



Chapter 9

Fast Approximate Max Flow in Undirected Graphs

Advanced Algorithms

SS 2019

Fabian Kuhn

Cut Sparsifiers

Last week: cut sparsifiers by [Benczúr, Karger; '02]

Given: graph $G = (V, E) \Rightarrow$ weighted graph $H = (V, E', w)$ with $E' \subseteq E$

- Such that **all cuts are preserved** up to a $(1 \pm \varepsilon)$ -factor & $|E'| = O\left(\frac{n \log n}{\varepsilon}\right)$
- Can be computed in time $\tilde{O}(m)$
- Also works for weighted graphs

Gives an immediate algorithms to approximate cut problems:

- First, compute a cut sparsifier, then solve the problem on the sparsifier
- Outputs an almost optimal cut on G , faster if running time depends on m
- Example: $(1 + \varepsilon)$ -approximate minimum s - t cut in time $\tilde{O}(n^{3/2}/\varepsilon^3)$
 - in undirected graphs...

What about the max flow problem?

- We can approximate the value of the maximum flow, but it is not clear how to construct a flow (the sparsifier does not contain most of the edges of G)

Max Flow with Cut Sparsifiers

- Benczúr and Karger give a way to use their cut sparsifier also for the undirected max flow problem

The main ideas are:

- Replace “important” edges in G by multiple parallel edges (capacity is divided evenly among the multiple edges replacing an original edge)
- “important” = small edge strength (strong edge connectivity k_e)
 - Small edge strength \Rightarrow large sampling probability in sparsifier algorithm
- This can be done such that the number of edges only grows by a constant factor and every edge e has sampling probability $\tilde{O}(n/m)$ in the sparsifier alg.
- One can then randomly partition the graph into $\tilde{O}(m/n)$ parts
 - All cuts are close to their expected size (same analysis as sparsifier analysis)
- We can then solve independent max flow problems in all the $\tilde{O}(n/m)$ parts and add the flows to get an $(1 - \varepsilon)$ -approximate max flow for G
- Allows to **turn an existing $O(m^{3/2})$ into an $\tilde{O}(m\sqrt{n}/\varepsilon)$ approximate alg.**
- **Today:** Main ideas of a faster (more involved) way of solving max flow

A slightly more general Problem

Given:

- Undirected graph $G = (V, E, c)$ with edge capacities $c_e > 0$
- Every node $v \in V$ has a **demand** $b_v \in \mathbb{R}$ s.t. $\sum_{v \in V} b_v = 0$

Goal: Find a flow f such that

$$\forall v \in V : f_{out}(v) - f_{in}(v) = b_v$$

$$\text{minimize } \max_{e \in E} \frac{|f_e|}{c_e}$$

How to solve max s - t flow:

1. Set $b_s = 1, b_t = -1, b_v = 0$ for $v \notin \{s, t\}$
2. Solve the above problem
3. Scale flow s.t. $\max_{e \in E} \frac{|f_e|}{c_e} = 1$

Matrix Representation of Max Flow

- Each edge $\{u, v\}$ added either as (u, v) or as (v, u) , flow f_e on edge $e = (u, v)$:
 - $f_e > 0$: flow from u to v , $f_e < 0$: flow from v to u
- B : node-edge incidence matrix (B is an $n \times m$ matrix)
 - Edge $e = (u, v)$: $B_{u,e} = +1, B_{v,e} = -1, B_{w,e} = 0$ for $w \notin \{u, v\}$

Valid flow: f is valid $\Leftrightarrow Bf = b$

Matrix Representation of Max Flow

- B : node-edge incidence matrix (B is an $n \times m$ matrix)
 - Edge $e = (u, v)$: $B_{u,e} = +1$, $B_{v,e} = -1$, $B_{w,e} = 0$ for $w \notin \{u, v\}$

Valid flow: f is valid $\Leftrightarrow Bf = b$

Capacity matrix:

Goal: minimize $\|C^{-1}f\|_{\infty}$ s.t. $Bf = b$

Dual Problem (Generalization of Min Cut)

For every cut $(S, V \setminus S)$:

- Capacity of cut c_S
- Total amount of flow across cut $(S, V \setminus S)$ is at least

$$b_S := \sum_{v \in S} b_v$$

Max flow min cut theorem:

$$\exists \text{ cut } S : \text{opt}(\mathbf{b}) = \frac{b_S}{c_S}$$

Dual Problem (Generalization of Min Cut)

Dual problem: maximum congested cut

- Vertex potentials $\mathbf{x} \in \mathbb{R}^n$, $x_v \in \mathbb{R}$
- Goal:

$$\max \mathbf{b}^\top \mathbf{x} \quad \text{s.t.} \quad \|CB^\top \mathbf{x}\|_1 \leq 1$$

- Example: consider a cut $(S, V \setminus S)$:
Vector \mathbf{x}_S is characteristic vector of set S ($x_v = 1 \Leftrightarrow v \in S$)

Dual Problem (Generalization of Min Cut)

Goal: vertex potentials $x_v \in \mathbb{R}$:

$$\max \mathbf{b}^\top \mathbf{x} \quad s.t. \quad \|CB^\top \mathbf{x}\|_1 \leq 1$$

Claim: Opt. solution of above problem has value $\text{opt}(\mathbf{b})$.

- We have seen that there exists \mathbf{x} with $\mathbf{b}^\top \mathbf{x} = \text{opt}(\mathbf{b})$

Congestion Approximator

A method that allows to get a good approximation of $\text{opt}(\mathbf{b})$

Definition: α -congestion approximator is a matrix $R \subseteq \mathbb{R}^{\ell \times n}$ s.t.

$$\forall \mathbf{b} \in \mathbb{R}^n: \|R\mathbf{b}\|_{\infty} \leq \text{opt}(\mathbf{b}) \leq \alpha \cdot \|R\mathbf{b}\|_{\infty}$$

Example 1:

- One row for each possible cut $(S, V \setminus S)$:

$$R_{S,v} = \frac{1}{c_S}$$

Congestion Approximator

A method that allows to get a good approximation of $\text{opt}(\mathbf{b})$

Definition: α -congestion approximator is a matrix $R \subseteq \mathbb{R}^{\ell \times n}$ s.t.

$$\forall \mathbf{b} \in \mathbb{R}^n: \|R\mathbf{b}\|_{\infty} \leq \text{opt}(\mathbf{b}) \leq \alpha \cdot \|R\mathbf{b}\|_{\infty}$$

Example 2:

- Assume T is a maximum weight spanning tree
- Add one row for each edge e of T , let S_e be the induced cut of e :

$$R_{e,v} = \frac{b_v}{c_{S_e}}$$

- Measures exactly the cost of routing the flow on the tree T
- Routing on the tree incurs at most a factor $m \Rightarrow \alpha = m$

Congestion Approximator

A method that allows to get a good approximation of $\text{opt}(\mathbf{b})$

Definition: α -congestion approximator is a matrix $R \subseteq \mathbb{R}^{\ell \times n}$ s.t.

$$\forall \mathbf{b} \in \mathbb{R}^n: \|R\mathbf{b}\|_{\infty} \leq \text{opt}(\mathbf{b}) \leq \alpha \cdot \|R\mathbf{b}\|_{\infty}$$

Example 3:

- Use all the trees of a low-congestion tree embedding
 - As considered in the lectures on May 17 and June 7
 - When picking a random tree, expected congestion of each edge is at most $O(\log n)$ times the congestion for an optimal solution of an arbitrary multicommodity flow problem
- Add one row for each tree T and each edge e of T , let S_e be the induced cut of e :

$$R_{e,v} = \frac{b_v}{c_{S_e}}$$

- Gives $\alpha = O(\log n)$

Congestion Approximator

A method that allows to get a good approximation of $\text{opt}(\mathbf{b})$

Definition: α -congestion approximator is a matrix $R \subseteq \mathbb{R}^{\ell \times n}$ s.t.

$$\forall \mathbf{b} \in \mathbb{R}^n: \|R\mathbf{b}\|_\infty \leq \text{opt}(\mathbf{b}) \leq \alpha \cdot \|R\mathbf{b}\|_\infty$$

