

Classical Algorithms for Maximum Flow

R. Kyng & M. Probst Gutenberg

Problem Set 7 — Thursday, April 27th

The exercises for this week will not count toward your grade, but you are highly encouraged to solve them all. This exercise sheet has exercises related to week 7. We encourage you to start early so you have time to go through everything.

To get feedback, you must hand in your solutions by 23:59 on May 4. Both hand-written and L^AT_EX solutions are acceptable, but we will only attempt to read legible text.

Exercise 1: Implementing Field Preservations for Cut-Link Tree Rotations

In Chapter 14, Section 14.3, we described the operations $\text{PLINK}(u, v)$ and $\text{PCUT}(u, v)$ by doing $O(\log n)$ tree rotations (in expectation). However, we omitted the details. Here we ask you to give the pseudo-code for the operation $\text{PTREEROTATION}(v, w)$ where it is assumed that on input v is the parent of w and the operation manipulates the tree over the path such that v and w change position as described in the script. For simplicity, you are allowed to assume that w is the left child of v in the treap \mathcal{T}_p , and that the nodes $\text{left}_{\mathcal{T}_p}(v)$, $\text{right}_{\mathcal{T}_p}(v)$, $\text{left}_{\mathcal{T}_p}(w)$, $\text{right}_{\mathcal{T}_p}(w)$ exist. Accompany your pseudo-code with a brief analysis that confirms that the run-time is indeed $O(1)$.

Solution The pseudocode is given in algorithm 1. Clearly, all updates in the algorithm (changing pointers and updating Δ_{\min} and Δ_{cost} values can be done in $O(1)$ hence the total runtime is $O(1)$.

Exercise 2: Max Flow in directed Graphs with Edge Capacities

Consider directed graph $G = (V, E, c)$ with arbitrary capacities $c \geq \mathbf{0}$.

Let $\mathbf{B} \in \mathbb{R}^{E \times V}$ be the edge vertex incidence matrix of the graph, i.e. if $e \in E$ and $(u, v) = e$ then $\mathbf{B}(e, u) = 1$ and $\mathbf{B}(e, v) = -1$.

We let $\chi_v \in \mathbb{R}^V$ denote the indicator of vertex v , i.e. $\chi_v(v) = 1$ and $\chi_v(u) = 0$ for $u \neq v$.

We let $s \in V$ denote the flow “source” and $t \in V$ the flow “sink”.

The maximum flow problem is given by

$$\begin{aligned} & \max_{\mathbf{f} \in \mathbb{R}^E, F \geq 0} F \\ \text{s.t. } & \mathbf{B}\mathbf{f} = F(-\chi_s + \chi_t) \\ & \mathbf{0} \leq \mathbf{f} \leq \mathbf{c} \end{aligned}$$

In the context of a given maximum flow problem, for a flow \mathbf{f} satisfying $\mathbf{B}\mathbf{f} = F(-\chi_s + \chi_t)$, we define $\text{val}(\mathbf{f}) = F$.

Let \mathbf{f}^* denote a feasible flow maximizing F , so that the maximum attainable flow value F is $\text{val}(\mathbf{f}^*)$.


