

Flows in One-Crossing-Minor-Free Graphs

*Erin Wolf Chambers*¹ *David Eppstein*²

¹Dept. of Mathematics and Computer Science, Saint Louis University

²Computer Science Department, University of California, Irvine

Abstract

We study the maximum flow problem in directed H -minor-free graphs where H can be drawn in the plane with one crossing. If a structural decomposition of the graph as a clique-sum of planar graphs and graphs of constant complexity is given, we show that a maximum flow can be computed in $O(n \log n)$ time. In particular, maximum flows in directed $K_{3,3}$ -minor-free graphs and directed K_5 -minor-free graphs can be computed in $O(n \log n)$ time without additional assumptions.

Submitted: August 2012	Reviewed: January 2013	Revised: February 2013	Accepted: February 2013	Final: February 2013
Published: March 2013				
Article type: Regular paper		Communicated by: S. Albers		

This material was presented in a preliminary form at ISAAC 2010. The work of David Eppstein was supported in part by NSF grants 0830403 and 1217322 and by the Office of Naval Research under grant N00014-08-1-1015.

E-mail addresses: echambe5@slu.edu (Erin Wolf Chambers) eppstein@ics.uci.edu (David Eppstein)

1 Introduction

Computing maximum flows is fundamental in algorithmic graph theory, and has many applications. Although flows can be computed in polynomial time for arbitrary graphs, it is of interest to find classes of graphs for which flows can be computed more quickly. Classes of note include planar graphs [2, 13, 19, 20, 23, 25, 27, 28, 31, 42], where the fastest current algorithms take $O(n \log \log n)$ time for undirected graphs [24] and $O(n \log n)$ for directed graphs [2]; graphs of bounded genus [4, 5], where the best running time is $O(n \log^2 n \log^2 C)$ for undirected graphs with integer capacities that sum to C ; graphs with small crossing number [21], where the best running time is $O(k^3 n \log n)$ for k crossings; and graphs of bounded treewidth [17], where the best running time is $O(n)$.

Planar graphs, graphs of bounded genus, and graphs of bounded treewidth are all examples of *minor-closed graph families*, families of graphs closed under edge contractions and edge deletions. According to the Robertson–Seymour graph minor theorem [38], any minor-closed graph family can be described as the X -minor-free graphs, graphs that do not have as a minor any member of a finite set X of non-members of the family; for instance, the planar graphs are exactly the $\{K_5, K_{3,3}\}$ -minor-free graphs [40]. In many cases the properties of a graph family are closely related to the properties of its excluded minors: for instance, the minor-closed graph families with bounded treewidth are exactly the families of X -minor-free graphs for which X includes at least one planar graph [35], and the families with bounded local treewidth (a functional relationship between the diameter of a graph and its treewidth) are exactly those for which X includes at least one *apex graph*, a graph that can be made planar by removing a single vertex [11]. If X includes a graph that can be drawn in the plane with a single pair of crossing edges, then the X -minor-free graphs have a structural decomposition as a *clique-sum* of smaller graphs that are either planar or have bounded treewidth [8, 36]. In this last case we say that the family of X -minor-free graphs is *one-crossing-minor-free*; families of this type include the $K_{3,3}$ -minor-free graphs and K_5 -minor-free graphs, since $K_{3,3}$ and K_5 are one-crossing graphs (Figure 1).

In this paper we consider flows in one-crossing-minor-free graph families. We provide $O(n \log n)$ algorithms to compute maximum flows in any directed H -minor-free graph, where H is a fixed one-crossing graph and where the structural decomposition of the graph is provided as part of the input. In the case of $K_{3,3}$ -minor-free graphs and K_5 -minor-free graphs, algorithms are known that can find such a decomposition efficiently [1, 33], and by combining our techniques with those known decomposition algorithms we provide an $O(n \log n)$ time algorithm for maximum flow in directed $K_{3,3}$ -minor-free and K_5 -minor-free graphs without requiring the decomposition to be part of the input.

Our main motivation for looking at flows in one-crossing-minor-free graphs is to try to make progress towards finding flows in arbitrary minor-closed graph families. Currently the best known algorithms for such families are those which are for arbitrary sparse graphs [16, 32, 10, 15], so this paper represents a first attempt at using known structural properties of minor-closed families to improve

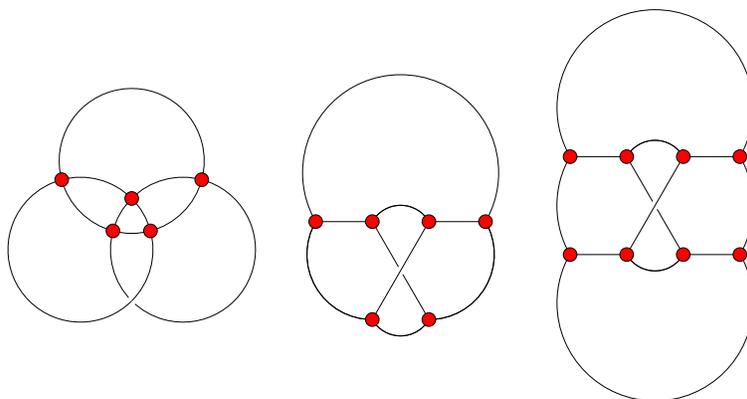


Figure 1: One-crossing drawings of K_5 (left), $K_{3,3}$ (center), and the Wagner graph (right).

the running time over these bounds. We note that other problems with fast planar algorithms have also generalized to minor-closed graph families, such as computing good separators [34], minimum spanning trees [30], and shortest paths [39]

Like one-crossing-minor-free graphs, graphs in more general minor-closed families have a structural decomposition in terms of bounded-genus surfaces, clique-sums, apexes (a constant number of vertices that can be adjacent to arbitrary subsets of each of the bounded-genus surfaces), and vortices (bounded-treewidth graphs glued into faces of the bounded-genus surfaces) [7, 37]. However, these decompositions are greatly simplified in the one-crossing-minor-free case: the surfaces are planes and there are no apexes or vortices [36]. To handle the general case, we would need to combine clique-sums, bounded-genus, apexes, and vortices. The problem of flows on bounded genus surfaces has been previously examined [4, 5] and the present work focuses on clique-sums, as these are the main feature in the structural decomposition for one-crossing-minor-free graphs. Recent work details a near linear time algorithm for multi-source multi-sink flows in planar graphs [3], which is equivalent to a 2-apex planar graph since one can simply add a source apex and connect to all sources and a sink apex and connect to all sinks. However, it remains unclear how to handle apexes in arbitrary bounded genus graphs and vortices in any setting.

