Hindawi Publishing Corporation
International Journal of Combinatorics
Volume 2013, Article ID 783710, 7 pages
http://dx.doi.org/10.1155/2013/783710

Research Article

Graphs with no $K_{3,3}$ Minor Containing a Fixed Edge

Donald K. Wagner

Mathematical, Computer, and Information Sciences Division, Office of Naval Research, Arlington, VA 22203, USA

Correspondence should be addressed to Donald K. Wagner; don.wagner@navy.mil

Received 30 November 2012; Accepted 6 February 2013

Academic Editor: Chính T. Hoang

It is well known that every cycle of a graph must intersect every cut in an even number of edges. For planar graphs, Ford and Fulkerson proved that, for any edge e, there exists a cycle containing e that intersects every minimal cut containing e in exactly two edges. The main result of this paper generalizes this result to any nonplanar graph G provided G does not have a $K_{3,3}$ minor containing the given edge e. Ford and Fulkerson used their result to provide an efficient algorithm for solving the maximum-flow problem on planar graphs. As a corollary to the main result of this paper, it is shown that the Ford-Fulkerson algorithm naturally extends to this more general class of graphs.

1. Introduction

This paper examines the structure of paths and cuts in a graph relative to a fixed edge. In particular, let G be a graph, and let e be an edge of G. Define an *e-path* of G to be a path P such that $P \cup \{e\}$ is a cycle of G. Define an *e-cut* of G to be a cut of G that contains e (in this paper, paths and cycles do not have repeated nodes and are equated with their edge sets. Also, cuts are minimal; i.e., no cut properly contains another.) Ford and Fulkerson [1] showed that if G is planar, then there exists an e-path that intersects every e-cut in exactly one edge. This Ford-Fulkerson property does not hold for graphs in general. Specifically, take $G = K_{3,3}$. Then, for any edge e of G and any e-path P, one can always find an e-cut that intersects P in more than one edge. The Ford-Fulkerson property, however, is not confined solely to planar graphs; in particular, if $G = K_5$, then it is easy to find an e-path, for any choice of e, that intersects every e-cut in exactly one edge.

One of the main goals of this paper is to extend the Ford-Fulkerson result to a larger class of graphs. Motivated by the $K_{3,3}$ example above, it is shown that if $K_{3,3}$ is excluded in the proper way, then this goal can be achieved. Below is the main result of the paper. Throughout the paper, n denotes the number of nodes of a graph, and m the number of edges.

Theorem 1. *Let G be a graph, and let e be an edge of G. If $G \setminus e$ is connected, and G does not have a $K_{3,3}$ minor containing e, then there exists an e-path of G that intersects every e-cut of G in exactly one edge. Moreover, such an e-path can be found in $O(m)$ time.*

Theorem 1 is then used to provide a very simple $O(n^2)$-time algorithm for the maximum-flow problem for graphs in this class. This is within a logarithmic factor of the fastest maximum-flow algorithm, namely, the recent algorithm due to Orlin [2].

The remainder of the paper is outlined as follows. The next section introduces a graph decomposition, which serves as a key ingredient for the proof of Theorem 1. Section 3 contains the proof of Theorem 1, and Section 4 applies Theorem 1 to the maximum-flow problem.

2. Graph Decomposition

This section describes a connectivity-based decomposition for graphs that do not have a $K_{3,3}$ minor containing a fixed edge. This decomposition was introduced in Wagner [3]. For the sake of completeness, the key results are presented and proved here.

The notion of connectivity used here is that of Tutte [4]. A *k-separation*, for a positive integer k, of a connected graph G is a partition $\{E_1, E_2\}$ of the edge set of G such that $|E_1| \geq$

$k \le |E_2|$ and the edge-induced subgraphs $G[E_1]$ and $G[E_2]$ have at most k nodes in common. A connected graph G is *k-connected*, for $k \ge 2$, if it does not have a k'-separation for any $k' < k$. A k-separation $\{E_1, E_2\}$ of a k-connected graph G is an *internal* k-separation if $|E_1| \ge k+1 \le |E_2|$.

The next theorem is a well-known result of Wagner [5].

Theorem 2. *Let G be a* 3*-connected graph. Then, G does not have a $K_{3,3}$ minor if and only if G is planar or isomorphic to K_5.*

It is sometimes more convenient to work with subdivisions rather than minors. A graph H is a *subdivision* of a graph K if it can be obtained from K by a sequence of the following operation: replace an edge xy by edges xz and yz, where z is a new node. If a graph G has a subgraph H that is a subdivision of a graph K, then G is said to have a *K subdivision*. It is well known and easy to prove that a 3-connected graph has a $K_{3,3}$ subdivision if and only if it has a $K_{3,3}$ minor.

If a graph H is a subdivision of a graph K, and x and y are nonadjacent nodes of K, then x and y are *independent* in H. The next lemma is due to Širáň [6].

