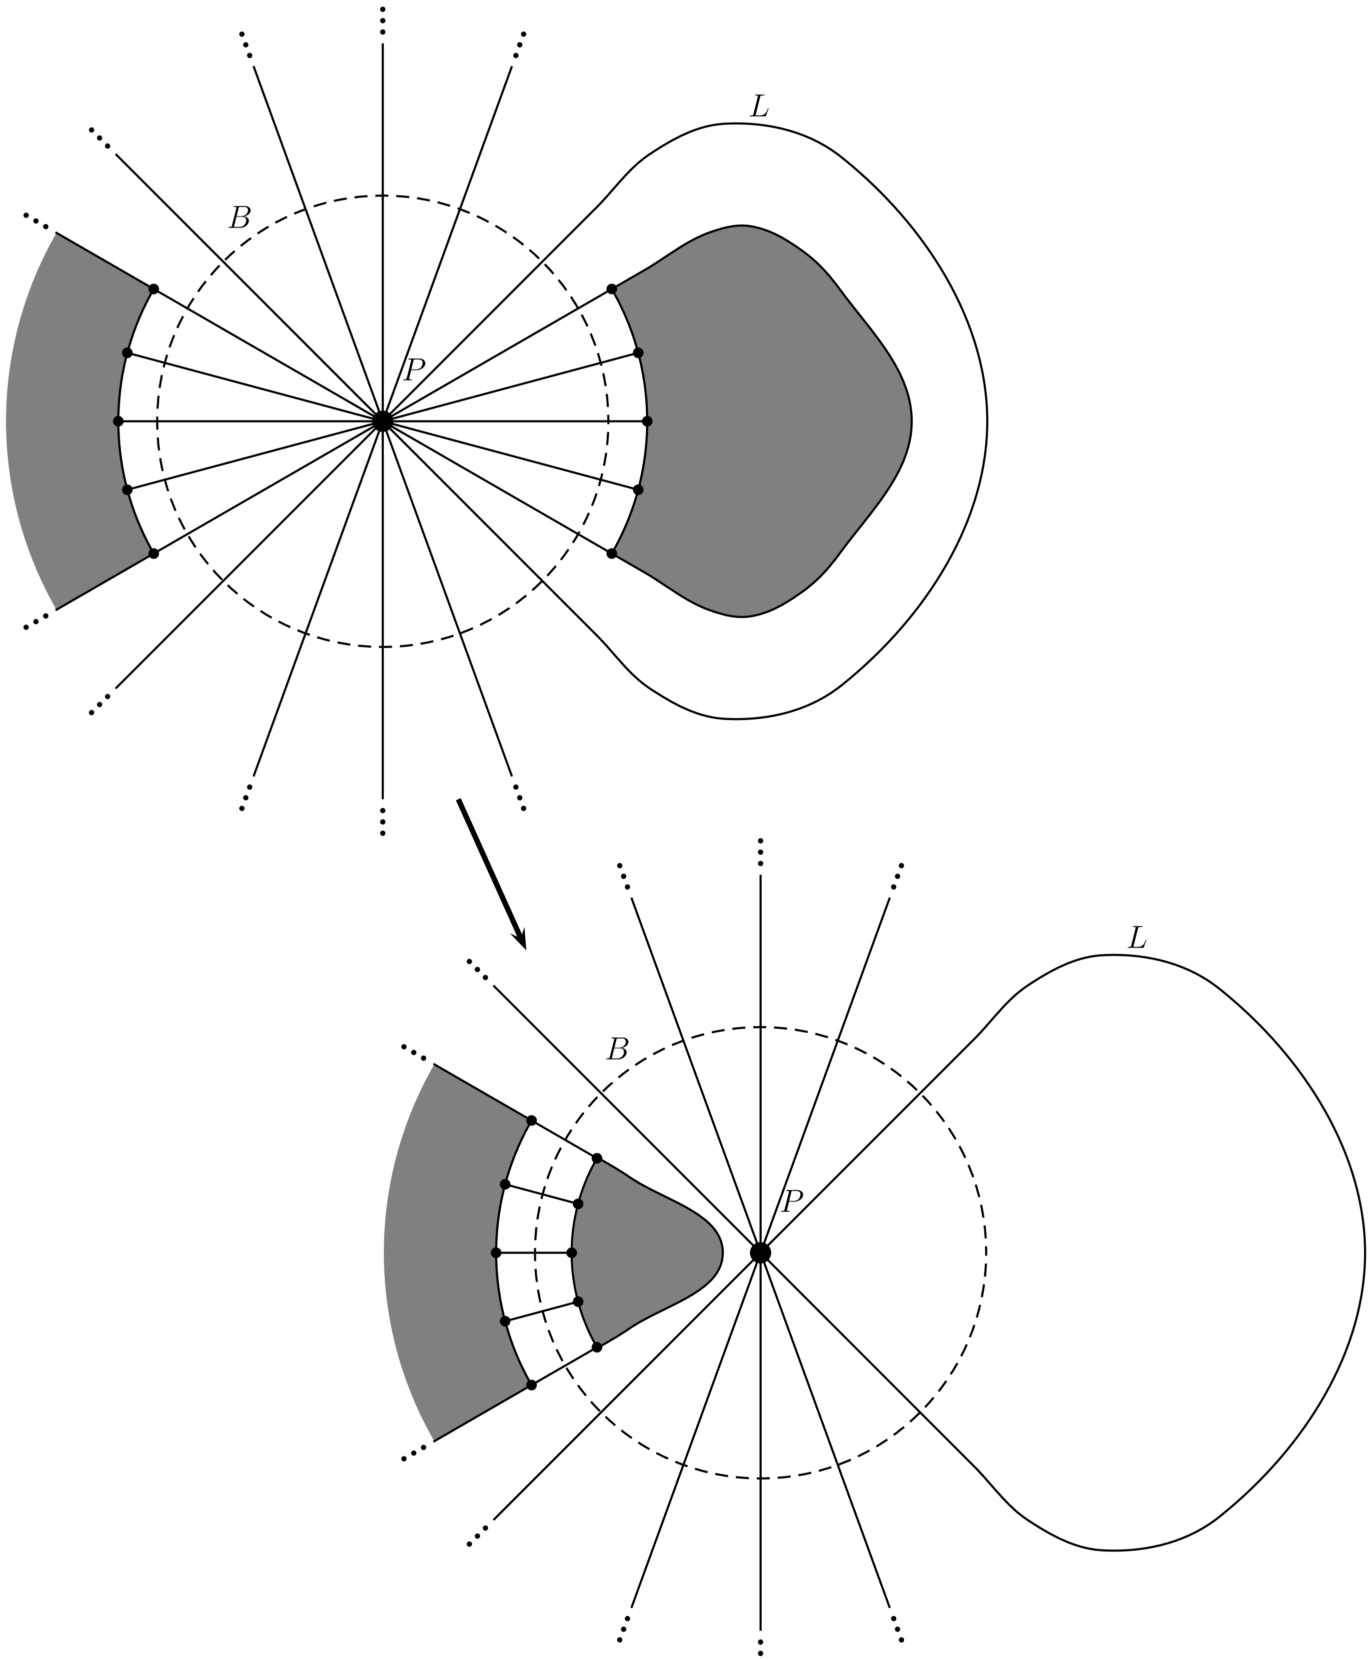

FIGURE 3.2. Untwisting the components at P.

And finally, let x be a point on ∂U which is contained in some edge of G. Let B_1 and B_2 be open discs about x and the antipode of x, respectively, whose intersections with $\Gamma(G)$ are straight line segments, s_1 and s_2, perpendicular to ∂U. Then in the projection of $\Gamma(G)$ in the plane with a crosscap, the projections of s_1 and s_2 form a single straight line segment. Hence the projection satisfies (PC4$'$), and by Lemma 3.2.7, we see that it results in an embedding of G in the plane with a crosscap.

For the right-to-left implication, suppose that G is a graph embedded in the plane with a crosscap P at the origin, labeled with polar coordinates. With the following projection, we map $G \backslash P$ to the unit disc model of the projective plane:

$$(r, \theta) \longmapsto \left(\frac{1}{1+r}, \theta \right). \tag{3.2.2}$$

The closure of the image of G then represents an embedding of G in the projective plane. □

We now prove three technical lemmas which ultimately yield a special disc in the projective plane. This disc provides the foundation and first layer of our eventual edge bipartition of projective plane graphs.

Lemma 3.2.9. *If G is a cubic, 2-connected projective plane graph which is not planar, and whose embedding in the unit disc model of the projective plane satisfies the following condition:*

(BDY1) The boundary of U contains at most one point (i.e. one pair of antipodal points) from each edge;

then there is a surface dual G^* (which may be a multigraph) of G whose embedding also satisfies condition (BDY1).

Proof. Let $G = (V, E)$ be a cubic, 2-connected projective plane graph which is not planar, and suppose that G satisfies condition (BDY1). Let $F_1, \ldots, F_k$ be the faces of G. Let $V' = v'_1, \ldots, v'_k$ be a set of points in the projective plane such that $v'_i \in F_i$, for all i. Since G is 2-connected and not planar, we know, by Lemma 3.2.2, that every edge of G is incident to two distinct faces. Let S be the set of edges e for which the following holds, where e is incident to F_i and F_j:

(1) There is a topological (v'_i, v'_j)-path α which avoids ∂U such that $|\alpha \cap G| = |\alpha \cap \mathring{e}| = 1$.

Clearly we may find a collection P_1 of topological paths which serve as the edges corresponding to S in a surface dual of G. For each edge e in $E - S$, we may also find a topological path α_e such that $\partial U \cap \alpha_e$ consists of exactly one point (i.e. one pair of antipodal points); thus, for each $e \in E - S$, let α_e be such a path. Let $P_2 = \{\alpha_e : e \in E - S\}$. Then $G^* = (V', P_1 \cup P_2)$ is our desired surface dual of G. $\square$

Lemma 3.2.10. *If G is a projective planar multigraph, then there is an embedding of G in the unit disc model of the projective plane such that the following hold:*

(1) The boundary of U contains at most one point (i.e. one pair of antipodal points) from each edge of G;

(2) $\partial U \cap V(G) = \emptyset$.

Proof. Let G be a projective planar multigraph. By Lemma 3.2.8, there is an embedding Γ of G in the plane with a crosscap P. Then the closure of the image of $\Gamma(G)$ under the projection 3.2.2 represents an embedding of G in the unit disc model of the projective plane such that the following hold:

(1) ∂U contains no vertices of G;

(2) ∂U contains at most one point (i.e. one pair of antipodal points) from each edge of G. $\square$

Lemma 3.2.11. *Let G be a 2-connected, non-planar, projective plane triangulation, such that, in the unit disc model of the projective plane, ∂U avoids $V(G)$ and contains at most one point (i.e. one pair of antipodal points) from each edge of G. Then there is a closed curve α in the projective plane which bounds a disc D such that the following hold:*

(1) $\alpha \cap V(G) = \emptyset$;

(2) $V(G) \subsetneq D$;

(3) $|\alpha \cap e| \in \{0, 2\}$ for every $e \in E(G)$;

(4) The graph induced by the edges which avoid α is connected.

Proof. Let G be a 2-connected, non-planar, projective plane triangulation, such that, in the unit disc model of the projective plane, ∂U avoids $V(G)$ and contains at most one point (i.e. one pair of antipodal points) from each edge of G. Let $\mathfrak{C}$ be a set of open discs B in U such that

(a) $\bigcup_{D\in\mathfrak{C}} D \supsetneq \partial U$;

(b) $\overline{B} \cap G$ avoids $V(G)$;

(c) If $\overline{B} \cap G$ is nonempty, then it is homeomorphic to the half-open unit interval and contains precisely one point of $G \cap \partial U$.

Since ∂U is compact, we may suppose, without loss of generality, that $\mathfrak{C}$ is finite. Let $D_1 = U - (\bigcup_{B\in\mathfrak{C}} B)$. Then D_1 is a closed disc and is bounded by a closed curve α_1. Notice that α_1 and D_1 satisfy conditions (1), (2), and (3) in the statement of the lemma. Let H_1 be the graph induced by the edges which are contained in D_1. If H_1 is connected, then condition (4) is satisfied, and the conclusion follows. See Figure 3.3.

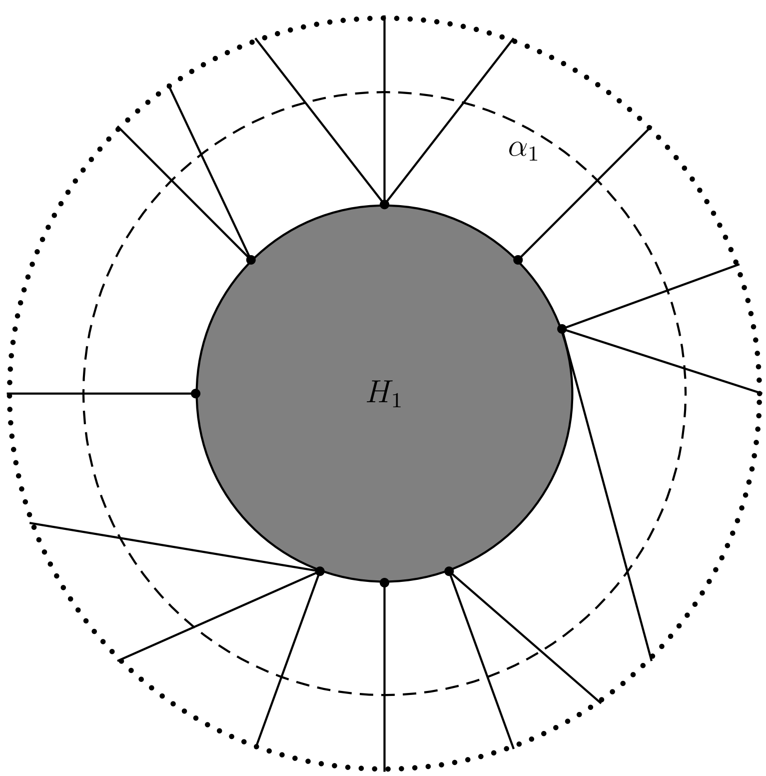

FIGURE 3.3. When H_1 is connected, the conclusion follows.