As an important tool in our results, we greatly simplify the *mimicking networks* of Hagerup et al [17] for multiterminal flow networks in the case of four terminals with a single source, leading to significantly reduced constant factors in the running time of algorithms that use these networks. Similar simplifications had been achieved in the undirected case [6] but to our knowledge our small directed mimicking network is novel.

2 Preliminaries

2.1 Flows and Cuts

A *flow network* is a directed graph $G = (V, E)$, where each (directed) edge $e \in E$ has an associated nonnegative real capacity c_e , along with two distinguished vertices s and t which are called the *source* and *sink*, respectively. A *flow* in this graph is a set of nonnegative real values f_e for each edge $e \in E$ such that $f_e \leq c_e$ for every $e \in E$ and $\sum_{uv \in E} f_{uv} = \sum_{vu \in E} f_{vu}$ for every $v \in V - \{s, t\}$. The *value* of the flow is the amount of net flow going from s to t , or $\sum_{sv \in E} f_{sv}$.

A *cut* in a flow network is a set of edges separating s from t ; the *capacity* of the cut is the sum of the capacities of all the edges that cross the cut in the direction from s to t . Clearly, the value of any flow is less than or equal to the value of any cut. The classic max-flow min-cut theorem states that the maximum possible flow from s to t is in fact equal to the minimum possible cut.

In our setting, we will need to know the maximum possible flow that can travel through a subgraph of the input graph; this will allow us to simplify the graph by removing the subgraph and replacing it with an equivalent (but much smaller) subgraph. To do this, we compute an *external flow*. In external flow networks, instead of a single source and sink, we have a set of terminals $Q = \{q_1, \dots, q_k\}$ where each terminal q_i has associated with it a number x_i ; we require that $\sum_i x_i = 0$. If x_i is positive then it is interpreted as a supply, and if x_i is negative it is interpreted as a demand. (It may also be the case that x_i is zero, in which case x_i carries neither a supply nor a demand.) A *realizable external flow* is a set of k values (x_1, \dots, x_k) for (q_1, \dots, q_k) , along with a flow f such that $\sum_{(q_i, v) \in E} f_{(q_i, v)} - \sum_{(v, q_i) \in E} f_{(v, q_i)} = x_i$ for all i . Basically, the flow remains balanced at every vertex in $V \setminus Q$, and the imbalance of flow at each vertex in $q_i \in Q$ is exactly x_i .

It will be helpful to define a special case of external flow networks, which we call *single-source external flow networks*. A single-source external flow network is, simply, an external flow network for which only q_1 may have a positive supply x_i ; every other terminal has $x_i \leq 0$ indicating that it is either a demand node or is inactive as a terminal. As before, a *realizable single-source external flow* is a realizable external flow subject to this constraint on the values of x_i .

We define $S \not\rightarrow T$, for sets of terminals S and T in an external flow network, to be the minimum value of a cut for which every terminal in S is on the source side of the cut and every terminal in T is on the sink side of the cut. We will further abbreviate the notation by writing strings of symbols instead of bracketed set notation for the sets of terminals on each side of the cut; e.g., $s \not\rightarrow abc$ should be interpreted as an abbreviation for $\{s\} \not\rightarrow \{a, b, c\}$.

2.2 Mimicking networks

Let G be an external flow network or single-source external flow network with a fixed specification of the edge capacities and a fixed ordered set of terminals Q , but where the supply and demand quantities x_i remain variable. A key ingre-

dient for our technique, as formalized by Hagerup et al. [17], is the concept of a *mimicking network*, a network H that has the same terminals as G and has the same realizable external flows, but may have many fewer edges and vertices than G . In what amounts to a form of separator based sparsification [12], Hagerup et al. solve flow problems on bounded-treewidth networks by repeatedly replacing subnetworks of the given network with smaller mimicking networks. Their construction of mimicking networks is based on an observation of Gale [14]:

Lemma 1 [17] *An external flow (x_1, \dots, x_k) is realizable in a network $G = (V, E)$ with terminals $Q = \{q_1, \dots, q_k\}$ if and only if the following relations are satisfied:*

$$\sum_{i=1}^k x_i = 0 \tag{1}$$

$$\sum_{q_i \in S} x_i \leq (S \not\rightarrow (Q \setminus S)), \text{ for all } S \subseteq Q \text{ with } \emptyset \neq S \neq Q. \tag{2}$$

Essentially, this means that in order to understand the possible flow patterns in a subnetwork with k terminals, one needs only to know the $(2^k - 2)$ minimum cut values from a nonempty subset of terminals to its nonempty complement. If two networks have the same minimum cut values for each subset then they behave the same with respect to flows. Based on this observation, Hagerup et al. show that there exists a replacement network of at most $2^{2^k - 2}$ vertices that behaves equivalently to any k -terminal network:

Lemma 2 [17] *Given any external flow network G with k terminals, there exists a flow network having at most $2^{2^k - 2}$ vertices which has the same external flow value as G .*

Specifically, the mimicking network of Lemma 2 can be constructed from G by finding a set of $2^k - 2$ minimum cuts, one for each partition of the terminals into two nonempty subsets, and by collapsing subsets of vertices in G into a single supervertex whenever all vertices in the subset are on the same side of every cut. Note that the mimicking network is not necessarily a minor of G , as the collapsed subsets need not form connected subgraphs of G .

However, the size of the mimicking networks formed by Lemma 2, while constant, is large. Our algorithms will involve external flow networks with up to four terminals, and if we applied Lemma 2 directly we might get as many as 16384 vertices in our mimicking networks. It is possible to reduce the number of vertices in the construction of Hagerup et al. from a power of two to a Dedekind number by requiring nested partitions of the terminals to have nested cuts, but this would still lead to 168 vertices for the mimicking network of a four-terminal network. We will describe in Section 3 a simpler mimicking network of Chaudhuri et al. [6] for external flows with at most three terminals, and a new mimicking network for single source external flows with at most four terminals that are both much smaller, requiring at most one nonterminal vertex.

