Chapter 2
Multicommodity Routing

Advanced Algorithms

SS 2019

Fabian Kuhn

The Multicommodity Flow Problem

Given:

- Directed graph $G = (V, E)$, each edge $e \in E$ has a capacity $c_e > 0$
- $k \geq 1$ source-destination pairs (s_i, t_i) with demand $d_i > 0$
  - these are the commodities

Goal:

- For each $i \in \{1, \dots, k\}$, compute an **s_i-t_i flow $f_i : E \to \mathbb{R}_{\geq 0}$** of value 1
  - Flow f_i needs to satisfy the usual flow constraints:
    - flow conservation for $v \notin \{s_i, t_i\}$
    - net flow leaving s_i has value 1, net flow entering t_i has value 1
- Minimize maximum edge congestion λ:

$$\boldsymbol{\lambda := \max_{e \in E} \frac{1}{c_e} \cdot \sum_{i=1}^{k} d_i \cdot f_i(e)}$$

Example: Multicommodity Flow

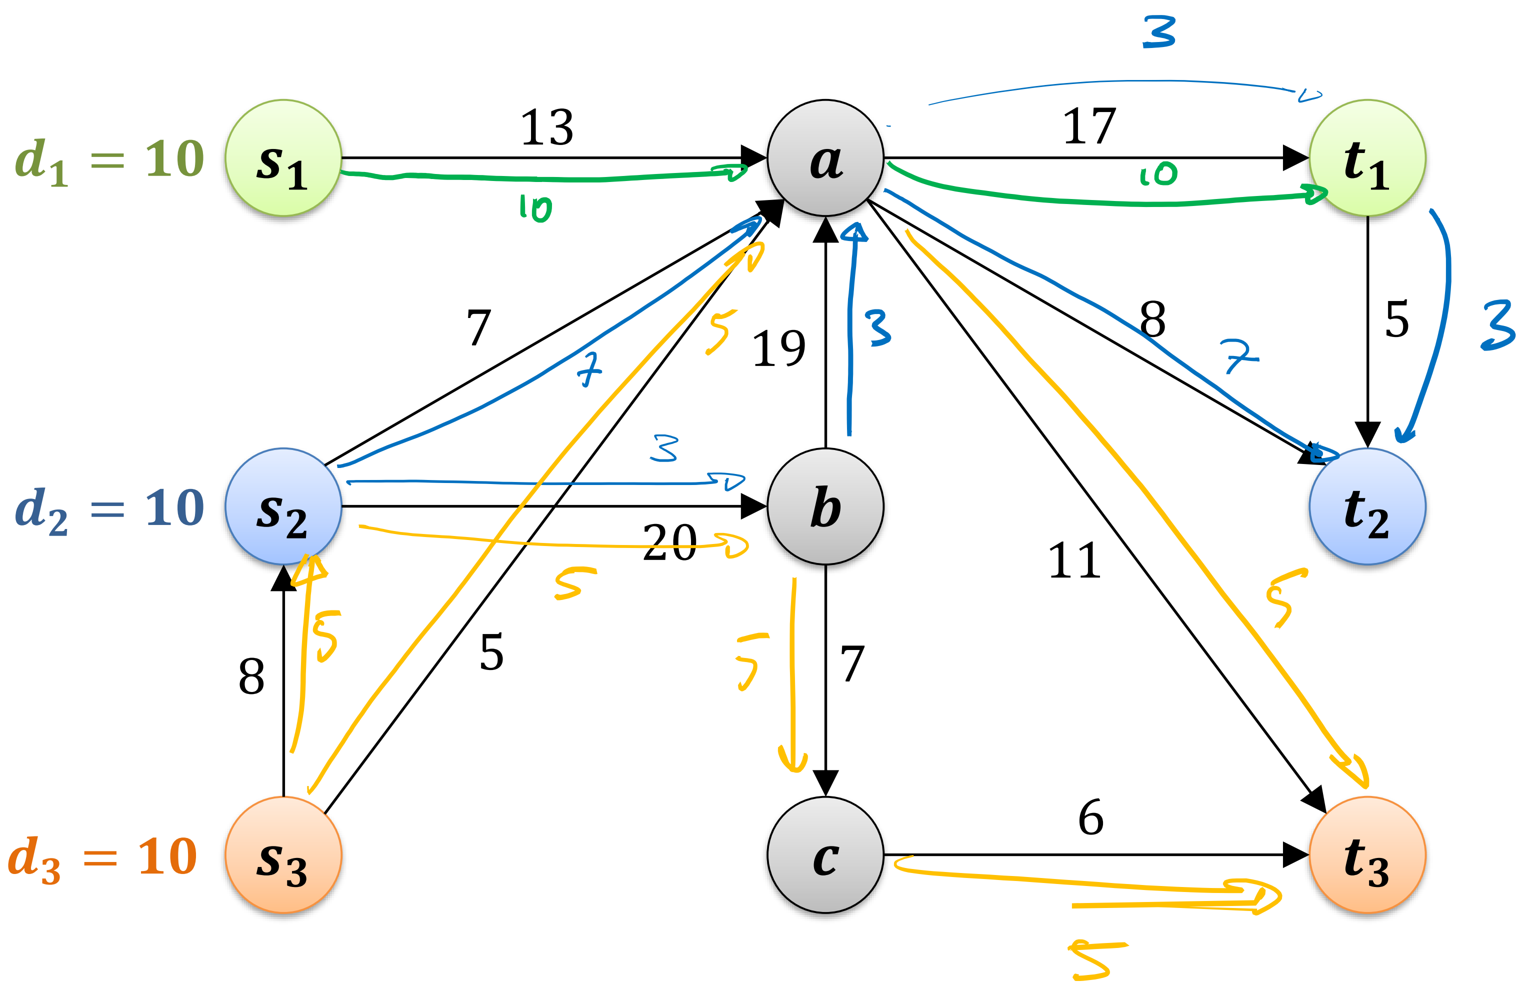

Multicommodity Flow as an LP $\quad G=(V,E)$

For $v \in V$: $\quad in(v)$: ~~edges into~~ v, $\quad out(v)$: edges out of v

---

