Feedback welcome! If you find any typos or anything is unclear or misleading, please email me and know!

For additional detail, see the companion paper, "Box-Simplex Games : Algorithms, Applications, and Algorithmic Graph Theory" on my website.

Box-Simplex Games

Algorithms, Applications, and Algorithmic Graph Theory

Part I

Aaron Sidford

Stanford University
Departments of Management Science & Engineering and Computer Science

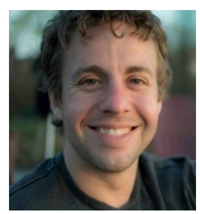

Contact Info:
- email: *sidford@stanford.edu*
- website: *www.aaronsidford.com*

Lecture Plan

Today and Tomorrow

- Box-simplex games
- Their structure
- Applications
- Algorithms

Why?

- **(Applications)** Continuous and combinatorial.
- **(Tools)** New optimization methods
- (**Reinforce**) Modifications of common methods

Friday

- Interior point methods

- Introduction of state-of-the-art method

The Problem

Input

- n-dimensional box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- m-dimensional simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\geq 0}^m \mid \|y\|_1 = 1\}$

Bounded vectors in $\mathbb{R}^n$

Probability distributions on m elements

Output:

- An approximate solution to

$$\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x, y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$$

Box-Simplex Game ℓ_1-ℓ_∞ **Game**

Key Motivating Questions

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\geq 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x, y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Question #1

How can we design efficient methods for solving box-simplex games?

Question #2

How can we leverage box-simplex solvers to solve continuous and combinatorial optimization problems?

Talk Plan (Today and Tomorrow)

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\geq 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x, y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Part 1
Structure of box-simplex games

- Primal and dual problems
- Approximate solutions
- Discuss state-of-the-art runtimes

Friday
Interior Point Methods

Part 2
Applications

- Box-constrained ℓ_∞-regression
- Linear programming
- Maximum cardinality bipartite matching
- Undirected maximum flow

Part 3
Algorithms

- ℓ_∞-Gradient Descent (constrained steepest descent)
- ℓ_1-Mirror Descent (multiplicative weights)
- Mirror prox and primal dual regularizers

Primal Problem

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \le 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\ge 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x,y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Lemma: $\max_{x \in \Delta^m} d^\top x = \max_{i \in [m]} d_i$ for all $d \in \mathbb{R}^m$ and therefore

$$f_{\max}(x) \stackrel{\text{def}}{=} \max_{y \in \Delta^m} f(x,y) = c^\top x + \max_{i \in [m]} [\boldsymbol{A}x - b]_i$$

Proof:

- Let $i_* \in \operatorname{argmax}_{i \in [m]} d_i$. Note that $\vec{1}_{i_*} \in \Delta^m$.
  - $\Rightarrow \max_{x \in \Delta^m} d^\top x \ge d^\top \vec{1}_{i_*} = d_{i_*} = \max_{i \in [m]} d_i$
- $d_i \le d_{i_*}$ and $x_i \ge 0$ for $x \in \Delta^m$ and $i \in [m]$
  - $\Rightarrow \max_{x \in \Delta^m} d^\top x = \max_{x \in \Delta^m} \sum_{i \in [m]} d_i x_i \le \max_{x \in \Delta^m} \sum_{i \in [m]} d_{i_*} x_i = d_{i_*} = \max_{i \in [m]} d_i$

Dual Problem

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}^m_{\geq 0} \mid \|y\|_1 = 1\}$
- $\max_{y \in \Delta^m} \min_{x \in B_\infty^n} f(x, y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Lemma: $\min_{x \in B_\infty^n} d^\top x = -\|d\|_1$ for all $d \in \mathbb{R}^n$ and therefore

$$f_{\min}(y) \stackrel{\text{def}}{=} \min_{x \in B_\infty^n} f(x, y) = -b^\top y - \|\boldsymbol{A}^\top y - b\|_1$$

Proof:

- Let $\operatorname{sign}(d) \in \mathbb{R}^n$ with $\operatorname{sign}(d)_i$ as 1 if $d_i > 0$, -1 if $d_i < 0$, and 0 othewise
  - $\Rightarrow \min_{x \in B_\infty^n} d^\top x \leq d^\top\big(-\operatorname{sign}(d)\big) = -\sum_{i \in [n]} |d_i| = -\|d\|_1$
- $|x_i| \leq 1$ for all $x \in B_\infty^n$ and $i \in [n]$
  - $\Rightarrow \min_{x \in B_\infty^n} d^\top x = \min_{x \in B_\infty^n} -\sum_{i \in [n]} |d_i||x_i| \geq \min_{x \in B_\infty^n} -\sum_{i \in [n]} |d_i| = -\|d\|_1$

Primal Dual Relationship

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\geq 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x, y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Primal Problem

- $\min_{x \in B_\infty^n} f_{\max}(x) = \max_{y \in \Delta^m} f(x, y)$
- $f_{\max}(x) = c^\top x + \max_{i \in [m]} [\boldsymbol{A} x - b]_i$