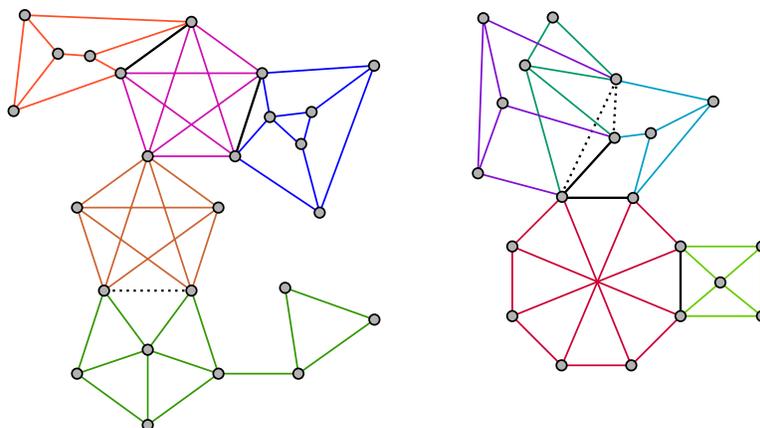


Figure 2: A $K_{3,3}$ -free graph expressed as a 2-sum of planar graphs and copies of K_5 (left), and a K_5 -free graph expressed as a 3-sum of planar graphs and copies of the Wagner graph (right). The edges forming the cliques of the clique-sum operations are shown as solid black (if they remain in the final graph) or dashed (if they were removed after the clique-sum operation); the colors of the remaining edges identify the subgraphs entering into the clique-sum operations.

2.3 Structure of minor free graphs

A *minor* of a graph G is a graph that can be formed from G by contracting and removing edges. A graph family F is *minor-closed* if every minor of a graph in F also belongs to F . If X is a finite set of graphs, the X -minor-free graphs are the graphs G such that no minor of G belongs to X ; the X -minor-free graphs are obviously a minor-closed graph family, and (much less obviously) the Robertson–Seymour graph minor theorem [38] states that every minor-closed graph family has this form. If F is the family of X -minor-free graphs, then X is the set of *forbidden minors* for F ; for instance, K_5 and $K_{3,3}$ are the forbidden minors for the planar graphs. We will abuse notation and abbreviate the $\{H\}$ -minor-free graphs (where H is a single graph) as the H -minor-free graphs.

A *clique-sum* is an operation that combines two graphs by identifying the vertices in two equal-sized cliques in the two graphs, and then possibly removing some of the edges of the cliques. A k -sum is a clique-sum where all cliques have at most k vertices. More generally, we will say that a given graph is a k -sum of a collection of more than two graphs if it can be formed by repeatedly replacing pairs of graphs of the collection by their k -sum. Clique-sums are closely related to vertex-connectivity of graphs: for instance, the decomposition of a biconnected graph into triconnected components, formalized in the SPQR-tree (described in Section 2.4), gives a representation of the graph as a 2-sum of triconnected graphs, cycles, and two-vertex multigraphs [9, 22, 29].

Besides their relation to connectivity, clique-sums have also played an important role in describing the structure of minor-closed graph families since the

proof by Wagner [40] that K_5 -minor-free graphs are exactly the graphs that can be formed by 3-sums of planar graphs and the 8-vertex nonplanar Wagner graph, the graph formed by adding four edges connecting opposite pairs of vertices in an 8-cycle (Figure 2, right). Wagner [41] and D. W. Hall [18] also proved that $K_{3,3}$ -minor-free graphs are exactly the 2-sums of planar graphs and the five-vertex complete graph K_5 (Figure 2, left). These two characterizations of H -minor-closed families can be generalized to any H -minor-free graph with the property that H can be drawn in the plane with only a single edge crossing, as $K_{3,3}$ and K_5 both can. In this case, the H -minor-free graphs can always be decomposed into 3-sums of planar graphs and graphs of bounded treewidth [36].

Algorithmically, a decomposition of a $K_{3,3}$ -minor-free graph into a 2-sum of planar graphs and K_5 can be computed in linear time [1]. Essentially, this is simply a matter of constructing the SPQR tree and verifying that all triconnected components are either planar or K_5 . As has been shown more recently, a decomposition of a K_5 -minor-free graph into a 3-sum of planar graphs and the Wagner graph can also be constructed in linear time [26, 33]. While decomposition of graphs in any H -minor-free family (where H is a one-crossing graph) into 3-sums of planar graphs and graphs of bounded treewidth have been directly considered [8], the current fastest algorithm to compute this decomposition is recent work by Grohe et. al which computes the graph minor decomposition for *any* fixed minor graph. However, the polynomial $O(n^2)$ is too high to lead to improved running times for flow algorithms, so we will assume that a decomposition has been given to us as part of the input. In case a future development leads to linear time algorithms for decomposition of one-crossing-minor-free graphs, this assumption may be removed.

2.4 SPQR trees

The SPQR tree of a graph is a formalized method of decomposing the graph into a 2-sum of simpler multigraphs, its *triconnected components* [9, 22, 29]. In order to have a self-contained exposition, we describe it briefly here. The SPQR tree is properly defined only for 2-connected graphs; to apply it to graphs that may not be 2-connected, they need to be first decomposed into 2-connected components and the SPQR tree construction applied separately to each such component.

A *cycle* is a 2-connected graph with n vertices and n edges; a *bond* is a multigraph with two vertices and three or more edges, each of which has the same two endpoints. Then an SPQR tree is a tree with the following properties:

- Each node t_i of the SPQR tree is associated with a multigraph G_i that is either a nontrivial 3-connected simple graph (an “R” node), a cycle (an “S” node), or a bond (a “P” node).
- Each tree edge $t_i t_j$ of the SPQR tree is associated with a multigraph edge e_{ij} in G_i and another multigraph edge e_{ji} in G_j . These two multigraph edges e_{ij} and e_{ji} are known as *virtual edges*.