Otherwise, we proceed by induction on the number of components of H_i. Suppose, then, that α_i is a closed curve in the projective plane which bounds a disc D_i; suppose also that α_i and D_i satisfy conditions (1), (2), and (3) in the statement of the lemma. Let H_i be the graph induced by the edges contained in D_i, and suppose that H_i has k components, with $k \geq 2$.

Let $b_1, \ldots, b_p$ be the points in $\alpha_i \cap G$. For each b_j, with $j \in \{1, \ldots, p\}$, there is an edge-segment, lying in D_i, with endpoints b_j and v_j, for some vertex $v_j \in V(G)$. And we know that each vertex of G lies in some component of H_i; we say that b_j is *connected to* the said component of H_i. Notice that $b_1, \ldots, b_p$ divides α_i up into p closed segments, which we call $S_1, \ldots, S_p$. For each S_m whose endpoints b_j and b_l are connected to distinct components of H_i, pick a point in S_m which avoids $\{b_1, \ldots, b_p\}$; let the resulting points be $z_1, \ldots, z_q$. Notice that $z_1, \ldots, z_q$ divide α_i up into q closed segments, which we call $T_1, \ldots, T_q$. Then the points in $\{b_1, \ldots, b_p\} \cap T_j$, for each $j \in \{1, \ldots, q\}$, are all connected to the same component of H_i. We say that a component of H_i is *represented* by T_j if the points of $\{b_1, \ldots, b_p\} \cap T_j$ are connected to that component. Since G is connected, we know that every component of H_i is represented by at least one of $T_1, \ldots, T_q$. See Figure 3.4.

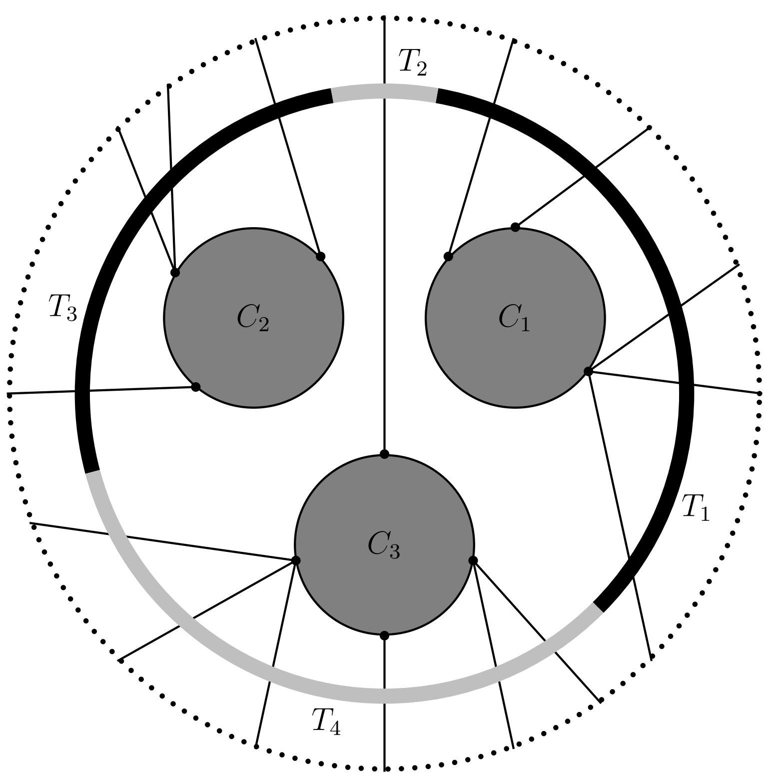

FIGURE 3.4. The components C_1, C_2, C_3 of H_i are represented by the segments T_1, T_2, T_3, T_4.

Part 1. We prove now that some component of H_i is represented by exactly one of $T_1, \ldots, T_q$. Note that, by induction, it suffices to prove the following: if C is a component of H_i which is represented by more than one of $T_1, \ldots, T_q$, then there is a component C' of H_i which is represented by fewer of $T_1, \ldots, T_q$ than is C.

Let C be some component of H_i. If C is represented by exactly one of $T_1, \ldots, T_q$, then the conclusion follows from Part 2 below. Otherwise, pick T_i, T_j, and a sequence of segments $T'_1, \ldots, T'_r$ such that the following hold:

(1) T_i and T_j are distinct representatives of C;

(2) None of $T'_2, \ldots, T'_{r-1}$ are representatives of C;

(3) $T'_1 = T_i$;

(4) $T'_r = T_j$;

(5) $|T'_l \cap T'_{l+1}| = 1$, for all $l \in \{1, \ldots, r-1\}$.

(In Figure 3.4, if we take C to be C_3, then a suitable sequence would be T_2, T_3, T_4.) Let C' be a component of H_i represented by at least one segment in $T'_2, \ldots, T'_{r-1}$. By the Jordan curve theorem, we know that all the segments which represent C' lie in $T'_2, \ldots, T'_{r-1}$. Then C' is represented by fewer of $T_1, \ldots, T_q$ than is C.

Thus there is a component of H_i which is represented by exactly one of $T_1, \ldots, T_q$.

Part 2. We will now produce a disc D_{i+1} and a closed curve α_{i+1} bounding D_{i+1} which satisfy conditions (1), (2), and (3) of the statement of the lemma, such that the graph induced by the edges contained in D_{i+1} has $k-1$ components. Let C be a component of H_i which is represented by exactly one of $T_1, \ldots, T_q$, say T_1.

We say that two discs D, D', in the unit disc model of the projective plane, are *antipodal* if $D \cap \partial U$ and $D' \cap \partial U$ are disjoint in U, but equal under the identification of U. Let $S_1, \ldots, S_s$ be the components of $G \backslash D_i$ whose closures meet T_1. Let X be the component of $G \backslash \alpha_i$ which contains C.

Let $\mathfrak{C}_1$ be a collection of open discs in the projective plane such that the following hold:

(1) $\bigcup_{B \in \mathfrak{C}_1} B$ is a disc in the projective plane;

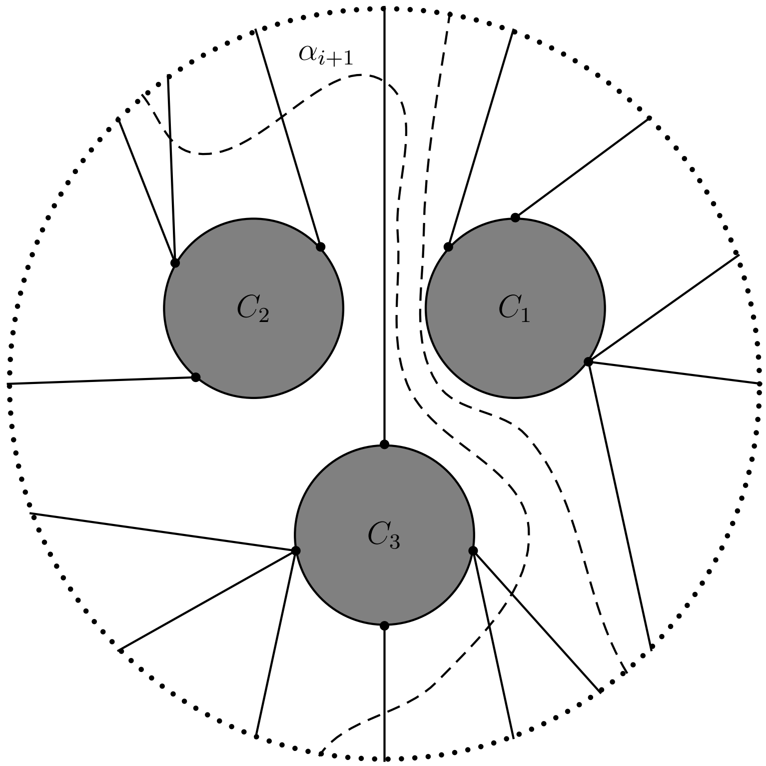

FIGURE 3.5. We aim to re-route α_i from Figure 3.4 to obtain the curve α_{i+1} shown here.

(2) $\bigcup_{B\in\mathfrak{C}_1} B \supsetneq \overline{S_1 \cup \cdots \cup S_s}$;

(3) $\overline{\bigcup_{B\in\mathfrak{C}_1} B} \cap G$ has precisely s components and avoids $V(G)$.

Since $\overline{S_1 \cup \cdots \cup S_s}$ is compact, we may assume, without loss of generality, that $\mathfrak{C}_1$ is finite. Let $\mathfrak{C}_2$ be a collection of open discs in the projective plane such that the following hold:

(1) $\bigcup_{B\in\mathfrak{C}_2} B$ is a disc in the projective plane;

(2) $\overline{X} \subsetneq \left(\bigcup_{B\in\mathfrak{C}_2} B\right) \bigcap G \subsetneq \overline{X} \cup S_1 \cup \cdots \cup S_s$.

Since the closure of the component of $G\backslash\alpha_i$ which contains C is compact, we may suppose, without loss of generality, that $\mathfrak{C}_2$ is finite. Let $\mathfrak{C}_3$ be a collection of open discs in the projective plane such that the following hold:

(1) $\bigcup_{B\in C_3} B$ is a disc in the projective plane;

(2) $\left(\bigcup_{B\in C_3} B\right) \bigcap G = \left(\bigcup_{B\in C_2} B\right) \bigcap G$;

(3) $\overline{\bigcup_{B\in C_2} B} \subsetneq \bigcup_{B\in C_3} B$;

(4) $\overline{\bigcup_{B\in C_3} B}$ avoids $V(G)$.

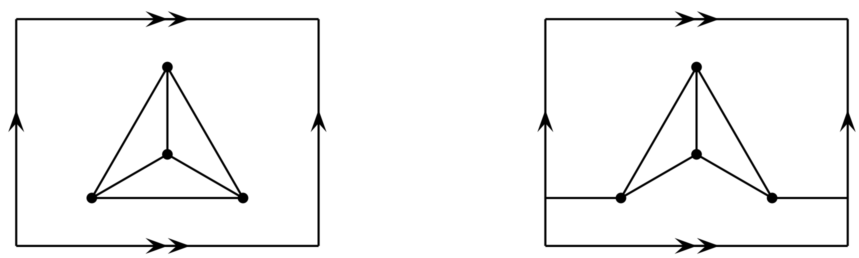

FIGURE 3.6. Toroidal embeddings of K_4, with a different number of faces.

Since $\overline{\bigcup_{B\in C_2} B}$ is compact, we may suppose, without loss of generality, that $\mathfrak{C}_3$ is finite. Let $D_{i+1} = \big(D_i \backslash \overline{\bigcup_{B\in\mathfrak{C}_2} B}\big) \bigcup \big(\bigcup_{B\in\mathfrak{C}_1\cup\mathfrak{C}_2} B\big)$. Then D_{i+1} is a disc in the projective plane. Let α_{i+1} be the closed curve which bounds D_{i+1}. Then D_{i+1} and α_{i+1} satisfy conditions (1), (2), and (3) in the statement of the lemma. And we see as well that the graph H_{i+1} induced by the edges contained in D_{i+1} has $k-1$ components. The conclusion follows by induction. $\square$

Duality in topological graph theory is a far less versatile concept than in matroid theory, since all duals are tied to a particular embedding. And any particular multigraph may admit numerous distinct embeddings, with variety in the number of faces. If there are embeddings of a multigraph G with distinct numbers of faces, then any surface duals of these embeddings will not be isomorphic, since they will have a different number of vertices. For example, we may embed K_4 in the torus in the two ways represented in Figure 3.6.

In Figure 3.6, we see that the embedding on the left has four faces, and the embedding on the right has three faces. Therefore the surface duals of these embeddings will have distinct numbers of vertices. Hence the surface duals of these embeddings are non-isomorphic, as graphs.

Given a multigraph G embedded in a surface, the following easy lemma allows us to speak of *the surface dual of* G, which is a multigraph, not embedded in a surface, isomorphic to all surface duals of G.