Lemma 3. *Let G be a* 3*-connected graph, and let $e = xy$ be an edge of G. If e is not contained in any $K_{3,3}$ minor of G, then for any $K_{3,3}$ subdivision H of G, x and y are independent degree-three nodes of H.*

In this paper, 2- and 3-separations play a crucial role, as do the related notions of 2- and 3-sums. First, consider a 2-separation $\{E_1, E_2\}$ of a 2-connected graph G. Let $\{p, q\} := V(G[E_1]) \cap V(G[E_2])$, and let $\{t\}$ be a set disjoint from $E(G)$. For $i \in \{1, 2\}$, define G_i to be the graph obtained from $G[E_i]$ by adding t as an edge joining p and q. Then, $\{G_1, G_2\}$ is a 2*-sum decomposition* of the graph G, and t is the *connecting* edge.

Now, let $\{E_1, E_2\}$ be an internal 3-separation of a 3-connected graph G. Let $\{x, y, z\} := V(G[E_1]) \cap V(G[E_2])$, and let S be the set of edges of G that have both end nodes in $\{x, y, z\}$. Let R be a set disjoint from $E(G)$ such that $|R \cup S| = 3$. For $i \in \{1, 2\}$, construct a graph G_i from $G[E_i \cup S]$ by adding the members of R in such a way that $T := R \cup S$ is a triangle of G_i and such that each edge of R has the same ends in G_1 as it does in G_2. Then, $\{G_1, G_2\}$ is a 3*-sum decomposition* of G, and T is the *connecting* triangle.

It is well known that if $\{G_1, G_2\}$ is a k-sum decomposition of a k-connected graph G, for $k \in \{2, 3\}$ then both G_1 and G_2 are k-connected and are isomorphic to proper minors of G.

Two special kinds of internal 3-separations are needed. Both are defined for a given 3-connected graph G relative to a fixed edge e.

First, let $\{E_1, E_2\}$ be a internal 3-separation of G. If both ends of e are in $V(G[E_1]) \cap V(G[E_2])$, then the 3-separation $\{E_1, E_2\}$ is said to be *straddled* by e. Observe that in this case, e is in the connecting triangle of the corresponding 3-sum decomposition. The notion of a straddling edge can be found in the work of Tseng and Truemper [7]. It is also related to the concept of "contractibility," which traces back to the work of Tutte [8].Specifically, an edge of a 3-connected graph is *contractible* if its contraction results in a 3-connected graph. It is easy to see that an edge is not contractible if and only if it straddles an internal 3-separation.

The second special internal 3-separation is as follows. Let $\{E_1, E_2\}$ be an internal 3-separation of G, and suppose E_1 has exactly seven edges, say $e, f_1, \ldots, f_6$. Suppose further that $\{e, f_1, f_2\}, \{e, f_3, f_4\}$, and $\{e, f_5, f_6\}$ are triangles of G such that no two of $\{f_1, \ldots, f_6\}$ are parallel. Then, $G[E_1]$ is a *crown*, and $\{E_1, E_2\}$ is a *crown* 3-separation of G with respect to e. Observe that the crown $G[E_1]$ has three nodes of degree two, which, by the 3-connectivity of G, constitute the set $V(G[E_1]) \cap V(G[E_2])$. It also has two nodes of degree four, which are the ends of e.

Let $\{E_1, E_2\}$ be a *crown* 3-separation of G with respect to e, and let $\{G_1, G_2\}$ be the corresponding 3-sum decomposition. If G_2 is planar, then G is said to be *crown-planar* with respect to e. Crown-planar graphs show up in the decomposition established in Theorem 5. In the context of Theorem 5, crown-planar graphs can alternatively be described as being obtained from a 3-connected planar graph by duplicating a degree-three node x, where the fixed edge e joins x to its twin.

Let G be a graph, and H a subgraph of G. Let P be a path of G, the end nodes of which are nodes of H and the internal nodes of which are not nodes of H. Then, the subgraph $H \cup P$ of G is said to be obtained from H by *adjoining* P, and P is an *adjoinable* path of G with respect to H.

Let G be a graph, and e an edge of G. Let H be a $K_{3,3}$ subdivision of G, and suppose that e joins two independent degree-three nodes of H. Since $K_{3,3}$ has nine edges, the graph H consists of nine paths, each of which is a subdivision of an edge of $K_{3,3}$. The six such paths that share an end with e are called the *principal* paths of H with respect to e; the remaining three paths are the *support* paths of H. The $K_{3,3}$ subdivision H of G is *good* (resp., *bad*) with respect to e if all six (resp., at most five) of the principal paths with respect to e consist of a single edge.

Lemma 4. *Let G be a* 3*-connected graph, and let e be an edge of G. Then, either (i) G has a $K_{3,3}$ minor that contains e, (ii) G has an internal* 3*-separation that is straddled by e, or (iii) every $K_{3,3}$ subdivision of G is good with respect to e.*

Proof. Let $e = xy$. Suppose that neither (i) nor (iii) holds. If G is planar or isomorphic to K_5, then (iii) holds vacuously, and so, Theorem 2 implies that G has a $K_{3,3}$ minor, and thus a $K_{3,3}$ subdivision. By Lemma 3, x and y are independent degree-three nodes in every $K_{3,3}$ subdivision of G. Since (iii) does not hold, there exists a $K_{3,3}$ subdivision of G, say H, in which some principal path with respect to e, say Q_1, has at least two edges. Let u denote the end node of Q_1 not in $\{x, y\}$. Let Q_2 denote the other principal path that has u as an end node, and let S_1 denote the support path that has u as an end node. Denote the other end node of S_1 by z. Consistent with the above, assume H and Q_1 are chosen so that the number of edges in S_1 is as small as possible.