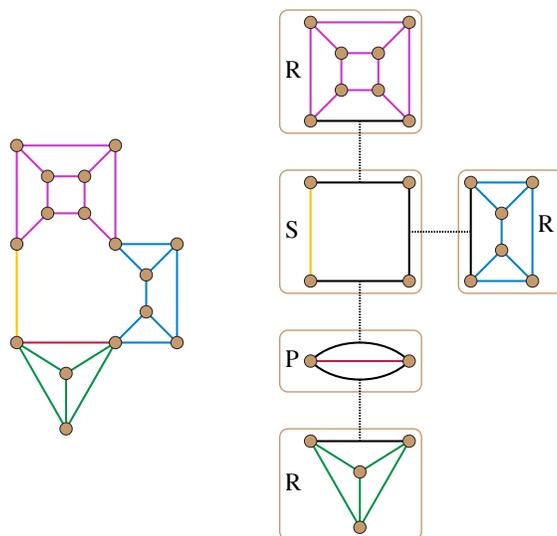


Figure 3: An graph (left) and its SPQR tree (right). Each tree node is shown with its corresponding triconnected component: a bond for the P node, a cycle for the S node, and triconnected graphs for the R nodes. The dashed black lines connect pairs of virtual edges (solid black) in these components; the original graph is formed by gluing together and then removing each pair of virtual edges.

- Each multigraph edge can be a virtual edge for at most one SPQR tree edge.
- Every SPQR tree edge connects two “R” nodes or two nodes of different types from each other.
- Each “P” node has at most one non-virtual edge.

A tree of this type represents a graph G that is formed by repeated 2-sum operations. Each 2-sum glues together the two designated virtual edges e_{ij} and e_{ji} for an SPQR tree edge and then deletes the glued-together edge. The order in which the 2-sums are performed does not affect the final result. Each edge that is not virtual, in the graph associated to each node, survives to become an edge in the overall graph. The graphs associated with the SPQR tree nodes are known as the *triconnected components* of G .

Conversely, any 2-connected graph with two trivial exceptions, the graphs K_1 and K_2 , has a unique representation as an SPQR tree; this representation can be constructed in linear time [9, 22]. In order to handle the two exceptions, we allow as a special case an SPQR tree with one “Q” node associated with one of these two exceptional graphs.

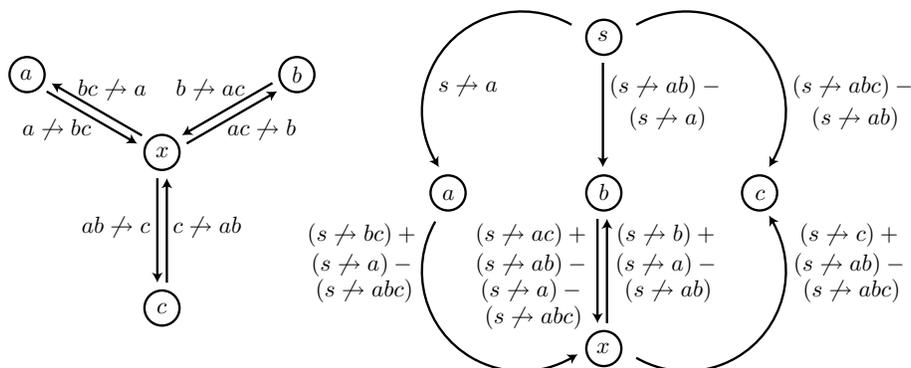


Figure 4: Mimicking networks for an external flow network with three terminals a , b , and c (left), and for a single-source external flow network with four terminals s , a , b , and c (right). In the single-source mimicking network, a , b , and c are permuted if necessary so that $s \not\rightarrow a \geq \max\{s \not\rightarrow b, s \not\rightarrow c\}$ and $s \not\rightarrow ab \geq s \not\rightarrow ac$.

3 Small mimicking networks

We show in this section that any three-terminal external flow network, and any four-terminal single-source external flow network, may be replaced by a mimicking network with at most one non-terminal vertex.

Lemma 3 *Let P , Q , and R be a partition of the terminals of any external flow network into three subsets. Then $(P \not\rightarrow Q \cup R) \leq (P \cup Q \not\rightarrow R) + (P \cup R \not\rightarrow Q)$ and $(P \cup Q \not\rightarrow R) \leq (P \not\rightarrow Q \cup R) + (Q \not\rightarrow P \cup R)$.*

Proof: If C_1 and C_2 are minimum cuts separating $P \cup Q$ from R and $P \cup R$ from Q , respectively, then $C_1 \cup C_2$ separates P from both Q and R , and has capacity at most $(P \cup Q \not\rightarrow R) + (P \cup R \not\rightarrow Q)$; therefore, the minimum cut value $P \not\rightarrow Q \cup R$ must be at most this capacity. The other inequality follows by a symmetric argument. \square

Lemma 4 (Chaudhuri et al. [6]) *Any external flow network with three terminals has a mimicking network with four vertices and six edges.*

Proof: Let the three terminals of flow network G be a , b , and c , and form a network that consists of these three terminals together with a fourth vertex d , in the form of a star $K_{1,3}$ with six edges connecting d in both directions to each terminal. For each terminal q , set the capacity of the edge from q to d to the minimum cut amount $q \not\rightarrow (\{a, b, c\} \setminus \{q\})$ as measured in G , and set the capacity of the edge from d to q to be $(\{a, b, c\} \setminus \{q\}) \not\rightarrow q$. It follows from Lemma 3 that each minimum cut separating one of the terminals from the other two consists of a single edge, and therefore that the capacity of this cut in the

mimicking network exactly matches the capacity of the corresponding cut in G . The correctness of the mimicking network then follows by Lemma 1. \square

It is tempting to try an even simpler network, a triangle with three vertices and six edges, in place of the four-vertex network of Lemma 4, but as Chaudhuri et al. observe this does not work: the triangle can only mimic networks for which $(a \not\rightarrow bc) + (b \not\rightarrow ac) + (c \not\rightarrow ab) = (ab \not\rightarrow c) + (ac \not\rightarrow b) + (bc \not\rightarrow a)$, and this equality is not necessarily true in other three-terminal external flow networks. However, in undirected flow networks, it is possible to form an undirected triangle network as a mimicking network, as shown in [6]: the capacity of the edge from a to b should be $\frac{1}{2}((a \not\rightarrow bc) + (b \not\rightarrow ac) - (c \not\rightarrow ab))$ and symmetrically for the other three edges.

