Practice Problems 8

Problem 1. Max-Flow Min-Cut: A Worked Example.

Consider the flow network $G = (V, E)$ shown in Fig. 1, with source s and sink t. Each edge is labeled with its capacity.

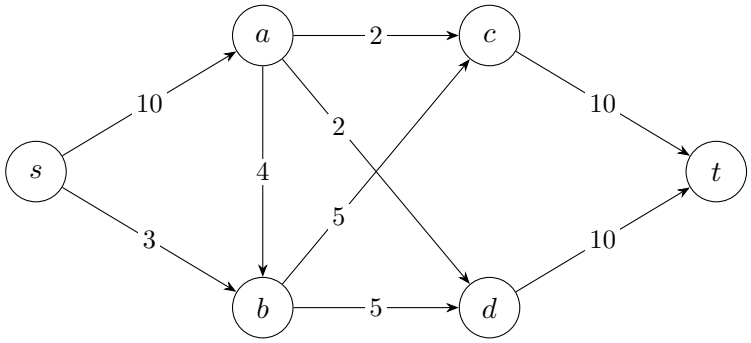


Figure 1: A flow network with source s and sink t.

(a) Run the Ford-Fulkerson algorithm on G, starting from the zero flow. List the augmenting paths you use, in order, together with the bottleneck capacity of each. Then draw the final flow f, labeling every edge with $f(e)/c(e)$, and state its value $|f|$.

(b) Draw the residual network G_f for your final flow. Let S be the set of vertices reachable from s by a directed path in G_f. Give S, and compute the capacity of the cut $(S, V \setminus S)$.

(c) Verify directly that your answers to (a) and (b) satisfy the max-flow min-cut theorem. Specifically, list the edges of G that cross the cut in the forward direction (from S to $V \setminus S$) and check that each is saturated, and list the edges that cross in the backward direction (from $V \setminus S$ to S) and check that each carries zero flow. Explain why these two conditions together force $|f|$ to equal the capacity of the cut.

(d) Your computation in part (b) is an instance of a general fact, which you will now prove. Let f be any flow admitting no augmenting path, and let

$$S \;=\; \{\, v \in V : \text{there is a directed path from } s \text{ to } v \text{ in } G_f \,\}.$$

Prove that the capacity of the cut $(S, V \setminus S)$ equals $|f|$. Then explain why this single fact proves *both* that f is a maximum flow and that $(S, V \setminus S)$ is a minimum cut.

Hint: First check that $s \in S$ and $t \notin S$, so that $(S, V \setminus S)$ really is an s-t cut. Then consider any edge $(u, v) \in E$ with $u \in S$ and $v \notin S$. What does the absence of a residual edge (u, v) in G_f tell you about $f(u, v)$? Ask the same question about an edge (v, u) with $v \notin S$ and $u \in S$, this time using the residual edge that runs *backwards*. Finally recall that $|f|$ equals the net flow across any s-t cut.

Problem 2. Multiple Sources and Multiple Sinks.

So far our networks have had a single source and a single sink. In many applications there are several of each. For example, several factories producing a commodity and several warehouses consuming it. Formally, a **multi-source multi-sink network** is a digraph $G = (V, E)$ with nonnegative edge capacities, together with two disjoint sets $\mathcal{S} = \{s_1, \ldots, s_p\} \subseteq V$ of *sources* and $\mathcal{T} = \{t_1, \ldots, t_q\} \subseteq V$ of *sinks*. As is standard, we assume that no edge enters a source and no edge leaves a sink. A **flow** is a function f obeying the usual capacity constraints $0 \le f(e) \le c(e)$, and obeying conservation at every vertex of $V \setminus (\mathcal{S} \cup \mathcal{T})$. Its **value** is the total flow out of the sources,

$$|f| \;=\; \sum_{i=1}^{p} \; \sum_{(s_i, v) \in E} f(s_i, v).$$

The goal is to find a flow of maximum value.

(a) A classmate proposes to solve the problem by running an ordinary single-source single-sink max-flow computation once for each pair (s_i, t_j) and adding up the results. Give a small network on which this overestimates the true maximum value, and explain in one sentence what the proposal ignores.

(b) Present a reduction that solves the problem using a single ordinary max-flow computation.

Hint: build from G a network G^* with one source and one sink, such that a maximum flow in G^* yields a maximum flow in G of the same value.

Prove that your reduction is correct. Be careful to argue in *both* directions: every flow in G of value k gives rise to a flow in G^* of value k, and every flow in G^* of value k restricts to a flow in G of value k. (For the second direction, check that conservation really does hold at each s_i in the restricted flow, and say why.)