Claim. If an adjoinable path of G with respect to H has one end that is an internal node of either Q_1 or Q_2, then the other end of the path is a node of $V(Q_1 \cup Q_2 \cup S_1)$.

Proof of Claim. If the other end of the path is not in $V(Q_1 \cup Q_2 \cup S_1)$, then it is easy to check that adjoining the path to H results in a graph that has a $K_{3,3}$ minor that contains e, a contradiction. *End of Claim.*

Observe that $\{Q_1 \cup Q_2 \cup \{e\}, E(H) - (Q_1 \cup Q_2 \cup \{e\})\}$ is an internal 3-separation of H straddled by e. Thus, either (ii) holds or there exists an adjoinable path R_1 of G, one end of which, say r_1, is a internal node of Q_1 (say) and the other end of which, say t_1, is not in $V(Q_1 \cup Q_2)$. By the Claim, t_1 is a node of S_1; if it is an internal node of S_1, then a contradiction to the choice of H is obtained by adjoining R_1 to H and deleting the internal nodes of the ur_1-subpath of subpath of Q_1. Thus, $t_1 = z$.

Observe that $\{Q_1 \cup Q_2 \cup S_1, E(H) - (Q_1 \cup Q_2 \cup S_1)\}$ is an internal 3-separation of H straddled by e. Thus, either (ii) holds or there exists an adjoinable path R_2 of G with respect to H, one end of which, say r_2, is in $V(Q_1 \cup Q_2 \cup S_1)$, the other end of which, say t_2, is not in $V(Q_1 \cup Q_2 \cup S_1)$, and neither end of which is in $\{x, y, z\}$. By the Claim, r_2 is a node of S_1 and t_2 is a node of P, where P is one of the principal or support paths of H not in $\{Q_1, Q_2, S_1\}$. Moreover, r_2 must equal u, for otherwise a contradiction to the choice of H is obtained by adjoining R_2 to H and deleting the internal nodes of the zt_2-subpath of P. By the Claim, R_1 and R_2 are node disjoint. Now, $H \cup R_1 \cup R_2$ has a $K_{3,3}$ minor that contains e, a contradiction. □

Theorem 5 below is the main result of the section.

Theorem 5. *Let G be a 3-connected graph, and let e be an edge of G. Then, either (i) G is planar, (ii) G is isomorphic to K_5, (iii) G has a $K_{3,3}$ minor that contains e, (iv) G has an internal 3-separation that is straddled by e, or (v) G is crown-planar with respect to e.*

Proof. Lemma 4 and Theorem 2 together imply that either one of (i)–(iv) holds, or every $K_{3,3}$ subdivision of G is good with respect to e. Assume that none of (i)–(iv) hold and let H denote a $K_{3,3}$ subdivision of G that is good with respect to e. Let $e = xy$, and let z denote the common end node of the three support paths of H. Let u, v, and w denote the remaining degree-three nodes of H. Let S_1, S_2, and S_3 denote the three support paths of H with respect to e, and without loss of generality, assume that the ends of S_1 are u and z.

Observe that $\{\{S_1, ux, uy, e\}, E(H) - \{S_1, ux, uy, e\}\}$ is an internal 3-separation of H straddled by e. Since (iv) does not hold, there exists an adjoinable path R_1 of G with respect to H, one end of which is in $V(S_1)$, the other end of which is in $V(S_2)$ (say), and neither end of which is equal to z. Similarly, there exists an adjoinable path R_2 of G with respect to H, one end of which is $V(S_3)$, the other end of which is in $V(S_1)$ (say), and neither end of which is equal to z.

Observe that $\{\{e, ux, uy, vx, vy, wx, wy\}, S_1 \cup S_2 \cup S_3 \cup R_1 \cup R_2\}$ is a crown 3-separation with respect to e of $H \cup R_1 \cup R_2$. Thus, either G has a crown 3-separation with respect to e or there exists an adjoinable path R_3 of G with respect to $H \cup R_1 \cup R_2$, one end of which is in $\{x, y\}$ and the other end of which, call it t, is in $V(S_1 \cup S_2 \cup S_3 \cup R_1 \cup R_2)$. If $t \neq z$, then observe that $H \cup R_1 \cup R_3 \cup R_3$ contains a bad $K_{3,3}$ subdivision with respect to e, a contradiction (note, if $t \in \{u, v, w\}$, then by the 3-connectivity of G, R_3 has at least two edges). Thus, $t = z$. It can now be checked that $H \cup R_1 \cup R_2 \cup R_3$ contains a $K_{3,3}$ minor containing e, a contradiction.