We also need mimicking networks for four terminals, but only for single-source external flows. In this case, the following variant of Lemma 1 is helpful.

Lemma 5 *A single-source external flow (x_1, \dots, x_k) is realizable in a network $G = (V, E)$ with terminals $Q = \{q_1, \dots, q_k\}$ if and only if the following relations are satisfied:*

$$\sum_{i=1}^k x_i = 0 \quad (3)$$

$$\sum_{q_i \in S} -x_i \leq (x_1 \not\rightarrow S), \text{ for all nonempty } S \subseteq Q \setminus \{x_1\}. \quad (4)$$

Proof: Augment G by a sink vertex t , add an edge of capacity $-x_i$ from each terminal q_i ($i > 1$) to t , and set the demand on t to x_1 . The result follows by the classical max-flow min-cut theorem: the desired single-terminal external flow exists if and only if the augmented graph has a two-terminal flow meeting the given demands, if and only if it has no cut with capacity less than x_1 . But any cut in the augmented graph has a capacity equal to its capacity in the original graph plus the sum of the demands for the terminals on the source side of the cut, and each inequality in the statement of the lemma can be restated as requiring the capacities of some of these cuts to be at least x_1 . So a cut with capacity less than x_1 exists, if and only if at least one of these inequalities is violated. \square

Lemma 6 *Let $P, Q, R,$ and S be four disjoint sets of terminals in an external flow network G . Then $(P \not\rightarrow Q \cup R \cup S) + (P \cup Q \cup R \not\rightarrow S) \leq (P \cup Q \not\rightarrow R \cup S) + (P \cup R \not\rightarrow Q \cup S)$.*

Proof: Let C_1 and C_2 be minimum cuts separating $P \cup Q$ from $R \cup S$ and $P \cup R$ from $Q \cup S$, respectively. Let A be the set of vertices of G on the source side of C_1 , and let B be the set of vertices of G on the source side of C_2 . We consider three subsets of the edges in $C_1 \cup C_2$: the edges from $A \cap B$ to $G \setminus (A \cap B)$, the edges from $A \cup B$ to $G \setminus (A \cup B)$, the edges in both $A \cap B$ to $G \setminus (A \cap B)$ and in $A \cup B$ to $G \setminus (A \cup B)$, and any remaining edges that go in either direction between $A \setminus B$ and $B \setminus A$.

The first of these three subsets of edges forms a cut separating P from $Q \cup R \cup S$, and the second of these three subsets of edges forms a cut separating $P \cup Q \cup R$ from S . We also note that any edges in the third set (and so in both the first and the second above) are counted twice in $|C_1| + |C_2|$, so the total value of $|C_1| + |C_2|$ is at least equal to $(P \not\rightarrow Q \cup R \cup S) + (P \cup Q \cup R \not\rightarrow S)$. \square

Lemma 7 *Any single-source external flow network with four terminals has a mimicking network with five vertices and seven edges.*

Proof: Given a single-source external flow network G , with terminals s, a, b , and c , we may assume without loss of generality (by permuting the final three terminals, if necessary) that $s \not\rightarrow a$ is at least as large as $s \not\rightarrow b$ and $s \not\rightarrow c$, and that $s \not\rightarrow ab$ is at least as large as $s \not\rightarrow ac$. We create a mimicking network with a single nonterminal vertex x and with the following seven edges:

- An edge from s to a with capacity $s \not\rightarrow a$.
- An edge from s to b with capacity $(s \not\rightarrow ab) - (s \not\rightarrow a)$.
- An edge from s to c with capacity $(s \not\rightarrow abc) - (s \not\rightarrow ab)$.
- An edge from a to x with capacity $(s \not\rightarrow bc) + (s \not\rightarrow a) - (s \not\rightarrow abc)$.
- An edge from b to x with capacity $(s \not\rightarrow ac) + (s \not\rightarrow ab) - (s \not\rightarrow a) - (s \not\rightarrow abc)$.
- An edge from x to b with capacity $(s \not\rightarrow b) + (s \not\rightarrow a) - (s \not\rightarrow ab)$.
- An edge from x to c with capacity $(s \not\rightarrow c) + (s \not\rightarrow ab) - (s \not\rightarrow abc)$.

Lemmas 3 and 6 can be used to show that all edge capacities in this network are non-negative. As we now verify, each of the seven minimum cuts from s to any subset of terminals in this network has the same capacity as the corresponding minimum cut in G :

- Every cut separating s from a must cut edge sa , so the cut that separates only that edge is the minimum cut, and has capacity $s \not\rightarrow a$.
- Every cut separating s from $\{a, b\}$ must cut edges sa and sb , so the cut that separates only those two edges is the minimum cut, and has capacity $s \not\rightarrow ab$.
- Every cut separating s from $\{a, b, c\}$ must cut edges sa, sb , and sc , so the cut that separates only those three edges is the minimum cut, and has capacity $s \not\rightarrow abc$.
- There are three minimal cuts from s to $\{b, c\}$: the cut $\{sa, sb, sc\}$, the cut $\{ax, sb, sc\}$, and the cut $\{xb, xc, sb, sc\}$. The first of these has capacity $(s \not\rightarrow abc) \geq (s \not\rightarrow bc)$, the second has capacity $s \not\rightarrow bc$, and the third is at least as large as the second by Lemma 3, so the minimum cut has capacity $s \not\rightarrow bc$.

- The minimal cuts from s to $\{a, c\}$ are the cut $\{sa, sb, sc\}$ with capacity $s \not\rightarrow abc$, the cut $\{sa, sc, bx\}$ with capacity $s \not\rightarrow ac$, and the cut $\{sa, sc, xc\}$ which is at least as large as the previous cut by Lemma 3. Therefore the minimum cut has capacity $s \not\rightarrow ac$.
- There are three minimal cuts from s to b : the cut $\{sa, sb\}$, the cut $\{ax, sb\}$, and the cut $\{xb, sb\}$. The first of these has capacity $(s \not\rightarrow ab) \geq (s \not\rightarrow b)$, the second has capacity at least equal to $s \not\rightarrow b$ by Lemma 6, and the third has capacity $s \not\rightarrow b$, so $s \not\rightarrow b$ is the minimum cut value.
- Among the cuts from s to c , the ones that cut off another terminal from s as well as c fall into one of the previous cases, and therefore have capacity at least $s \not\rightarrow c$. There are two remaining minimal cases: the cut $\{sc, ax, bx\}$ with capacity $(s \not\rightarrow bc) - (s \not\rightarrow abc) + (s \not\rightarrow ac)$, at least as large as $s \not\rightarrow c$ by Lemma 6, and the cut $\{sc, xc\}$ with capacity $s \not\rightarrow c$. Therefore, the minimum cut value is $s \not\rightarrow c$.