Lemma 3.2.12. *Let G be a multigraph embedded in a surface. Then all surface duals of G are isomorphic, as multigraphs.*

Proof. Let G be a multigraph embedded in a surface. Let G_1 and G_2 be two surface duals of G, and let f be the number of faces of G. We define $\zeta : V(G_1) \longrightarrow V(G_2)$ as such: if F is the unique face of G corresponding to a vertex v, then $\zeta(v)$ is the unique vertex of G_2 corresponding to F. Clearly ζ is a one-to-one correspondence. Let x and y be distinct vertices of G_1, and let F_x and F_y be the faces of G corresponding to x and y, respectively. We must prove the following:

(1) G_1 has k loops at x if and only if G_2 has k loops at $\zeta(x)$;

(2) x and y are adjacent via an edge of multiplicity k if and only if $\zeta(x)$ and $\zeta(y)$ are adjacent via an edge of multiplicity k.

Proof of (1). For the left-to-right implication, suppose that G_1 has k loops at x. Then the number of edges of G whose interiors lie in $\overset{\circ}{\overline{F_x}}$ is precisely k. Then since $\zeta(x)$ is the vertex of G_2 corresponding to F_x, we know that G_2 contains k loops at $\zeta(x)$. The right-to-left implication is similar. This concludes the proof of (1).

Proof of (2). For the left-to-right implication, suppose that x and y are adjacent via an edge of multiplicity k in G_1. Then precisely k edges lie in $\partial F_x \cap \partial F_y$. Then clearly $\zeta(x)$ and $\zeta(y)$ are adjacent in G_2 via an edge of multiplicity k. The right-to-left implication is similar. This concludes the proof of (2). $\square$

The next two lemmas describe simple facts about graph duality.

Lemma 3.2.13. *Let G be a projective plane multigraph, and let G^* be a surface dual of G. Let e be an edge in G whose set of endpoints is P, and let e^* be the edge of G^* corresponding, via duality, to e. If e is a loop, and F is the face of G^* in which P lies, then $(e^*)^\circ$ lies in $\overset{\circ}{\overline{F}}$. If e is a non-loop edge in G, and F and F' are the two faces of G^* in which the vertices of P lie, then e^* lies in $\partial F \cap \partial F'$.*

Proof. Let G be a projective plane multigraph, and let G^* be a surface dual of G. Let e be an edge in G whose set of endpoints is P, and let e^* be the edge of G^* corresponding, via duality, to e.

For the first part, suppose that e is a loop at v, and F is the face of G^* in which v lies. Notice that e is a closed curve which meets G^* only at some point in $(e^*)^\circ$. Therefore e^* lies on the boundary of precisely one face of G^*, namely F. Therefore, for any point $x \in (e^*)^\circ$, we may find a disc containing x which consists entirely of limit points of F. Therefore $(e^*)^\circ \subsetneq \overset{\circ}{\overline{F}}$.

For the second part, suppose that e is a non-loop edge, and F and F' are the two faces of G^* in which the vertices of P lie. By the definition of surface dual, we know that $e^* \subseteq \partial F$ and $e^* \subseteq \partial F'$. Therefore $e^* \subseteq \partial F \cap \partial F'$. $\square$

Lemma 3.2.14. *Let G be a projective plane multigraph, and let G^* be a surface dual of G. Let B_1 and B_2 be distinct blocks of G such that $E(B_1) \cap \partial F$ and $E(B_2) \cap F$ are nonempty, for some face F of G. Let $E^*_{B_1}$ and $E^*_{B_2}$ be the edges of G^* corresponding, via duality, to $E(B_1)$ and $E(B_2)$, respectively. Then $E^*_{B_1}$ and $E^*_{B_2}$ lie in distinct blocks of G^*.*

Proof. Let G' be a projective plane multigraph, and let G^* be a surface dual of G'. We know that there is a connected, projective plane multigraph G such that the following hold:

(1) There is a one-to-one correspondence ζ between $\{B : B \text{ is a block of } G'\}$ and $\{B : B \text{ is a block of } G\}$ such that B is isomorphic to $\zeta(B)$ for every block B of G';

(2) G^* is a surface dual of G.

Thus it suffices to prove the result for G. Let B_1 and B_2 be distinct blocks of G such that $E(B_1) \cap \partial F$ and $E(B_2) \cap F$ are nonempty, for some face F of G. Let $E^*_{B_1}$ and $E^*_{B_2}$ be the edges of G^* corresponding, via duality, to $E(B_1)$ and $E(B_2)$, respectively. Let v^* be the vertex of G^* corresponding, via duality, to F. If either of B_1, B_2 consists of a single edge, say B_1, then $E^*_{B_1}$ is a single edge with v^* as an endpoint; in this case, we see that $E^*_{B_1}$ and $E^*_{B_2}$ lie in distinct blocks of G^*.

Suppose, therefore, that both B_1 and B_2 contain more than one edge. And suppose, en route to a contradiction, that there is a $(V(E^*_{B_1}), V(E^*_{B_2}))$-path P in G^* which avoids v^*. From P we obtain a sequence $(F_1, \ldots, F_k)$ of faces of G such that $F \notin \{F_1, \ldots, F_k\}$ and $\partial F_i \cap \partial F_{i+1}$ share an edge when $i \in \{1, \ldots, k-1\}$. Let $F_P = \overline{F_1 \cup \cdots \cup F_k}$. Notice that F and F_P are disjoint and nonempty. Furthermore, notice that for any point $x \in \partial F_P$, the set $F_P \backslash x$ is connected. Let $v_1 \in V(B_1) \backslash V(B_2)$ and $v_2 \in V(B_2) \backslash V(B_1)$ be distinct vertices which are incident to F. Since ∂F_P consists entirely of edges of G, we know that there are two internally disjoint (v_1, v_2)-paths in $G \cap \partial F_P$; but since B_1

and B_2 are distinct blocks, we know that there is a set $S \subseteq V(G) \backslash \{v_1, v_2\}$ of at most one vertex which separates v_1 and v_2. This is a contradiction. Hence there is no such path P, in which case we see that $E^*_{B_1}$ and $E^*_{B_2}$ lie in distinct blocks of G^*. □

The next lemma is another easy and-well known fact about graph duality. It is the central premise of the proof-techniques in the main results of this chapter.

Lemma 3.2.15. *Let G be a connected multigraph embedded on the projective plane, and let G^* be a surface dual of G. Let X be a subset of edges of G, and let X^* be the edges of G^* corresponding, by duality, to the edges X. Then G/X is isomorphic, as a multigraph, to the surface dual of $G^* \backslash X^*$.*

Proof. Let G be a multigraph embedded on the projective plane, and let G^* be a surface dual of G. Let X be a set of edges of G, and let X^* be the edges of G^* corresponding, by duality, to the edges X. If X is empty, then X^* is empty, and the conclusion follows.

We proceed by induction. Suppose that $|X| = k$, and suppose that for any set Y of at most $k-1$ edges of G, the graph G/Y is isomorphic, as a multigraph, to the surface dual of $G^* \backslash Y^*$, where Y^* is the set of edges of G^* corresponding to Y. Let $e \in X$, and let e^* be the edge of G^* corresponding, by duality, to e. Then $G/(X-e)$ and the surface dual of $G^* \backslash (X^* - e^*)$ are isomorphic, as multigraphs. Then, using Lemma 3.2.13, we know there is a one-to-one correspondence f between the vertices of $G/(X-e)$ and the faces of $G^* \backslash (X^* - e^*)$ such that the following hold:

(1) there are m loops at vertex x in $G/(X-e)$ if and only if, in $G^* \backslash (X^* - e^*)$, the interiors of m edges lie in $\overset{\circ}{\overline{f(x)}}$; and

(2) there is an edge of multiplicity m between x and y, in G/X, if and only if m edges of $G^* \backslash X^*$ lie in $\partial f(x) \cap \partial f(y)$.

To show that G/X is graphically isomorphic to $G^* \backslash X^*$, it suffices (using Lemma 3.2.13 once again) to construct a one-to-one correspondence f' between the vertices of G/X and the faces of $G^* \backslash X^*$ such that the following hold:

(1) there are m loops at vertex x in G/X if and only if, in $G^* \backslash X^*$, the interiors of m edges lie in $\overset{\circ}{\overline{f'(x)}}$; and

(2) there is an edge of multiplicity m between x and y, in G/X, if and only if m edges of $G^*\backslash X^*$ lie in $\partial f'(x) \cap \partial f'(y)$.

Case 1. Suppose that e is a loop in $G/(X-e)$ at vertex v. Let F be the unique face of $G^*\backslash(X^*-e^*)$ such that $f(v)=F$. Then by Lemma 3.2.13, the set $(e^*)^\circ$ lies in $\overset{\circ}{\overline{F}}$. In this case, we let $f'(v)$ be the face of $G^*\backslash(X^*)$ equal to $F \cup e^*$, and we let $f'|_{V(G/X)-v} = f|_{V(G/(X-e))-v}$.

Case 2. Suppose that e is not a loop and has endpoints u, v. Let $w \in V(G/X)$ be the vertex to which u and v are contracted. Let F_u and F_v be the unique faces of $G^*\backslash(X^*-e^*)$ such that $f(u)=F_u$ and $f(v)=F_v$. Then, by Lemma 3.2.13, the edge e^* lies in $(\partial F_u) \cap (\partial F_v)$. In this case, we let $f'(w)$ be the face of $G^*\backslash X^*$ equal to $F_u \cup F_v \cup e^*$, and we let $f'|_{V(G/X)-w} = f|_{V(G/(x-e))-\{u,v\}}$.
$\square$

3.3 The Case of the Projective Plane – Main Result

We are now ready to prove our main result on projective planar graphs.

Theorem 3.3.1. *For any projective planar graph G, there is a bipartition $\{X, Y\}$ of $E(G)$ such that G/X and G/Y have tree-width at most three.*

Proof. Let G be a projective planar graph. If G is planar, the conclusion follows from Corollary 1.2.6. Suppose, therefore, that G is not planar. By Theorem 3.2.1, it suffices to find a bipartition $\{X, Y\}$ of $E(G)$ such that G/X and G/Y contain no minor isomorphic to K_5, M_6, M_8, and M_{10}.

By Lemmas 3.2.6 and 3.2.4, we may suppose that G is cubic and 2-connected. By Lemma 3.2.9, we have an embedding of G and a surface dual G^* such that the following hold:

(1) ∂U avoids $V(G)$ and $V(G^*)$;

(2) ∂U contains at most one point (i.e. one pair of antipodal points) of each edge in G and G^*.

Since G is cubic and 2-connected, we know that G^* is a 2-connected triangulation. Using the closed curve α and the disc D from Lemma 3.2.11, let G^*_{-1} be the subgraph of G^* induced by the edges that meet α.

Let V_0 consist of the vertices of G^* which lie on the boundary of the face of $G^*\backslash E(G^*_{-1})$ that contains α. See Figure 3.7. For $i \in \{1, 2, 3, \ldots\}$, let V_i consist of the vertices of G^* that are a distance i from V_0. For $i \in \{0, 1, 2, \ldots\}$, let G_i^* be the graph induced by the edges e which satisfy one of the following:

(1) both endpoints of e lie in V_i; or

(2) e has endpoints in V_i and V_{i+1}.

We call the graphs G_i^* *distance layers*. Let $X^* = \bigcup_{k \geq 0} E(G_{2k-1}^*)$, and let $Y^* = \bigcup_{k \geq 0} G_{2k}^*$. Let X and Y consist of the edges of G which correspond, via duality, to the edges in X^* and Y^*, respectively. Then $\{X, Y\}$ is a bipartition of the edges of G. We now prove that G/X and G/Y have no minors isomorphic to K_5, M_6, M_8, and M_{10}. By Lemma 3.2.15, we know that G/X and G/Y are isomorphic to the surface duals of $G^* \backslash X^*$ and $G^* \backslash Y^*$, respectively.

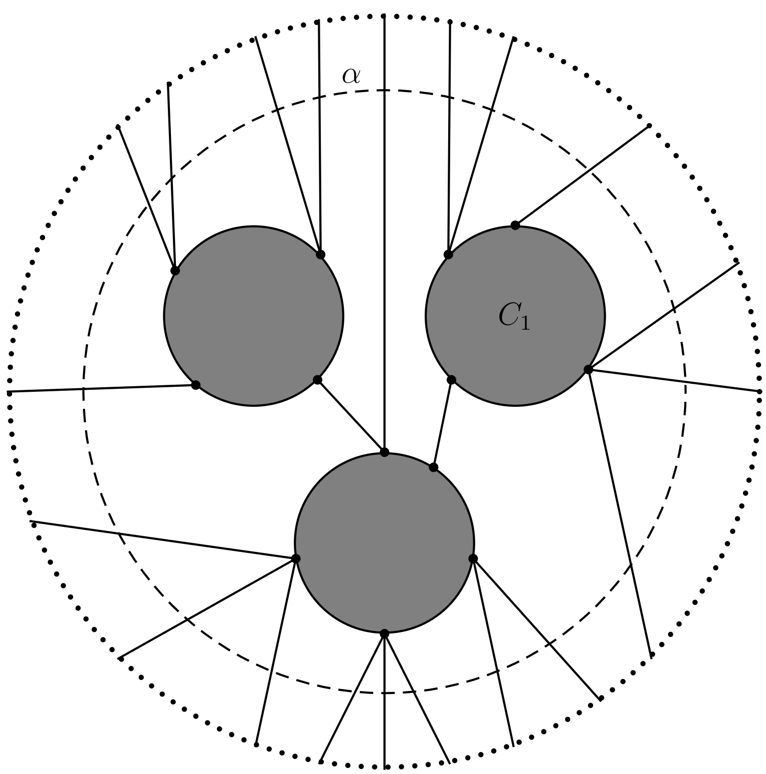

FIGURE 3.7. The boundary of the shaded regions contains V_0. Note that the graph represented here, G^*, is a triangulation; for clarity, only some representative edges are shown. See Figure 3.8 for a closer look at an example of what C_1 might be.

