

In these next two lectures we are going to talk about an important algorithmic problem called the *Network Flow Problem*. Network flow is important because it can be used to express a wide variety of different kinds of problems. So, by developing good algorithms for solving network flow, we immediately will get algorithms for solving many other problems as well. In Operations Research there are entire courses devoted to network flow and its variants. Topics in today's lecture include:

- The definition of the network flow problem
- The basic Ford-Fulkerson algorithm
- The maxflow-mincut theorem
- The bipartite matching problem

1 The Network Flow Problem

We begin with a definition of the problem. We are given a directed graph G , a start node s , and a sink node t . Each edge e in G has an associated non-negative *capacity* $c(e)$, where for all non-edges it is implicitly assumed that the capacity is 0. For example, consider the graph in Figure 1 below.

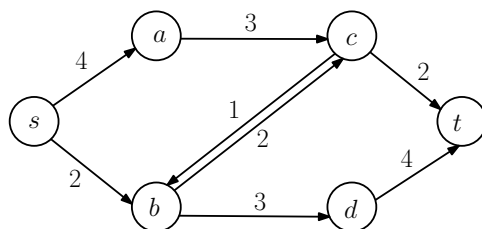


Figure 1: A network flow graph.

Our goal is to push as much *flow* as possible from s to t in the graph. The rules are that no edge can have flow exceeding its capacity, and for any vertex except for s and t , the flow *in* to the vertex must equal the flow *out* from the vertex. That is,

Capacity constraint: On any edge e we have $f(e) \leq c(e)$.

Flow conservation: For any vertex $v \notin \{s, t\}$, flow in equals flow out: $\sum_u f(u, v) = \sum_u f(v, u)$.

Subject to these constraints, we want to maximize the total flow into t . For instance, imagine we want to route message traffic from the source to the sink, and the capacities tell us how much bandwidth we're allowed on each edge.

E.g., in the above graph, what is the maximum flow from s to t ? Answer: 5. Using “capacity[flow]” notation, the positive flow looks as in Figure 2. Note that the flow can split and rejoin itself.

How can you see that the above flow was really maximum? Notice, this flow saturates the $a \rightarrow c$ and $s \rightarrow b$ edges, and, if you remove these, you disconnect t from s . In other words, the graph has an “ s - t cut” of size 5 (a set of edges of total capacity 5 such that if you remove them, this disconnects the sink from the source). The point is that any unit of flow going from s to t must take up at least 1 unit of capacity in these pipes. So, we know we're optimal.

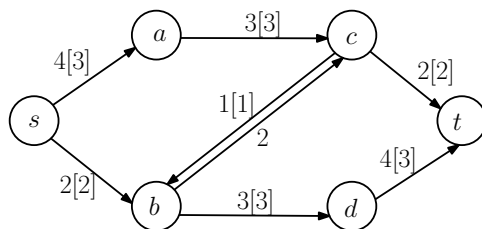


Figure 2: A network flow graph with positive flow shown using “capacity[flow]” notation.

We just argued that in general, the maximum s - t flow \leq the capacity of the minimum s - t cut.¹ An important property of flows, that we will prove as a byproduct of analyzing an algorithm for finding them, is that the maximum s - t flow is in fact *equal* to the capacity of the minimum s - t cut. This is called the *Maxflow-Mincut Theorem*. In fact, the algorithm will find a flow of some value k and a cut of capacity k , which will be proofs that both are optimal!

To describe the algorithm and analysis, it will help to be a bit more formal about a few of these quantities.

Definition 1 An s - t **cut** is a set of edges whose removal disconnects t from s . Or, formally, a cut is a partition of the vertex set into two pieces A and B where $s \in A$ and $t \in B$. (The edges of the cut are then all edges going from A to B).

Definition 2 The **capacity** of a cut (A, B) is the sum of capacities of edges in the cut. Or, in the formal viewpoint, it is the sum of capacities of all edges going from A to B . (Don’t include the edges from B to A .)

