

Problem: An ℓ_1 -Interior Point Method for Maximum Flow

Consider an undirected graph $G = (V, E)$ with capacities $\mathbf{c} \in \mathbb{R}_+^E$, as usual with n vertices and m edges. We will assume $m > 10$ and that the entries of \mathbf{c} are positive integers and $\|\mathbf{c}\|_1 \leq m^{10}$. Let $s, t \in V$ be a source and sink vertex respectively. Assume we are also given the maximum flow value $0 < F \leq m^{10}$ (F is an integer) that can be routed between s and t .

Now we add a new special edge from \tilde{e} between s and t with no capacity limit. We then define the extended edge set $\tilde{E} = E \cup \{\tilde{e}\}$. We pick the orientation of \tilde{e} such that $\mathbf{1}_{\tilde{e}}^\top \mathbf{f} = \mathbf{f}(\tilde{e})$ measures the amount of flow going from s to t on edge \tilde{e} . Let \mathbf{B} be the edge-vertex incidence matrix of E . Ultimately, we want to find a flow $\mathbf{f} \in \mathbb{R}^E$ that routes F units of flow from s to t without using the edge \tilde{e} . We will do this by setting up a linear program that asks us to find a flow $\mathbf{f} \in \mathbb{R}^{\tilde{E}}$ that routes F units of flow from s to t without using the edge \tilde{e} .

We can now write this as a linear program that asks us to minimize the flow on \tilde{e} , subject to the constraints $\mathbf{B}\mathbf{f} = F(-\mathbf{1}_s + \mathbf{1}_t)$ and for all $e \in E$ that $-\mathbf{c}(e) \leq \mathbf{f}(e) \leq \mathbf{c}(e)$.

$$\begin{aligned} & \min_{\mathbf{f} \in \mathbb{R}^{\tilde{E}}} \mathbf{f}(\tilde{e}) & (1) \\ \text{s.t. } & \mathbf{B}\mathbf{f} = F(-\mathbf{1}_s + \mathbf{1}_t) \\ & \text{and for all } e \in E \\ & -\mathbf{c}(e) \leq \mathbf{f}(e) \leq \mathbf{c}(e) \end{aligned}$$

Now we are going to consider a *barrier program* which is an variant of the above program but using barrier functions that make it suitable for optimizing with an interior point method. We will consider an inequality constraint set that also constrains $\mathbf{f}(\tilde{e})$. We denote this capacity constrained flow set by

$$\mathcal{I} = \left\{ \mathbf{f} \in \mathbb{R}^{\tilde{E}} : \mathbf{f}(\tilde{e}) > 0 \text{ and for all } e \in E. -\mathbf{c}(e) < \mathbf{f}(e) < \mathbf{c}(e) \right\}$$

Then, we define a barrier $B : \mathcal{I} \rightarrow \mathbb{R}$ by

$$B(\mathbf{f}) = \sum_{e \in E} -\log \left(1 - \frac{\mathbf{f}(e)}{\mathbf{c}(e)} \right) - \log \left(1 + \frac{\mathbf{f}(e)}{\mathbf{c}(e)} \right)$$

We further define our overall potential $\Phi(\mathbf{f}) = 10m \log(\mathbf{f}(\tilde{e})) + B(\mathbf{f})$. We can now introduce our barrier program:

$$\begin{aligned} & \min_{\mathbf{f} \in \mathcal{I}} \Phi(\mathbf{f}) & (2) \\ \text{s.t. } & \mathbf{B}\mathbf{f} = F(-\mathbf{1}_s + \mathbf{1}_t) \end{aligned}$$

Part A: The Potential Function: Initialization and End-Goal. We would now like to begin understanding the Programs (1) and (2) better. First, we will try understand why solutions to Program (1) give us a maximum flow, and second we will investigate how a feasible solution \mathbf{f} for Program (2) with a small potential value $\Phi(\mathbf{f})$ give us an approximate maximum flow. Finally, we will see that there is an easy way to get a reasonable feasible starting point \mathbf{f}_0 for Program (2). Later we will see how to improve this solution to get a maximum flow.