We now examine the structure of the surface duals of $G^* \backslash X^*$ and $G^* \backslash Y^*$. We know that $G^* \backslash X^* = G_0^* \cup G_2^* \cup G_4^* \cup \cdots$ and $G^* \backslash Y^* = G_{-1}^* \cup G_1^* \cup G_3^* \cup \cdots$. We know as well that G_i^* and G_j^* are vertex-

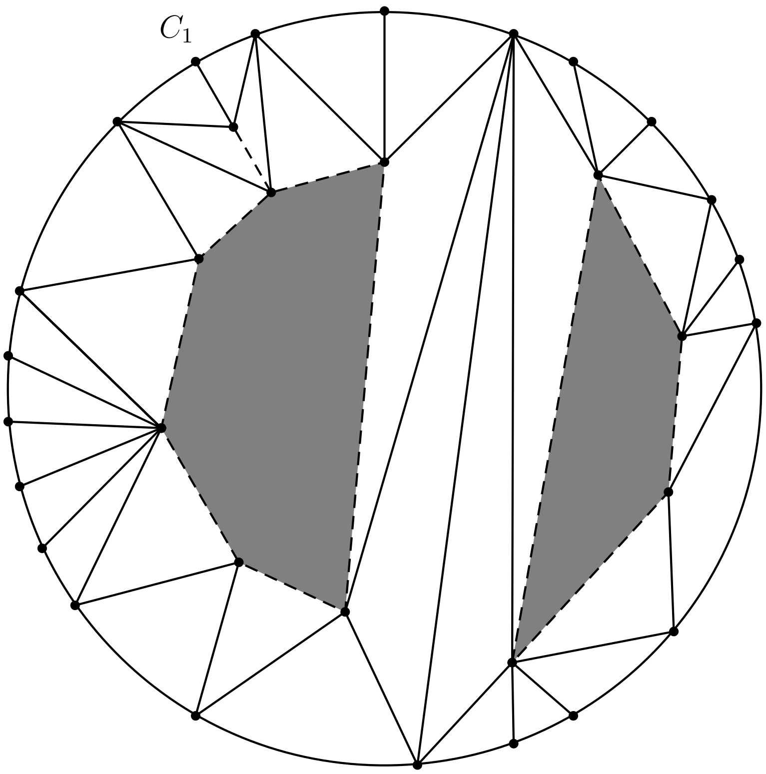

FIGURE 3.8. A closer look at C_1. The shaded regions represent components of $G_1^* \cup G_2^* \cup G_3^* \cup \cdots$ which lie in the disc bounded by C_1.

disjoint if $|i - j| > 1$. Then each connected component of $G^* \backslash X^*$ is a connected component of some G_{2i}^*, with $i \geq 0$; and each connected component of $G^* \backslash Y^*$ is a connected component of some G_{2i-1}^*, with $i \geq 0$.

By Lemma 3.2.14, we can conclude the following:

(1) No block B of the surface dual of $G^* \backslash X^*$ contains edges corresponding, via duality, to two distinct blocks of $G^* \backslash X^*$;

(2) No block B of the surface dual of $G^* \backslash Y^*$ contains edges corresponding, via duality, to two distinct blocks of $G^* \backslash Y^*$.

Hence, using Lemma 3.2.3, it suffices to show that for each block B of $G^* \backslash X^*$ and $G^* \backslash Y^*$, the surface dual of B contains no minor isomorphic to K_5, M_6, M_8, and M_{10}. We now prove the following:

3.3.2. *The surface duals of $G^* \backslash X^*$ and $G^* \backslash Y^*$ contain no K_5- and no M_8-minor.*

Since $E(G^*_{-1})$ is in X^*, we know that $G^*\backslash X^*$ is embedded in the disc D. Therefore we can find a surface dual of $G^*\backslash X^*$ that lies in D; and since D is a disc, this surface dual is planar. Therefore, since K_5 and M_8 are not planar, the surface dual of $G^*\backslash X^*$ contains no K_5- and no M_8-minor.

Let F_{-1} be the face of $G^*\backslash E(G^*_0)$ in which $\big(E(G^*_0)\big)^\circ$ lies. Then F_{-1} is a face of $G^*\backslash Y^*$. Let f_{-1} be the vertex of the surface dual of $G^*\backslash Y^*$ corresponding to F_{-1}, and suppose, without loss of generality, that $f_{-1} \in D$.

We first prove that G^*_{-1} has only one face. Suppose that two edges uv and uw in G^*_{-1} are adjacent in the ordering induced by α. Since G^* is a triangulation, we know that vw is an edge of G^*_0. Let α' be a component of $\alpha\backslash G^*$ with endpoints $\alpha \cap uv$ and $\alpha \cap uw$. Then α' lies in F_{-1}. Since α' was chosen arbitrarily, we know that all components of $\alpha\backslash G^*$ lie in F_{-1}. Therefore G^*_{-1} lies in $\overline{\mathring{F}_{-1}}$. Hence G^*_{-1} has only one face. See Figure 3.9. Hence the surface dual of G^*_{-1} has only one vertex, namely f_{-1}, and $|E(G^*_{-1})|$ loops.

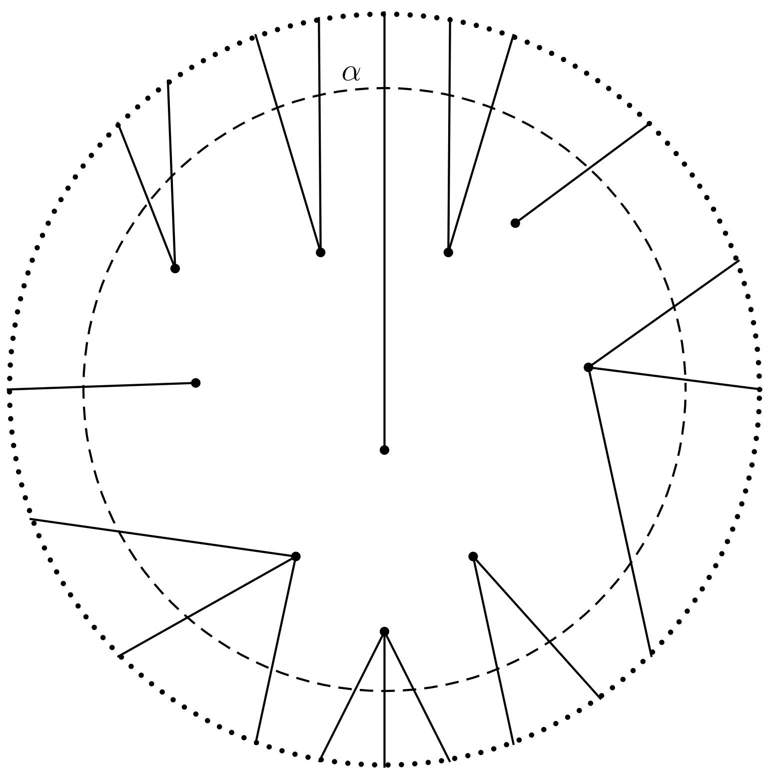

FIGURE 3.9. G^*_{-1} has only one face.

Notice that $G^*\backslash(Y^* \cup G^*_{-1})$ lies in D. Therefore we can find a surface dual of $G^*\backslash(Y^* \cup G^*_{-1})$ that lies in D, ensuring that the vertex corresponding to the face containing F_{-1} is f_{-1}; and since D is a disc, we know that this surface dual is planar. Therefore the surface dual of $G^*\backslash Y^*$ is a planar graph containing a vertex f_{-1}, to which we add $|E(G^*_{-1}|$ loops at f_{-1}. Therefore the surface dual of $G^*\backslash Y^*$ contains no K_5- and no M_8-minor. This concludes the proof of statement 3.3.2.

The following definition will be useful: for each $i \in \{0, 1, 2, \ldots\}$, an *internal face* of G^*_i is a face of G^*_i which avoids α. We now prove the following:

3.3.3. *For each $i \in \{0, 1, 2, \ldots\}$, every (V_i, V_i)-edge e that is incident to an internal face of G^*_i lies in a 3-cycle of G^*_i.*

Let $i \in \{0, 1, 2, \ldots\}$ and let uv be a (V_i, V_i)-edge that is incident to an internal face F of G^*_i. Then there is a vertex w and an edge uw, such that $w \neq v$ and uw is incident to F. Since G^* is a triangulation, we know that there is a triangular face F' of G^* such that the following hold:

(1) uv is incident to F';

(2) $F' \subseteq F$.

Let w' be the unique vertex incident to F' and distinct from u and v. Since w' is adjacent to u, we know that $w' \in V_i \cup V_{i+1}$. Therefore uw' and vw' are edges of G^*_i. Then uvw' is a 3-cycle of G^*_i containing uv. This concludes the proof of statement 3.3.3.

Using the definition of G^*_i, with $i \in \{0, 1, 2, \ldots\}$, and statement 3.3.3, we see that each connected component of G^*_i, with $i \in \{0, 1, 2, \ldots\}$, is of the form described by the following construction:

(1) Let T be a tree;

(2) For each $v \in V(T)$, let P_v be a graph consisting of either an edge or a cycle with at least three edges, such that P_u and P_v are disjoint when $u \neq v$;

(3) For each edge e in T, with endpoints u and v, let $x_{e,u}$ and $x_{e,v}$ be arbitrary vertices of P_u and P_v, respectively;

(4) Let L' be the graph formed from the graph $\bigcup_{v \in V(T)} P_v$ by identifying, for every $uv \in E(T)$, the vertices x_u and x_v;

(5) Embed L' in D;

(6) Let L be the graph (embedded in D) formed from L' by doing the following, for each cycle C in L_0:

  (a) Let D_C be the closed disc bounded by C;

  (b) Let Z_C be a finite collection of points in $\mathring{D}_C$;

  (c) Embed any number of $\big(V(C), V(C)\big)$-edges in $\mathring{D}_C \backslash Z_C$;

  (d) Embed $\big(V(C), Z_C\big)$-edges in such a way that each $\big(V(C), V(C)\big)$-edge lies in a 3-cycle.

See Figures 3.10 and 3.11 for illustrations of the construction.

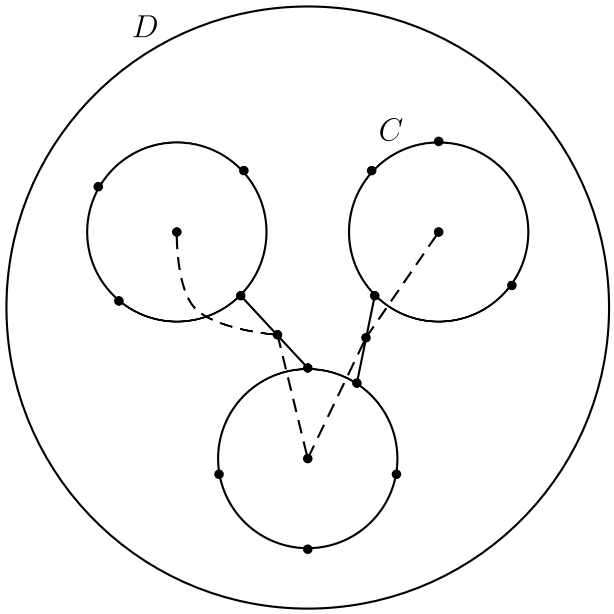

FIGURE 3.10. We see here the first five steps in the construction. The tree T (here, a path) is represented with dashed edges.

