28 Oct 2024

Approximate Edge-Disjoint Paths (Congestion Minimization)

Input: Graph $G = (V, E)$ (directed)

Capacities $c(e) \in \mathbb{N}$

Source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$

Each s_i must route 1 unit of flow to t_i along a single path.

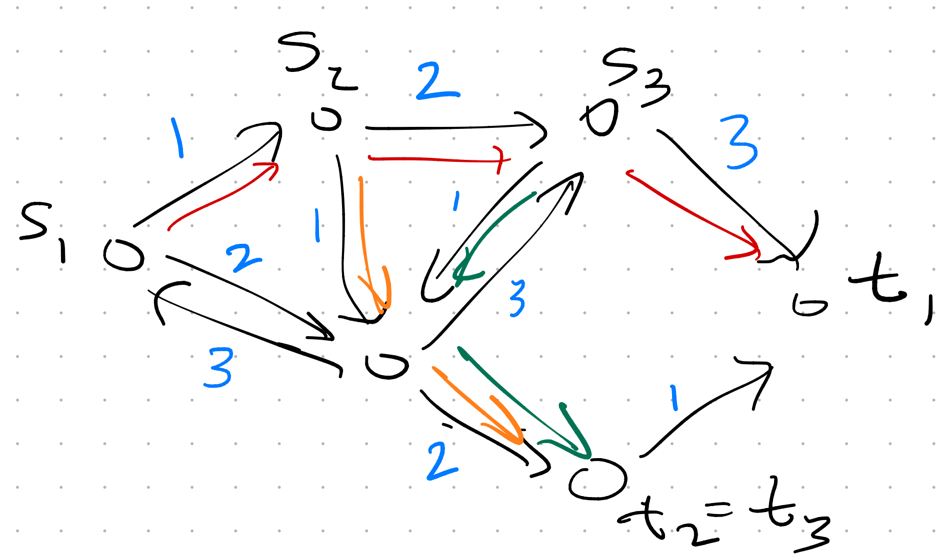

Solution is $Q = (P_1, \ldots, P_k) \in \prod_{i=1}^{k} \mathcal{P}_{(s_i, t_i)}$.

Load on edge e is $n_Q(e) = \# \{ i \mid e \in P_i \}$.

Congestion of e is

$$\mathrm{cong}_Q(e) = \max \left\{ \frac{n_Q(e)}{c(e)}, 1 \right\}.$$

"Minimize the amount by which the network slows down due to its most congested edge."

$$\min_Q \max_e \left\{ \mathrm{cong}_Q(e) \right\}.$$

Plan for congestion minimization:

we (humankind) know how to solve multicommodity flow. It's just the "use only one path per commodity" causing difficulty.

So solve (fractional) MCF and round fractional solution to an integer one.

$$\begin{aligned} \max \quad & \sum_Q y_Q \\ \text{s.t.} \quad & \sum_Q n_Q(e)\, y_Q \le c(e) \quad \forall e \\ & y_Q \ge 0 \quad \forall Q \end{aligned}$$

<u>Lemma</u>. Let OPT_G denote the optimal value of $\min_Q \max_e \{ \mathrm{cong}_Q(e) \}$. Let Λ = opt. of MCF LP.

Then
$$\min\{1, \Lambda\} \cdot OPT_G \ge 1.$$

Proof. $\min\{1, \Lambda\} \cdot \mathrm{OPT}_G \geq 1$ means:

$$\mathrm{OPT}_G \geq 1 \quad \text{and} \quad \Lambda \cdot \mathrm{OPT}_G \geq 1$$

$\forall Q \ \forall e \ \ \mathrm{cong}_Q(e) \geq 1$

$\therefore \min_Q \max_e \{\mathrm{cong}_Q(e)\} \geq 1$

QED

proof below

Given Q^* such that $\max_e \{\mathrm{cong}_{Q^*}(e)\} = \mathrm{OPT}_G$

define LP solution y as

$$y_Q = \begin{cases} \mathrm{OPT}_G^{-1} & \text{if } Q = Q^* \\ 0 & \text{o.w.} \end{cases}$$

Then y_Q is feasible for LP because

$$\sum_Q n_Q(e)\, y_Q = \mathrm{OPT}_G^{-1} \cdot n_{Q^*}(e)$$

$$\leq \left(\mathrm{cong}_{Q^*}(e)\right)^{-1} \cdot n_{Q^*}(e) \leq c(e)$$

And LP objective at y_Q is OPT_G^{-1}.