Since all of these cuts have the same value as the corresponding cuts in G , it follows by Lemma 5 that the network constructed as above is a valid mimicking network for G . \square

The mimicking networks described in Lemmas 4 and 7 are shown in Figure 4. We note that it is important for our techniques that the mimicking network for a three-terminal network is planar and has its three terminals in a single face of its planar embedding (more strongly, in fact, it is outerplanar); however, we do not rely on the planarity of the four-terminal mimicking network.

4 Flow algorithm

Recall first that any H -minor free graph G (where H is a one-crossing graph) can be decomposed into 3-sums of planar and bounded treewidth graphs; we refer to these smaller graphs as *components* of the clique-sum decomposition. In order to handle the case that three or more components are glued together by a 3-sum at a single shared clique, we may represent the decomposition as a 2-colored tree, in which the vertices of one color represent the components of the decomposition and the vertices of the other color represent the cliques at which they are glued together. We may identify two distinguished components of the decomposition, one containing s and another containing t , where s and t are the two terminals of our given flow problem; if s or t is part of a clique on which multiple components are glued, then the distinguished component containing that terminal may be chosen arbitrarily among all components containing it. The two distinguished components are necessarily connected by a path in the clique-sum decomposition tree (possibly a path of length zero).

The vertices of G that belong to the cliques of the clique-sum decomposition, in effect, will form terminals in each component, so any flows into and out of components of the decomposition can be treated as an external flow

computations. Our algorithm will iteratively replace each component in the decomposition with a mimicking network of constant size. However care is needed in ordering these computations, as some components may have more than a constant number of vertices that belong to cliques of the decomposition and we can only use mimicking networks for components that have a constant number of terminals. We perform this replacement in two phases: in the first phase, the components that are replaced do not belong to the path from s to t of the decomposition tree, and in the second phase the replaced components lie along this path.

Refining the decomposition. For technical reasons it is necessary to ensure that the planar components in our decomposition are only glued together by 3-sums along their faces, which we achieve by refining the given clique-sum decomposition tree. To do so, we make the following changes, each of which replaces one of the components in the decomposition by a clique-sum of smaller components:

- If any component of the decomposition is not biconnected, replace it with the 1-sum of its biconnected components, glued together by 1-sums at its articulation vertices.
- If any component of the decomposition is biconnected but not triconnected, find an SPQR tree representing it as a 2-sum of its triconnected components [9, 22, 29], as outlined in Section 2.4.
- At this stage of refinement, each component is triconnected, and in particular each planar component has a unique planar embedding. Within each planar component, find all the triangles at which 3-sums are glued, and compare them against the list of triangular faces of the graph. The gluing triangles that are not faces are separating triangles of the component, and may be used to partition each planar component into a 3-sum of smaller planar components.

All these refinements may be performed in linear time, and after this stage each gluing triangle in each planar component of the decomposition forms a face of the component.

Simplification phase I: components off the terminal path. Next, we deal with all components that are *not* on the s to t path in the clique-sum decomposition tree. Our algorithm iteratively finds a component C_i that is a leaf in the clique-sum decomposition tree; let C_j be the component to which C_i is glued by a clique-sum, one step closer to the s - t path. Then C_i is connected to C_j and the rest of the graph via at most three vertices, so we can treat the flow calculation within that component as an external flow problem with at most three terminals. Since the subgraph is either planar or of bounded treewidth, we can in $O(n \log n)$ time compute the min-cut values of all 6 possible partitions of its terminals, and by Lemma 4, we can replace the component C_i

with a mimicking network C'_i of at most four vertices. We then split into cases according to how C_i and C_j fit into the rest of the clique-sum decomposition tree:

- If C_j is a bounded-treewidth component, then C'_i can be glued directly into C_j by performing the clique-sum that connects these two components, forming a decomposition tree with one fewer component. The treewidth of the merged component is the maximum of 3 and the treewidth of C_j (using the fact that the mimicking network for C_i is at worst an oriented form of K_4), so after any number of repetitions of steps of this type the treewidth remains bounded.
- If C_j is a planar component, and C_i connects to C_j via one or two terminals or via a triangle of three terminals that is not used for any other 3-sum in the decomposition tree, then again C'_i can be immediately glued into C_j forming a decomposition tree with one fewer component. Due to the planar structure of our three-terminal mimicking networks and due to the fact that the gluing triangle is a face of C_j , this step preserves both the planarity of C_j and the property that all gluing triangles in the new larger planar component are still faces of their component.
- If C_j is a planar component and the gluing triangle of C_i is shared with another four-vertex replacement network C'_h , then C'_h and C'_i may be merged into a single mimicking network. Then, if this merged component is the only component that uses that gluing triangle, it may be glued into C_j as in the previous case.
- In the remaining cases, we replace C_i by C'_i in the decomposition tree but do not merge it with C_j .

After all simplifications of this type have been performed, the initial decomposition tree will have all leaves iteratively removed until none remain; therefore, it must have the structure of a path of components connecting s to t , where each of the gluing triangles in this path may also be connected to a single mimicking network off the path.