$\min \lambda$

$\forall i \in \{1,\dots,k\}: \quad \forall v \notin \{s_i, t_i\} \quad \sum_{e \in in(v)} f_i(e) = \sum_{e \in out(v)} f_i(e)$

$$\sum_{e \in out(s_i)} f_i(e) - \sum_{e \in in(s_i)} f_i(e) = 1$$

$$\sum_{e \in in(t_i)} f_i(e) - \sum_{e \in out(t_i)} f_i(e) = 1$$

$\forall e \in E: \quad \sum_{i=1}^{k} d_i f_i(e) \le \lambda \cdot c_e, \quad \lambda \ge 0, \quad f_i(e) \ge 0 \quad (\forall i, e)$

The Multicommodity Routing Problem

Goal:

- For each $i \in \{1, \dots, k\}$, compute an s_i-t_i path P_i
- Minimize maximum edge congestion λ:

$$\lambda := \max_{e \in E} \frac{1}{c_e} \cdot \sum_{i : e \in P_i} d_i$$

- The same as the multicommodity flow problem, however, each of the flows has to be routed on a single path

Remark: For the routing problem, we assume that for a constant $\alpha > 0$,

$$\forall i \in \{1, \dots, k\}, \forall e \in E : \quad d_i \leq \alpha \cdot c_e$$

can be phrased ~~as an~~ ILP:

require $f_i(e) \in \{0, 1\}$

Example: Multicommodity Routing

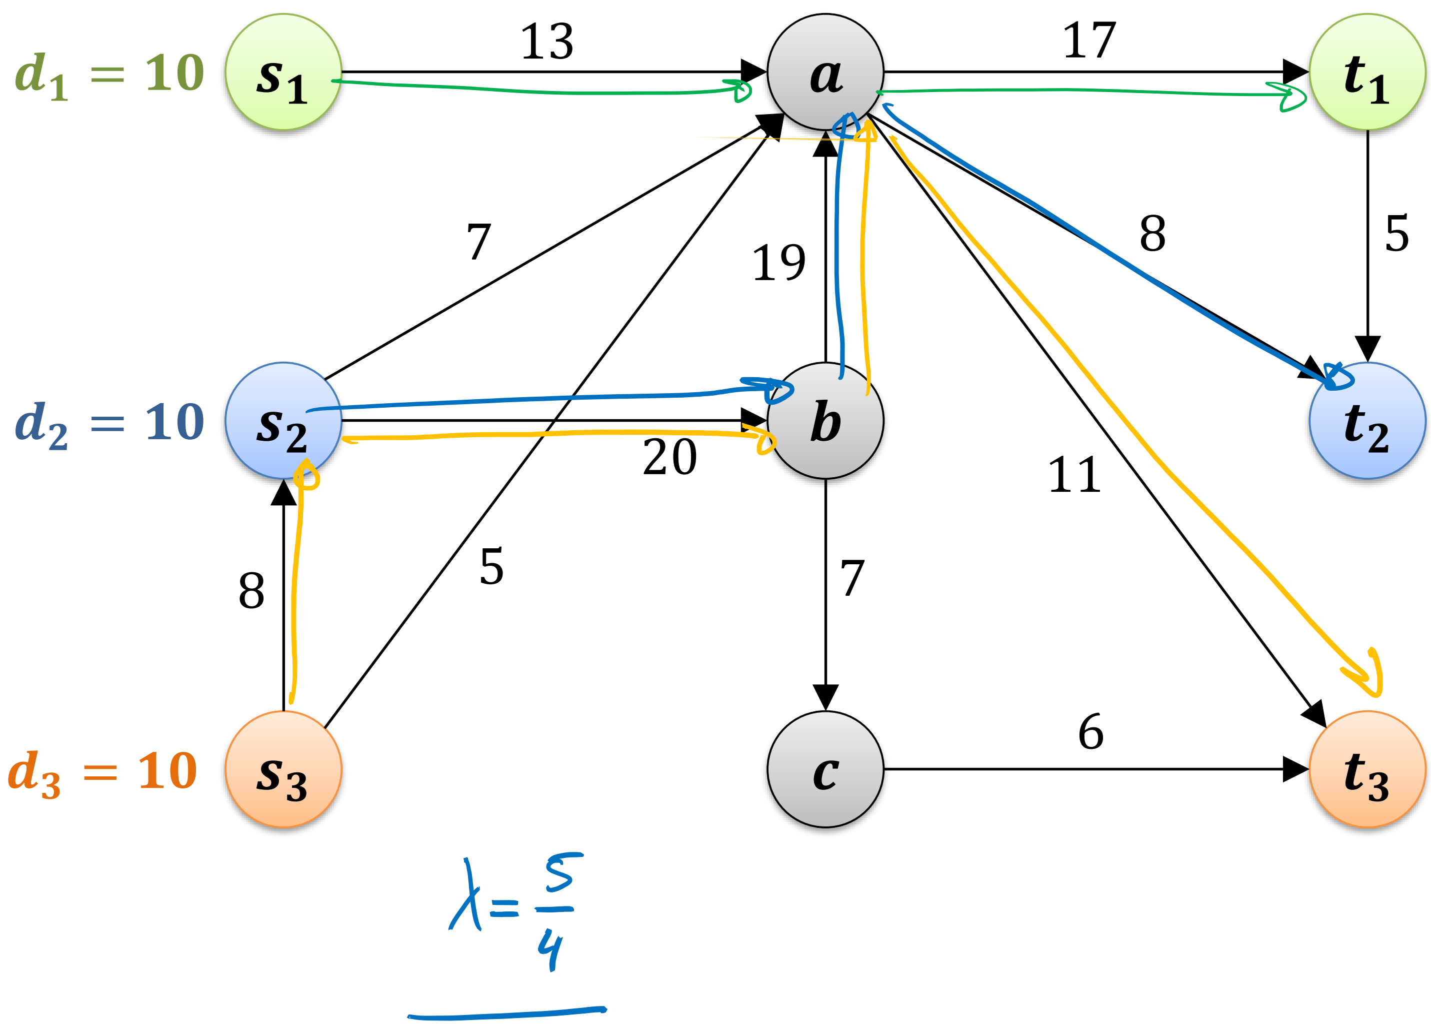

Rounding the Multicommodity Flow LP

Let's start with a simpler problem:

- For each of the k source-destination pairs (s_i, t_i), we are given a collection $\mathcal{P}_i = \{P_{i,1}, \dots, P_{i,\ell_i}\}$ of s_i-t_i paths
- s_i and t_i have to be connected by one of the paths in $\mathcal{P}_i$

Integer Linear Program:

Variables $x_{i,j} \in \{0,1\}$ $x_{i,j} = 1 \iff$ choose path $P_{i,j} \in \mathcal{P}_i$ for comm. i

min λ

$\forall i \in \{1, \dots, k\}: \quad \sum_{j=1}^{\ell_i} x_{i,j} = 1$

$\forall e \in E: \quad \sum_{i=1}^{k} \sum_{j: e \in P_{i,j}} x_{i,j} \cdot d_i \leq \lambda \cdot c_e$

$\lambda \geq 0$

Rounding the Multicommodity Flow LP

Let's start with a simpler problem:

- For each of the k source-destination pairs (s_i, t_i), we are given a collection $\mathcal{P}_i = \{P_{i,1}, \dots, P_{i,\ell_i}\}$ of s_i-t_i paths
- s_i and t_i have to be connected by one of the paths in $\mathcal{P}_i$

LP Relaxation:

relax condition $x_{i,j} \in \{0,1\} \rightarrow x_{i,j} \in [0,1]$

assume that λ^* opt. LP solution

Rounding the Multicommodity Flow LP

- For each of the k source-destination pairs (s_i, t_i), we are given a collection $\mathcal{P}_i = \{P_{i,1}, \dots, P_{i,\ell_i}\}$ of s_i-t_i paths

Randomized Rounding:

For each $i \in \{1, \dots, k\}$: value $x_{i,j} \in [0,1]$ for path $P_{i,j} \in \mathcal{P}_i$

$$\sum x_{i,j} = 1$$

→ pick path from $\mathcal{P}_i$ according to distr. given by $x_{i,j}$

rand. var. $X_{i,j} = 1 \iff$ rand. rounding picks path $P_{i,j}$ $\quad \mathbb{E}[X_{i,j}] = x_{i,j}$

