NETWORK FLOWS AND THE MAX-FLOW MIN-CUT THEOREM

AL STAPLES-MOORE

Abstract. The Max-Flow Min-Cut Theorem is an elementary theorem within the field of network flows, but it has some surprising implications in graph theory. We define network flows, prove the Max-Flow Min-Cut Theorem, and show that this theorem implies Menger's and König's Theorems.

Contents

1. Introduction to Network Flows

Graph theory provides a framework for discussing systems in which it is possible to travel between discrete vertices. If we extend a directed graph to a network flow by assigning a capacity and a flow value to every edge, then this flow can be used to model any number of systems in which a resource travels from one point to another, e.g. the spread of data in a network, traffic along roads, water in pipes, and so on.

1.1. Basic Definitions.

We model a network as a directed graph with a weight at every edge. For this paper, all graphs considered will be simple and finite.

Definition 1.1. A *network* N is a directed graph $G = (V, E)$ with a mapping $w : E \to \mathbb{R}$ that assigns a weight to each edge. The function w is called the *weight function* of N.

To define a flow on a network $N = (G, w)$, it is necessary to introduce the following additional features:

(a) Two vertices of G are given special names: a *source* s and a *sink* t.
(b) We extend w to a function $c : V \times V \to \mathbb{R}$ as follows:

$$c(u, v) = \begin{cases} w(\overline{uv}) & \text{if } \overline{uv} \in E \\ 0 & \text{otherwise.} \end{cases}$$

c is called the *capacity function* of N, and represents how much data can flow along that edge.

Definition 1.2. A *flow* on a network N with source s, sink t, and capacity function c is a function $f : V \times V \to \mathbb{R}$ such that:

(a) $f(u, v) < c(u, v) \forall u, v \in V$.
(b) For every vertex v not equal to s or t, $\sum_{u \in V} f(u, v) = \sum_{w \in V} f(v, w)$.
(c) $\sum_{u \in V} f(s, u) \geq 0$ and $\sum_{u \in V} f(u, t) \geq 0$.

The *value* of a flow is simply $\sum_{u \in V} f(s, u) \geq 0$, that is, the total data leaving the source.

In other words, the flow does not exceed the capacity on any edge, and at every vertex (other than the source and sink) the quantity of data entering equals the quantity of data leaving. The source has a non-negative amount of data leaving it, and the sink has a non-negative amount of data entering it.
For the purposes of this paper, we will assume that c and f take only non-negative, integral values.

1.2. **Additional Tools.** It is useful in talking about flows to define an edge's residual capacity as the difference between an edge's maximum allowed data flow and the amount of flow actually passing through it. We can then extend this definition to define the residual capacity along a path from s to t.

Definition 1.3. The *residual capacity* of an edge $\overline{uv}$ is $c_f(u, v) = c(u, v) - f(u, v)$.

Definition 1.4. Let p be a simple path from s to t. The *residual capacity* of p is $c_f(p) = \min\{c_f(u, v) : \overline{uv} \in p\}$. If $c_f(p) > 0$, then p is called an *augmenting path.*

We call a flow maximal between s and t if no legal flow on that network has a greater flow value. An augmenting path for a given flow is a path along which more data could flow. In other words, if there is an augmenting path, the flow is not maximal for that graph.

Lemma 1.5. *Let f be a flow on a network $N = (G, c)$. Then f is maximal only if f has no augmenting paths.*

Proof. Let f be a flow and suppose that an augmenting path p exists with residual capacity $c_f(p)$. Then we can define a new flow f' by adding $c_f(p)$ additional flow along each edge in f. By the definition of residual capacity, f' satisfies the capacity restraints for each edge in p, and has $c_f(p)$, a positive number, more net flow than f. Therefore f is not maximal. □

Lastly, we define a vertex cut in the context of a network.

Definition 1.6. A *vertex cut* of a flow f on a network $N = (G, w)$ with graph $G = (V, E)$ is a partition of V into disjoint sets S and T such that $s \in S$, $t \in T$, and $S \cup T = V$. The *net flow* of a cut (S, T) is $f(S, T) = \sum_{u \in S, v \in T} f(u, v)$, and the *capacity* is $c(S, T) = \sum_{u \in S, v \in T} c(u, v)$

Since for all vertices u and v, $f(u, v)$ is non-negative, it is clear that for any vertex cut (S, T), $c(S, T)$ is bounded below by zero. Since the set of all cuts for a given graph is finite, there exists a vertex cut of minimal capacity.

2. The Max-Flow Min-Cut Theorem

Fix a network (G, c). The Max-Flow Min-Cut Theorem states that the cut of minimum capacity vertex cut of a network N is equal to the maximal flow that could travel along that network. To prove this, we begin with the following lemma.