Example 4:

- Add one row for each tree T and each edge e of T of a low-congestion tree embedding, let S_e be the induced cut of e
- The construction required $\tilde{O}(m)$ trees $\Rightarrow R$ has $\tilde{O}(mn)$ rows
- Can be improved by first computing a cut sparsifier
 - Now, we only need $\tilde{O}(n)$ trees $\Rightarrow R$ has $\tilde{O}(n^2)$ rows
- In [Sherman; 2013], a recursive variant of this is described:
 - Based on a construction of [Mardy; 2010]
 - Instead of trees, embed into more complicated structures (needs less of them)
 - Gives a congestion approximator R with $n^{1+o(1)}$ rows that can be computed in time $m \cdot n^{o(1)}$ and with $\alpha = n^{o(1)}$

An Unconstrained Version of the Problem

- Assume that an α -congestion approximator R with $\leq n^2/2$ rows is given

Using it, we can turn max flow into an unconstrained optimization problem:

$$\min_{\text{flow } \mathbf{f}} \gamma(\mathbf{f}) := \|C^{-1}\mathbf{f}\|_{\infty} + 2\alpha \cdot \|R(\mathbf{b} - B\mathbf{f})\|_{\infty}$$

Theorem 1:

There is an algorithm `AlmostRoute`(\mathbf{b}, ε) that returns a flow f for which

$$\gamma(\mathbf{f}) \leq (1 + \varepsilon) \cdot \text{opt}(\mathbf{b}).$$

The algorithm requires $O\left(\frac{\alpha^2 \cdot \log \alpha \cdot \log n}{\varepsilon^3}\right)$ iterations that require time $\tilde{O}(m)$ plus a multiplication by R and by R^{\top} .

An Unconstrained Version of the Problem



$$\min_{\text{flow } \mathbf{f}} \gamma(\mathbf{f}) := \|\mathbf{C}^{-1}\mathbf{f}\|_{\infty} + 2\alpha \cdot \|\mathbf{R}(\mathbf{b} - \mathbf{B}\mathbf{f})\|_{\infty}$$

Theorem 2: There is an algorithm that computes a valid $(1 + \varepsilon)$ -approximate flow that applies `AlmostRoute` $O(\log n)$ times.

An Unconstrained Version of the Problem



$$\min_{\text{flow } \mathbf{f}} \gamma(\mathbf{f}) := \|\mathbf{C}^{-1}\mathbf{f}\|_{\infty} + 2\alpha \cdot \|\mathbf{R}(\mathbf{b} - \mathbf{B}\mathbf{f})\|_{\infty}$$

Theorem 2: There is an algorithm that computes a valid $(1 + \varepsilon)$ -approximate flow that applies `AlmostRoute` $O(\log n)$ times.

A Differentiable Objective Function

Softmax function (on a vector $\mathbf{x} \in \mathbb{R}^d$):

$$\text{lmax}(\mathbf{x}) := \ln \left(\sum_{i=1}^d (e^{x_i} + e^{-x_i}) \right)$$

Properties of softmax:

$$\|\mathbf{x}\|_{\infty} \leq \text{lmax}(\mathbf{x}) \leq \|\mathbf{x}\|_{\infty} + \ln(2d)$$

$$\|\nabla \text{lmax}(\mathbf{x})\|_1 \leq 1$$

$$\nabla \text{lmax}(\mathbf{x})^{\top} \mathbf{x} \geq \text{lmax}(\mathbf{x}) - \ln(2d)$$

$$\|\nabla \text{lmax}(\mathbf{x}) - \nabla \text{lmax}(\mathbf{y})\|_1 \leq \|\mathbf{x} - \mathbf{y}\|_{\infty}$$

A Differentiable Objective Function

Softmax function (on a vector $\mathbf{x} \in \mathbb{R}^d$):

$$\text{lmax}(\mathbf{x}) := \ln \left(\sum_{i=1}^d (e^{x_i} + e^{-x_i}) \right)$$

Replace

$$\gamma(\mathbf{f}) := \|C^{-1}\mathbf{f}\|_{\infty} + 2\alpha \cdot \|R(\mathbf{b} - B\mathbf{f})\|_{\infty}$$