Dual Problem

- $\max_{y \in \Delta^m} f_{\min}(y) = \min_{x \in B_\infty^n} f(x, y)$
- $f_{\min}(y) = -b^\top y - \|\boldsymbol{A}^\top y - b\|_1$

Comparison

- Trivially: $f_{\max}(x) \geq f_{\min}(y)$ (weak duality)
- Interestingly: $\min_{x \in B_\infty^n} f_{\max}(x) = \max_{y \in \Delta^m} f_{\min}(y)$ (strong duality)

We will prove algorithmically later

Approximate Solutions

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\geq 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x, y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Primal Problem

- $\min_{x \in B_\infty^n} f_{\max}(x) = \max_{y \in \Delta^m} f(x, y)$
- $f_{\max}(x) = c^\top x + \max_{i \in [m]} [\boldsymbol{A} x - b]_i$

Dual Problem

- $\max_{y \in \Delta^m} f_{\min}(y) = \min_{x \in B_\infty^n} f(x, y)$
- $f_{\min}(y) = -b^\top y - \|\boldsymbol{A}^\top y - b\|_1$

Approximate Solutions

- Let $x_* \in \operatorname{argmin}_{x \in B_\infty^n} f_{\max}(x)$ and $y_* \in \operatorname{argmax}_{y \in \Delta^m y \in \Delta^m} f_{\min}(y)$
- ϵ-approximate primal solution: $x_\epsilon \in B_\infty^n$ with $f_{\max}(x_\epsilon) \leq f_{\max}(x_*) + \epsilon$
- ϵ-approximate dual solution: $y_\epsilon \in \Delta^m$ with $f_{\min}(y_\epsilon) \geq f_{\min}(y_*) - \epsilon$
- ϵ-approximate *(primal-dual) saddle point (or equilibrium)*: $(x_\epsilon, y_\epsilon) \in B_\infty^n \times \Delta^m$

$$f_{\max}(x_\epsilon) - f_{\min}(y_\epsilon) \leq \epsilon$$

Equilibrium

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \le 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\ge 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x,y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Primal Problem

- $\min_{x \in B_\infty^n} f_{\max}(x) = \max_{y \in \Delta^m} f(x,y)$
- $f_{\max}(x) = c^\top x + \max_{i \in [m]} [\boldsymbol{A}x - b]_i$

Dual Problem

- $\max_{y \in \Delta^m} f_{\min}(y) = \min_{x \in B_\infty^n} f(x,y)$
- $f_{\min}(y) = -b^\top y - \|\boldsymbol{A}^\top y - b\|_1$

ϵ-approximate (primal-dual) saddle point (or equilibrium)

- **Definition**: $(x_\epsilon, y_\epsilon) \in B_\infty^n \times \Delta^m$ and $f_{\max}(x_\epsilon) - f_{\min}(y_\epsilon) \le \epsilon$
- **Duality gap**: $\text{gap}(x_\epsilon, y_\epsilon) = f_{\max}(x_\epsilon) - f_{\min}(y_\epsilon)$
  - Total $f(x_\epsilon, y_\epsilon)$ change by best responses: $= f_{\max}(x_\epsilon) - f(x_\epsilon, y_\epsilon) + [f(x_\epsilon, y_\epsilon) - f_{\min}(y_\epsilon)]$
  - Sum of x_ϵ and y_ϵ suboptimality: $= f_{\max}(x_\epsilon) - f(x_\epsilon, y_\epsilon) + [f(x_\epsilon, y_\epsilon) - f_{\min}(y_\epsilon)]$

Don't need x_* and y_* to compute!

State-of-the-art

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\geq 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x,y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Theorem: there is a method which can compute an ϵ-approximate saddle point in time $\tilde{O}\left(\mathrm{nnz}(\boldsymbol{A}) \|\boldsymbol{A}\|_{\mathrm{op},\infty} / \epsilon\right)$

Nearly linear time algorithm

Notation

- $\mathrm{nnz}(\boldsymbol{A}) \stackrel{\text{def}}{=} n + m +$ number of nonzero entries in $\boldsymbol{A}$

size of the input

- $\|\boldsymbol{A}\|_{\mathrm{op},\infty} \stackrel{\text{def}}{=} \sup_{x \neq 0} \frac{\|\boldsymbol{A}x\|_\infty}{\|x\|_\infty} =$ max ℓ_1 norm of row of $\boldsymbol{A}$

"ℓ_∞ operator norm" bounds up to constant how suboptimal primal/dual solutions which just optimize b and c are

- $\tilde{O}(\cdot)$ hides logarithmic factors in $\mathrm{nnz}(\boldsymbol{A})$, $\|\boldsymbol{A}\|_{\mathrm{op},\infty}/\epsilon$

First-order method

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{x \in \mathbb{R}^m_{\geq 0} \mid \|x\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x,y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Theorem: there is a method which solves box-simplex games to accuracy ϵ in time $\tilde{O}\left(\mathrm{nnz}(\boldsymbol{A})\|\boldsymbol{A}\|_{\mathrm{op},\infty}/\epsilon\right)$.