We see that each block of G_i^*, with $i \in \{1, 2, \ldots\}$, that contains more than one edge is contained in some component of $G_0^* \backslash \big\{E(P_v) : v \in V(T)$, and P_v is an edge$\big\}$. Therefore, it suffices for us to prove the following:

3.3.4. *The surface dual of an arbitrary component of* $G_0^* \backslash \big\{E(P_v) : v \in V(T)$, *and* P_v *is an edge*$\big\}$ *contains no* M_6- *and no* M_{10}-*minor.*

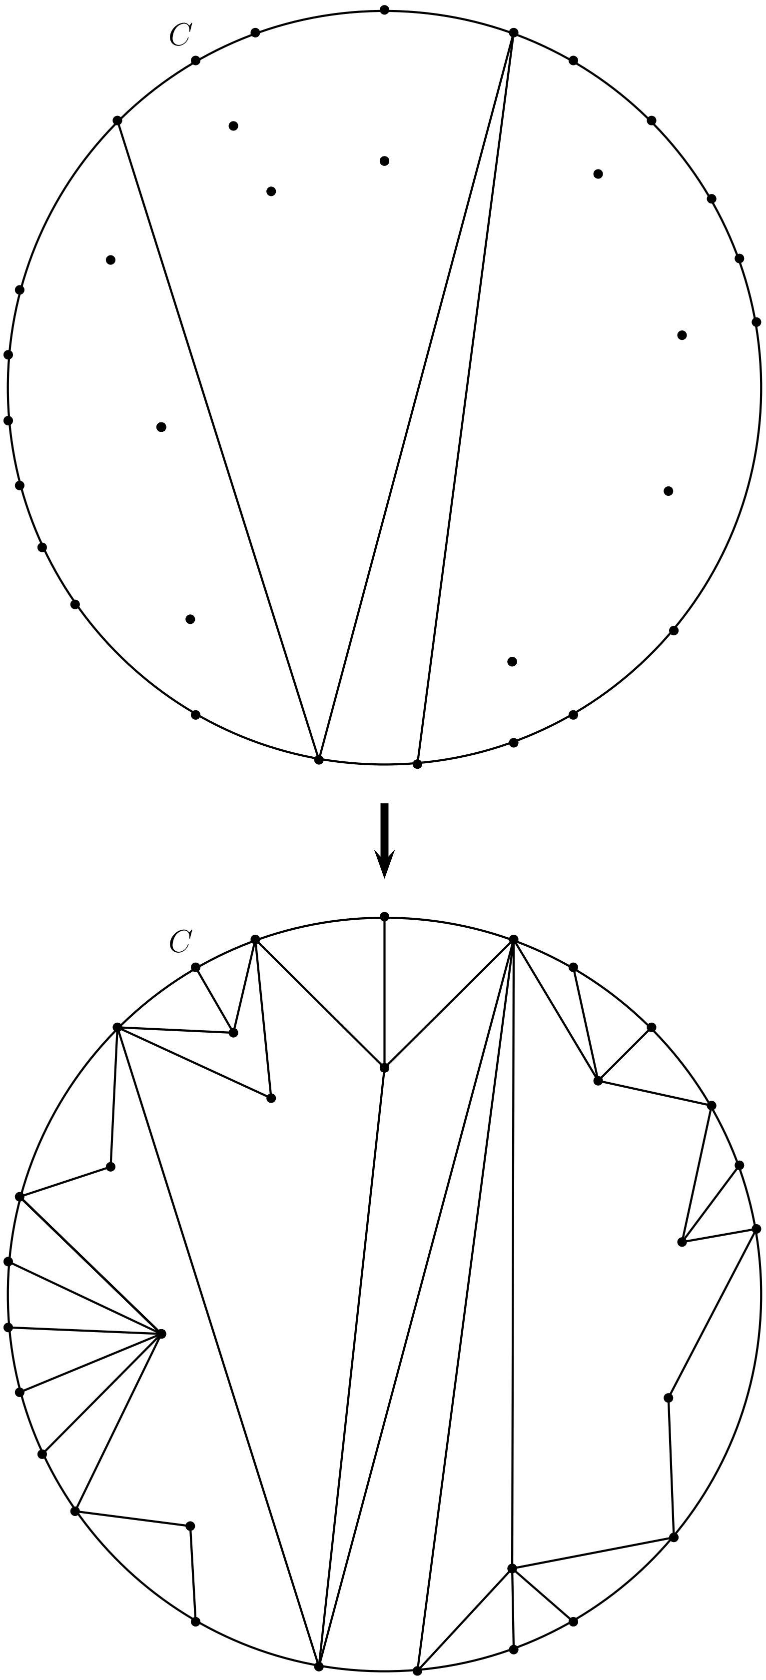

FIGURE 3.11. We see here step six of the construction. The figure on top shows steps (6a), (6b), and (6c). The figure on bottom shows step (6d).

Let F_0 be the face of G_0^* which contains α. We know that the boundary of F_0 may contain cycles of G_0^*; and we know that the blocks of G_0^* (other than the blocks which consist of single edges) are those subgraphs of G_0^* which are contained in the closed discs bounded by those cycles. Let C be a cycle contained in the boundary of F_0. (Then $C = P_v$ for some $v \in V(T)$.) Let D_1 be the closed disc bounded by C. We consider now the edges of G_0^* embedded in D_1. Let J consist of all chords of C. (Note that all such chords are embedded in D_1.) Let $\{v_1, v_2, \ldots, v_k\} = V_1 \cap D_1$; and for each $i \in \{1, \ldots, k\}$, let S_i consist of all (V_0, v_i)-edges. Let $K = \{S_1, \ldots, S_k\}$. Then K consists of all (V_0, V_1)-edges which lie in D_1. We know that every edge of G_0^* embedded in D_1 is in $C \cup J \cup K$. Then $G^*[C \cup J \cup K]$ is an arbitrary component of $G_0^* \backslash \{E(P_v) : v \in V(T), \text{ and } P_v \text{ is an edge}\}$. We must therefore show that the surface dual of $G^*[C \cup J \cup K]$ contains no M_6- and no M_{10}-minor.

We now construct a surface dual of $G^*[C \cup J \cup K]$. Let f_0 be a point in $D \cap D_1$. Notice that the surface dual of $G^*[C \cup J]$ is a tree, with all its leaves identified (notice that these leaves all lie outside of D_1); let R_0 be such a multigraph, embedded in D_1. Without loss of generality, we may assume the following:

(1) f_0 is the vertex of R_0 to which the leaves were identified;

(2) R_0 avoids $\{v_1, \ldots, v_k\}$.

Given R_0, we obtain a sequence $R_1, \ldots, R_k$ of multigraphs by performing the following inductive process on R_{i-1}:

(1) Let F_1^i be the face of $G^*[C \cup J \cup \{S_1 \cup \cdots \cup S_{i-1}\}]$ in which v_i lies;

(2) Let f_i be the vertex of R_{i-1} corresponding to F_1^i;

(3) Let $(n_1, n_2, \ldots, n_l)$ be the rotation scheme at f_i of R_{i-1};

(4) Delete f_i (and all the edges incident to it) from R_{i-1} and embed a cycle C_i of length $|S_i|$ (possibly a loop or doubled edge) in the resulting face such that C_i bounds a closed disc which avoids $\{v_{i+1}, \ldots, v_k\} \cup (R_{i-1} \backslash f_i)$;

(5) For each $j \in \{1, \ldots, l\}$, embed the appropriate $\big(V(C_i), n_j\big)$-edge.

Thusly we obtain R_k, which is a surface dual of $G^*[C \cup J \cup K]$. If M_6 or M_{10} is to be a minor of R_k, then R_k must contain two disjoint, nested cycles Z_1 and Z_2, and three pairwise disjoint (Z_1, Z_2)-paths. See Figure 3.12 .

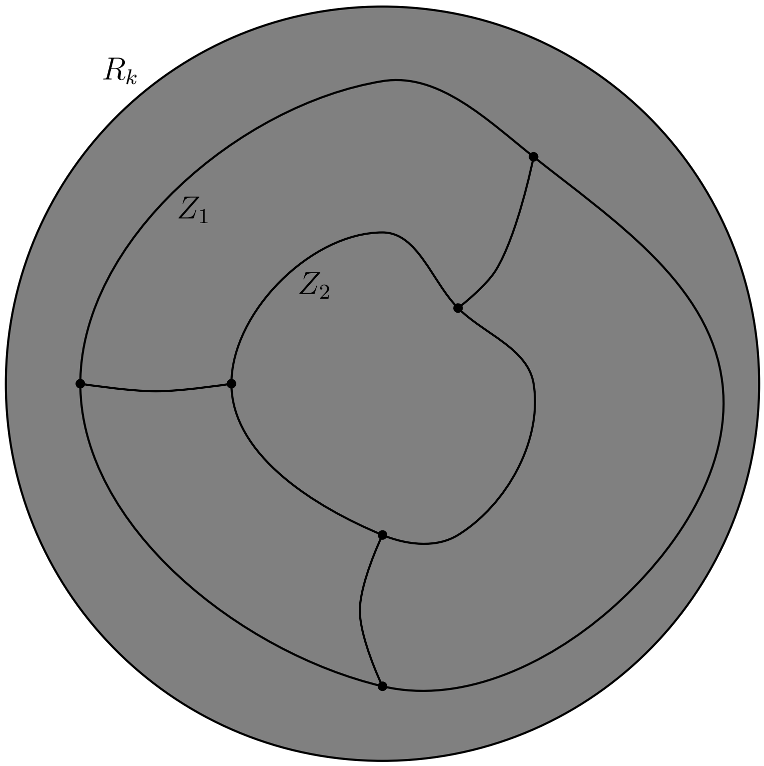

FIGURE 3.12. When M_6 or M_{10} is a minor of R_k, then R_k contains this substructure.

We prove now that R_k contains no such cycles. Let Z_1 and Z_2 be disjoint nested cycles in R_k, and suppose that Z_2 is contained in the disc bounded by Z_1. Notice that cycles in R_k are of two types: those which contain f_0 and those which do not contain f_0. In step (4) of the construction of R_i from R_{i-1}, we specified that the closed disc bounded by C_i avoid $\{v_{i+1}, \ldots, v_k\}$. Therefore no pair of $C_1, \ldots, C_k$ are nested. Since $\{v_1, \ldots, v_k\}$ are pairwise non-adjacent in G_0^*, we know that no two distinct C_i and C_j (with $i \neq j$) share an edge. Therefore, every cycle in R_k which avoids f_0 is in $\{C_1, \ldots, C_k\}$. Therefore $Z_2 = C_i$ for some $i \in \{1, \ldots, k\}$, and Z_1 contains f_0. Then $Z_1 \backslash f_0$ is a path in the tree-like structure $R_k \backslash f_0$, and we know that there are not two disjoint $(Z_2, Z_1 \backslash f_0)$-paths in $R_k \backslash f_0$. Therefore three pairwise disjoint (Z_1, Z_2)-paths do not exist in R_k. Hence R_k contains no M_6- and no M_{10}-minor. This concludes the proof of statement 3.3.4. Thus concludes the proof of Theorem 3.3.1. $\square$

3.4 The Case of the Torus – Introductory Results

In this section, we need a few specialized definitions. Suppose that G is a plane multigraph with a 3-cycle xyz. We define the *interior of* xyz, notated $int(xyz)$, as the subgraph of G induced by the edges whose interiors lie in the interior of the disc bounded by xyz. Note that no edge of xyz lies in $int(xyz)$. We define the *exterior of* xyz, notated $ext(xyz)$, as the subgraph induced by the edges whose interiors lie outside of the disc bounded by xyz.

The next lemma is an easy and well-known fact about representativity. The curve it provides will allow us to find a suitable set of edges, whose deletion produces a suitable planar graph.

Lemma 3.4.1. *If G is a graph embedded in a surface, then there is a closed, noncontractible curve α such that $|\alpha \cap H| = rep(G)$ and $\alpha \cap G \subseteq V(G)$.*

Proof. Let G be a graph embedded in a surface, and let α' be a closed, contractible curve in the surface such that $|\alpha \cap G| = rep(G)$, and let $\{x_1, \ldots, x_k\} = \alpha \cap G$ Let $F_1 \ldots, F_k, F_{k+1} = F_1$ be the ordered list of faces of G which meet α, such that $x_i \in \partial F_i \cap \partial F_{i+1}$. Then $\partial F_i \cap \partial F_{i+1} \cap V(G)$ is nonempty, for every $i \in \{1, \ldots, k\}$.

For every $x_i \in \{x_1, \ldots, x_k\} \backslash V(G)$, let B_i be a closed disc containing x_i such that $B_i \cap G$ avoids $V(G)$ and is homeomorphic to the closed unit interval; let $\mathfrak{B}$ be the set of these B_i's. For every $B_i \in \mathfrak{B}$, let α_i be a curve such that the following hold:

(1) The endpoints of α_i are the same as the endpoints of the curve $\alpha' \cap B_i$;

(2) $\alpha_i \cap G$ consists precisely of a single vertex of G;

(3) $\alpha_i \cup (B_i \cap \alpha')$ is a contractible curve.

We construct our desired curve α by replacing, in α', for every $B_i \in \mathfrak{B}$, the subsegment $B_i \cap \alpha'$ with the curve α_i. $\square$

We will need the following well-known theorem of Tutte [22], which, notably, implies that 4-connected planar graphs are Hamiltonian.

Theorem 3.4.2. *Let G be a plane graph, and let e and f be distinct edges of G such that e is not a cut-edge and such that e and f both lie in some cycle which is contained in the boundary of some face. Then G contains a cycle C which satisfies the following:*

(1) e and f are edges of C;

(2) Every C-bridge has at most three vertices of attachment;

(3) If a C-bridge B shares a vertex with a cycle C' such that C' contains e and is contained in the boundary of some face, then B has exactly two vertices of attachment.

We will need the following lemma, which comes from [20].

Lemma 3.4.3. *Let G be a 4-connected plane graph. If C is a cycle of G of length at least four, and $int(C)$ is 3-connected, then $int(C)$ has a Hamilton cycle containing any three edges of C.*

We will also use the following lemma, which is a slight modification of Lemma 2.3 in [7].

Lemma 3.4.4. *Let G be a 4-connected plane triangulation with the cycle xyz as the boundary of the infinite face, and suppose that z has degree greater than three. Then G has an edge partition $\{A, B\}$ such that the following hold:*

(1) Each of $G[A], G[B]$ contains xyz and is outerplanar;

(2) Every path in A from x or y to z uses xz or yz;

(3) B has no path between any of x, y, z except those contained in xyz;

(4) A contains every edge of G which has an endpoint in $\{x, y\}$.

