Today

- Probabilistic Subtree Embeddings
- Min-Congestion (Oblivious) routing
- Tree-Based Oblivious Routing
- $O(\log n)$-competitive

An <u>α-distortion probabilistic tree embedding</u> of (V,d) is a distribution over trees Υ containing V s.t. $\forall u,v \in V$

1) $d(u,v) \le d_T(u,v)$ $\quad \forall T \in \Upsilon$

2) $\mathbb{E}_{T\sim\Upsilon}[d_T(u,v)] \le \alpha \cdot d(u,v)$

$\forall$ over trees w/ $\Pr > 0$

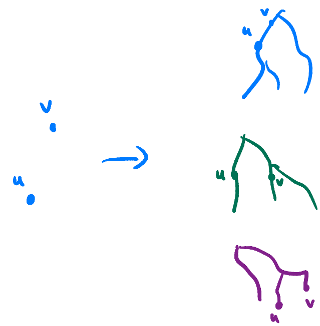

"FRT"

<u>Theorem:</u> every metric has a (Poly-time sampleable) $O(\log n)$-distortion tree embedding

Given tree $T=(V,E)$, let $T(u,v)$ be the unique (simple) $u \rightsquigarrow v$ path in T

let $T\{e\}$ be $S \subseteq V$ s.t. $\delta(e) = \{e\}$

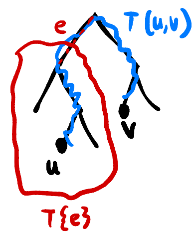

Dual LP

$\max \langle c, x\rangle$ s.t. $Ax \le b$ $\quad$ ($\min \langle -c, x\rangle$)

(Primal)

$\min \langle b, \lambda\rangle$ s.t. $A^T\lambda = c^T$, $\lambda \ge 0$ $\quad$ ($\max \langle -b, \lambda\rangle$)

(Dual)

$D' \le P'$ $\qquad P', D' = -P, -D$

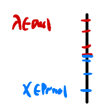

<u>Strong Duality:</u> If Primal feasible + bounded then $P = D$

Will use Subtree embeddings's result from last Class also works but takes longer to describe

Probabilistic Subtree Embeddings

Fact: Given edge-weighted graph $G=(V,E,w)$, $\exists$ a distribution Υ over spanning trees of G s.t. $\mathbb{E}_{T\sim\Upsilon}[d_T(u,v)] \le \tilde{O}(\log n)\cdot d_G(u,v)$ $\forall u,v\in V$ where d_T gives shortest paths $\in T$ w.r.t. w

recall

Note: $d_G(u,v)\le d_T(u,v)$ $\forall T\in\Upsilon$ b/c T is a spanning tree

Stronger than result from last class (up to loglog) b/c (1) spanning trees and (2) T on same vertex set as G

PTE Corollary: Given edge-weighted graph $G=(V,E,w)$, $\exists$ a spanning tree $T\subseteq G$ s.t.

$$\sum_{\{u,v\}\in E} d_T(u,v) \le \tilde{O}(\log n)\cdot \sum_{\{u,v\}\in E} w(\{u,v\})$$

Note: this says $\exists$ a spanning tree that, on average, distorts edges $\le \tilde{O}(\log n)$

Proof basically by LoE

Let Υ be distribution from fact so by LoE have

$$\mathbb{E}_{T\sim\Upsilon}\left[\sum_{\{u,v\}\in E} d_T(u,v)\right] \le \tilde{O}(\log n)\cdot \sum_{\{u,v\}\in E} w(\{u,v\})$$

So by averaging, must $\exists$ at least 1 spanning tree T s.t.

$$\sum_{\{u,v\}\in E} d_T(u,v) \le \tilde{O}(\log n)\cdot \sum_{\{u,v\}\in E} w(\{u,v\})$$

So far: all graphs are $\approx$ trees wrt cuts and distances

Today: all graphs are $\approx$ trees wrt routing

Given graph $G=(V,E)$, a <u>demand</u> is a function $D: V\times V \to \{0,1\}$

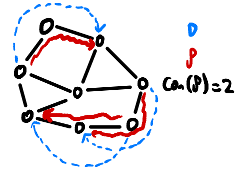

Paths $\mathcal{P}$ <u>route</u> unit demand D if $\forall u,v$ s.t. $D(u,v)=1$, $\exists$ a $u \rightsquigarrow v$ path $\in \mathcal{P}$