1. Is Program (1) a convex optimization program? Is Program (2) a convex program?
2. Prove that any optimal solution \mathbf{f}^* for the Program (1) routes F units of flow from s to t on the edges of E and has $\mathbf{f}(\tilde{e}) = 0$.
3. Prove that if $\mathbf{B}\mathbf{f} = F(-\mathbf{1}_s + \mathbf{1}_t)$, and $\mathbf{f} \in \mathcal{I}$, and $\Phi(\mathbf{f}) \leq -10m \log m$ then $\mathbf{f}(\tilde{e}) \leq 1/m$.
4. Prove that for $\mathbf{f}_0 = F\mathbf{1}_{\tilde{e}}$ we have $\mathbf{B}\mathbf{f}_0 = F(-\mathbf{1}_s + \mathbf{1}_t)$, and $\mathbf{f}_0 \in \mathcal{I}$, and $\Phi(\mathbf{f}_0) \leq 100m \log m$.

Part B: IPM Progress Using Updates. In Part A of this homework problem, we saw how to find a starting point for minimizing $\Phi(\mathbf{f})$ in Program (2). Now, we would like to understand how to improve the solution. To do this, we will think about Taylor series approximations of Φ .

1. Given $\mathbf{f} \in \mathcal{I}$, define $\mathbf{l} \in \mathbb{R}^{\tilde{E}}$ by $\mathbf{l}(\tilde{e}) = 1/\mathbf{f}(\tilde{e})$ and

$$\mathbf{l}_f(e) = \frac{1}{\min(c(e) - \mathbf{f}(e), c(e) + \mathbf{f}(e))} \quad (3)$$

for all $e \in E$. We define a corresponding diagonal matrix $\mathbf{L}_f \in \mathbb{R}^{\tilde{E} \times \tilde{E}}$ by $\mathbf{L}_f(e, e) = \mathbf{l}_f(e)$ for all $e \in \tilde{E}$. Prove that if $\|\mathbf{L}_f \boldsymbol{\delta}\|_\infty \leq 1/2$ then $\mathbf{f} + \boldsymbol{\delta} \in \mathcal{I}$ and

$$\Phi(\mathbf{f} + \boldsymbol{\delta}) \leq \Phi(\mathbf{f}) + \nabla_{\mathbf{f}} \Phi(\mathbf{f})^\top \boldsymbol{\delta} + 4 \|\mathbf{L}_f \boldsymbol{\delta}\|_1^2.$$

2. Prove that if for some $\kappa > 1$ we have $\|\mathbf{L}_f \boldsymbol{\delta}\|_1 \leq \kappa$ and $\nabla_{\mathbf{f}} \Phi(\mathbf{f})^\top \boldsymbol{\delta} = -1$ then

$$\Phi(\mathbf{f} + \frac{1}{8\kappa^2} \boldsymbol{\delta}) \leq \Phi(\mathbf{f}) - \frac{1}{16\kappa^2}.$$

Part C: The Update. In Part B, we saw a set of sufficient conditions for an update $\boldsymbol{\delta}$ to improve the potential function value. Of course, if we use this to optimization $\Phi(\mathbf{f})$ by moving from \mathbf{f} to $\mathbf{f} + \boldsymbol{\delta}$, then if we require that $\mathbf{B}\mathbf{f} = F(-\mathbf{1}_s + \mathbf{1}_t)$ and $\mathbf{B}(\mathbf{f} + \boldsymbol{\delta}) = F(-\mathbf{1}_s + \mathbf{1}_t)$, then we must have $\mathbf{B}\boldsymbol{\delta} = \mathbf{0}$, i.e. the update should be a circulation. In this part, we will show that in fact a good circulation $\boldsymbol{\delta}$ exists for lowering the value of Φ , and we can phrase the task of finding $\boldsymbol{\delta}$ as an optimization problem. Finally, we will put everything together and argue that if we can find the update circulations, we can compute a maximum flow in a reasonable amount of time.