By

$$\phi(\mathbf{f}) := \text{lmax}(C^{-1}\mathbf{f}) + \text{lmax}(2\alpha \cdot R(\mathbf{b} - B\mathbf{f}))$$

AlmostRoute(\mathbf{b}, ε)

- Initialize $\mathbf{f} = 0$, scale \mathbf{b} so $2\alpha \cdot \|\mathbf{R}\mathbf{b}\|_\infty = 16\varepsilon^{-1} \ln n$
- Repeat:
 - While $\phi(\mathbf{f}) < 16\varepsilon^{-1} \ln n$, scale \mathbf{f} and \mathbf{b} up by 17/16
 - Set $\delta := \|C \cdot \nabla\phi(\mathbf{f})\|_1$
 - If $\delta \geq \varepsilon/4$, set $f_e := f_e - \frac{\delta}{1+4\alpha^2} \cdot \text{sgn}\left(\left(\nabla\phi(\mathbf{f})\right)_e\right) \cdot c_e$
 - Otherwise, terminate and output \mathbf{f} after undoing all scalings

Also, output vertex potentials $\mathbf{x} := \mathbf{R}^\top \cdot \nabla \max(2\alpha \cdot \mathbf{R}(\mathbf{b} - \mathbf{B}\mathbf{f}))$

AlmostRoute(\mathbf{b}, ε)

Lemma: When AlmostRoute(\mathbf{b}, ε) terminates, we have

$$\phi(\mathbf{f}) \leq (1 + \varepsilon) \cdot \frac{\mathbf{b}^\top \mathbf{x}}{\|C\mathbf{B}^\top \mathbf{x}\|_1}$$

AlmostRoute(\mathbf{b}, ε)

Lemma: When AlmostRoute(\mathbf{b}, ε) terminates, we have

$$\phi(\mathbf{f}) \leq (1 + \varepsilon) \cdot \frac{\mathbf{b}^\top \mathbf{x}}{\|C\mathbf{B}^\top \mathbf{x}\|_1}$$

AlmostRoute(\mathbf{b}, ε)

Lemma: When AlmostRoute(\mathbf{b}, ε) terminates, we have

$$\phi(\mathbf{f}) \leq (1 + \varepsilon) \cdot \frac{\mathbf{b}^\top \mathbf{x}}{\|C\mathbf{B}^\top \mathbf{x}\|_1}$$

AlmostRoute(\mathbf{b}, ε)

Lemma: When AlmostRoute(\mathbf{b}, ε) terminates, we have

$$\phi(\mathbf{f}) \leq (1 + \varepsilon) \cdot \frac{\mathbf{b}^\top \mathbf{x}}{\|C\mathbf{B}^\top \mathbf{x}\|_1}$$

AlmostRoute(\mathbf{b}, ε)

Lemma: The number of iterations of AlmostRoute(\mathbf{b}, ε) is at most

$$O\left(\frac{\alpha^2 \cdot \log \alpha \cdot \log n}{\varepsilon^3}\right).$$

- Initialize $\mathbf{f} = 0$, scale \mathbf{b} so $2\alpha \cdot \|\mathbf{R}\mathbf{b}\|_\infty = 16\varepsilon^{-1} \ln n$
- Repeat:
 - While $\phi(\mathbf{f}) < 16\varepsilon^{-1} \ln n$, scale \mathbf{f} and \mathbf{b} up by 17/16
 - Set $\delta := \|C \cdot \nabla\phi(\mathbf{f})\|_1$
 - If $\delta \geq \varepsilon/4$, set $f_e := f_e - \frac{\delta}{1+4\alpha^2} \cdot \text{sgn}\left(\left(\nabla\phi(\mathbf{f})\right)_e\right) \cdot c_e$
 - Otherwise, terminate

AlmostRoute(\mathbf{b}, ε)

Lemma: The number of iterations of AlmostRoute(\mathbf{b}, ε) is at most

$$O\left(\frac{\alpha^2 \cdot \log \alpha \cdot \log n}{\varepsilon^3}\right).$$