Finally, it needs to be shown that if G has a crown 3-separation with respect to e, and none of (i)–(iv) hold, then G is crown-planar with respect to e. To see this, let $\{G_1, G_2\}$ be the 3-sum decomposition of G corresponding to the crown 3-separation with respect to e, where $e \in E(G_1)$. Then, it suffices to show that G_2 is planar. If this is not the case, then by Theorem 2, G_2 is either isomorphic to K_5 or has a $K_{3,3}$ subdivision. In either case, it is straightforward to see that G has a $K_{3,3}$ subdivision for which e does not join two independent nodes, contradicting Lemma 3. □

The next result is from Wagner [3]. It shows that if G is a simple 2-connected graph having an edge e that is not contained in a $K_{3,3}$ minor, then the number of edges of G is bounded $5n - 12$. The proof is a straightforward induction using 2- and 3-sum decompositions, together with Theorem 5 and the well-known fact that any planar graph has at most $3n - 6$ edges.

Lemma 6. *Let G be a simple 2-connected graph having at least three nodes. If, for some edge e, G does not have a $K_{3,3}$ minor containing e, then G has at most $5n - 12$ edges.*

3. Admissible Paths

This section presents a proof of Theorem 1. This section begins with two lemmas that relate an e-cut of a graph to that of a member of a k-sum decomposition of the graph.

Lemma 7. *Let G be a 2-connected graph, and let e be an edge of G. Suppose that $\{E_1, E_2\}$ is a 2-separation of G with $e \in E_1$. Let $\{G_1, G_2\}$ be the corresponding 2-sum decomposition, and let t denote the connecting edge. Let D be an e-cut of G. Then, either D or $D - E_2 \cup \{t\}$ is an e-cut of G_1. Moreover, in the latter case, $D - E_1 \cup \{t\}$ is a t-cut of G_2.*

Proof. Let p and q denote the nodes common to $G[E_1]$ and $G[E_2]$. Let $\{X, Y\}$ denote the node partition of $V(G)$ corresponding to the e-cut D of G.

First, it is shown that $D \cap E_2 \neq \emptyset$ if and only if $p \in X$ (say) and $q \in Y$. To this end, suppose that $p \in X$ and $q \in Y$. Observe, there exists a pq-path in $G[E_2]$, and any such path must contain an edge from D. Thus, $D \cap E_2 \neq \emptyset$. Now, suppose that $D \cap E_2 \neq \emptyset$, and let $f \in D \cap E_2$. Since $\{e, f\} \subseteq D$, there exists two paths, say P and Q, each of which joins an end node of f to an end node of e and such that $V(P) \subseteq X$ (say) and $V(Q) \subseteq Y$. Since $e \in E_1$ and $f \in E_2$, P must go through p (say) and Q through q. Thus, $p \in X$ and $q \in Y$.

Now, define $X_1 := X \cap V(G_1)$ and $Y_1 := Y \cap V(G_1)$. Then, X_1 and Y_1 are both nonempty since they each contain an end of e. Also, they are disjoint and their union equals $V(G_1)$. By the previous paragraph, it can be seen that the set of edges of G_1 that have exactly one end in X_1 is D (if $\{p, q\} \subseteq X_1$ (say)) or $D - E_2 \cup \{t\}$ (if $p \in X_1$ (say) and $q \in Y_1$). Thus, the conclusion that either D or $D - E_2 \cup \{t\}$ is an e-cut of G_1 follows provided the subgraphs $G_1[X_1]$ and $G_1[Y_1]$ are connected. To this end, let u and v be nodes of X_1, and let R be a uv-path in $G[X]$. If $R \subseteq E_1$, then R is a path of $G_1[X_1]$. Suppose that this is not the case. Then, $\{p, q\} \subseteq V(R)$. Moreover, the edges of $R \cap E_2$ constitute a pq-subpath S of R. Thus, replacing, in R, the subpath S by the edge t yields a uv-path in G_1. Thus, $G_1[X_1]$ is connected. Similarly, $G_1[Y_1]$ is connected.

Now, suppose that $D - E_2 \cup \{t\}$ is an e-cut of G_1. Showing that $D - E_1 \cup \{t\}$ is a t-cut of G_2 is done in a manner similar to the above. Specifically, define $X_2 := X \cap V(G_2)$ and $Y_2 := Y \cap V(G_2)$. Then, X_2 and Y_2 are nonempty since $p \in X_2$ and $q \in Y_2$. Also, they are disjoint and their union equals $V(G_2)$. Moreover, the set of edges of G_2 that have exactly one end in X_2 is precisely $D - E_1 \cup \{t\}$. Thus, the conclusion that $D - E_1 \cup \{t\}$ is a t-cut of G_2 follows provided the subgraphs $G_2[X_2]$ and $G_2[Y_2]$ are connected. To this end, let u and v be nodes of X_2, and let R be a uv-path in $G[X]$. Since $q \in Y_2$, $q \notin V(R)$. Therefore, $R \subseteq E_2$. Thus, $G_2[X_2]$ is connected. Similarly, $G_2[Y_2]$ is connected. □