$Y_{e,i}$: contribution of comm. i to edge e

$Y_{e,i} = d_i$ ⟷ edge e is in the path chosen for comm. i by rounding

congestion of edge e → $Y_e = \sum_{i=1}^{k} Y_{e,i}$ $\quad \mathbb{E}[Y_e] \le \lambda^*$

Rounding the Multicommodity Flow LP

- For each of the k source-destination pairs (s_i, t_i), we are given a collection $\mathcal{P}_i = \{P_{i,1}, \dots, P_{i,\ell_i}\}$ of s_i-t_i paths

Randomized Rounding:

- Random variables Y_e for all $e \in E$:

$$Y_e := \sum_{i=1}^{k} Y_{e,i}\,, \qquad \text{where } Y_{e,i} := \frac{d_i}{c_e} \cdot \sum_{j: e \in P_{i,j}} X_{i,j}$$

$$\mathbb{E}[Y_e] = \sum_{i=1}^{k} \frac{d_i}{c_e} \sum_{j: e \in P_{i,j}} x_{i,j} \le \lambda^*$$

$$\mathbb{E}[Y_{e,i}] = \frac{d_i}{c_e} \cdot \sum_{j: e \in P_{i,j}} x_{i,j}$$

to show that max. edge cong. $\le \beta \cdot \lambda^*$

$\hookrightarrow \forall e \in E: \; Y_e \le \beta \cdot \lambda^* \quad \rightarrow \mathbb{P}(Y_e > \beta \cdot \lambda^*) = ? \quad \le \alpha$

$Y_{e,i}$ are indep. for diff. i, $\; Y_{e,i} \in \{0, \frac{d_i}{c_e}\}$

Chernoff Bounds

Theorem: Let $X_1, \dots, X_n$ be independent random variables and let $a_1, \dots, a_n$ be positive numbers such that $0 < a_i \leq A$ for all i. Assume that each variable X_i can take values 0 or a_i such that $\mathbb{P}(X_i = a_i) = p_i$. Define $X := X_1 + \dots + X_n$ and let μ be chosen such that $\mu \geq \mathbb{E}[X] = \sum_{i=1}^{n} p_i \cdot a_i$.
Then, for all $\varepsilon > 0$, it holds that

$$\mathbb{P}(X \geq (1+\varepsilon)\cdot\mu) \leq \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\mu/A} \leq e^{-\frac{\varepsilon^2}{2+\varepsilon}\cdot\frac{\mu}{A}}$$

$$\mathbb{P}(X \leq (1-\varepsilon)\cdot\mu) \leq \left(\frac{e^{-\varepsilon}}{(1-\varepsilon)^{1-\varepsilon}}\right)^{\mu/A} \leq e^{-\frac{\varepsilon^2}{2A}\cdot\mu}$$

need to assume that $\mu \leq \mathbb{E}[X]$

Rounding the Multicommodity Flow LP

- For each of the k source-destination pairs (s_i, t_i), we are given a collection $\mathcal{P}_i = \{P_{i,1}, \dots, P_{i,\ell_i}\}$ of s_i-t_i paths

Randomized Rounding:

- Random variables Y_e for all $e \in E$:

$$Y_e := \sum_{i=1}^{k} Y_{e,i}\,, \qquad \text{where } Y_{e,i} := \frac{d_i}{c_e} \cdot \sum_{j:e\in P_{i,j}} X_{i,j}$$

  - $Y_{e,i}$ can take values $\frac{d_i}{c_e} \leq \alpha$ or 0, $\mathbb{E}[Y_e] \leq \lambda^*$
  - $Y_{e,i}$ are independent for different i

- **Chernoff Bound:**

$$\forall e \in E:\ \mathbb{P}\big(Y_e \geq (1+\varepsilon)\cdot\lambda^*\big) \leq \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\lambda^*/\alpha}$$

Rounding the Multicommodity Flow LP

Theorem: After randomized rounding, with probability at least $1 - {}^{1}/_{n}$, the maximum edge congestion λ is upper bounded by

$$\lambda \le O\left(\frac{\log n}{\log\log n}\right) \cdot \max\{1, \lambda^*\}$$

Proof: set: $\lambda' := \max\{1, \lambda^*\}$

$$\forall e \in E:\ \mathbb{P}(Y_e \ge (1+\varepsilon)\cdot \lambda^*) \le \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\lambda^*/\alpha}$$