The <u>congestion</u> of e wrt $\mathcal{P}$ is $\text{Con}_{\mathcal{P}}(e) := |\{P \in \mathcal{P} : e \in P\}|$ and of $\mathcal{P}$ is $\text{Con}(\mathcal{P}) := \max_e \text{Con}_{\mathcal{P}}(e)$

<u>Min-Congestion Routing Problem:</u> given demand D, find paths $\mathcal{P}_D$ routing D, minimizing $\text{Con}(\mathcal{P})$

Let OPT_D be $\min_{\mathcal{P} : \mathcal{P} \text{ routes } D} \text{Con}(\mathcal{P})$

notice not same as $\mathbb{E}[\text{Con}(\mathcal{P}_D)]$

Say randomized algorithm is α-competitive if $\mathbb{E}\left[\text{Con}_{\mathcal{P}_D}(e)\right] \le \alpha \cdot \text{OPT}_D \quad \forall D, e \in E$

1-Competitive Randomized Algorithm

Variable $x_P \sim$ Probability $P=(s,\dots,t)$ sampled as $s \rightsquigarrow t$ path in $\mathcal{P}_D$

Min z s.t.

$$z \ge \sum_{P : e \in P} x_P \quad \forall e \in E$$

$$\sum_{P=(s,\dots,t)} x_P = 1 \quad \forall s,t \text{ s.t. } D(s,t)=1$$

$$x_P \in [0,1] \quad \forall P$$

Easy to see an LP relaxation of IP and $\exists$ poly-time separation oracle so poly-time solvable

Problem w/ this approach: have to resolve an LP every time a new demand

Solution: oblivious algorithms

Say randomized algorithm is <u>oblivious</u> if $\forall s,t \in V$ $\exists$ distribution $\mathcal{P}_{st}$ over $s \rightsquigarrow t$ paths s.t.

$$\mathcal{P}_D = \{P_{st} \sim \mathcal{P}_{st} : D(s,t)=1\}$$

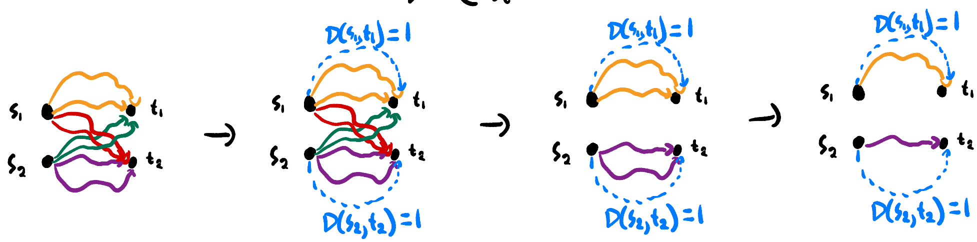

Can see $\exists$ n-node graphs s.t. any deterministic oblivious algorithm $\Omega(\sqrt{n})$-competitive (see later in notes), hence motivation for randomness to get poly-log competitiveness

<u>Theorem:</u> every graph has an $\tilde{O}(\log n)$-competitive randomized oblivious routing algorithm

Tree-Based Oblivious Routing

Say oblivious routing algorithm is <u>tree-based</u> if $\exists$ a distribution over spanning trees Υ s.t.

$$\mathcal{P}_{st} := T(s,t) \text{ for } T \sim \Upsilon$$

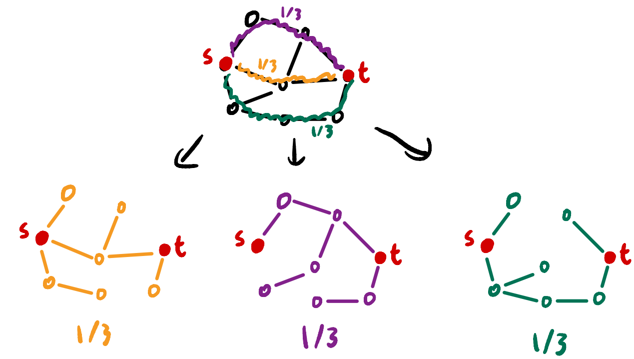

May seem restrictive but very nice to analyze as per the following claim

Given spanning tree $T \subseteq G$, let the <u>load on e</u> be $L_T(e) := |\delta_G(T\{e\})|$

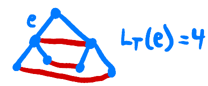

Given distribution Υ over spanning trees, let the load on e be $L_\Upsilon(e) := \mathbb{E}_{T\sim\Upsilon}[L_T(e)]$

and let $L(\Upsilon) := \max_e L_\Upsilon(e)$

<u>Claim:</u> any tree-based oblivious routing scheme w/ tree distribution Υ is $L(\Upsilon)$-competitive