Proof. Let G be a 4-connected plane triangulation with the cycle xyz as the boundary of the infinite face, and suppose that z has degree greater than three. Let C be the cycle of G which bounds the infinite face of $G - z$. Let v_x be the unique neighbor of x in C which is not y, and let v_y be the unique neighbor of y in C which is not x. Let $D = \{v_x x, v_y y, xy\}$. Since G is 4-connected, we know that $v_x \neq v_y$. And since z has more than three neighbors in G, we know that either $|V(G-z)| = 3$ or $|E(C) \geq 4$. Therefore, whether trivially or by Lemma 3.4.3, we know that $G - z$ has a Hamiltonian cycle H containing D.

Let X consist of the edges of H together with $E(xyz)$ and the edges of G lying in the disc bounded by H. Let Y consist of $E(xyz)$ together with the edges of G whose interiors lie outside of the disc bounded by H.

Notice that every vertex of $G[A]$ is incident to the face of $G[A]$ which lies outside of the disc bounded by H and inside of the disc bounded by xyz; therefore $G[A]$ is outerplanar. Notice that every vertex of $G[B]$ is contained in the face of $G[B]$ which contains the disc bounded by H; therefore $G[B]$ is outerplanar. Hence condition (1) of the lemma holds. Since A contains no edges incident to z except for xz and yz, we know that every path in A from x or y to z uses xz or yz. Hence condition (2) of the lemma holds. Since the only edges in B which have an endpoint in $\{x, y\}$ are $\{xy, xz, yz\}$, we know that B has no path between any of x, y, z except those contained in xyz. Hence condition (3) of the lemma holds. And finally, since every edge with an endpoint in $\{x, y\}$ is either in xyz, in H, or contained in the disc bounded by H, we know that A contains all such edges. Hence condition (4) of the lemma holds. □

The following theorem is a strengthening of Theorem 2.2 from [7].

Theorem 3.4.5. *If G is a plane graph and v is a vertex of G, then the edges of G can be bipartitioned into $\{S, T\}$ such that $G[S]$ and $G[T]$ are series-parallel and all the edges incident to v lie in S.*

Proof. Let $\mathcal{S}$ consist of all triples (G, e, f) satisfying the following three conditions:

(1) G is a plane graph;

(2) The edges e and f of G are distinct, incident, and co-facial;

(3) G does not admit the desired edge partition with respect to v, where v is taken to be the vertex shared by e and f.

Let $(G, e, f) \in \mathcal{S}$ be a triple such that G has the fewest vertices. Without loss of generality, we may suppose that G is a triangulation, and via a stereographic projection, we may suppose that e and f lie on the infinite face. Clearly $|V(G)| > 4$.

Case 1. Suppose that G is 4-connected. Then Theorem 3.4.2 yields a Hamiltonian cycle H which contains e and f. Let v be the vertex shared by e and f. Since e, f, and v lie on the boundary of the infinite face, we see that all edges incident to v lie either on H or in the disc bounded by H. Let S be the graph induced by H and the edges which lie inside of the disc bounded by H. Let T

be the graph induced by the edges which lie outside of the disc bounded by H. Then S and T are outerplanar and therefore series-parallel, and all the edges incident to v lie in S.

Case 2. Suppose that G is not 4-connected. Then G has a separating triangle. Let xyz be a separating triangle such that $int(xyz)$ is minimal with respect to number of vertices. If e and f are contained in xyz, then xyz is the boundary of the infinite face, and $int(xyz)$ is not connected; this contradicts the fact that the interior of a separating triangle in a (simple) plane triangulation is connected. Therefore one of e, f is not contained in xyz. Furthermore, notice that xyz contains at most one edge on the boundary of the infinite face. Therefore, without loss of generality, we may assume that if xyz contains an edge of the infinite face, then $e = xy$. Let $I = xyz \cup int(xyz)$, and let $E = xyz \cup ext(xyz)$. Notice that e and f lie in E. By minimality, we know that E has an edge partition $\{S', T'\}$ into two series-parallel graphs such that the edges incident to v are contained in S'.

If there is only one vertex u in the interior of xyz, then we obtain the desired partition of the edges of G by letting $S = S' \cup \{ux, uy\}$ and $T = T' \cup uz$. Since the property of being series-parallel is retained under doubling edges and subdividing edges, this partition is as desired.

Suppose, then, that there is more than one vertex in $int(xyz)$. Since the $int(xyz)$ is minimal with respect to vertices, we know that I has no separating triangles. Hence I is 4-connected, and the degrees of x, y, and z in I are each greater than three. Thus I has an edge partition $\{A, B\}$ as specified in Lemma 3.4.4. Let $S = S' \cup (A \backslash \{xy, xz, yz\})$, and let $T = T' \cup (B \backslash \{xy, xz, yz\})$. Since A contains all edges of I with an endpoint (with respect to $G[I]$) in $\{x, y\}$, we know that any edges of $G[I]$ which are incident to v lie in S. Recall that A contains no paths between x or y and z except xz and yz. Therefore we may obtain S from S' by doubling and subdividing edges and by adding leaves to x and y. Therefore S is series-parallel. Recall also that B contains no paths between any of x, y, z except those contained in xyz. Therefore we may obtain T from T' by adding leaves to z. Therefore T is series-parallel. $\square$

3.5 The Case of the Torus – Main Result

We now prove our main result on toroidal graphs.

Theorem 3.5.1. *If G is a toroidal graph, then there is a bipartition $\{E_1, E_2\}$ of $E(G)$ such that $tw(G/E_1) \leq 3$ and $tw(G/E_2) \leq 4$.*

Proof. Let G be a toroidal graph. By Lemmas 3.2.6 and 3.2.4, we may suppose that G is cubic and 2-connected. Our first goal is to find a suitable set Z of edges which satisfies a few properties, most notably that $G \backslash Z$ is planar. If there is a closed, non-contractible curve on the torus which meets G at precisely one interior point of one edge e, then let $Z = \{e\}$. Otherwise, we must look more closely to find Z. Suppose that G admits no such curve on the torus. Then the closure of each face of G is a disc.

By Lemma 3.4.1, there exists a closed, homotopically nontrivial curve Γ on the torus such that $\Gamma \cap G \subset V(G)$ and $|\Gamma \cap G| = rep(G)$. We now prove the following:

3.5.2. *No face of G contains more than one connected component of $\Gamma \backslash G$.*

Let F be a face of G, and assume that F contains two distinct connected components A and A' of $\Gamma \backslash G$. Note that for every point x in $\Gamma \cap G$, every open disc about x intersects at least two distinct faces; for otherwise we could very easily find a new curve homotopically equivalent to Γ which intersects G one fewer times than does Γ. Therefore, since F is a disc, we know that $\Gamma \backslash (A \cup A')$ consists of two connected components B and B' and at least one of $B \cup F$ and $B' \cup F$ is homotopically nontrivial. Without loss of generality, assume that $B \cup F$ is homotopically nontrivial. Let Γ' be the curve consisting of B and a segment whose endpoints are the two vertices in $\overline{B} \cap F$ and whose interior lies in the interior of F. Then Γ' is a homotopically nontrivial curve which hits G only at vertices and which hits G fewer times than does Γ. This contradicts the minimality of Γ. Therefore the statement of 3.5.2 holds.

Since every face of G is disc, and since $|\Gamma \cap G| = rep(G)$, statement 3.5.2 implies that $|\Gamma \cap e| \leq 1$, for every $e \in E(G)$.

Let $v_1, v_2, \ldots, v_k$ be the vertices contained in Γ. The rotation scheme at v_i, for $i \in \{1, \ldots, k\}$, will be as follows: $e_i, \Gamma, e_i', e_i'', \Gamma$, where e_i, e_i', and e_i'' are the three edges incident to v_i.

We will now alter Γ slightly to produce a curve Γ' which will yield a desirable set of edges; a set whose deletion produces a planar graph. For each v_i, let B_i be an open disc such that the following hold:

(1) $v_i \in B_i$;

(2) $B_i \cap G$ is an open star;

(3) $\overline{B_i \cap G}$ equals $\overline{B_i} \cap G$ and contains none of $V(G) - v_i$.

Let Γ' be the curve obtained from Γ by replacing, for every $i \in \{1, \ldots, k\}$, the path $\overline{B_i} \cap \Gamma$ with the path on the boundary of B_i which intersects e_i. Then $\Gamma' \cap G$ consists of precisely one interior point from e_i, for each $i \in \{1, \ldots, k\}$. Let $Z = \{e_1, \ldots, e_k\}$. Note that $|\Gamma' \cap G| = rep(G)$. We now prove the following:

3.5.3. *The edges of Z are pairwise non-adjacent.*

Assume, en route to a contradiction, that e_i and e_j are adjacent at vertex v, with $i \neq j$. Then by the statement of 3.5.2 and the fact that $|\Gamma' \cap e_i| = |\Gamma' \cap e_j| = 1$, it follows that $|i - j| = 1$. Let B be an open disc such that the following hold:

(1) $v \in B$;

(2) $\Gamma' \cap G$ is an open star;

(3) $\overline{B \cap G}$ equals $\overline{B} \cap G$ and contains none of $V(G) - v$.

Let Γ'' be the curve obtained from Γ' by replacing the path $\Gamma' \cap B$ with the path on the boundary of B which intersects neither e_i nor e_j. Then Γ'' is a noncontractible curve and $|\Gamma'' \cap G| < |\Gamma' \cap G|$. This contradicts the minimality of $|\Gamma' \cap G|$. Hence the statement of 3.5.3 holds.

Therefore $G \backslash Z$ is planar and the face created by the deletion of Z is bounded by two cycles S and T. Let G' be $G \backslash Z$ with the vertices of degree two suppressed, and embed G' in the plane such that S and T are facial cycles. Let H be a plane dual of G', and let s and t be the vertices of H associated with the faces S and T of G'. Since G' is cubic, we know that H is a triangulation.

Using Theorem 3.4.5, we bipartition $E(H)$ into E'_1 and E'_2 such that all the edges incident to s lie in E'_1, and $H \backslash E'_1$ and $H \backslash E'_2$ are series parallel. Let E_T consist of the edges incident to t, and let

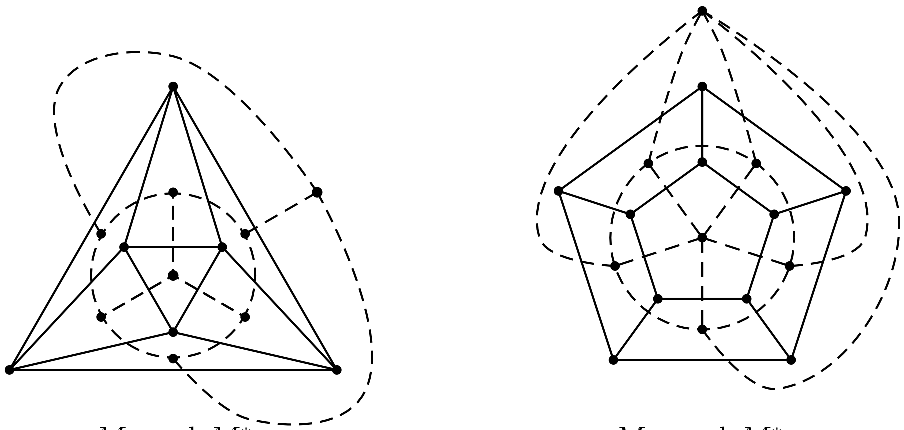

M_6 and M_6^* $\qquad$ M_{10} and M_{10}^*

FIGURE 3.13. M_6, M_{10}, and their plane duals.

$E_1 = E_1' \backslash E_T$ and $E_2 = E_2' \cup E_T$. Then $H \backslash E_2$ is a subgraph of $H \backslash E_2'$ and therefore is series-parallel. Therefore, by Lemma 3.2.15, we know that G/E_2 is series-parallel.

We see that $H \backslash E_1$ can be obtained from $H \backslash E_1'$ by adding a vertex to $V(H \backslash E_1')$ and adding any number of edges between that vertex and the rest of the vertices in $H \backslash E_1'$. We know that the unique plane dual of M_6 is the cube, and the unique plane dual of M_{10} is the double 5-wheel. See Figure 3.13.

Furthermore, both the cube minus any vertex and the double 5-wheel minus any vertex contain a K_4-minor. Therefore, since $H \backslash E_1'$ contains no K_4-minor, we know that $H \backslash E_1$ contains no minor isomorphic to the cube or the double 5-wheel. Hence, by Lemma 3.2.15, we know that G'/E_1 has no minor isomorphic to M_6 or M_{10}. And by planarity, we know that G'/E_1 has no minor isomorphic to K_5 or M_8. Hence $tw(G'/E_2) \leq 2$ and $tw(G'/E_1) \leq 3$. Furthermore, $S \subseteq E_1$ and $T \subseteq E_2$.

Let u be the vertex to which S is contracted in G'/E_1, and let v be the vertex to which T is contracted in G'/E_2. Then we can construct a tree decomposition of $G/(E_1 \cup Z)$, with width at most four, by adding u to every bag in the tree decomposition of G'/E_1. And we can construct a tree decomposition of G/E_2, with width at most three, by adding v to every bag in the tree decomposition of G'/E_2. Hence $\{E_1 \cup Z, E_2\}$ is the desired bipartition of $E(G)$. $\square$

Chapter 4

Structure of Cubic, Internally 4-Connected Graphs

4.1 Introduction

In this chapter, we deal only with graphs. Let B be a G-bridge of H, and let A be the set of attachments of B. For notational ease, in this chapter, we define $H\backslash B$ to be $H\backslash E(B)\backslash(V(B)\backslash A)$.