Simplification phase II: components along the terminal path. To complete the algorithm we will perform a similar sequence of replacements along the path from s to t . However, we must use a slightly different technique, since we may now need to deal with four terminals at a time: s and the three vertices in the 3-sum connecting the first two components in the s - t path. Additionally, after some of the replacements in this final stage of the algorithm, we will have to compute flows in networks that are neither planar nor of bounded treewidth, but that can be made planar by the removal of a constant number of edges; a result of Hochstein and Weihe [21] allows us to compute minimum cuts in these graphs in near-linear time. In this stage of the algorithm, as long as s and t belong to different components, we perform the following steps:

- Let C_i be the component containing s and T_i be the set of vertices in the clique-sum connecting C_i to the next component in the s – t path.
- Compute the minimum cut amounts between s and each nonempty subset of T_i , using the algorithm of Hochstein and Weihe [21] if C_i has been formed by adding a constant number of edges to a planar component, and the algorithm of Hagerup et al. [17] if C_i is a bounded-treewidth component.
- Replace C_i by a mimicking network C'_i for single-source external flows (as described in Lemma 7) using the computed cut amounts.
- Glue C'_i , and (if it exists) the other mimicking network sharing the same gluing triangle, into the next component in the s – t path, forming a path with one fewer component. If the next component was planar, it becomes a graph formed from a planar graph by adding a constant number of edges, while if it had bounded treewidth, its new treewidth is again bounded by the maximum of its old treewidth and 3.

Eventually this simplification will leave us with a planar or bounded treewidth graph containing both s and t , and we can compute the maximum flow between them directly in near linear time.

Reversing the simplifications and constructing a flow. We then reverse the sequence of simplifications we have performed, replacing each mimicking network with the larger network it replaced; for each such replacement we perform a single flow computation to find valid flow amounts forming the same external flow as in the mimicking network. At the end of this reversal of the simplification process, we will have a correct maximum flow in the original network that we were given as input.

We summarize our results as a theorem.

Theorem 1 *If we are given as input a decomposition of a directed flow network into a 3-clique-sum of planar and bounded-treewidth components, then in $O(n \log n)$ time we may compute a maximum flow between any two terminals s and t of the network.*

Corollary 2 *For any fixed one-crossing graph H , maximum flows in directed H -minor-free flow networks may be computed in $O(n \log n)$ time once a clique-sum decomposition of the network has been found.*

In the case of K_5 -free and $K_{3,3}$ -free graphs, we can find a clique-sum decomposition in $O(n)$ time [1, 33]. The case of $K_{3,3}$ -free graphs is particularly simple, since it uses only 2-sums, and therefore involves simpler mimicking networks.

Corollary 3 *We may compute maximum flows between any two terminals in a K_5 -free or $K_{3,3}$ -free directed flow network in $O(n \log n)$ time.*

5 Conclusions

We have shown how to find flow in near linear time when the input graph is a clique-sum of planar and bounded tree-width graphs by using mimicking networks to iteratively simplify the graph. This technique allows us to use known near linear algorithms in each bounded tree-width or planar component of the decomposition.

There is no added generality in considering 4-sums of planar graphs in place of 3-sums (any 4-sum involving a planar graph can be rearranged into a combination of 3-sums), but our methods immediately generalize to 2-sums of bounded-genus graphs and bounded-treewidth graphs. Flow computation in 3-sums of bounded genus graphs is more problematic due to the possible existence of non-facial non-separating triangles.

The larger goal, however, is computing flow quickly in more arbitrary minor-free families of graphs. Since flow can be computed efficiently in bounded genus and bounded tree-width graphs, the primary remaining open questions are those of computing flow in graphs with vortices or apexes, since these are the relevant building blocks for more general minor free families.

References

- [1] T. Asano. An approach to the subgraph homeomorphism problem. *Theoretical Computer Science*, 38(2–3):249–267, 1985. doi:10.1016/0304-3975(85)90222-1.
- [2] G. Borradaile and P. Klein. An $O(n \log n)$ algorithm for maximum st -flow in a directed planar graph. *Journal of the ACM*, 56(2):Art. No. 9, 2009. doi:10.1145/1502793.1502798.
- [3] G. Borradaile, P. N. Klein, S. Mozes, Y. Nussbaum, and C. Wulff-Nilsen. Multiple-source multiple-sink maximum flow in directed planar graphs in near-linear time. In *Proc. 52nd IEEE Symp. on Foundations of Comp. Sci. (FOCS 2011)*, pages 170–179, 2011.
- [4] E. W. Chambers, J. Erickson, and A. Nayyeri. Homology flows, cohomology cuts. In *Proc. 41st ACM Symposium on Theory of Computing (STOC 2009)*, pages 273–282, 2009. doi:10.1145/1536414.1536453.
- [5] E. W. Chambers, J. Erickson, and A. Nayyeri. Minimum cuts and shortest homologous cycles. In *Proc. 25th ACM Symposium on Computational Geometry (SoCG 2009)*, pages 377–385, 2009. doi:10.1145/1542362.1542426.
- [6] S. Chaudhuri, K. V. Subrahmanyam, F. Wagner, and C. D. Zaroliagis. Computing mimicking networks. *Algorithmica*, 26(1):31–49, 2000. doi:10.1007/s004539910003.
- [7] E. D. Demaine, M. Hajiaghayi, and K. ichi Kawarabayashi. Algorithmic Graph Minor Theory: Decomposition, Approximation, and Coloring. In *Proc. 46th IEEE Symp. on Foundations of Computer Science (FOCS 2005)*, pages 637–646, 2005. URL: http://erikdemaine.org/papers/Decomposition_FOCS2005/.
- [8] E. D. Demaine, M. Hajiaghayi, and D. M. Thilikos. 1.5-Approximation for treewidth of graphs excluding a graph with one crossing as a minor. In *Proc. 5th International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX 2002)*, volume 2462 of *Lecture Notes in Computer Science*, pages 67–80. Springer-Verlag, 2002. doi:10.1007/3-540-45753-4_8.
- [9] G. Di Battista and R. Tamassia. Incremental planarity testing. In *Proc. 30th Annual Symposium on Foundations of Computer Science (FOCS 1989)*, pages 436–441, 1989. doi:10.1109/SFCS.1989.63515.
- [10] S. I. Diatch and D. A. Spielman. Faster lossy generalized flow via interior point algorithms. In *Proc. 40th ACM Symp. Theory Comput.*, pages 451–460, 2008.

