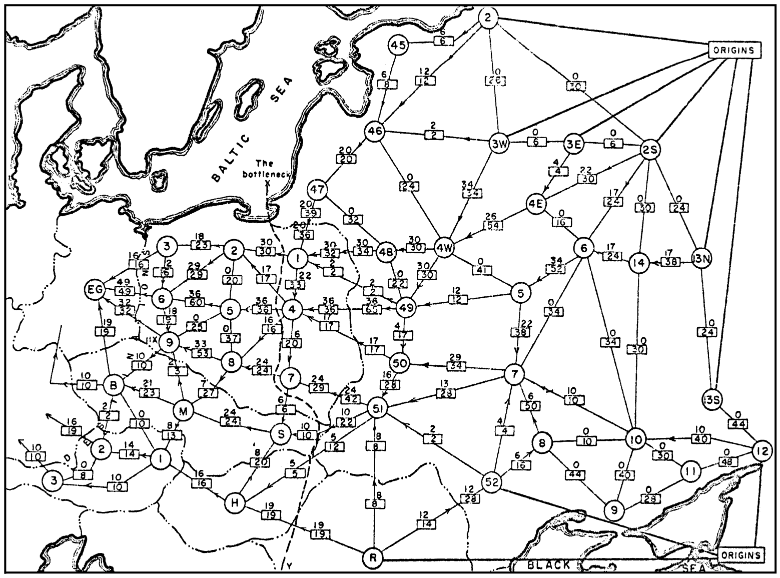

15 Maximum Flows and Minimum Cuts (November 10)

In the mid-1950s, Air Force researchers T. E. Harris and F. S. Ross published a classified report studying the rail network that linked the Soviet Union to its satellite countries in Eastern Europe. The network was modeled as a graph with 44 vertices, representing geographic regions, and 105 edges, representing links between those regions in the rail network. Each edge was given a weight, representing the rate at which material could be shipped from one region to the next. Essentially by trial and error, they determined both the maximum amount of stuff that could be moved from Russia into Europe, as well as the cheapest way to disrupt the network by removing links (or in less abstract terms, blowing up train tracks), which they called 'the bottleneck'. Their results (including the figure at the top of the page) were only declassified in 1999.[1]

This one of the first recorded applications of the *maximum flow* and *minimum cut* problems, which are defined as follows. Let $G = (V, E)$ be a directed graph, and let s and t be special vertices of G called the *source* and *target*. As in the previous lectures, we use $u \to v$ to denote the directed edge from vertex u to vertex v.

15.1 Flows

An **(s, t)-flow** (or just **flow** if the source and target are clear from context) is a function $f : E \to \mathbb{R}_{\geq 0}$ that satisfies the following *balance constraint* for all vertices v except (possibly) s and t:

$$\sum_u f(u \to v) = \sum_w f(v \to w).$$

In English, the total flow into any vertex is equal to the total flow out of that vertex. (To keep the notation simple, we assume here that $f(u \to v) = 0$ if there is no edge $u \to v$ in the graph.) The *value* of the flow f is defined as the excess flow out of the source vertex s:

$$\boxed{|f| = \sum_w f(s \to w) - \sum_u f(u \to s)}$$

[1] Both the fiigure and the story were taken from Alexander Schrijver's fascinating survey 'On the history of combinatorial optimization (till 1960)'.

It's not hard to prove that $|f|$ is also equal to the excess flow *into* the target vertex t. First we observe that

$$\sum_v \left(\sum_w f(v \to w) - \sum_u f(u \to v) \right) = \sum_v \sum_w f(v \to w) - \sum_v \sum_u f(u \to v) = 0$$

because both summations count the total flow across all edges. On the other hand, the balance constraint implies that

$$\begin{aligned} &\sum_v \left(\sum_w f(v \to w) - \sum_u f(u \to v) \right) \\ &\qquad = \left(\sum_w f(t \to w) - \sum_u f(u \to s) \right) + \left(\sum_w f(t \to w) - \sum_u f(u \to t) \right) \\ &\qquad = |f| + \left(\sum_w f(t \to w) - \sum_u f(u \to t) \right). \end{aligned}$$

It follows that

$$|f| = \sum_u f(u \to t) - \sum_w f(t \to w).$$

Now suppose each edge e in our graph has a non-negative *capacity* $c(e)$. We say that a flow f is *subject to* the capacity function c if $f(e) \le c(e)$ for all e. Most of the time we will consider only flows that are subject to some fixed capacity function c. We say that a flow f *saturates* an edge e if $f(e) = c(e)$, and *avoids* an edge e if $f(e) = 0$. The *maximum flow problem* is to compute an (s,t)-flow in a given directed graph, subject to a given capacity function, whose value is as large as possible.

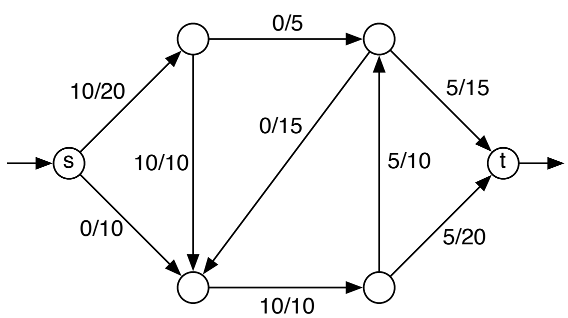

An (s,t)-flow with value 10. Each edge is labeled with its flow/capacity.