If G and H are two graphs, then a *topological embedding of* G *in* H is a one-to-one map from G to H. If f is a topological embedding of G in H, then the f-image of an edge is called a *branch*; and a vertex v in H is a *branch vertex* if v is the f-image of a vertex in G. Notice that each branch has branch vertices as its endpoints. We will often refer to a branch by its endpoints, as in, "the branch xy." We will often speak of connected subsets of branches. If xy is a branch, then an *xy-segment* is a connected subset of xy. An *open xy-segment* is an xy-segment which is homeomorphic to the open unit interval. A *closed xy-segment* is an xy-segment which is homeomorphic to the closed unit interval. A *half-open xy-segment* is an xy-segment which is homeomorphic to the half-open unit interval. If a and b are distinct points on the branch xy, then we will often let (a, b) refer to the open xy-segment with endpoints a and b. We will often let $[a, b]$ refer to the closed xy-segment with endpoints a and b. And we will often let $(a, b]$ and $[a, b)$ refer to the two appropriate half-open xy-segments with endpoints a and b.

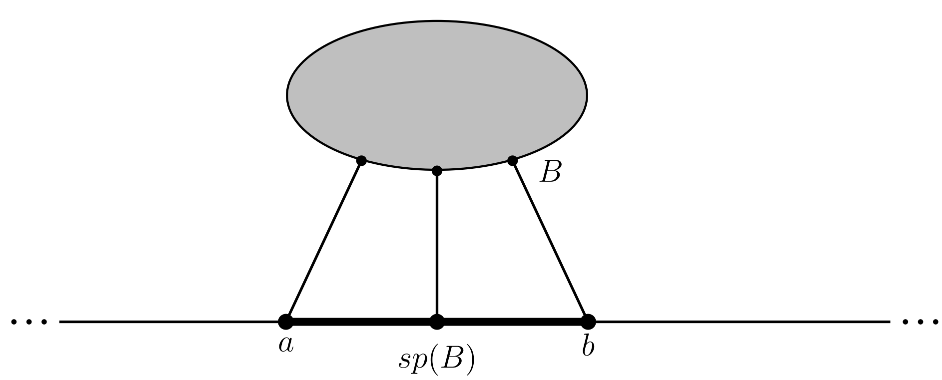

FIGURE 4.1. The attachments of the bridge B lie on a single branch. The span of B is shown in bold.

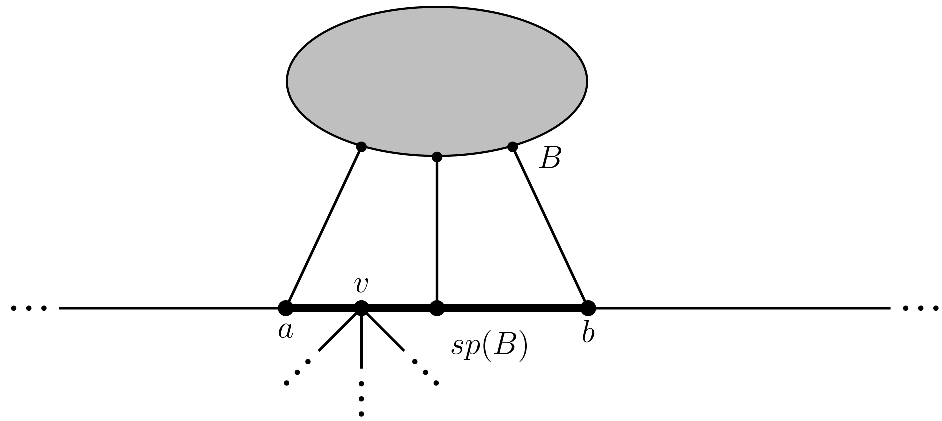

FIGURE 4.2. The attachments of the bridge B lie on two distinct, incident branches. The span of B is shown in bold.

Let B be an $f(G)$-bridge which has exactly two attachments, say a and b, neither of which is a branch vertex, such that the branches on which a and b lie are either equal or incident. Then B determines a closed, connected topological subspace of $f(G)$ called the *span of* B, notated $sp(B)$, which we define as follows: if both attachments of B lie on the same $f(G)$-branch, then $sp(B)$ is the closed subsegment of that branch with a and b as endpoints; otherwise a and b lie on distinct branches, in which case the branches are incident, at v say, and we define $sp(B)$ to be the union of the closed branch-segments $[a, v]$ and $[v, b]$. See Figures 4.1 and 4.2.

Let B be any $f(G)$-bridge in H. Notice that f yields an embedding of G in $H \backslash B$. Clearly B contains an (a, b)-path for any pair $\{a, b\}$ of distinct attachments of B. Then if P is an (a, b)-path in B for some pair $\{a, b\}$ of attachments of B, then P is an $f(G)$-bridge of $(H \backslash B) \cup P$; if a and b lie on equal or incident branches, then we may consider the span of P. If every pair $\{a, b\}$ of attachments of B lies on exactly one branch or two incident branches, then we define the *span of* B (denoted $sp(B)$) as follows: for each pair $\{a_i, b_i\}$ of attachments of B which lies on either exactly one branch or two incident branches, let P_i be an (a_i, b_i)-path in B; let $sp(B)$ be the union of the spans $sp(P_i)$ over all such pairs $\{a_i, b_i\}$.

If xy is a branch, then an attachment a which lies in the interior of xy is an (xy, x)-*incoming attachment* if there is a bridge B such that a is an attachment of B and $x \in sp(B)$. For example, in Figure 4.2, if the vertex b is on the branch vw, then b is a (vw, v)-incoming attachment.

Two bridges are said to *overlap* if their spans intersect in more than one (topological) point. Using the notion of overlapping, we define an equivalence relation, in which two bridges B, B' are equivalent if the following holds: there is a sequence of bridges $B = B_1, B_2, \ldots, B_k = B'$ such that B_i and B_{i+1} overlap, for $i \in \{1, \ldots, k-1\}$. We call the resulting equivalence classes *clusters*, and we refer to the cluster containing B as the *cluster closure of* B, denoted $cl(B)$. We define the *span of a cluster* to be the union of the spans of the bridges in the cluster. A bridge or cluster B *spans a point* x *completely* if x lies in the interior of the span of B. A bridge or cluster B *spans a set* S *of points completely* if B spans completely every point in S.

We define two operations on a graph G:

(O1) Subdivide two non-adjacent edges and add an edge between the newly-created pair of vertices.

(O2) For three edges A, B, C which form a path and not a triangle, subdivide A and C once, and perform $2n$ subdivisions on B, where $n \geq 1$. Name the new vertices $v_0, v_1, v_2, \ldots, v_{2n}, v_{2n+1}$, respectively, where v_1 is adjacent to an endpoint of A, and v_{2n} is adjacent to an endpoint of C, and $v_1, \ldots, v_{2n}$ form a path. Add edges $v_0v_2, v_{2n-1}v_{2n+1}$, and v_iv_{i+3} for all $i \in \{1, 3, 5 \ldots, 2n-3\}$. See Figure 4.3

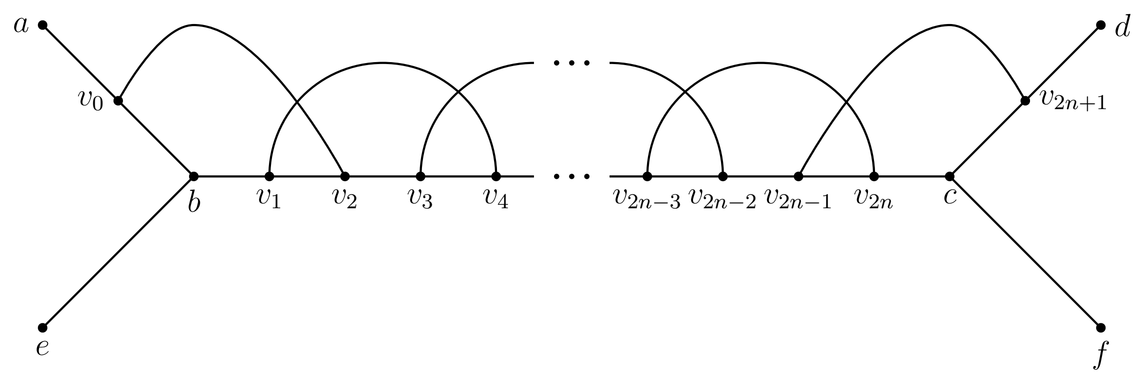

FIGURE 4.3. The operation (O2).

In 1968, Kotzig [12] proved that every cubic, internally 4-connected graph with more than eight vertices can be constructed from the cube by repeated instances of (O1). We prove here a stronger result involving topological embeddings, and Kotzig's result follows as a corollary.

4.2 Results

The following lemma proves that our operations maintain internal 4-connectedness.

Lemma 4.2.1. *If G is a cubic, internally 4-connected graph, and G' is a graph obtained from G via an instance of (O1) or (O2), then G' is internally 4-connected.*

Proof. Let G be a cubic, internally 4-connected graph.

Case 1. Suppose that G' is a graph obtained from G via an instance of (O1), such that G contains non-adjacent edges ab and cd, and G' contains edges av_1, v_1b, cv_2, v_2d, v_1v_2.

Suppose, en route to a contradiction, that $\{x, y, z\}$ is a nonverticial 3-separation in G'. Clearly, if neither of v_1, v_2 is in $\{x, y, z\}$, then $\{x, y, z\}$ is a nonverticial 3-separation of G; this is a contradiction. If precisely one of v_1, v_2 is in $\{x, y, z\}$, say, $x = v_1$, then one of $\{a, y, z\}$,$\{b, y, z\}$ is a nonverticial 3-separation of G; this is a contradiction. If both of v_1, v_2 are in $\{x, y, z\}$, say, $v_1 = x$ and $v_2 = y$, then one of $\{a, c, z\}$, $\{a, d, z\}$, $\{b, c, z\}$, $\{b, d, z\}$ is a nonverticial 3-separation of G; this is a contradiction. Therefore G' has no such 3-separation. Hence G' is internally 4-connected.

Case 2. Suppose that G' is a graph obtained from G via an instance of (O2), with vertices labeled as in Figure 4.3. We can view G as being embedded in G', and therefore speak of the branches ab, bc, and cd. Suppose, en route to a contradiction, that $\{x, y, z\}$ describes a nonverticial 3-separation in G'. Clearly, if no v_i, for $i \in \{0, \ldots, 2n + 1\}$, is in $\{x, y, z\}$, then $\{x, y, z\}$ is a nonverticial 3-separation of G; this is a contradiction. If v_i, for $i \in \{0, \ldots, 2n + 1\}$, is in $\{x, y, z\}$, say $x = v_i$, then one of $\{a, y, z\}$, $\{b, y, z\}$, $\{c, y, z\}$, $\{d, y, z\}$ is a nonverticial 3-separation of G; this is a contradiction. If precisely two or three of x, y, z meet $\{v_0, v_1, \ldots, v_{2n+1}\}$, then we can similarly find a nonverticial 3-separation of G which contradicts the internal 4-connectedness of G. Therefore G' has no nonverticial 3-separations, and hence G' is internally 4-connected. □

We now prove our main result, that if we embed one internally 4-connected graph in another internally 4-connected graph, then the embedding admits an instance of (O1) or (O2); that is, one of the bridges of the embedding will contain an instance of (O1) or (O2).

Theorem 4.2.2. *Let G and H be non-isomorphic, cubic, internally 4-connected graphs, and let f be a topological embedding of G in H. Then there is a cubic, internally 4-connected graph G' and* *a topological embedding f' of G' in H such that G' is obtained from G via one instance of (O1) or (O2), and $f'|_G = f$.*

Proof. Let G and H be non-isomorphic, cubic, internally 4-connected graphs, and let f be a topological embedding of G in H.

Case 1. Suppose that there is an $f(G)$-bridge B which has attachments on nonadjacent branches of $f(G)$; let $\{a, b\}$ be a pair of such attachments. Let e_a and e_b be the edges of G corresponding to the branches on which a and b, respectively, lie. We know that B contains an (a, b)-path P which contains no attachments except a and b. Let G' be the graph obtained from G by performing (O1) on the edges e_a and e_b. Then we may view G as a subspace of G', and therefore we may define the topological embedding $f' : G' \longrightarrow H$ as follows:

(1) $f'|_G = f$;

(2) f' maps the edge $(f^{-1}(a))(f^{-1}(b))$ to the path P.

Then by Lemma 4.2.1, we know that G' is internally 4-connected; and by definition, we know that $f'|_G = f$.

Case 2. Suppose that there is no $f(G)$-bridge which has attachments on nonadjacent branches of $f(G)$. Then since G is cubic, we know that for any $f(G)$-bridge B, the attachments of B either lie on a pair of adjacent branches or a triple of branches which share a single branch vertex.