Lemma 8. *Let G be a 3-connected graph, and let e be an edge of G. Suppose that $\{E_1, E_2\}$ is a 3-separation of G that is straddled by e. Let $\{G_1, G_2\}$ be the corresponding 3-sum decomposition, and let T denote the connecting triangle. Let D be an e-cut of G. Then, $D - E_2 \cup \{t, e\}$ is an e-cut of G_1 for some $t \in T - \{e\}$.*

Proof. Let x, y, and z denote the nodes common to $G[E_1]$ and $G[E_2]$. Let $T := \{e, t, t'\}$ with $e = xy$ and $t = yz$. Let $\{X, Y\}$ denote the node partition of $V(G)$ corresponding to the e-cut D of G. Without loss of generality, assume $\{x, z\} \subseteq X$ and $y \in Y$. Let $X_1 := X \cap V(G_1)$ and $Y_1 := Y \cap V(G_1)$. Then, X_1 and Y_1 are both nonempty since they each contain an end of e. Also, they are disjoint and their union equals $V(G_1)$. Moreover, the set of edges of G_1 that have exactly one end in X_1 is precisely $D - E_2 \cup \{t, e\}$. The result now follows provided the subgraphs $G_1[X_1]$ and $G_1[Y_1]$ are connected. To this end, let u and v be nodes of X_1, and let R be a uv-path in $G[X]$. If $R \subseteq E_1$, then R is a path of $G_1[X_1]$. Suppose that this is not the case. Then, $\{x, z\} \subseteq V(R)$. Moreover, the edges of $R \cap E_2$ constitute an xz-subpath S of P. Thus, replacing, in R, the subpath S by the edge t' yields a uv-path in G_1. Thus $G[X_1]$ is connected. Similarly, $G[Y_1]$ is connected. □

Let G be a graph, and let e be an edge of G. An e-path of G is called *admissible* if it intersects every e-cut in exactly one edge. The above two lemmas are now used to show how admissible paths relate to k-sums.

Lemma 9. *Let G be a 2-connected graph, and let e be an edge of G. Suppose that G has a 2-separation $\{E_1, E_2\}$ with $e \in E_1$. Let $\{G_1, G_2\}$ be the corresponding 2-sum, and let t be the connecting edge. Suppose that P_1 is an admissible e-path of G_1, and P_2 is an admissible t-path of G_2. If $t \notin P_1$, P_1 is an admissible e-path of G, and if $t \in P_1$, then $P_1 - \{t\} \cup P_2$ is an admissible e-path of G.*

Proof. Let D be an e-cut of G. By Lemma 7, either D or $D - E_2 \cup \{t\}$ is an e-cut of G_1.

First, consider the case that D is an e-cut of G_1. Since D is an e-cut of both G and G_1, it must be that $D \subseteq E_1$. In particular, $t \notin D$. Since P_1 is an admissible e-path of G_1, it follows that $|D \cap (P_1 - \{t\})| = 1$. Thus, if $t \notin P_1$, then P_1 is an admissible e-path of G, and if $t \in P_1$, then $P_1 - \{t\} \cup P_2$ is an admissible e-path of G.

Now, assume that $D - E_2 \cup \{t\}$ is an e-cut of G_1. Since P_1 is an admissible e-path of G_1, $|(D - E_2 \cup \{t\}) \cap P_1| = 1$. If $t \notin P_1$, then $P_1 \subseteq E_1$, from which it follows $|D \cap P_1| = 1$, implying that P_1 is an admissible e-path of G. Thus, assume $t \in P_1$. Since P_1 is an admissible e-path of G_1, $(D - E_2 \cup \{t\}) \cap P_1 = \{t\}$. Therefore, $(D \cap E_1) \cap (P_1 - \{t\}) = \emptyset$. Since P_2 is an admissible t-path of G, by Lemma 7, $|(D - E_1 \cup \{t\}) \cap P_2| = 1$. It follows that $|(D \cap P_1 - \{t\}) \cap P_2| = 1$, implying that $P_1 - \{t\} \cap P_2$ is an admissible e-path of G. □

Lemma 10. *Let G be a 3-connected graph, and let e be an edge of G. Suppose that G has an internal 3-separation straddled by e. Let $\{G_1, G_2\}$ be the corresponding 3-sum, and let T be the connecting triangle. Let P be an admissible e-path of G_1 that is edge disjoint from $T - \{e\}$. Then, P is an admissible e-path of G.*

Proof. Since P is edge disjoint from $T - \{e\}$, P is a path of G. Let D be an e-cut of G. By Lemma 8, $D - E_2 \cup \{t, e\}$ is an e-cut of G_1 for some $t \in T - \{e\}$. By assumption, $|(D - E_2 \cup \{t, e\}) \cap P| = 1$, which implies $|D \cap P| = 1$, as required. □

Below is the re-statement of Theorem 1 in terms of admissible paths.

Theorem 11. *Let G be a graph, and let e be an edge of G. If $G \backslash e$ is connected, and G does not have a $K_{3,3}$ minor containing e, then G has an admissible e-path. Moreover, such an e-path can be found in $O(m)$ time in general and in $O(n)$ time if G is 2-connected and simple.*

Proof. The proof is by a series of reductions.