Fix a demand D, a $T \in \Upsilon$ and $e \in T$, let $D(T\{e\}) := \sum_{u \in T\{e\}} \sum_{v \notin T\{e\}} D(u,v) + D(v,u)$

Have $OPT_D \geq D(T\{e\}) / L_T(e)$ b/c routing D requires $\geq D(T\{e\})$ paths to cross $T\{e\}$

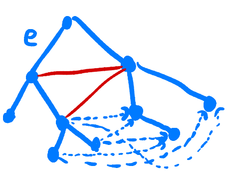

$$\left.\begin{array}{l} L_T(e) = 3 \\ D(T\{e\}) = 6 \end{array}\right] \rightarrow OPT_D \geq 2$$

$E \setminus T$

So $\quad Con(\mathcal{P}_D(e)) = \sum_{\substack{T\in\Upsilon \\ \text{s.t. } e\in T}} Pr(T) \cdot D(T\{e\}) \leq \sum_{\substack{T\in\Upsilon \\ \text{s.t. } e\in T}} Pr(T) \cdot L_T(e) \cdot OPT_D = OPT_D \cdot L_\Upsilon(e)$

So $\forall e, D$ have $Con(\mathcal{P}_D(e)) \leq L_\Upsilon(e) \cdot OPT_D \leq L(\Upsilon) \cdot OPT_D$ so get $L(\Upsilon)$-competitive

So now showing theorem just requires Υ w/ load $\leq O(\log n)$

Finding Low Load Spanning Tree Distributions

Claim: Given graph $G=(V,E)$, $\exists$ a distribution Υ over spanning trees s.t. $L(\Upsilon) \leq \tilde{O}(\log n)$

Can exactly capture problem of finding Υ minimizing $L(\Upsilon)$ w/ an LP as follows

Variable X_T for $\Pr_{\Upsilon}(T)$ and z for $L(\Upsilon)$

$$\min z \text{ s.t.}$$

$$\sum_T X_T = 1$$

$$\underbrace{\sum_{T \ni e} X_T \cdot L_T(e)}_{L_\Upsilon(e)} \leq z \quad \forall e \in E$$

$$X_T \geq 0 \quad \forall \text{ STs } T$$

Can solve in poly-time via ellipsoid + MWU so all algorithmic

Not clear why $\exists$ good solution; show by taking dual; see derivation later

Variable w_e $\forall e \in E$ and Γ

$$\max \Gamma \text{ s.t.}$$

$$\sum_{\{u,v\} \in E} d_T(u,v) \geq \Gamma \cdot \sum_{e \in E} w_e \quad \forall \text{ STs } T$$

length of $T(u,v)$ w/ weights w

$$w_e \geq 0 \quad \forall e$$

Let P, D be optimal primal, dual values

WTS $P \leq \tilde{O}(\log n)$

By strong duality $P = D$ so wts $D \leq \tilde{O}(\log n)$

Consider any feasible dual solution (w, D)

By PTE corollary, $\exists$ spanning tree T s.t. $\sum_{\{u,v\} \in E} d_T(u,v) \leq \tilde{O}(\log n) \cdot \sum_e w_e$

By corresponding constraint $\in$ dual LP for T, know $D \cdot \sum_e w_e \leq \sum_{\{u,v\} \in E} d_T(u,v) \leq \tilde{O}(\log n) \cdot \sum_e w_e$

So $D \leq \tilde{O}(\log n)$

$\tilde{O}(\log n)$-competitive OR follows by 2 claims

Derivation of Dual

Primal

Variable x_T for $\Pr_{\gamma}(T)$ and z for $L(\gamma)$

Min z s.t.

$$\sum_T x_T = 1$$

$$\underbrace{\sum_{T \ni e} x_T \cdot L_T(e)}_{L_\gamma(e)} \le z \quad \forall e \in E$$

$$x_T \ge 0 \quad \forall \text{ STs } T$$

$\longleftrightarrow$

Max $-z$ s.t.

$$\sum_T x_T \le 1 \quad (\alpha)$$

$$-\sum_T x_T \le -1 \quad (\beta)$$

$$\sum_{T \ni e} x_T \cdot L_T(e) - z \le 0 \quad \forall e \in E \quad (w_e)$$

$$-x_T \le 0 \quad \forall \text{ STs } T \quad (y_T)$$

Dual

Min $\alpha - \beta$ s.t.