Algorithm 1: PTreeRotation(v,w)

```
/* Recall we assume  $left_{\mathcal{T}_p}(v) = w$ ; and  $w$  and  $v$  both have two children.      */
/* We first change the parent pointer for the parent that supports the subtree
   (if there is one).                                                                */
if  $parent_{\mathcal{T}_p}(v) \neq NULL$  then
   $p \leftarrow parent_{\mathcal{T}_p}(v)$ 
  if  $left_{\mathcal{T}_p}(p) = v$  then
     $left_{\mathcal{T}_p}(p) \leftarrow w$ 
  end
  else
     $right_{\mathcal{T}_p}(p) \leftarrow w$ 
  end
end
/* Compute roots of the subtrees rooted at the children of  $v$  and  $w$ .                */
 $a \leftarrow left_{\mathcal{T}_p}(w)$ .
 $b \leftarrow right_{\mathcal{T}_p}(w)$ .
 $c \leftarrow right_{\mathcal{T}_p}(v)$ .
/* Change the tree structure according to the rotation.                            */
 $right_{\mathcal{T}_p}(w) \leftarrow v$ .
 $left_{\mathcal{T}_p}(v) \leftarrow b$ .
 $\Delta c(v) \leftarrow \Delta cost(v)$ .
 $\Delta c(w) \leftarrow \Delta min(w) + \Delta cost(w)$ .
 $\Delta m(b) \leftarrow \Delta min(w) + \Delta min(b)$ .
 $\Delta m(c) \leftarrow \Delta min(c)$ .
/* Recompute the fields at  $v$  and  $w$ .                                              */
 $\Delta min(w) \leftarrow \Delta min(v)$ . ; // The old minimum of  $v$  is now the minimum of  $w$ .
 $\Delta cost(w) \leftarrow \Delta c(w)$ .; // We set the cost of  $w$  to be the pre-computed cost.
/* The new minimum of  $v$  is the minimum in subtrees rooted at  $b$  or  $c$  or the
   element  $v$  itself.                                                                */
 $\Delta min(v) \leftarrow \min\{\Delta m(b), \Delta m(c), \Delta c(v)\}$ .
 $\Delta cost(v) \leftarrow \Delta c(v) - \Delta min(v)$ .
```

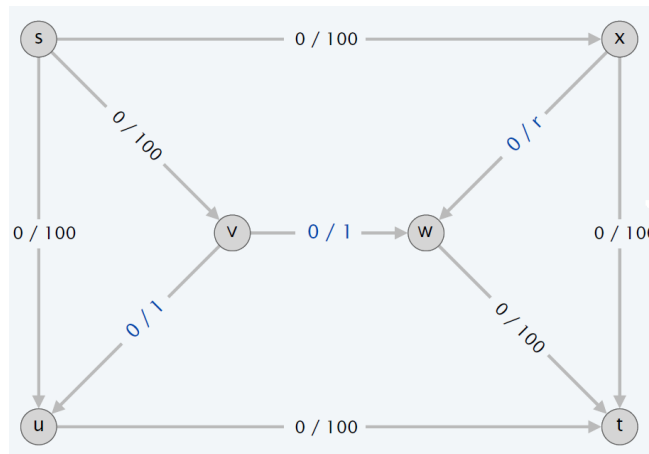
Exercise 2.A: Convergence of Ford-Fulkerson

Show that the Ford-Fulkerson algorithm may not terminate; moreover, it may converge a value not equal to the value of the maximum flow.

Hint: You might use the graph below with the given capacities, where $r = \frac{\sqrt{5}-1}{2}$ (which implies that $r^2 = 1 - r$).




Solution. Consider the graph below with the given capacities.¹ The goal is to find a sequence of augmenting paths such that the algorithm converges to a value which is much smaller than the value of maximum flow.




First, we augment Path 1: s, v, w, t .


¹All the figures used in this exercise are borrowed from the slides of the course Theory of Algorithms by Kevin Wayne, Princeton University.




Then, we augment Path 2: s, x, w, v, u, t .




Then, we augment Path 3: s, v, w, x, t .




Then, we augment Path 4: s, x, w, v, u, t .




Then, we augment Path 5: s, u, v, w, t .




Then, we augment Path 6: s, x, w, v, u, t .




Then, we augment Path 7: s, v, w, x, t .



Then, we augment Path 8: s, x, w, v, u, t .



Then, we augment Path 9: s, u, v, w, t .



After augmenting Path 1 the flow is 1.

After augmenting Path 5 the flow is $1 + 2r + 2r^2$.

After augmenting Path 9 the flow is $1 + 2r + 2r^2 + 2r^3 + 2r^4$.

Using the given sequence of augmenting paths, after $(1 + 4k)$ -th such path, the value of the flow is equal to

$$F = 1 + 2 \sum_{i=1}^k r^i \leq 1 + 2 \sum_{i=1}^{\infty} r^i = 3 + 2r \leq 5. \quad (1)$$

We observe that the value of maximum flow is equal to $200 + 1 = 201$. Therefore, we can conclude that the Ford-Fulkerson algorithm may not terminate; moreover, it may converge a value not equal to the value of the maximum flow.