(I) *Reduction to the Simple 2-Connected Case.* Since $G \backslash e$ is connected, every e-path of G is contained in the same block of G as e. Therefore, the search for an admissible e-path of G can be restricted to this block. Computing the blocks of G and identifying the one containing e can be done in $O(m)$ time; see Tarjan [9]. Also, since any e-path uses at most one edge from any parallel class, any parallel edges can be deleted, which requires $O(m)$ time. Now, by Lemma 6, m is $O(n)$.

(II) *Reduction to the 3-Connected Case.* The next step is to show that the admissible-path problem on G can be reduced in linear time to solve a sequence of admissible-path problems, where each problem in the sequence is defined on a graph that is 3-connected and does not contain a $K_{3,3}$ minor using a specified edge, and such that the total size of this sequence is linear in the size of G.

(IIa) *Admissible Paths and 2-Sum Decompositions.* The first step in defining this sequence is to examine the relationship between an instance of the admissible-path problem and a 2-sum decomposition in the underlying graph. So, suppose that G is not 3-connected, and let $\{E_1, E_2\}$ be a 2-separation of G with $e \in E_1$. Let $\{G_1, G_2\}$ be the corresponding 2-sum decomposition, and let $t = pq$ be the connecting edge. It is straightforward to verify that G_1 (resp., G_2) is 2-connected and does not have a $K_{3,3}$ minor containing e (resp., t). By Lemma 9, if P_1 is an admissible e-path of G_1, and P_2 is an admissible t-path of G_2, then an admissible e-path P of G is equal to either P_1, if $t \notin P_1$, and $P_1 - \{t\} \cup P_2$, otherwise.

(IIb) *The Reduction Procedure*. To turn the above relationship between admissible paths and 2-sum decomposition into a computationally efficient algorithm requires two straightforward ideas. First, one chooses the 2-sum decomposition judiciously, and second one applies this judicious choice recursively. Tutte [4] and Hopcroft and Tarjan [10] showed that one can always find a 2-separation $\{E_1, E_2\}$ of G such that $e \in E_1$, and in the resulting 2-sum decomposition $\{G_1, G_2\}$, G_2 is either 3-connected, a cycle on three or more edges, or a bond (i.e., the planar dual of a cycle) on three or more edges. Applying this choice of 2-sum decomposition recursively reduces the admissible-path problem on G to solve a sequence of admissible-path problems on a collection of graphs, say $\{H_1, \ldots, H_j\}$, every member of which is either 3-connected, a cycle, or a bond. Moreover, Hopcroft and Tarjan [10] showed that the sequence of 2-separations necessary to generate $\{H_1, \ldots, H_j\}$ can be found in $O(n)$ time and that the size of the collection, that is, $\sum_{i=1}^{j} |V(H_i)|$, is $O(n)$. Observe that an admissible path on a cycle or bond can be trivially found in linear time (in the size of the cycle or bond). Thus, in $O(n)$ time, the admissible-path problem on G can be reduced to solve a sequence of admissible-path problems, each of which is on a graph that is 3-connected and does not have a $K_{3,3}$ minor using a specified edge. Moreover, the total size of the graphs in the sequence is $O(n)$. Thus, it suffices to prove the theorem assuming G is 3-connected.

(III) *Reduction to the Planar Case*. Assume that G is 3-connected. By Theorem 5, G either is planar, isomorphic to K_5, crown-planar with respect to e or has an internal 3-separation straddled by e. These cases are considered one at a time. As a first step, it is shown that one can recognize which case is applicable in $O(n)$ time. Clearly, recognizing if G is isomorphic to K_5 can be done in constant time. Also, it is well known that planarity can be recognized in $O(n)$ time; see, for example, Hopcroft and Tarjan [11]. Determining whether G has an internal 3-separation straddled by e can be done by simply first deleting the ends of e and then determining if the resulting graph has a cut vertex; the latter can be done in $O(n)$ time using algorithm of Tarjan [9]. The only other possibility for G is that it is crown-planar with respect to e.

(IIIa) *The Base Cases*. This subcase considers these cases when G is either planar, isomorphic to K_5, or crown-planar with respect to e. For each of these three cases, it is shown how to find an admissible path of G in $O(n)$ time.

First, assume that G is planar. Then, as observed by Ford and Fulkerson [1], it is straightforward to find an admissible e-path of G. This is done as follows. Embed G in the plane, so that e is on the infinite face, and then delete e. Then, there exist two e-paths of G the union of which defines the outer face of $G \setminus e$. Ford and Fulkerson [1] showed that each of these paths is an admissible e-path of G (this is also easily seen by planar duality). Observe that if $G \setminus e$ is 2-connected, these two admissible e-paths are internally node disjoint, a fact that will be used later. Finding an embedding of a planar graph can be done in $O(n)$ time [11, 12], and so for planar graphs, it follows that the two admissible e-paths can be found in $O(n)$ time.