$\boxed{\lambda = \max_e Y_e}$

$$\mathbb{P}(Y_e \ge (1+\varepsilon)\cdot\lambda') \overset{!}{\le} \frac{1}{n^3}$$

$$\mathbb{P}(Y_e \ge (1+\varepsilon)\lambda') \le \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\lambda'/\alpha} = e^{\underbrace{\frac{\lambda'}{\alpha}(\varepsilon - (1+\varepsilon)\ln(1+\varepsilon))}_{\le -3\cdot\ln n}}$$

choose $1+\varepsilon = c\cdot\frac{\ln n}{\ln\ln n}$

$$\frac{\lambda'}{\alpha}(\varepsilon - (1+\varepsilon)\ln(1+\varepsilon)) \le \frac{\lambda'}{\alpha}\cdot\Big(c\frac{\ln n}{\ln\ln n} - \overbrace{c\frac{\ln n}{\ln\ln n}\cdot \underbrace{\ln\Big(c\frac{\ln n}{\ln\ln n}\Big)}_{\Theta(\log\log n)}}^{(c-o(1))\cdot\ln(n)}\Big)$$

$$\mathbb{P}(\text{some } Y_e \ge (1+\varepsilon)\lambda') \le |E|\cdot\frac{1}{n^3} \le \frac{1}{n}$$

then: $\lambda \le (1+\varepsilon)\lambda'$ with prob. $1 - \frac{1}{n}$

Proofing the Chernoff Bound

- $X_i \in \{0, a_i\}, \quad 0 < a_i \leq A, \quad \mathbb{P}(X_i = a_i) = p_i,$
- $X = X_1 + \cdots + X_n, \quad \mu \geq \mathbb{E}[X] = \sum_{i=1}^{n} a_i \cdot p_i$

Chernoff Bound:

$$\mathbb{P}(X \geq (1+\varepsilon)\cdot\mu) \leq \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\mu/A}$$

Let's start with some useful tools:

- Markov inequality:

$$\text{For } Z \geq 0: \mathbb{P}(Z \geq z) \leq \mathbb{E}[Z]/z$$

- Linearity of expectation:

$$\mathbb{E}[X+Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$

- For independent rand. var.:

$$\mathbb{E}[X \cdot Y] = \mathbb{E}[X] \cdot \mathbb{E}[Y]$$

- For all $x \in \mathbb{R}$:

$$(1+x) \leq e^{x}$$

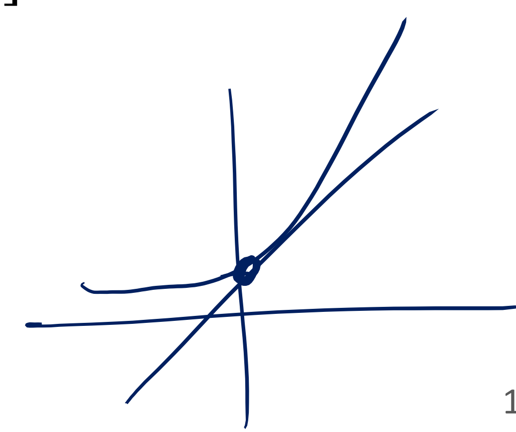

Proofing the Chernoff Bound

e^{tx}

- $X_i \in \{0, a_i\}, \quad 0 < a_i \le A, \quad \mathbb{P}(X_i = a_i) = p_i,$
- $X = X_1 + \dots + X_n, \quad \mu \ge \mathbb{E}[X] = \sum_{i=1}^{n} a_i \cdot p_i$

$1 + x \le e^x$

assume $a_i \ge 1$

Chernoff Bound:

$$\mathbb{P}(X \ge (1+\varepsilon)\cdot\mu) \le \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\mu/A}$$

Proof:

$$\mathbb{P}(X \ge (1+\varepsilon)\mu) \overset{\forall t>0}{=} \mathbb{P}\left(e^{tX} \ge e^{t(1+\varepsilon)\mu}\right) \overset{\text{Markov}}{\le} \frac{\mathbb{E}[e^{tX}]}{e^{t(1+\varepsilon)\mu}}$$

$$\mathbb{E}[e^{tX}] = \mathbb{E}\left[e^{t\sum_{i=1}^{n} X_i}\right] = \mathbb{E}\left[\prod_{i=1}^{n} e^{tX_i}\right] \overset{\text{indep. of } X_i}{=} \prod_{i=1}^{n} \mathbb{E}\left[e^{tX_i}\right] = \prod_{i=1}^{n}\left(p_i e^{ta_i} + (1-p_i)\right)$$

$$= \prod_{i=1}^{n}\left(1 + \underbrace{p_i(e^{ta_i}-1)}_{x}\right) \le \prod_{i=1}^{n} e^{p_i(e^{ta_i}-1)} \underset{(a_i \le A)}{\le} e^{\sum p_i (e^{tA}-1)} \underset{(a_i \le A)}{\le} e^{(e^{tA}-1)\mu/A}$$