So $\Lambda \geq \mathrm{OPT}_G^{-1}$, QED.

Suppose given LP solution y^* achieving value Λ.

It means $\Lambda^{-1} y^*$ is a non-neg vector whose coordinates sum to 1.

$\Longrightarrow$ probability distrib $\mathcal{D}(y^*)$ on $\prod_{i=1}^{k} \mathcal{P}(s_i, t_i)$.

Let $\mathcal{D}_i(y^*)$ denote marginal distribution of path P_i when $(P_1, \dots, P_k) = Q$ is sampled from $\mathcal{D}(y^*)$.

Raghavan–Thompson: Output $\hat{Q} = (\hat{P}_1, \dots, \hat{P}_k)$ by drawing each $\hat{P}_i$ independently from $\mathcal{D}_i(y^*)$.

Analysis. Focus on one edge e.

Try using Chernoff bound to bound from above $\Pr(e \text{ is too congested})$.

The random variable we will analyze is $\underbrace{\dfrac{\min\{1, \Lambda\}}{c_e}}_{\text{normalizing factor}} \cdot \underbrace{n_{\hat{Q}}(e)}_{\text{sum of indep. RV's}}$.

Let $$X_i = \begin{cases} \frac{\min\{1,\Lambda\}}{c(e)} & \text{if } e \in \hat{P}_i \\ 0 & \text{o.w.} \end{cases}$$

Observe

1. $$X_1 + \dots + X_k = \frac{\min\{1,\Lambda\}}{c_e} \cdot n_{\hat{Q}}(e)$$

2. $$0 \le X_i \le 1$$

3. $$\mathbb{E}[X_1 + \dots + X_k]$$
$$= \frac{\min\{1,\Lambda\}}{c(e)} \cdot \sum_{i=1}^{k} \Pr(e \in \hat{P}_i)$$

$(P_1, \dots, P_k) \sim \mathcal{D}(y^*)$

$$= \frac{\min\{1,\Lambda\}}{c(e)} \cdot \sum_{i=1}^{k} \Pr(e \in P_i)$$

$$= \frac{\min\{1,\Lambda\}}{c(e)} \cdot \mathbb{E}_{\mathcal{D}(y^*)}[n_Q(e)].$$

$$= \frac{\min\{1,\Lambda\}}{c(e)} \cdot \Lambda^{-1} \cdot \sum_Q n_Q(e) y^*_Q$$

$$\leq \frac{\min\{1,\Lambda\}}{\cancel{c(e)}} \cdot \Lambda^{-1} \cdot \cancel{c(e)}$$

$$\leq \min\{1,\Lambda\} \cdot \Lambda^{-1} \leq 1 .$$

From Chernoff:

$$\Pr\left(X_1 + \dots + X_k > \frac{3 \log(2m)}{\ln \log(2m)}\right) < \frac{1}{2m} .$$

Union bound:

$$\Pr\left(\exists \text{ edge with } \frac{\min(1,\Lambda)}{c(e)} \cdot n_{\hat{Q}}(e) > \frac{3 \log(2m)}{\ln \log(2m)}\right)$$

$$< \frac{1}{2} .$$

With prob $\geq \frac{1}{2}$, all edges obey

$$\min(1,\Lambda) \cdot \frac{n_{\hat{Q}}(e)}{c(e)} \leq \frac{3 \log(2m)}{\ln \log(2m)} .$$

Also

$$\min(1,\Lambda) \cdot 1 \leq \frac{3 \log(2m)}{\ln \log(2m)} .$$

$$\forall e \quad \min\left(1, \Lambda\right) \cdot \mathrm{cong}_{\hat{Q}}(e) \leq \frac{3 \log(2m)}{\ln \log(2m)}$$

w. prob. $\geq \frac{1}{2}$

$$\max_e \left\{ \mathrm{cong}_{\hat{Q}}(e) \right\} \leq \frac{3 \log(2m)}{\ln \log(2m)} \cdot \frac{1}{\min(1, \Lambda)}$$

$$\mathrm{OPT}_G \geq \frac{1}{\min(1, \Lambda)}$$

So max-congestion of $\hat{Q}$ exceeds OPT_G by $\leq \frac{3 \log(2m)}{\ln \log(2m)}$.