Second, consider the case that G is isomorphic to K_5 or is crown-planar with respect to e. In each of these cases, there exist two internally node-disjoint e-paths in G, each of which contains exactly two edges. Since, in general, every e-cut must intersect every e-path in an odd number of edges, it must be that each of these e-paths is admissible. Thus, finding these two admissible e-paths can be done in $O(n)$ time.

(IIIb) 3-*Sum Decompositions*. The final step is to analyze when G has an internal 3-separation straddled by e. Consider the graph $G \setminus \{x, y\}$, where x and y are the ends of e. Since G has an internal 3-separation, $G \setminus \{x, y\}$ has a cut node. Conversely, each cut node of $G \setminus \{x, y\}$ corresponds to an internal 3-separation of G straddled by e. Now, choose a block of $G \setminus \{x, y\}$ that has exactly one cut node in common with rest of $G \setminus \{x, y\}$, and let $\{E_1, E_2\}$ be the corresponding internal 3-separation of G. Let $\{G_1, G_2\}$ be the associated 3-sum decomposition of G, and let T be the connecting triangle. Then, by the choice the cut node, G_1 (say) does not have an internal 3-separation. Thus, by Theorem 5, G_1 is planar, isomorphic to K_5, or crown planar with respect to e. Moreover, since G_1 is 3-connected, $G_1 \setminus e$ is 2-connected. Thus, by Case (IIIa), G_1 has two internally node-disjoint admissible e-paths. Since these two paths are internally node disjoint, one of them, call it P, must be edge disjoint from $T - \{e\}$. Thus, by Lemma 10, P is an admissible e-path of G. Finding the appropriate internal 3-separation easily reduces to finding the cut nodes of $G \setminus \{x, y\}$ and so requires $O(n)$ time using Tarjan [9]. Once the appropriate internal 3-separation is identified, finding the admissible path requires $O(n)$ time by Case (IIIa). □

Theorem 11 can be strengthened as follows: if $G \setminus e$ is 2-connected (resp., 2-edge-connected), then there exist two internally node disjoint (resp., edge-disjoint) admissible e-paths. The proof of this requires a bit more work but follows the same line of reasoning.

Theorem 11 provides a sufficient condition for the existence of an admissible e-path. It is, however, not a necessary condition. For example, take the graph $K_{3,3}$, and let e be any edge. Now, add a new edge that creates a triangle T containing e. Then, $T - \{e\}$ is an e-path that intersects every e-cut in exactly one edge. It would be interesting to determine if one excludes triangles containing e, whether the condition is also necessary.

4. The Ford-Fulkerson Algorithm

Let G be a graph, and let x and y be distinguished nodes of G. Assume G has the edge $e = xy$, and consider an instance of the minimum e-cut problem (equivalently, the maximum-flow problem) defined on G. Ford and Fulkerson [1] provided a very simple algorithm for solving this problem provided that G is planar. This section shows that the Ford-Fulkerson algorithm extends virtually unchanged provided that G does not have a $K_{3,3}$ minor containing e; the running time of the algorithm is shown to be $O(n^2)$. This is within a logarithmic

```
Algorithm ford-fulkerson;
begin
    H := G;
    D := {e};
    while D does not contain an e-cut of H do
        begin
            set P to be an admissible e-path of H;
            ϵ := min{c_f | f ∈ P};
            c_f := c_f − ϵ for f ∈ P;
            choose f ∈ P such that c_f = 0 and set H ← H \ f;
            D ← D ∪ {f};
        end;
end;
```

Algorithm 1

factor of the fastest maximum-flow algorithm, namely the recent (and more complicated) algorithm due Orlin [2]. The only faster algorithm for graphs in this class is the $O(n)$-time algorithm in Wagner [3].

Consider the following property for the graph G.

Property A. For any subset X of edges with $e \in X$, if $G \setminus X$ has an e-path, then it has an admissible e-path.

Algorithm 1 is the Ford-Fulkerson algorithm, stated in terms of solving the minimum e-cut problem. An instance of the minimum e-cut problem is specified by (G, e, c), where c is a vector of nonnegative edge capacities.

The next theorem shows that the algorithm is correct provided Property A holds. In particular, Theorem 11, then implies that the algorithm works for any instance (G, e, c) provided G does not have a $K_{3,3}$ minor containing e.

Theorem 12. *Assume that G satisfies Property A. Then, Algorithm 1 correctly computes a minimum e-cut of* (G, e, u). *In particular, at termination, the set D contains a unique e-cut of G, and it is a minimum e-cut of G.*

Proof. Each execution of the *while* loop chooses an admissible e-path. Since, by assumption, G satisfies Property A, this step is well defined. Each execution of the *while* loop adds exactly one edge to D, and therefore eventually, D will contain an e-cut of G, and so the algorithm will terminate.

Assume that the *while* loop executes t times, and index the instantiations of H, D, P, f, c, and ϵ by $1, \ldots, t$. So, H_t denotes the instantiation of H at termination of the algorithm.

