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 $\boldsymbol{s_i}$**-**$\boldsymbol{t_i}$ **flow** $\boldsymbol{f_i : E \to \mathbb{R}_{\geq 0}}$ 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

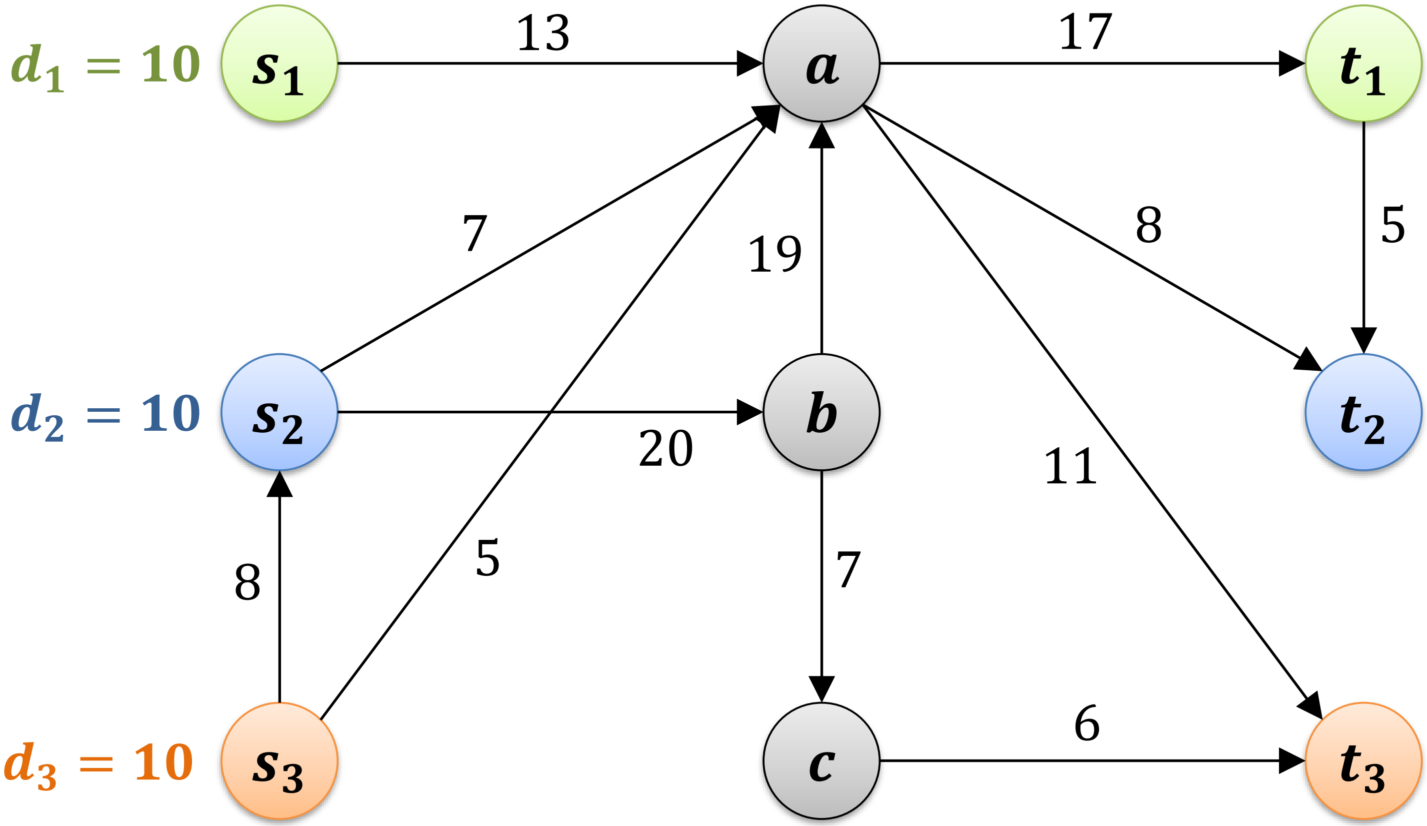

Multicommodity Flow as an LP

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:

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

- 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 $\alpha > 0$,

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

Example: Multicommodity Routing

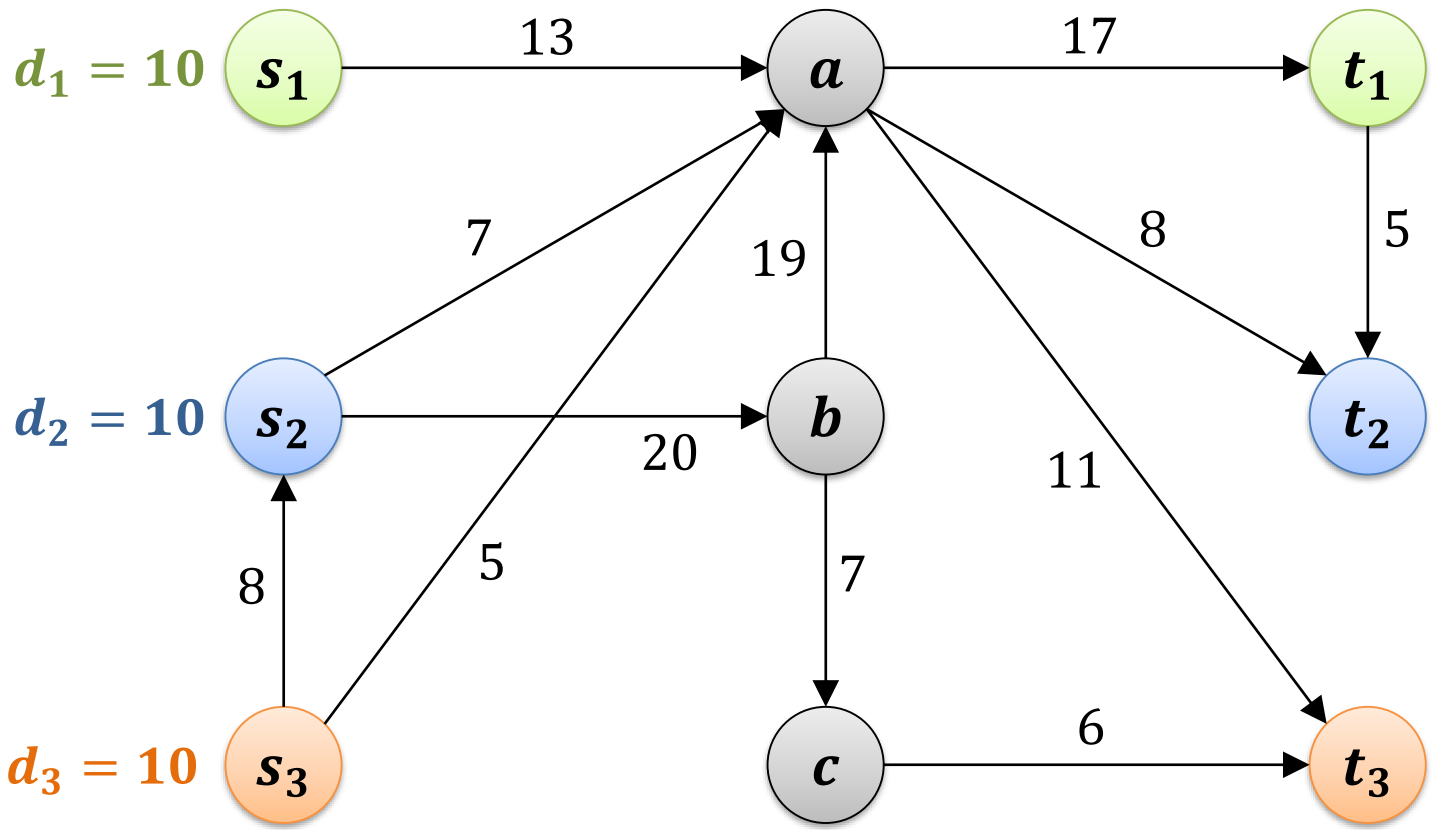

Rounding the Multicommodity Flow LP

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:

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:

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:

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 $e \in E$:

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

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 $X \coloneqq X_1 + \cdots + X_n$ 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

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

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

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:**

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

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

$$\lambda \leq O\left(\frac{\log n}{\log\log n}\right) \cdot \lambda^*.$$

Proof:

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

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

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

Chernoff Bound:

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

Proof:

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:

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

Proof:

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:

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

Proof:

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:

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 $f_1, \dots, f_t$ and a collection of s-t paths $P_1, \dots, P_t$ such that

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

Proof: Inductively construct $P_1, \dots, P_t$ (and corresponding flows $f_1, \dots, f_t$)

- 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

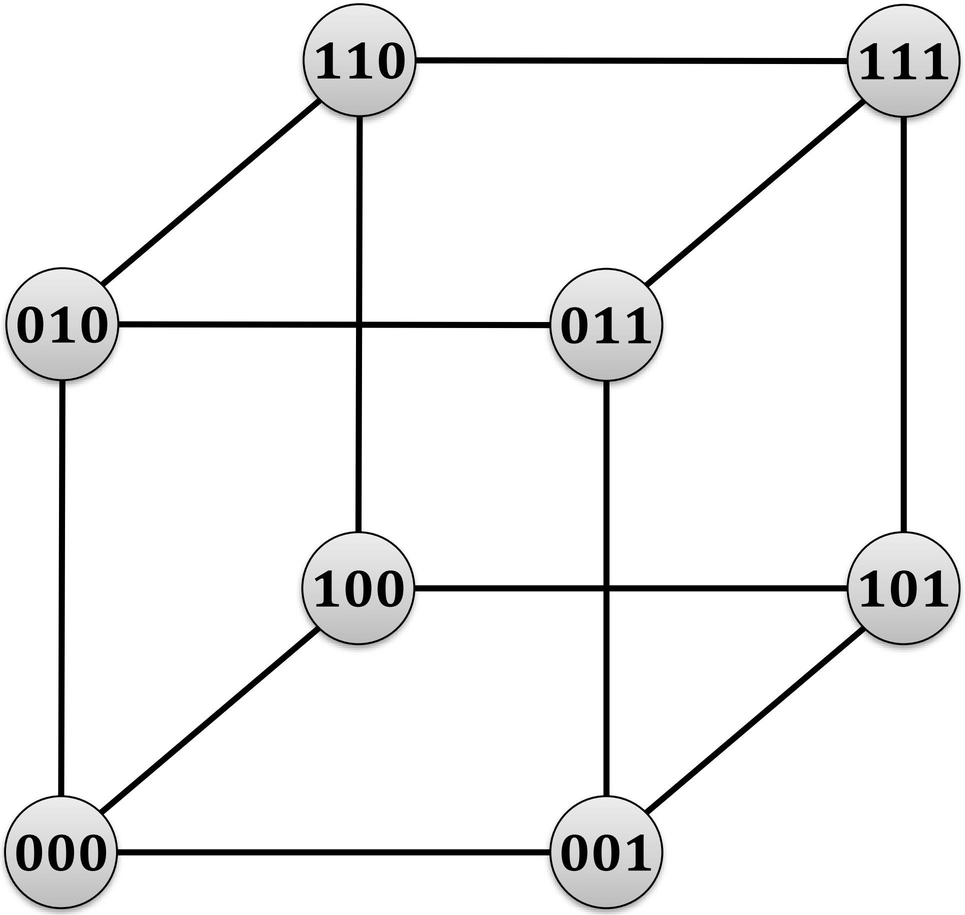

Routing on the Hypercube

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

Valiant's Trick

Analyzing Bit Fixing with Valiant's Trick

Analyzing Bit Fixing with Valiant's Trick