1. Consider the following optimization program:

$$\begin{aligned} \min_{\boldsymbol{\delta} \in \mathbb{R}^{\tilde{E}}} \quad & \|\mathbf{L}_f \boldsymbol{\delta}\|_1 & (4) \\ \text{s.t.} \quad & \mathbf{B}\boldsymbol{\delta} = \mathbf{0} \\ & \nabla_{\mathbf{f}} \Phi(\mathbf{f})^\top \boldsymbol{\delta} = -1 \end{aligned}$$

Is this a convex optimization program?

2. Prove that the value γ^* of Program (4) is bounded by $\gamma^* \leq O(1)$.

Hint: Consider vectors of the form $\tilde{\boldsymbol{\delta}} = \alpha(\mathbf{f}^ - \mathbf{f})$ for some scalar α , where \mathbf{f}^* is an optimal solution to Program (1).*

3. Prove that there exists an optimal solution δ^* for Program (4) which is supported on a single simple cycle.
4. Suppose you have access to the following subroutines:
 - An algorithm that computes a probabilistic low stretch spanning tree of any graph G with m edges in time $\tilde{O}(m)$. Recall that a probabilistic low stretch tree is a random sub-tree T s.t. for every edge of G , the expected stretch of the edge w.r.t. T is $\tilde{O}(1)$.
 - An algorithm $\text{ROUNDFLOW}(\mathbf{f})$ that given a flow $\mathbf{f} \in \mathcal{I}$ with $\mathbf{B}\mathbf{f} = F(-\mathbf{1}_s + \mathbf{1}_t)$ returns an *integral* flow¹ $\hat{\mathbf{f}}$ with $-\mathbf{c}(e) \leq \hat{\mathbf{f}}(e) \leq \mathbf{c}(e)$, $\hat{\mathbf{f}}(\bar{e}) = 0$ and $\mathbf{B}\hat{\mathbf{f}} = \hat{F}(-\mathbf{1}_s + \mathbf{1}_t)$ where $\hat{F} \geq F - \mathbf{f}(\bar{e}) - 10$. The routine takes time $T_{\text{ROUNDFLOW}} = \tilde{O}(m)$.

Describe a randomized algorithm for solving our undirected maximum flow problem w.h.p. in time $\tilde{O}(m^2)$ using the IPM, low-stretch spanning trees, the ROUNDFLOW algorithms, and Ford-Fulkerson with basic path augmentation (i.e. no blocking flows). You should not need to use other algorithms such as Dinic's max flow algorithm or data structures like Link-Cut trees.

Part D: Stability. In the previous problem, we started analyzing the update Program (4), and we will start to see why this might be possible, by showing that an update δ with small norm $\|\mathbf{L}_f \delta\|_1$ can only cause the lengths $\mathbf{l}_{f+\delta}$ to change significantly in very few entries.

1. Consider an update δ such that $\|\mathbf{L}_f \delta\|_1 \leq 1/2$, and define $\mathbf{s} = \mathbf{L}_f^{-1} \mathbf{l}_{f+\delta}$. Let $\mathcal{U} = \{e : |\mathbf{s}(e) - 1| > 1/m^{o(1)}\}$. Prove that $|\mathcal{U}| \leq m^{o(1)}$.
2. Bonus: Analyze stability of the lengths during a sequence of updates. *Hint: a bucketing scheme may help.*
3. Bonus: Analyze stability of the gradient. What is the right notion of approximation?

Bonus Exercise: Tree data structures

We want to consider a simplified model of our data structures to understand the basics of how we can use a link-cut tree-based data structure to maintain our flow during the IPM iterations.