$\sum p_i \le \mu/A$

Proofing the Chernoff Bound

- $X_i \in \{0, a_i\}, \quad 0 < a_i \le A, \quad \mathbb{P}(X_i = a_i) = p_i,$
- $X = X_1 + \cdots + X_n, \quad \mu \ge \mathbb{E}[X] = \sum_{i=1}^{n} a_i \cdot p_i$

Chernoff Bound:

$$\mathbb{P}(X \ge (1+\varepsilon)\cdot\mu) \le \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\mu/A}$$

Proof:

$$\mathbb{P}(X \ge (1+\varepsilon)\mu) \le \frac{\mathbb{E}[e^{tX}]}{e^{t(1+\varepsilon)\mu}} \le e^{(e^{tA}-1)\mu - t(1+\varepsilon)\mu}$$

choose t s.t. $(e^{tA} - 1 - t(1+\varepsilon))$ is min.

$$A \cdot e^{tA} = 1+\varepsilon \implies tA = \ln\left(\frac{1+\varepsilon}{A}\right) \implies t = \frac{\ln\left(\frac{1+\varepsilon}{A}\right)}{A}$$

$$\mathbb{P}(X \ge (1+\varepsilon)\mu) \le e^{\left(\frac{1+\varepsilon}{A} - \frac{A}{A} - \frac{(1+\varepsilon)\ln\left(\frac{1+\varepsilon}{A}\right)}{A}\right)\mu} = \left(\frac{e^{1+\varepsilon-A}}{\left(\frac{1+\varepsilon}{A}\right)^{1+\varepsilon}}\right)^{\mu/A}$$

not quite what we wanted.

Proofing the Chernoff Bound

$e^x = \sum_{i=0}^{\infty} \frac{x^i}{i!}$

- $X_i \in \{0, a_i\}, \quad 0 < a_i \leq A, \quad \mathbb{P}(X_i = a_i) = p_i,$
- $X = X_1 + \cdots + X_n, \quad \mu \geq \mathbb{E}[X] = \sum_{i=1}^{n} a_i \cdot p_i$

Chernoff Bound:

$$\mathbb{P}(X \geq (1+\varepsilon)\cdot\mu) \leq \left(\frac{e^{\varepsilon}}{(1+\varepsilon)^{1+\varepsilon}}\right)^{\mu/A}$$

Proof:

$$\mathbb{E}[e^{tX}] \leq \prod_{i=1}^{n} e^{p_i(e^{ta_i}-1)} = e^{\sum_{i=1}^{n} p_i(e^{ta_i}-1)} \leq e^{\sum_{i=1}^{n} \frac{a_i p_i}{A}(e^{tA}-1)} \leq e^{(e^{tA}-1)\cdot\frac{\mu}{A}}$$

$$e^{ta_i} - 1 \leq \frac{a_i}{A}(e^{tA}-1) \xrightarrow{a=a_i} A(e^{ta}-1) \leq a(e^{tA}-1)$$

$$A\left(ta + \sum_{i=2}^{\infty}\frac{(ta)^i}{i!}\right) \leq a\left(tA + \sum_{i=2}^{\infty}\frac{(tA)^i}{i!}\right)$$

$$a^{i-1} \leq A^{i-1}$$

$$\sum_{i=2}^{\infty}\frac{(ta)^i \cdot A}{i!} \leq \sum_{i=2}^{\infty}\frac{(tA)^i a}{i!}$$

$$\underline{i \geq 2}$$

$$a^i A \leq A^i \cdot a$$

$$\underline{0 \leq a \leq A}$$

Multicommodity Routing: The General Case

- What if the possible paths $\mathcal{P}_i$ for commodity i are not given?
  - Using all exponentially many possible paths is not feasible

We can reduce to the rounding problem with fixed paths:

1. Solve the multicommodity flow LP
   - Returns a valid flow of value 1 for each commodity

2. Compute a set of paths $\mathcal{P}_i$ for each $i \in \{1, \dots, k\}$ such that the flow f_i corresponds to a probability distribution on the paths in $\mathcal{P}_i$
   - Using flow decomposition, one can always find a collection $\mathcal{P}_i$ of at most m paths

3. Round as before by using the path sets $\mathcal{P}_i$

Flow Decomposition

Flow Decomposition Lemma:

$\forall e : f(e) \geq \sum_{i=1}^{h} f_i(e)$

Let $G = (V, E)$ be a directed network with edge capacities $c_e > 0$, let $s, t \in V$, and let f be a flow in the network. Then there is a collection of feasible flows $f_1, \dots, f_h$ and a collection of s-t paths $P_1, \dots, P_h$ such that

- The number of paths is $h \leq |E|$
- The value of f is equal to the sum of the values of $f_1, \dots, f_h$
- Flow f_i sends positive flow only on the edges of P_i

Proof: Inductively construct $P_1, \dots, P_h$ (and corresponding flows $f_1, \dots, f_h$)