15.2 Cuts

An **(s,t)-cut** (or just **cut** if the source and target are clear from context) is a partition of the vertices into disjoint subsets S and T—meaning $S \cup T = V$ and $S \cap T = \varnothing$—where $s \in S$ and $t \in T$.

If we have a capacity function $c : E \to \mathbb{R}_{\ge 0}$, the *cost* of a cut is the sum of the capacities of the edges that start in S and end in T:

$$\boxed{\|S,T\| = \sum_{v \in S} \sum_{w \in T} c(v \to w).}$$

(Again, if $v \to w$ is not an edge in the graph, we assume $c(v \to w) = 0$.) Notice that the definition is asymmetric; edges that start in T and end in S are unimportant. The *minimum cut* problem is to compute an (s,t)-cut whose cost, relative to a given capacity function, is as large as possible.

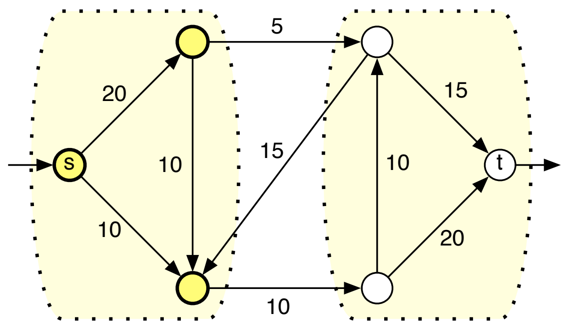

An (s,t)-cut with cost 15. Each edge is labeled with its capacity.

Intuitively, the minimum cut is the cheapest way to disrupt all flow from s to t. Indeed, it is not hard to show that the value of *any* (s,t)-flow that is subject to c is bounded about by the cost of *any* (s,t)-cut. Choose your favorite flow f and your favorite cut (S,T), and then follow the bouncing equal signs:

$$\begin{aligned}
|f| &= \sum_w f(s \to w) - \sum_u f(u \to s) && \text{by definition} \\
&= \sum_{v \in S} \left(\sum_w f(v \to w) - \sum_u f(u \to v) \right) && \text{by the balance constraint} \\
&= \sum_{v \in S} \left(\sum_{w \in T} f(v \to w) - \sum_{u \in T} f(u \to v) \right) && \text{removing duplicate edges} \\
&\le \sum_{v \in S} \sum_{w \in T} f(v \to w) && \text{since } f(u \to v) \ge 0 \\
&\le \sum_{v \in S} \sum_{w \in T} c(v \to w) && \text{since } f(u \to v) \le c(v \to w) \\
&= \|S,T\| && \text{by definition}
\end{aligned}$$

Our derivation actually implies the following stronger observation: $|f| = \|S,T\|$ **if and only if f saturates every edge from S to T and avoids every edge from T to S.** Moreover, if we have a flow f and a cut (S,T) that satisfies this equality condition, f must be a maximum cut, and (S,T) must be a minimum flow.

15.3 The Max-Flow Min-Cut Theorem

Surprisingly, for any weighted directed graph, there is always a flow f and a cut (S,T) that satisfy the equality condition. This is the famous *max-flow min-cut theorem*:

The value of the maximum flow is equal to the cost of the minimum cut.

The rest of this section gives a proof of this theorem; we will eventually turn this proof into an algorithm.

Fix a graph G, vertices s and t, and a capacity function $c : E \to \mathbb{R}_{\geq 0}$. The proof will be easier if we assume that the capacity function is *reduced*: For any vertices u and v, either $c(u \to v) = 0$ or $c(v \to u) = 0$, or equivalently, if an edge appears in G, then its reversal does not. This assumption is easy to enforce. Whenever an edge $u \to v$ and its reversal $v \to u$ are both the graph, replace the edge $u \to v$ with a path $u \to x \to v$ of length two, where x is a new vertex and $c(u \to x) = c(x \to v) = c(u \to v)$. The modified graph has the same maximum flow value and minimum cut cost as the original graph.

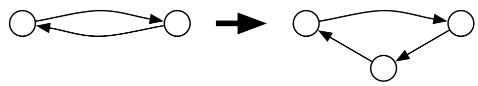

Enforcing the one-direction assumption.

Let f be a flow subject to c. We define a new capacity function $c_f : V \times V \to \mathbb{R}$, called the *residual capacity*, as follows:

$$c_f(u \to v) = \begin{cases} c(u \to v) - f(u \to v) & \text{if } u \to v \in E \\ f(v \to u) & \text{if } v \to u \in E \\ 0 & \text{otherwise} \end{cases}.$$

Since $f \geq 0$ and $f \leq c$, the residual capacities are always non-negative. It is possible to have $c_f(u \to v) > 0$ even if $u \to v$ is not an edge in the original graph G. Thus, we define the *residual graph* $G_f = (V, E_f)$, where E_f is the set of edges whose non-zero residual capacity is positive. Notice that the residual capacities are *not* necessarily reduced; it is quite possible to have both $c_f(u \to v) > 0$ and $c_f(v \to u) > 0$.