Exercise 2.B: Iterative Refinement for Maximum Flow

Suppose we have an algorithm FLOWREFINE, which given a maximum flow instance $G = (V, E, \mathbf{c})$ with source $s \in V$ and sink $t \in V$ returns a feasible s - t flow $\tilde{\mathbf{f}}$, i.e. $\mathbf{B}\tilde{\mathbf{f}} = F(-\chi_s + \chi_t)$ for some F , and $\mathbf{0} \leq \tilde{\mathbf{f}} \leq \mathbf{c}$, and $\tilde{\mathbf{f}}$ is guaranteed to route at least half the maximum flow, i.e. $F = \text{val}(\tilde{\mathbf{f}}) \geq 0.5 \text{val}(\mathbf{f}^*)$.

Suppose that the running time of FLOWREFINE is $O(|E|^c)$ for some constant $c \geq 1$.

Explain how we can use FLOWREFINE to find a flow $\hat{\mathbf{f}}$ that routes at least $(1 - \epsilon) \text{val}(\mathbf{f}^*)$ in time $O(|E|^c \log(1/\epsilon))$.

Solution. Assume that we are given a maximum flow instance $G = (V, E, \mathbf{c})$ with source $s \in V$ and sink $t \in V$. Consider the flow $\mathbf{f} = \mathbf{0}$. We run the algorithm FLOWREFINE on graph G . We know that, it will return a feasible s - t flow \mathbf{f}_1 such that $\text{val}(\mathbf{f}_1) \geq 0.5 \text{val}(\mathbf{f}_G^*)$, where \mathbf{f}_G^* is a maximum flow in G . Let $G_{\mathbf{f}_1}$ be the residual graph of \mathbf{f}_1 and set $\mathbf{f} = \mathbf{f} + \mathbf{f}_1$. Let's run FLOWREFINE this time on graph $G_{\mathbf{f}_1}$. It should return a flow \mathbf{f}_2 . Now, we set $\mathbf{f} = \mathbf{f} + \mathbf{f}_2$. Thus, we get

$$\begin{aligned} \text{val}(\mathbf{f}) &= \text{val}(\mathbf{f}_1) + \text{val}(\mathbf{f}_2) \\ &\geq \text{val}(\mathbf{f}_1) + \frac{1}{2} \text{val}(\mathbf{f}_{G_{\mathbf{f}_1}}^*) \\ &= \text{val}(\mathbf{f}_1) + \frac{1}{2} \text{val}(\mathbf{f}_G^*) - \frac{1}{2} \text{val}(\mathbf{f}_1) \\ &= \frac{1}{2} \text{val}(\mathbf{f}_1) + \frac{1}{2} \text{val}(\mathbf{f}_G^*) \\ &\geq \frac{1}{4} \text{val}(\mathbf{f}_G^*) + \frac{1}{2} \text{val}(\mathbf{f}_G^*) \\ &\geq \frac{3}{4} \text{val}(\mathbf{f}_G^*) \end{aligned}$$

where we used $\text{val}(\mathbf{f}_{G_{\mathbf{f}_1}}^*) + \text{val}(\mathbf{f}_1) = \text{val}(\mathbf{f}_G^*)$ based on Lemmas 2.6 and 2.7 from Lecture 10.

By repeating the same argument, we can conclude that after $k = \log(1/\epsilon)$ iterations

$$\text{val}(\mathbf{f}) \geq \left(1 - \frac{1}{2^k}\right) \text{val}(\mathbf{f}_G^*) = (1 - \epsilon) \text{val}(\mathbf{f}_G^*)$$

We know that the running time of `FLOWREFINE` is $O(|E|^c)$ for some constant $c \geq 1$ and one can generate a residual graph in $O(|E|)$. Therefore, the above algorithm runs in time $O(|E|^c \log(1/\epsilon))$.