$$-\sum_e w_e = -1 \quad (z)$$

$$\alpha - \beta - y_T + \sum_{e \in T} w_e \cdot L_T(e) = 0 \quad (x_T) \quad \forall \text{ STs } T$$

$$\alpha, \beta, y \ge 0$$

$\Gamma = \beta - \alpha$,
$y_T \ge 0$ but
o/w doesn't matter

$\Leftrightarrow$

Max Γ s.t.

$$\sum_e w_e = 1$$

$$\sum_{e \in T} w_e L_T(e) \ge \Gamma \quad \forall \text{ STs } T$$

$$w_e \ge 0 \quad \forall e$$

But $\sum_{e \in T} w_e L_T(e) = \sum_{e \in T} w_e \sum_{\{u,v\} \in E} \mathbb{1}\left(\{u,v\} \in \delta(T \setminus \{e\})\right) = \sum_{\{u,v\} \in E} \sum_{e \in T(u,v)} w_e = \sum_{\{u,v\} \in E} d_T(u,v)$

Switch $\sum$ order:

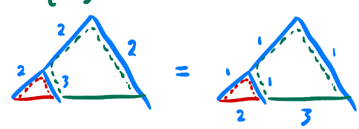

So final dual has variable $w_e \; \forall e \in E$ and Γ and is

Max Γ s.t.

$$\sum_e w_e = 1$$

$$\sum_{\{u,v\} \in E} d_T(u,v) \ge \Gamma \quad \forall \text{ STs } T$$

$$w_e \ge 0 \quad \forall e \in E$$

LHS feasible for RHS

$w \leftarrow \frac{1}{\sum_e w_e} \cdot w$ feasible w/ same Γ

Max Γ s.t.

$$\sum_{\{u,v\} \in E} d_T(u,v) \ge \Gamma \cdot \sum_{e \in E} w_e \quad \forall \text{ STs } T$$

$$w_e \ge 0 \quad \forall e \in E$$

$\Omega(\sqrt{n})$-Competitiveness of Deterministic Oblivious Routing

Notice a deterministic algorithm is oblivious if $\forall s,t \in V$ $\exists$ $s \rightsquigarrow t$ path P s.t. $\mathcal{P}_D = \{P_{st} : D(s,t)=1\}$ $\forall D$

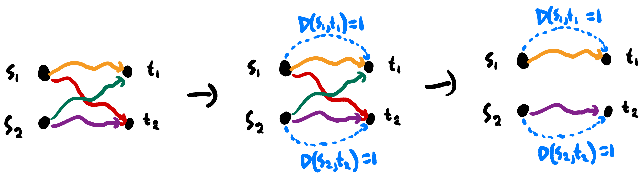

Consider the following graph

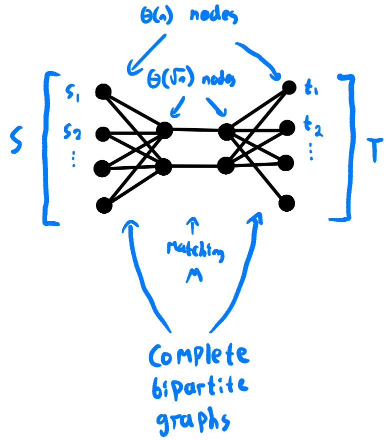

Fix the deterministic oblivious routing algorithm and let P_{st} be its $s \rightsquigarrow t$ path

Let $P_i := P_{s_i t_i}$

For $e \in M$, let $\mathcal{P}_e := \{P_i : e \in P_i\}$ and let $D_e(s,t) := \begin{cases} 1 & \text{if } s_i = s, t_i = t, e \in P_i \\ 0 & \text{o/w} \end{cases}$ be routed demand

So we have $\text{con}(\mathcal{P}_{D_e}) \geq |\mathcal{P}_e|$

By averaging, $\exists$ $\bar{e}$ s.t. $|\mathcal{P}_{\bar{e}}| \geq \Omega(\sqrt{n})$

For simplicity suppose $|\mathcal{P}_{\bar{e}}| = \Theta(\sqrt{n})$ (if not, just drop paths from $\mathcal{P}_{\bar{e}}$)

So $\text{con}(\mathcal{P}_{D_{\bar{e}}}) \geq \Omega(\sqrt{n})$

OTOH $\forall$ D that only send from S to T and $\sum_{u,v} D(u,v) \leq O(\sqrt{n})$,

$\exists$ $\mathcal{P}_D$ routing D w/ $\text{con}(\mathcal{P}_D) = O(1)$

Thus, the algorithm is $\Omega(\sqrt{n})$-competitive