Lemma 2.1. *Let f be a flow on a network (G, c) with net flow v and let C be a vertex cut (S, T) with capacity k. Then $v \leq k$.*

Proof. Define $P = \{\overline{xy} : x \in S, y \in T\}$, the set of edges from S to T. As the flow along each edge cannot exceed the capacity along that edge for all edges in the network flow, in particular this is true for each edge in P. It therefore follows that the net flow over C $f(S, T) \leq k$. □

The previous lemma is a result that follows easily from defintions, yet it shows a fundamental connection between vertex cut size and max flow. The following theorem strengthens this connection.

Theorem 2.2. *For any network (G, c), the value of the maximal flow is equal to the minimum-capacity cut.*

Proof. We will begin with an arbitrary flow on the network. We will then define an augmentation process which, when taken to completion, will result in a maximum flow. Finally, we show how this process defines a minimum-capacity cut, and moreover that the maximal flow and the minimal cut have the same value.

Begin with any flow f. Since the zero-flow, the flow where every edge carries value zero, is valid for any network, such a flow must exist. Define the digraph D_f on the vertex set V with edge set $E' = \{\overline{uv} | c_f(u, v) > 0$. Suppose there is a path p from s to t within D_f. In G, p is a path along which every edge could carry more flow, and therefore p is an augmenting path. Let m be the minimum of $c_f(p)$, the residual capacity of the forward-oriented edges from s to t, and the set $\{f(u, v) | \overline{uv}$ is a backwards-oriented edge from s to $t\}$. $m > 0$ by the construction of D_f. Then increasing the flow of each forward-oriented edge by m and decreasing the flow of each backwards-oriented edge by m preserves the non-negativity of flows and the capacity restraint along each edge. This augmentation also increases the flow from s to t by the positive quantity m.

Denote the augmented flow by f'. Since p is not an augmenting path in f', as augmenting the flow by m either made one of the forward-oriented edges have maximum capacity flow or one of the backwards-oriented edges have zero flow.

Repeat the above process until no augmenting paths remain. As the number of total paths from s to t is finite in any finite graph, this task will finish in a finite number of steps. In the resulting digraph D_f, denote the set of vertices reachable from s as R and the set of vertices unreachable from s as U. It is clear that s is in R, and as no paths from s to t yet remain in D_f, t is in U. For each edge from a vertex in R to a vertex in U, each forward-oriented edge must be at full capacity and each backwards-oriented edge must have zero flow. Thus the augmented flow f' is a maximum flow, the cut (R, U) is a minimum cut, and moreover the flow of f' equals the capacity of (R, U), proving the claim. □

3. Applications of the Max-Flow Min-Cut Theorem

The Max-Flow Min-Cut Theorem is a fundamental result within the field of network flows, but it can also be used to show some profound theorems in graph theory.

3.1. Menger's Theorem.

There are multiple versions of Menger's Theorem, which all link the number of disjoint paths between two vertices to the size of some sort of minimum cut to separate those vertices. The theorem below is the vertex cut version for undirected graphs.

Definition 3.1. Let G be an undirected graph (V, E) and let u, v be vertices in V. Then a set $W \subseteq V$ is called an *vertex cut* of u and v if the removal of W separates u and v.

For all graphs, for each pair of vertices (u, v), it is clear that a minimal edge cut exists. Also note that this definition of vertex cut is easily reconcilable with the definition we used earlier – just let S be the union of the subgraph containing s separated by W and W itself and let T be the subgraph containing t separated by W.

Theorem 3.2. *For any finite undirected graph $G = (V, E)$ with vertices x and y, the minimum vertex cut of x and y is equal to the number of pairwise internally-disjoint paths (that is, the number of paths that pairwise share no edges) from x to y.*

Proof. Denote a maximal set of pairwise internally-disjoint paths from x to y as P, with $|P| = n$. Define a flow f on G as follows:

(a) Let x be the source and y be the sink
(b) Extend G to a network N by the capacity function $c(u, v) = 1$ for each edge $\overline{uv} \in E$.
(c) The flow along an edge e in E is 1 if $e \in p$ for some path $p \in P$ and e is forward-oriented from x to y, and 0 otherwise.
(d) If a vertex v is not x or y and it is part of a path p in P, erase all edges entering and leaving v that are not in p.

First, we must show that this flow must satisfy the capacity restraint and the conversion-of-flow restraint. The capacity restraint is trivially satisfied, as every edge has capacity 1 and flow 0 or 1. The conversion-of-flow restraint is also satisfied: pick any v in V such that v is not x or y. If v is not a member of any path in P, then no flow passes through it. If v is in at least one member of P, then it is in exactly one, as the paths in P are pairwise internally-disjoint. Denote the path containing v as p. Within p, for every edge $\overline{uv}$ entering v there is another path $\overline{vw}$ leaving v, with $f(u, v) = f(v, w) = 1$. Thus, for any vertex that is not a source or sink, the flow entering that vertex is equal to the flow leaving that vertex, satisfying the conversion-of-flow restraint.

Now we have a flow from s to t. Our next step will be to show that this flow is maximal. Suppose there is an augmenting path q from s to t. All edges are at full capacity, so any augmenting path cannot share any edge with any of the existing paths. Furthermore, q cannot pass through a non-source non-sink vertex belonging to a path in P, as we erased these edges in our construction of f. Therefore q is a path from s to t with no internal vertices in common with any member of P, and q is not in P. However, P was constructed to be a maximal set of internally-disjoint paths from s to t, a contradiction. Therefore f is maximal.