- For details, see, e.g., mins 17:00 – 29:50 of https://www.youtube.com/watch?v=zgutyzA9JM4&t=1020s

Application to Multicommodity Routing

- Decompose flow of each commodity $i \in \{1, \dots, k\}$
- Value of flow on each path is used as sampling probability

Oblivious Routing

- An "online" version of the multicommodity routing problem
- Decide for each source-destination request independently on which path to route it
  - For each $s, t \in V$, there is a probability distribution on s-t paths
  - If a message is sent from s to t, a path is chosen according to this distribution
- **Goal:** Be competitive with best multicommodity flow solution
- In this lecture, we will look at a very specific example:
  permutation routing on the d-dimensional hypercube
- **Permutation routing:**
  each node is source and destination of exactly one routing request
- **Hypercube $\boldsymbol{Q = (V, E)}$:**
  $V = \{0,1\}^d$, edge between u and v if Hamming distance $= 1$

Hypercube

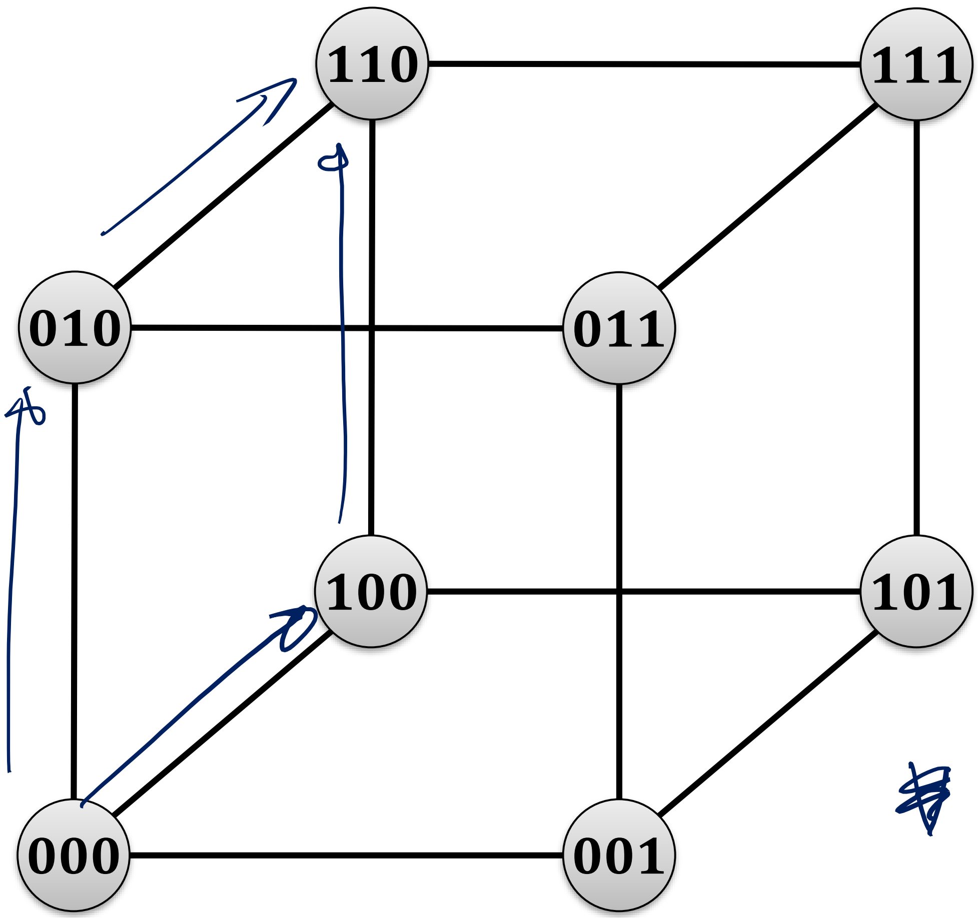

$V = \{0,1\}^d$

Routing on the Hypercube

Bit Fixing Algorithm:

- Fix "wrong" bits from left to right
- Example: $00101100 \longrightarrow 10010110$

$$\rightarrow \mathbf{1}0101100 \rightarrow 10\mathbf{0}01100 \rightarrow 100\mathbf{1}1100 \rightarrow 1001\mathbf{0}100 \rightarrow 100101\mathbf{1}0$$

Permutation Routing:

- Assumption: d-dimensional hypercube $Q = (V, E), n = |V|$
- $n = 2^d$ routing requests (s_i, t_i) (each of demand 1)
- Each node $v \in V$ is source s_i and destination t_i for exactly one request
  - Within these assumptions, requests are given in a worst-case manner
- Round-based model, ≤ 1 message per edge and round
  - In each round, every node can forward one message on each of its edges