Definition 3 It will also be mathematically convenient for any edge (u, v) to define $f(v, u) = -f(u, v)$. This is called **skew-symmetry**. (We will think of flowing 1 unit on the edge from u to v as equivalently flowing -1 units on a back-edge from v to u .)

The skew-symmetry convention makes it especially easy to add two flows together. For instance, if we have one flow with 1 unit on the edge (c, b) and another flow with 2 units on the edge (b, c) , then adding them edge by edge does the right thing, resulting in a net flow of 1 unit from b to c . In fact, let’s now formally define the sum of two flows. If f and g are flows, then $h = f + g$ is defined as $h(u, v) = f(u, v) + g(u, v)$ for all pairs (u, v) . Notice that if f and g satisfy flow-conservation and skew-symmetry, then h does too. In fact, using skew-symmetry, if we wanted we could rewrite the flow conservation condition as $\forall v \notin \{s, t\}, \sum_u f(u, v) = 0$, since the total flow *out* of a node will always be the negative of the total flow *into* a node.

How can we find a maximum flow and prove it is correct? Here’s a very natural strategy: find a path from s to t and push as much flow on it as possible. Then look at the leftover capacities (an important issue will be how exactly we define this, but we will get to it in a minute) and repeat. Continue until there is no longer any path with capacity left to push any additional flow on. Of course, we need to *prove* that this works: that we can’t somehow end up at a suboptimal solution by making bad choices along the way. This approach, with the correct definition of “leftover capacity”, is called the Ford-Fulkerson algorithm.

¹This immediately implies the value of *any* s - t flow \leq maximum s - t flow \leq minimum s - t cut \leq *any* s - t cut.

2 The Ford-Fulkerson algorithm

The Ford-Fulkerson algorithm² is simply the following: while there exists an $s \rightarrow t$ path P of positive *residual capacity* (defined below), push the maximum possible flow along P . By the way, these paths P are called *augmenting paths*, because you use them to augment the existing flow.

Residual capacity is just the capacity left over given the existing flow, where we will use skew-symmetry to capture the notion that if we push f units of flow on an edge (u, v) , this *increases* our ability to push flow on the back-edge (v, u) by f .

Definition 4 Given a flow f in graph G , the **residual capacity** $c_f(u, v)$ is defined as $c_f(u, v) = c(u, v) - f(u, v)$, where recall that by skew-symmetry we have $f(v, u) = -f(u, v)$.

For example, given the flow in Figure 2, the edge (s, a) has residual capacity 1. The back-edge (a, s) has residual capacity 3, because its original capacity was 0 and we have $f(a, s) = -3$.

Definition 5 Given a flow f in graph G , the **residual graph** G_f is the directed graph with all edges of positive residual capacity, each one labeled by its residual capacity. Note: this may include back-edges of the original graph G .

Let's do an example. Consider the graph in Figure 1 and suppose we push two units of flow on the path $s \rightarrow b \rightarrow c \rightarrow t$. We then end up with the following residual graph:

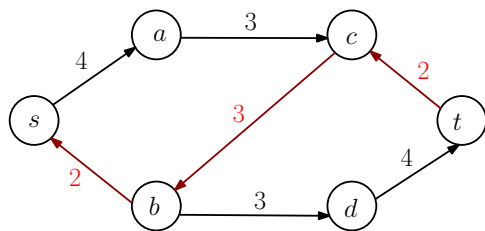


Figure 3: Residual graph resulting from pushing 2 units of flow along the path $s-b-c-t$ in the graph in Figure 1. The red edges denote the changes.

If we continue running Ford-Fulkerson, we see that in this graph the only path we can use to augment the existing flow is the path $s \rightarrow a \rightarrow c \rightarrow b \rightarrow d \rightarrow t$. Pushing the maximum 3 units on this path we then get the next residual graph, shown in Figure 4.