Assume a min-ratio cycle data structure to support the following operations:

- $\text{INITIALIZE}(G, \hat{\mathbf{g}}^{(0)}, \hat{\mathbf{L}}^{(0)})$: initially, the data structure is given the underlying graph G and the initial approximate gradients and lengths on the edges of G . This also initializes an associated dynamic spanning tree T , whose edges will always be a subset of the edges of G .
- $\text{UPDATE}(\hat{\mathbf{g}}^{(t)}, \hat{\mathbf{L}}^{(t)})$: the t -th update replaces the current gradient and lengths by $\hat{\mathbf{g}}^{(t)}$ and $\hat{\mathbf{L}}^{(t)}$ (assume a sparse representation can be used, and only specify entries that should be updated). This also returns a list of edges E_{add} that need to be added to T and edges E_{delete} that need to be removed from T . It is guaranteed that T will remain a spanning tree after the update. The amortized number of edges of T that change per update is $m^{o(1)}$.
- $\text{QUERY}()$: returns an off-tree edge e (w.r.t. the current tree T) and a scaling $\alpha \in \mathbb{R}$ such that the cycle given by sending α units of flow forward along e and backwards around its tree path in T (approximately) minimizes (4) with respect to $\hat{\mathbf{g}}^{(t)}$ and $\hat{\mathbf{L}}^{(t)}$.
- Assume that if $\hat{\mathbf{g}}^{(t)} \approx \mathbf{g}^{(t)}$ and $\hat{\mathbf{L}}^{(t)} \approx \mathbf{L}^{(t)}$ (for the lengths, this means coordinate-wise $1 \pm 1/m^{o(1)}$ factor approximation), then the tree cycle for edge e associated with T has value $m^{o(1)}$ in Program (4) and still satisfies the gradient condition (this is slightly unrealistic, but very close the truth).

¹An integral flow is a flow with integer entries.

For this exercise, you should ignore gradient maintenance: Just assume the gradient vector is maintained correctly and for free by the data structure.

Assume we have another data structure `MAINTAINTREE` which can support the following operations:

- `INITIALIZE(V, T)` : Given a spanning tree T with edges F on vertex set V , initialize the data structure. The edges should have an orientation (but this is not part of the relevant to the property of being a spanning tree). Maintain a vector $\mathbf{t} \in \mathbb{R}^F$ of real numbers, one for each edge in T . Initially all entries of \mathbf{t} are set to zero. Time is $\tilde{O}(m)$.
- `UPDATE($E_{\text{add}}, E_{\text{delete}}$)` : Remove edges E_{delete} from T and add edges E_{add} to T . The edges must maintain that T is a spanning tree of V . This changes vector \mathbf{t} : new entries are set to zero. Time per edge is $\tilde{O}(1)$.
- `QUERYMAX(u, v)` : Given vertices $u, v \in V$, returns the maximum entry in \mathbf{t} (index and value) on the path between u and v . Time per query is $\tilde{O}(1)$.
- `ADDUNSIGNED(u, v, δ)` : Given vertices $u, v \in V$, add $\delta \in \mathbb{R}$ to every entry of \mathbf{t} corresponding to edges on the u - v path in T . Time for the operation is $\tilde{O}(1)$.
- `ADDSIGNED(u, v, δ)` : Given vertices $u, v \in V$, add $\pm\delta$ to every entry of \mathbf{t} corresponding to edges on the u - v path in T : Add $+\delta$ to the edge if the tree path traverses the edge in the direction of its orientation. Add $-\delta$ to the edge if the tree path traverses the edge in the opposite direction of its orientation. Time for the operation is $\tilde{O}(1)$.
- You may also assume operations for `QUERYSUMSIGNED(u, v)` and `QUERYSUMUNSIGNED(u, v)` : (self-explanatory), but I don't think you should need them. Time for the operation is $\tilde{O}(1)$.

The goals of the exercise:

- Describe how to maintain the solution flow vector $\mathbf{f}^{(t)}$ during the IPM under updates given by tree cycle flows. Overall time should be $m^{1+o(1)}$.
- Describe how to maintain the lengths $\hat{\mathbf{L}}^{(t)} \approx \mathbf{L}^{(t)}$ with a coordinate-wise $1 \pm 1/m^{o(1)}$ factor approximation. Overall time should be $m^{1+o(1)}$.