It is first shown that D_t contains a unique e-cut of G. This is equivalent to showing that H_t has exactly two components. Suppose this is not the case; that is, suppose H_t has at least three components. Since H_t was obtained from H_{t-1} by deleting exactly one edge, H_{t-1} must have at least two components. Let j denote the least index such that H_j has at least two components. Observe that, in H_j, the ends of e are in the same component, for otherwise the algorithm would have terminated at $j < t$, contradicting the definition of t. By definition, H_{j-1} is connected. Since H_j is obtained from H_{j-1} by deleting the edge f_{j-1}, this edge must be a cut edge of H_{j-1}. From the algorithm, f_{j-1} is contained in the e-path P_{j-1}. But this is impossible since no path, starting and ending in the same component of a graph, can contain a cut edge of the graph. This D_t contains a unique e-cut of G; denote this e-cut by D^*.

Now, consider (H_2, e, c^2), that is, the minimum e-cut problem that results after one execution of the *while* loop. The graph H_2 has one less edge than H_1. Applying Algorithm 1 to (H_2, e, c^2) will produce the set $D_t - \{f_1\}$, which, by induction, will contain a unique e-cut of H_2 that is also a minimum e-cut of (H_2, e, c^2); denote this e-cut by D^{**}. Observe that (G, e, u) is obtained from (H_2, e, u^2) first adding the edge f_1 with a capacity of zero and then adding ϵ_1 to the capacity of every edge of P_1. Consider these steps individually. Adding an edge of capacity zero effectively leaves the minimum e-cut problem unchanged. In particular, either $D^{**} \cup \{f_1\}$ or D^{**} is minimum e-cut for the resulting minimum e-cut problem, depending on whether the ends of f_1 are in different components of $H_2 \setminus D^{**}$ or not. In either case, the capacity of the resultant minimum e-cut is unchanged. The second step also effectively does not change the minimum e-cut problem. In particular, by the algorithm, P_1 is an admissible e-path. Therefore, adding ϵ_1 to each edge of P_1 adds ϵ_1 to the capacity of each e-cut. Consequently, either $D^{**} \cup \{f_1\}$ or D^{**} (again, depending on whether the ends of f_1 are in different components of $H_2 \setminus D^{**}$ or not) is a minimum e-cut of (G, e, u). In either case, this minimum e-cut of (G, e, u) is contained in D_t and, thus, by the uniqueness demonstrated in the previous paragraph, is equal to D^*, as required. □

Corollary. *If G does not have a* $K_{3,3}$ *minor containing e, then the complexity of Algorithm 1 is* $O(n^2)$.

Proof. By Theorem 11, finding the first admissible e-path requires $O(m)$ time. Also, note that the initial step in this first admissible-path computation reduces the graph to a simple 2-connected graph, and so by Lemma 6, the resulting number of edges is $O(n)$. Thus, each subsequent admissible-path computation requires $O(n)$ time. Since the algorithm deletes one edge every iteration, it requires at most $O(n)$ iterations. Thus, the algorithm requires $O(n^2)$ time. □

As a final note, Algorithm 1 can be dualized, using planar path-cut duality, to an algorithm for finding a shortest e-path in a graph. This dual version of Algorithm 1 can, in fact, be shown to be equivalent to Dijkstra's shortest-path algorithm.

References

[1] L. R. Ford and D. R. Fulkerson, "Maximal flow through a network," *Canadian Journal of Mathematics*, vol. 8, pp. 399–404, 1956.

[2] J. B. Orlin, "Max flows in $O(nm)$ time, or better," Working paper, 2012.

[3] D. K. Wagner, "$K_{3,3}$ minors and the maximum-flow problem," *Algorithmic Operations Research*, vol. 3, no. 1, pp. 30–42, 2008.

[4] W. T. Tutte, *Graph Theory*, Addison-Wesley, Reading, Mass, USA, 1984.

[5] K. Wagner, "Über eine Erweiterung eines Satzes von Kuratowski," *Deutsche Mathematik*, vol. 2, pp. 280–285, 1937.

[6] J. Širáň, "Edges and Kuratowski subgraphs of nonplanar graphs," *Mathematische Nachrichten*, vol. 113, pp. 187–190, 1983.

[7] F. T. Tseng and K. Truemper, "A decomposition of the matroids with the max-flow min-cut property," *Discrete Applied Mathematics*, vol. 15, no. 2-3, pp. 329–364, 1986.

[8] W. T. Tutte, "A theory of 3-connected graphs," *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen A*, vol. 64, pp. 441–455, 1961.

[9] R. E. Tarjan, "Depth-first search and linear graph algorithms," *SIAM Journal on Computing*, vol. 1, no. 2, pp. 146–160, 1972.

[10] J. E. Hopcroft and R. E. Tarjan, "Dividing a graph into triconnected components," *SIAM Journal on Computing*, vol. 2, pp. 135–158, 1973.

[11] J. E. Hopcroft and R. E. Tarjan, "Efficient planarity testing," *Journal of the Association for Computing Machinery*, vol. 21, pp. 549–568, 1974.

[12] K. Mehlhorn and P. Mutzel, "On the embedding phase of the Hopcroft and Tarjan planarity testing algorithm," *Algorithmica*, vol. 16, no. 2, pp. 233–242, 1996.