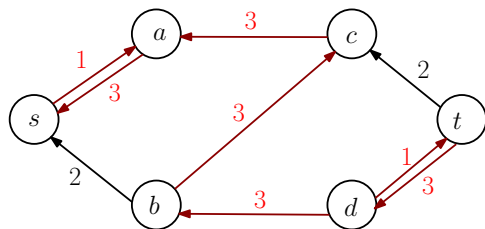


Figure 4: Residual graph resulting from pushing 3 units of flow along the path $s-a-c-b-d-t$ in the graph in Figure 3. Again, the red edges denote the changes.

²The algorithm is due to Lester R. Ford Jr. and Delbert R. Fulkerson, then at the RAND corporation working on figuring out the capacity of the Soviet rail network. They didn't want to find the max-flow, they wanted the min-cut. As we will soon see, the two problems are inextricably entwined. They presented their algorithm in a RAND report of 1955.

At this point there is no longer a path from s to t so we are done.

We can think of Ford-Fulkerson as at each step finding a new flow (along the augmenting path) and adding it to the existing flow, using the definition of adding flows we gave before. The definition of residual capacity ensures that the flow found by Ford-Fulkerson is *legal* (doesn't exceed the capacity constraints in the original graph). We now need to prove that in fact it is *maximum*. We'll worry about the number of iterations it takes and how to improve that later.

Note that one nice property of the residual graph is that it means that at each step we are left with same type of problem we started with. So, to implement Ford-Fulkerson, we can use any black-box path-finding method (e.g., DFS).

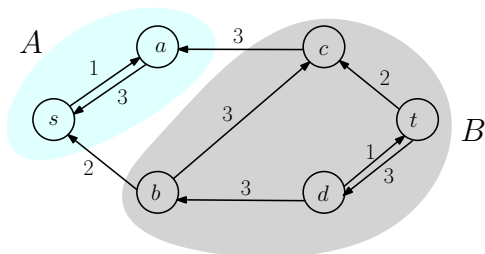
2.1 The Analysis

For now, let us assume that all the capacities are integers. If the maximum flow value is F , then the algorithm makes at most F iterations, since each iteration pushes at least one more unit of flow from s to t . We can implement each iteration in time $O(m+n)$ using DFS. So we get the following result.

Theorem 6 *If the given graph G has integer capacities, Ford-Fulkerson terminates in time $O(F(m+n))$ where F is the value of the maximum s - t flow.*

Theorem 7 *When it terminates, the Ford-Fulkerson algorithm outputs a maximum flow.*

Proof: Let's look at the final residual graph. This graph must have s and t disconnected by definition of the algorithm. Let A be the component containing s and B be the rest. Let c be the capacity of the (A, B) cut in the *original* graph — so we know we can't do better than c .



The claim is that we in fact *did* find a flow of value c (which therefore implies it is maximum). Here's why: let's look at what happens to the residual capacity of the (A, B) cut after each iteration of the algorithm. Say in some iteration we found a path with k units of flow. Then, even if the path zig-zagged between A and B , every time we went from A to B we added k to the flow from A to B and subtracted k from the residual capacity of the (A, B) cut, and every time we went from B to A we took away k from this flow and added k to the residual capacity of the cut³; moreover, we must have gone from A to B *exactly* one more time than we went from B to A . So, the residual capacity of this cut went down by exactly k . So, the drop in capacity is equal to the increase in flow. Since at the end the residual capacity is zero (remember how we defined A and B) this means the total flow is equal to c .

So, we've found a flow of value *equal* to the capacity of this cut. We know we can't do better, so this must be a max flow, and (A, B) must be a minimum cut. ■

Notice that in the above argument we actually proved the nonobvious *maxflow-mincut* theorem:

³This is where we use the fact that if we flow k units on the edge (u, v) , then in addition to reducing the residual capacity of the (u, v) edge by k we also *add* k to the residual capacity of the back-edge (v, u) .

Theorem 8 In any graph G , for any two vertices s and t , the maximum flow from s to t equals the capacity of the minimum (s, t) -cut.

We have also proven the *integral-flow theorem*: if all capacities are integers, then there is a maximum flow in which all flows are integers. This seems obvious, but it turns out to have some nice and non-obvious implications.