Bad Example for Bit Fixing Algorithm

$n = 2^d \qquad d = \log n$

$\frac{\sqrt{n}}{d}$

instances where we need time $\frac{\sqrt{n}}{\log n}$

$2^{d/2} = \sqrt{n}$

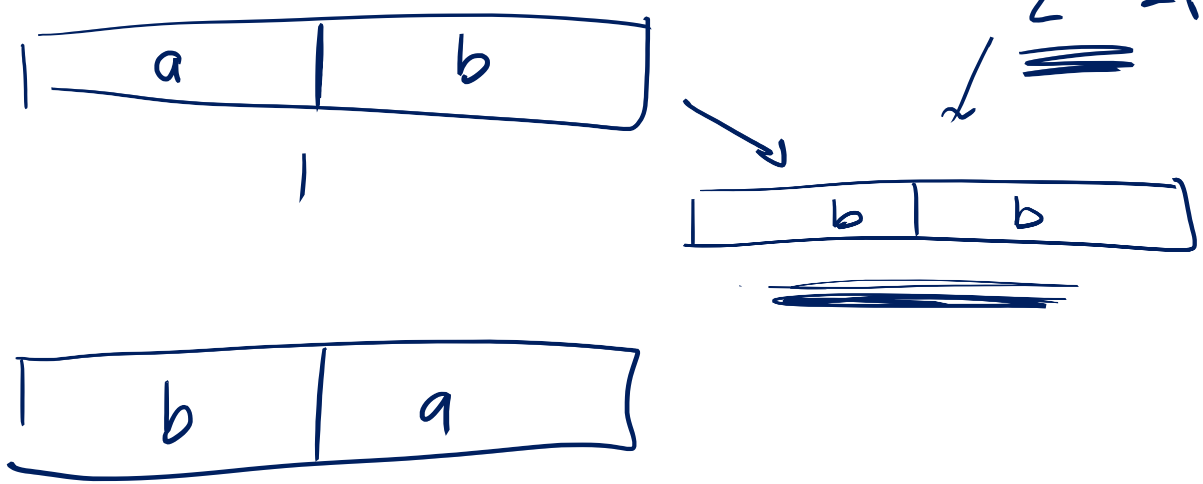

Valiant's Trick

s_i sends to random node r_i first

$$s_i \xrightarrow{\text{bit fixing}} r_i \xrightarrow{\text{bit fixing}} t_i$$

r_i: unif. at random

Analyzing Bit Fixing with Valiant's Trick

$|P_i| \le d$

time: $|P_i| + \underline{\text{delay}}$

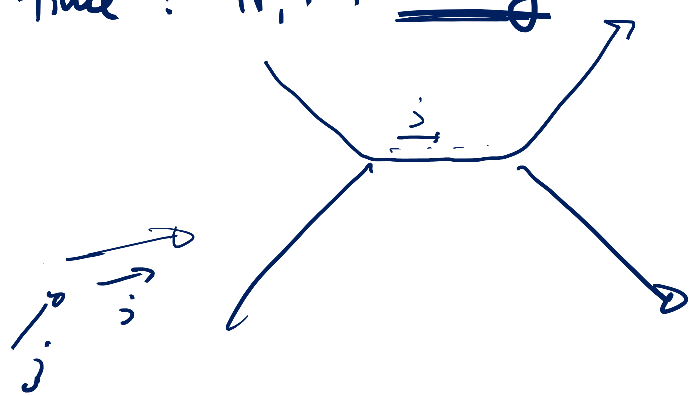

$X_{ij} = 1 \iff$ path P_i & P_j intersect

delay of $i \le X_i = \sum_{j \neq i} X_{ij}$

$\mathbb{E}[X_{ij}] = ?$ need $\mathbb{E}[X_i]$

Y_e: # paths going through e

$X_i \le \sum_{e \in P_i} Y_e$

$\mathbb{E}[X_i] \le \sum_{e \in P_i} \mathbb{E}[Y_e]$

Analyzing Bit Fixing with Valiant's Trick

what is $\mathbb{E}[Y_e]$?

$Y = \sum_{e \in E} Y_e$

total path length (over all n paths)

$\mathbb{E}[Y_e] = \mathbb{E}[Y_{e'}]$

$\mathbb{E}[Y]$ = exp. total path length

$= n \cdot \mathbb{E}[\text{length of single path}] = n \cdot \frac{d}{2}$

$\mathbb{E}[Y_e] = \frac{1}{|E|} \cdot \mathbb{E}[Y] = \frac{1}{n \cdot d} \cdot \frac{nd}{2} = \frac{1}{2}$

~~Chernoff~~

$\mathbb{E}[X_i] \leq \sum_{e \in P_i} \mathbb{E}[Y_e] \leq d/2$

$\mathbb{P}(X_i \geq 3d) \leq 2^{-3d}$

routing completes in $O(d)$ steps w.h.p.