- [11] D. Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27:275–291, 2000. [arXiv:math.CO/9907126](#), [doi:10.1007/s004530010020](#).
- [12] D. Eppstein, Z. Galil, G. F. Italiano, and T. H. Spencer. Separator based sparsification for dynamic planar graph algorithms. In *Proc. 25th ACM Symp. on Theory of Computing (STOC 1993)*, pages 208–217, 1993. [doi:10.1145/167088.167159](#).
- [13] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8(399–404), 1956.
- [14] D. Gale. A theorem on flows in networks. *Pacific Journal of Mathematics*, 7:1073–1082, 1957.
- [15] A. V. Goldberg and S. Rao. Beyond the flow decomposition barrier. *J. ACM*, 45(5):783–797, 1998.
- [16] A. V. Goldberg and R. E. Tarjan. A new approach to the maximum-flow problem. *J. Assoc. Comput. Mach.*, 35(4):921–940, 1988. [doi:10.1145/48014.61051](#).
- [17] T. Hagerup, J. Katajainen, N. Nishimura, and P. Ragde. Characterizing multiterminal flow networks and computing flows in networks of small treewidth. *Journal of Computer and System Sciences*, 57(3):366–375, 1998. [doi:10.1006/jcss.1998.1592](#).
- [18] D. W. Hall. A note on primitive skew curves. *Bulletin of the American Mathematical Society*, 49(12):935–936, 1943. [doi:10.1090/S0002-9904-1943-08065-2](#).
- [19] R. Hassin. Maximum flow in (s, t) planar networks. *Information Processing Letters*, 13(3):107, 1981. URL: <http://www.math.tau.ac.il/~hassin/planarst.pdf>.
- [20] R. Hassin and D. B. Johnson. An $O(n \log^2 n)$ algorithm for maximum flow in undirected planar networks. *SIAM J. Comput.*, 14(3):612–624, 1985. [doi:10.1137/0214045](#).
- [21] J. M. Hochstein and K. Weihe. Maximum s - t -flow with k crossings in $O(k^3 n \log n)$ time. In *Proc. 18th ACM–SIAM Symposium on Discrete Algorithms (SODA 2007)*, pages 843–847, 2007. URL: <http://portal.acm.org/citation.cfm?id=1283383.1283473>.
- [22] J. Hopcroft and R. Tarjan. Dividing a graph into triconnected components. *SIAM J. Comput.*, 2(3):135–158, 1973. [doi:10.1137/0202012](#).
- [23] A. Itai and Y. Shiloach. Maximum flow in planar networks. *SIAM J. Comput.*, 8(2):135–150, 1979. [doi:10.1137/0208012](#).

- [24] G. F. Italiano, Y. Nussbaum, P. Sankowski, and C. Wulff-Nilsen. Improved algorithms for min cut and max flow in undirected planar graphs. In *Proceedings of the 43rd annual ACM Symposium on Theory of Computing*, STOC '11, pages 313–322, New York, NY, USA, 2011. ACM. doi:10.1145/1993636.1993679.
- [25] D. B. Johnson and S. Venkatesan. Using divide and conquer to find flows in directed planar networks in $O(n^{3/2} \log n)$ time. In *Proc 20th Allerton Conference on Communication, Control, and Computing*, pages 898–905, 1982.
- [26] A. Kézdy and P. McGuinness. Sequential and parallel algorithms to find a K_5 minor. In *Proc. 3rd ACM–SIAM Symposium on Discrete Algorithms (SODA 1992)*, pages 345–356, 1992. URL: <http://portal.acm.org/citation.cfm?id=139404.139475>.
- [27] S. Khuller and J. Naor. Flow in planar graphs with vertex capacities. *Algorithmica*, 11(3):200–225, 1994. doi:10.1007/BF01240733.
- [28] S. Khuller, J. Naor, and P. Klein. The lattice structure of flow in planar graphs. *SIAM J. Discrete Math.*, 6(3):477–490, 1993. doi:10.1137/0406038.
- [29] S. Mac Lane. A structural characterization of planar combinatorial graphs. *Duke Mathematical Journal*, 3(3):460–472, 1937. doi:10.1215/S0012-7094-37-00336-3.
- [30] M. Mares. Two linear time algorithms for mst on minor closed graph classes. *Archivum Mathematicum*, 40:315–320, 2004.
- [31] G. L. Miller and J. Naor. Flow in planar graphs with multiple sources and sinks. *SIAM J. Comput.*, 24(5):1002–1017, 1995. doi:10.1137/S0097539789162997.
- [32] J. B. Orlin. A faster strongly polynomial minimum cost flow algorithm. *Oper. Res.*, 41(2):338–350, 1993.
- [33] B. Reed and Z. Li. Optimization and recognition for K_5 -minor free graphs in linear time. In *Proc. 8th Latin American Theoretical Informatics Symposium (LATIN 2008)*, volume 4957 of *Lecture Notes in Computer Science*, pages 206–215. Springer-Verlag, 2008. doi:10.1007/978-3-540-78773-0_18.
- [34] B. Reed and D. Wood. Fast separation in a graph with an excluded minor. In S. Felsner, editor, *DMTCS Proceedings, Discrete Mathematics and Theoretical Computer Science*, volume AE, pages 45–50, 2005.
- [35] N. Robertson and P. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.

- [36] N. Robertson and P. Seymour. Excluding a graph with one crossing. In N. Robertson and P. Seymour, editors, *Graph Structure Theory: Proc. AMS-IMS-SIAM Joint Summer Research Conference on Graph Minors*, volume 147 of *Contemporary Mathematics*, pages 669–675. American Mathematical Society, 1993.
- [37] N. Robertson and P. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003. doi:10.1016/S0095-8956(03)00042-X.
- [38] N. Robertson and P. Seymour. Graph Minors. XX. Wagner’s conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004. doi:10.1016/j.jctb.2004.08.001.
- [39] S. Tazari and M. Müller-Hannemann. Shortest paths in linear time on minor-closed graph classes, with an application to steiner tree approximation. *Discrete Appl. Math.*, 157(4):673–684, Feb. 2009. doi:10.1016/j.dam.2008.08.002.
- [40] K. Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114(1):570–590, December 1937. doi:10.1007/BF01594196.
- [41] K. Wagner. Über eine Erweiterung des Satzes von Kuratowski. *Deutsche Mathematik*, 2:280–285, 1937.
- [42] K. Weihe. Maximum (s, t) -flows in planar networks in $O(|V| \log |V|)$ time. *Journal of Computer and System Sciences*, 55(3):454–476, 1997. doi:10.1006/jcss.1997.1538.