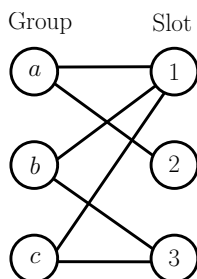
Technically, we just proved Theorem 8 only for integer capacities. What if the capacities are not integers? Firstly, if the capacities are rationals, then choose the smallest integer N such that $N \cdot c(u, v)$ is an integer for all edges (u, v) . It is easy to see that at each step we send at least $1/N$ amount of flow, and hence the number of iterations is at most NF , where F is the value of the maximum s - t flow. (One can argue this by observing that scaling up all capacities by N will make all capacities integers, whence we can apply our above argument.) And hence we get Theorem 8 for rational capacities as well.

What if the capacities are irrational? In this case Ford-Fulkerson may not terminate. And the solution it converges to (in the limit) may not even be the max-flow! (See here, here.) But the maxflow-minicut Theorem 8 still holds, even with irrational capacities. There are several ways to prove this; here's one. Suppose not, and suppose there is some flow network with the maxflow being $\epsilon > 0$ smaller than the mincut. Choose integer N such that $\frac{1}{N} \leq \frac{\epsilon}{2m}$, and round all capacities down to the nearest integer multiple of $1/N$. The mincut with these new edge capacities may have fallen by $m/N \leq \epsilon/2$, and the maxflow could be the same as the original flow network, but still there would be a gap of $\epsilon/2$ between maxflow and mincut in this rational-capacity network. But this is not possible, because used Ford-Fulkerson to prove maxflow-minicut for rational capacities in the previous paragraph.

In the next lecture we will look at methods for reducing the number of iterations the algorithm can take. For now, let's see how we can use an algorithm for the max flow problem to solve other problems as well: that is, how we can *reduce* other problems to the one we now know how to solve.

3 Bipartite Matching

Say we wanted to be more sophisticated about assigning groups to homework presentation slots. We could ask each group to list the slots acceptable to them, and then write this as a bipartite graph by drawing an edge between a group and a slot if that slot is acceptable to that group. For example:



This is an example of a **bipartite graph**: a graph with two sides L and R such that all edges go between L and R . A **matching** is a set of edges with no endpoints in common. What we want here in assigning groups to time slots is a **perfect matching**: a matching that connects every point in L with a point in R . For example, what is a perfect matching in the bipartite graph above?

More generally (say there is no perfect matching) we want a **maximum matching**: a matching with the maximum possible number of edges. We can solve this as follows:

Bipartite Matching:

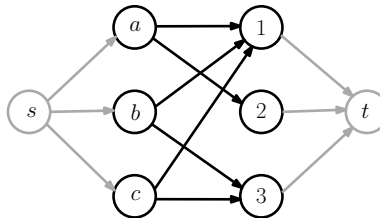
1. Set up a fake “start” node s connected to all vertices in L . Connect all vertices in R to a fake “sink” node t . Orient all edges left-to-right and give each a capacity of 1.

2. Find a max flow from s to t using Ford-Fulkerson.
3. Output the edges between L and R containing nonzero flow as the desired matching.

This finds a legal matching because edges from R to t have capacity 1, so the flow can't use two edges *into* the same node, and similarly the edges from s to L have capacity 1, so you can't have flow on two edges *leaving* the same node in L . It's a *maximum* matching because any matching gives you a flow of the same value: just connect s to the heads of those edges and connect the tails of those edges to t . (So if there was a better matching, we wouldn't be at a maximum flow).

What about the number of iterations of path-finding? This is at most the number of edges in the matching since each augmenting path gives us one new edge.

Let's run the algorithm on the above example. We first build this flow network.



Then we use Ford-Fulkerson. Say we start by pushing flow on s - a - 1 - t and s - c - 3 - t , thereby matching a to 1 and c to 3. These are bad choices, since with these choices b cannot be matched. But the augmenting path s - b - 1 - a - 2 - t *automatically* undoes them as it improves the flow!

Matchings come up in many different problems like matching up suppliers to customers, or cell-phones to cell-stations when you have overlapping cells. They are also a basic part of other algorithmic problems.