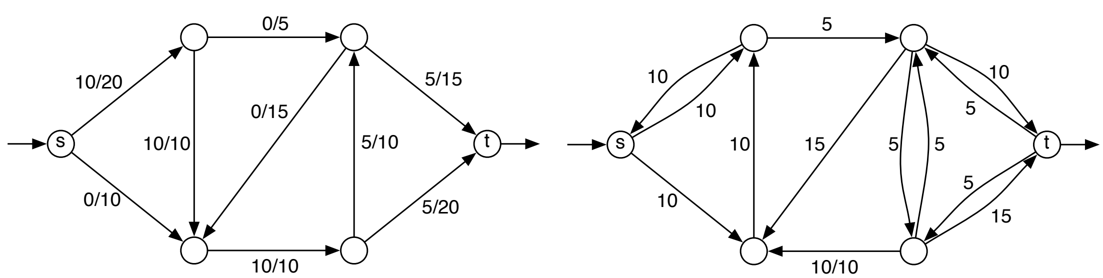

A flow f in a weighted graph G and its residual graph G_f.

Suppose there is no path from the source s to the target t in the residual graph G_f. Let S be the set of vertices that are reachable from s in G_f, and let $T = V \setminus S$. The partition (S, T) is clearly an (s, t)-cut. For every vertex $u \in S$ and $v \in T$, we have

$$c_f(u \to v) = (c(u \to v) - f(u \to v)) + f(v \to u) = 0,$$

which implies that $c(u \to v) - f(u \to v) = 0$ and $f(v \to u) = 0$. In other words, our flow f saturates every edge from S to T and avoids every edge from T to S. It follows that $|f| = \|S, T\|$. Moreover, f is a maximum flow and (S, T) is a minimum cut.

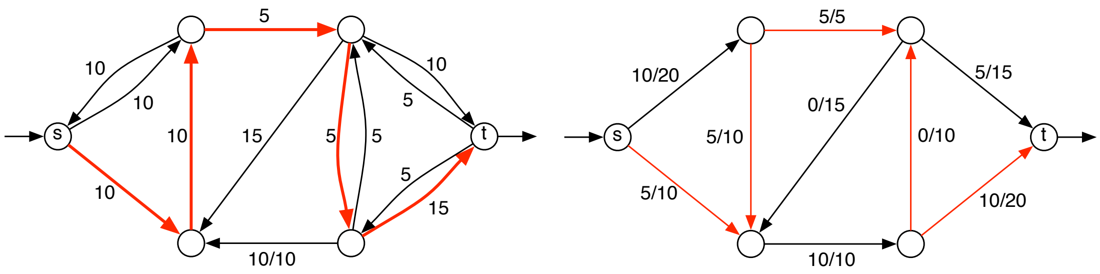

An augmenting path in G_f with value $F = 5$ and the augmented flow f'.

On the other hand, suppose there is a path $s = v_0 \to v_1 \to \cdots \to v_r = t$ in G_f. We refer to $v_0 \to v_1 \to \cdots \to v_r$ as an *augmenting path*. Let $F = \min_i c_f(v_i \to v_{i+1})$ denote the maximum amount of flow that we can push through the augmenting path in G_f. We define a new flow function $f' : E \to \mathbb{R}$ as follows:

$$f'(u \to v) = \begin{cases} f(u \to v) + F & \text{if } u \to v \text{ is in the augmenting path} \\ f(u \to v) - F & \text{if } v \to u \text{ is in the augmenting path} \\ f(u \to v) & \text{otherwise} \end{cases}$$

To prove this is a legal flow function subject to the original capacities c, we need to verify that $f' \geq 0$ and $f' \leq c$. Consider an edge $u \to v$ in G. If $u \to v$ is in the augmenting path, then $f'(u \to v) > f(u \to v) \geq 0$ and

$$\begin{aligned} f'(u \to v) &= f(u \to v) + F && \text{by definition of } f' \\ &\leq f(u \to v) + c_f(u \to v) && \text{by definition of } F \\ &= f(u \to v) + c(u \to v) - f(u \to v) && \text{by definition of } c_f \\ &= c(u \to v) && \text{since } c \text{ is reduced.} \end{aligned}$$

On the other hand, if the reversal $v \to u$ is in the augmenting path, then $f'(u \to v) < f(u \to v) \leq c(u \to v)$ and

$$\begin{aligned} f'(u \to v) &= f(u \to v) - F && \text{by definition of } f' \\ &\geq f(u \to v) - c_f(v \to u) && \text{by definition of } F \\ &= f(u \to v) - f(u \to v) && \text{by definition of } c_f \\ &= 0 && \text{since } c \text{ is reduced.} \end{aligned}$$

Finally, we observe that (without loss of generality) only the first edge in the augmenting path leaves s, so $|f'| = |f| + F > 0$. In other words, f is *not* a maximum flow.

This completes the proof!