Case 2a. Suppose that every $f(G)$-bridge spans a branch vertex, and suppose that there are two paths P_1, P_2 in H which satisfy the following conditions:

(1) P_1 and P_2 are disjoint subgraphs of $f(G)$-bridges;

(2) The endpoints of both paths are attachments of $f(G)$-bridges;

(3) There is exactly one $f(G)$-branch which contains endpoints of both P_1 and P_2;

(4) The span of P_1 contains exactly one endpoint of P_2.

Let G' be the graph obtained from G by performing (O2) with $n = 1$, on the edges corresponding to the branches on which the attachments of P_1 and P_2 lie. Let f' be an embedding of G' in H such that the following hold:

(1) $f'|_G = f$;

(2) f' maps $G' \backslash G$ to P_1 and P_2.

Then by Lemma 4.2.1, we know that G' is internally 4-connected; and by definition, we know that $f'|_G = f$.

Case 2b. Suppose that every $f(G)$-bridge spans a branch vertex, and suppose that no two paths in H satisfy the conditions of Case 2a. We will here derive a contradiction. We know that for any branch ab, there is a (topological) point x on ab such that the (ab, a)-incoming attachments lie on the closed ab-segment $[a, x]$ and the (ab, b)-incoming attachments lie on the closed ab-segment $[x, b]$. Then for any branch vertex v which is spanned by an $f(G)$-bridge, we can find a nonverticial 3-separation in the following way: Consider the induced subgraph $S = H[v \cup N(v) \cup A_v]$, where A_v is the set of all attachments of bridges which span v. We see that $(H[E(S)], H[E(H)\backslash E(S)])$ is a nonverticial 3-separation of H; this is a contradiction.

Case 2c. Suppose that there is an $f(G)$-bridge B whose attachments lie wholly on one branch vw_1. We prove now that there is a branch which is spanned by $cl(B)$. We see that $cl(B)$ must span one of v, w_1 completely to preserve internal 4-connectivity. Suppose, without loss of generality, that $cl(B)$ spans v completely. Let vw_2 and vw_3 be the other two branches incident to v. Suppose, en route to a contradiction, that $cl(B)$ spans none of vw_1, vw_2, vw_3 completely. Then there are unique attachments x, x', x'' of $cl(B)$ such that the following hold:

(1) The closed vw_1-segment $[x, v]$ is the smallest such segment containing every (vw_1, v)-incoming attachment;

(2) The closed vw_2-segment $[x', v]$ is the smallest such segment containing every (vw_2, v)-incoming attachment;

(3) The closed vw_3-segment $[x'', v]$ is the smallest such segment containing every (vw_3, v)-incoming attachment.

Then $\{x, x', x''\}$ is a nonverticial 3-separation of H; this is a contradiction. Therefore $cl(B)$ spans one of vw_1, vw_2, vw_3. Suppose, without loss of generality, that $cl(B)$ spans vw_1. Let B_0 and B_0' be the unique bridges which satisfy the following:

(1) $v \in sp(B_0)$ and $w_1 \in sp(B_0')$;

(2) $vw_1 \cap sp(B_0) \supseteq vw_1 \cap sp(D)$, for every bridge D which spans v;

(3) $vw_1 \cap sp(B_0') \supseteq vw_1 \cap sp(D)$, for every bridge D which spans w_1.

Since we forbade an instance of (O2) (by completing Case 2a), we know that $vw \backslash (sp(B_0) \cup sp(B_0'))$ is nonempty. We will construct a sequence $B_1, B_2, \ldots, B_k = B_0'$ of bridges in $cl(B)$ in the following way: Given $B_0, \ldots, B_i$, let B_{i+1} be a bridge of $cl(B)$ such that the following hold:

(1) $sp(B_i) \subsetneq vw_1$;

(2) $sp(B_i) \backslash \big(sp(B_0) \cup \cdots \cup sp(B_i)\big)$ is nonempty;

(3) $sp(B_i) \supseteq sp(D)$, for every bridge D which satisfies conditions (1) and (2);

(4) If $sp(B_i) \cap sp(B_0')$ is nonempty, then B_i is the final bridge in the sequence.

Let G' be the graph obtained from G by performing (O2) with $n = k + 1$ on the appropriate three edges. Let f' be an embedding of G' in H such that $f'|_G = f$ and $G' \backslash G$ gets mapped to appropriate paths in $B_0, B_1, \ldots, B_k$. $\square$

The following corollary allows us to ignore topological embeddings and speak only of minors.

Corollary 4.2.3. *Let G and H be non-isomorphic, cubic, internally 4-connected graphs, and assume that H contains G as a minor. Then there is a graph G' which is a minor of H and which arises from G via an application of (O1).*

Proof. Assume conditions of corollary. Cubicity ensures that H contains G as a topological minor. Hence there is an embedding f of G in H. Theorem 4.2.2 yields a cubic, internally 4-connected graph M and an embedding of M in H. Further, M is obtained from G via an instance of (O1) or (O2). If M arises via (O1), we are done.

Therefore assume that M arises from G via one application of (O2), as labeled in Figure 4.3.

Case 1. Assume that n is even. We define f' to be f with the following modification:

$f'(bc)$ is the path $u_0, u_1, \ldots, u_k$, where $u_0 = b, u_k = c$, and if i is even then $u_i = v_{j+1}$, where $v_j = u_{i+1}$, and if i is odd, then $u_i = v_{j+3}$, where $v_j = u_{i-1}$.

Figure 4.4 shows the image of G in solid, and the $f'(G)$-bridge in dashed.

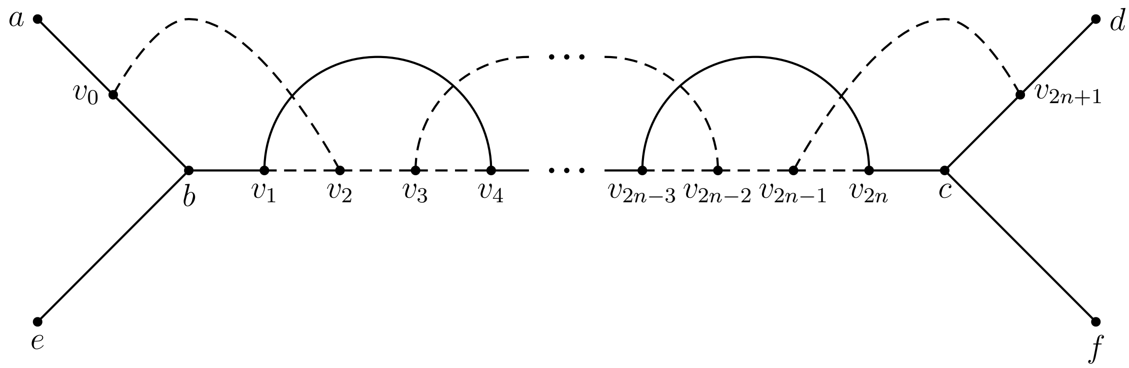

FIGURE 4.4.

We see then that if we delete all edges v_iv_{i+1} where $i \in \{1, 3, 5, \dots, 2n-1\}$, and if we contract all but one edge in the remaining dashed (v_0, v_{2n+1})-path in Figure 4.4, then we obtain a graph G' which is a minor of H and which arises from G via an application of (O1).

Case 2. Assume that n is odd. We define f' to be f with the following modifications:

(1) $f'(c) = v_{2n+1}$;

(2) $f'(cf)$ is the path induced by v_{2n+1}, c, f;

(3) $f'(cd) = v_{2n+1}d$;

(4) $f(bc)$ is the path induced by $u_0, u_1, \dots, u_k$, where $u_0 = b, u_k = v_{2n+1}$, and if i is even, then $u_i = v_{j+1}$, where $v_j = u_i$, and if i is odd, then $u_i = v_{j+3}$, where $v_j = u_i$.

We see then that if we delete all edges v_iv_{i+1} where $i \in \{1, 3, 5, \dots, 2n-1\}$, and if we contract all but one edge in the dashed (v_0, c)-path in Figure 4.5 below, then we obtain a graph G' which is a minor of H and which arises from G via an application of (O1). $\square$

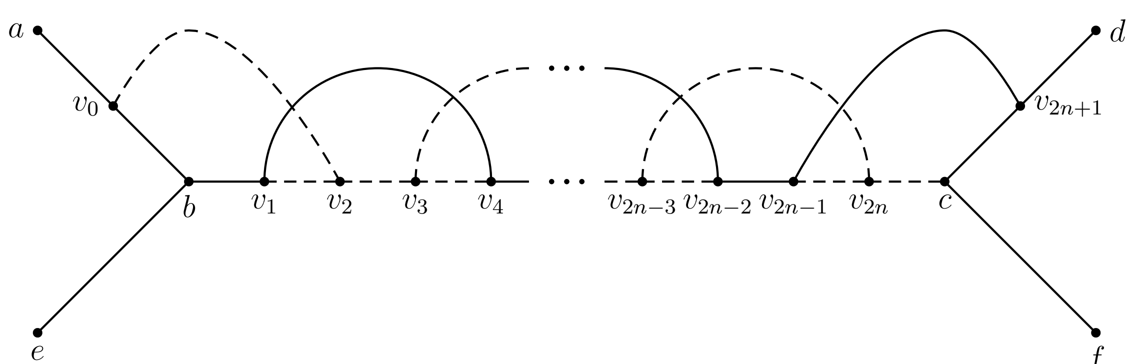

FIGURE 4.5.

References

[1] S. Arnborg, D. Corneil, and A. Proskurowski, Forbidden minors characterization of partial 3-trees, *Discrete Math* **80** (1990), 1–19.

[2] S. Arnborg and A. Proskurowski, Characterization and recognition of partial 3-trees, *SIAM J. Alg. Discrete Math.* **7** (1986), 305–314.

[3] T. Böhme, H. Broersma, and H. Tuinstra, A note on a conjecture concerning tree-partitioning 3-regular graphs, *Tatra Mt. Math. Publ.*, **18** (1999), 15–21.

[4] G. Chartrand, D. Geller, and S. Hedetniemi, Graphs with forbidden subgraphs, *J. Combin. Theory Ser. B* **10** (1971) 305–314.

[5] E. D. Demaine, M. Hajiaghayi, and B. Mohar, Approximation algorithms via contraction decomposition, *18th Symposium on Discrete Algorithms (SODA)*, ACM-SIAM, January 2007, 278–287.

[6] R. Diestel, *Graph Theory, Third Edition*, Graduate Texts in Mathematics, **173**, Springer-Verlag, Heidelberg, 2005.

[7] G. Ding, B. Oporowski, D. Sanders, and D. Vertigan, Surfaces, tree-width, clique-minors, and partitions, *J. Combin. Theory Ser. B* **79** (2000), no. 2, 221–246.

[8] D. Gonçalves, Edge partition of planar graphs into two outerplanar graphs, *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, ACM, New York, 2005, 504–512.

[9] R. Halin, S-functions for graphs, *J. Geometry* **8** (1976), 171–186.

[10] F. Jaeger, On vertex-induced forests in cubic graphs, Proceedings of the Fifth Southeastern Conference on Combinatorics, Graph Theory and Computing (Florida Atlantic Univ., Boca Raton, Fla., 1974), pp. 501–512. Congressus Numerantium, No. X, *Utilitas Math.*, Winnipeg, Man., 1974.

[11] K.S. Kedlaya, Outerplanar partitions of planar graphs, *J. Combin. Theory Ser. B* **67** (1996), 238–248.

[12] A. Kotzig, Regularly connected trivalent graphs without non-trivial cuts of cardinality 3, *Acta Fac. Rerum Natur. Univ. Comenian. Math.* **21** (1968), 1–14.

[13] A. Nakamoto and S. Negami, Diagonal flips in graphs on closed surfaces with specified face size distributions, *Yokohama Math. J.* **49** (2002), 171–180.

[14] C. St. J. A. Nash-Williams, Edge-disjoint spanning trees of finite graphs, *J. London Math. Soc.* **36** (1961), 445–450.

[15] S. Negami, Diagonal flips in pseudo-triangulations on closed surfaces, *Discrete Math* **240** (2001), 187–196.

[16] S. Negami, Diagonal flips of triangulations on surfaces, a survey. Proceedings of the 10th Workshop on Topological Graph Theory (Yokohama, 1998), *Yokohama Math. J.* **47** (1999), Special Issue, 1–40.

[17] B. Oporowski, private communication, 2003.

[18] J. Oxley, private communication, 2002.

[19] N. Robertson and P.D. Seymour, Graph Minors, II. Algorithmic aspects of tree-width, *J. Algorithms* **7** (1986), 309–322.

[20] D.P. Sanders, On Hamiltonian cycles in certain planar graphs, *J. Graph Theory* **21** (1996), 43–50.

[21] W. T. Tutte, On the problem of decomposing a graph into n connected factors, *J. London Math. Soc.* **36** (1961), 221–230.

[22] W. T. Tutte, A theorem on planar graphs, *Trans. Amer. Math. Soc.* **82** (1956), 99–116.

Vita

Evan Morgan was born in July of 1980 in Portland, Oregon. He began and finished his undergraduate studies at Lawrence University in Appleton, Wisconsin, in June 2002, majoring in mathematics and minoring in computer science and English. He earned a Master of Science degree in mathematics from Louisiana State University in May 2004. He is currently a candidate for the degree of Doctor of Philosophy in mathematics at Louisiana State University.