Since f was constructed to have flow 1 along each of its n pairwise internally-disjoint paths from x to y, the net flow of F is simply n. By the Max-Flow Min-Cut Theorem, the maximum flow from x to y is equal to the size of the minimal vertex cut of x and y, so the minimal vertex cut of x and y must be of size n. Thus the number of pairwise-internally disjoint paths is equal to the size of the minimum vertex cut, proving the claim. □

3.2. **König's Theorem.** Bipartite graphs are an important subclass of graphs with a number of applications. A graph is bipartite if it can be separated into two subsets such that no edges travel within a single subset. König's Theorem relates the size of a matching, a set of pairs of connected vertices from the two subsets, to the size of a vertex cover.

Definition 3.3. For any bipartite graph G with vertex sets X and Y, a *matching* is a set of disjoint pairs (x, y) such that $x \in X$, $y \in Y$, and the edge $\overline{xy}$ is in G. Since X and Y are finite, the cardinality of the set of such pairs is bounded above, so there is a maximal matching.

Definition 3.4. Let $G = (V, E)$ be any graph. A *vertex cover* is a subset W of V such that every edge in E has at least one vertex in W. Since the size of a cover is bounded below by 1 for any graph with at least one edge, every graph has a minimal cover.

Theorem 3.5. *For any finite bipartite graph G, the number of edges in a maximal matching equals the number of vertices in a minimal vertex cover.*

Proof. We will first extend G to a network, adding a source and a sink. We will then see that, in our new network, a maximal flow corresponds to a maximal matching, and a minimum cut corresponds to a minimum cover. From here, the min-cut max-flow theorem implies the desired result.

Let X and Y be a bipartite separation of the vertices of G. Starting with this graph, construct a digraph $G' = (V', E')$ where V' has all the vertices of V as well as a source s and a sink t. E' consists of all the edges in E, as well as new edges leading from s to every vertex in X, and also edges leading from each vertex in Y into t. Assign capacity values to edges as follows: give infinite capacity to each edge in E' that was originally in E (i.e., each edge from X to Y). To each of the newly added edges, give capacity 1.

Given a matching of cardinality k, it is easy to find a flow of value k. Simply push a flow of value 1 along the paths $\overline{sx}, \overline{xy}, \overline{yt}$, where (x, y) is one of the matched pairs. Likewise, any flow f must have a corresponding matching with cardinality equal to the flow's value. Thus, a maximal flow in G' corresponds to a maximal-cardinality matching in G.

Let W be a covering in G with r vertices, and let $W(X)$ and $W(Y)$ be subsets of W consisting of the vertices of W in X and Y, respectively. Next, let X' be

$X - W(X)$ and Y' be $Y - W(Y)$. The fact that W covers G implies that there is no edge from X' to Y'. Let S be the union of s, $W(Y)$, and Y' and let T be the union of t, $W(X)$, and Y'. Consider the cut (S,T) of G'; its cardinality is r, the cardinality of W. Thus, any vertex covering defines a vertex cut with equal value. Likewise, consider any vertex cut (S,T) in G' with finite value r. Edges from X to Y have infinite capacity, so each edge from S to T must either go from s to X or from Y to t, which have capacity 1. Since the cut had value r, (S,T) has r arcs.

Let W be the union of the set of vertices x in X such that $\overline{sx}$ is in (S,T) and the set of vertices y in Y such that $\overline{yt}$ is in (S,T); clearly, W has r vertices. For every covering in G there is a corresponding cut in G'. So a minimum cut corresponds to a minimum covering in G.

Thus far, we have shown that a maximal flow corresponds to a maximum-cardinality matching, and a minimum cut corresponds to a minimum vertex cover. By the max-flow min-cut theorem, a minimum cut is equal in value to a maximal flow. Therefore, by transitivity, the cardinality of a maximal matching in G is equal in value to that of a minimum covering. $\square$

Acknowledgments. It is a pleasure to thank my mentor, Alex Wright, for introducing me to this topic and assisting me during the writing process. I would also like to thank the University of Chicago Math Department for organizing the VIGRE REU program, and the National Science Foundation for making this program possible.

References

[1] W. D. Wallis. A Beginner's Guide to Graph Theory. Birkhäuser Boston. 2007.
[2] S. C. Locke. Menger's Theorems and Max-Flow-Min-Cut. http://math.fau.edu/locke/Menger.htm
[3] William T. Trotter. Max-Flow Min-Cut and Its LP Relatives. http://people.math.gatech.edu/ trotter/Section2-LP.htm
[4] V. K. Balakrishnan. Network Optimization. Chapman & Hall, London. 1995.
[5] S. Joseph. The Max-Flow Min-Cut Theorem. www.math.uri.edu/ eaton/maxflowmincut.pdf