- **First order method**: only access objective by evaluating the function and computing the gradient, $\nabla f(x,y) = (\boldsymbol{A}^\top y + c, \boldsymbol{A}x - b)$
- **Note:** only need b, c, and matrix vector multiplies.
  - Can compute in parallel $\tilde{O}(1)$ depth and $O(\mathrm{nnz}(\boldsymbol{A}))$ work.
- **The method for this theorem?**
  - First order method + matrix vector multiplies with $|\boldsymbol{A}|$
  - Parallel with $\tilde{O}(1)$ depth

Entrywise absolute value

History and More State-of-the-art

First Order Methods

- $\tilde{O}\left(\text{nnz}(\boldsymbol{A})\|\boldsymbol{A}\|_{\text{op},\infty}/\epsilon\right)$
  - First in [S17]
  - Later variants (influencing this presentation [J**S**T19,C**S**T21,AJJ**S**T21]
- Prior state of the art
  - $\tilde{O}(\text{nnz}(\boldsymbol{A})\|\boldsymbol{A}\|_{\text{op},\infty}^2/\epsilon^2)$ – folklore / [S13, KLO**S**14] (influencing this presentation)
  - $\tilde{O}\left(\text{nnz}(\boldsymbol{A})\sqrt{n}\|\boldsymbol{A}\|_{\text{op},\infty}/\epsilon\right)$ - AGD and smoothing.
- [**S**T18] alternative approach and improvements in sparse case

Interior Point Methods

- [CLS19,B20] $\tilde{O}(\max\{m,n\}^{\omega})$ where $\omega < 2.373$ is fast matrix multiplication constant
- [BLLS**S**SW21] $\tilde{O}(mn + \min\{m,n\}^{2.5})$
- [L**S**14,L**S**15] $\tilde{O}(\text{nnz}(\boldsymbol{A})\sqrt{\min\{m,n\}} + \min\{m,n\}^{2.5})$

$w < 2.373$ is current fast matrix multiplication (FMM) constant [W13]

Talk Plan (Today and Tomorrow)

- Box: $B_\infty^n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq 1\}$
- Simplex: $\Delta^m \stackrel{\text{def}}{=} \{y \in \mathbb{R}_{\geq 0}^m \mid \|y\|_1 = 1\}$
- $\min_{x \in B_\infty^n} \max_{y \in \Delta^m} f(x,y) \stackrel{\text{def}}{=} y^\top \boldsymbol{A} x + c^\top x - b^\top y$

Part 1
Structure of box-simplex games

- Primal and dual problems
- Approximate solutions
- Discuss state-of-the-art runtimes

Friday
Interior Point Methods

Part 2
Applications

- Box-constrained ℓ_∞-regression
- Linear programming
- Maximum cardinality bipartite matching
- Undirected maximum flow

Part 3
Algorithms

- ℓ_∞-Gradient Descent (constrained steepest descent)
- ℓ_1-Mirror Descent (multiplicative weights)
- Mirror prox and primal dual regularizers

Problem #1: Box-constrained ℓ_∞-Regression

Box-constrained ℓ_∞ Regressioon

- **Input**: matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and vector $b \in \mathbb{R}^m$
- **Problem**: $\mathrm{OPT}_\infty = \min_{x \in B_\infty^n} \|\boldsymbol{A}x - b\|_\infty$
- **Goal**: find ϵ-additive approximation, i.e. $x_\epsilon \in B_\infty^n$ with $\|\boldsymbol{A}x_\epsilon - b\|_\infty \leq \mathrm{OPT}_\infty + \epsilon$

Claim: can compute in $\tilde{O}(\mathrm{nnz}(\boldsymbol{A})\|\boldsymbol{A}\|_{op,\infty}\epsilon^{-1})$

Proof:

- $\|\boldsymbol{A}x - b\|_\infty = \max_{i \in [m]}[\max\{[\boldsymbol{A}x - b]_i, -[\boldsymbol{A}x - b]_i\}] = \max_{y \in \Delta^{2m}} y^\top \begin{pmatrix} \boldsymbol{A}x - b \\ -(\boldsymbol{A}x - b) \end{pmatrix}$
- New matrix has same $\|\cdot\|_{op,\infty}$ and just double nnz

Problem #2: Linear Programming

<u>Approximate Linear Programming</u>

- **Input**: $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and $\epsilon, \delta, D \geq 0$
- **Problem**: $\mathrm{OPT}_{\mathrm{lp}} = \min_{x \in \mathbb{R}^n \mid \boldsymbol{A}x \geq b} c^\top x$
- **Promise**: $\exists x_*^{\mathrm{lp}} \in \operatorname{argmin}_{x \in \mathbb{R}^n \mid \boldsymbol{A}x \geq b} c^\top x$ with $\left\| x_*^{\mathrm{lp}} \right\|_\infty \leq D$
- **Goal**: find $x_{\epsilon,\delta}$ with $c^\top x_{\epsilon,\delta} \leq \mathrm{OPT}_{\mathrm{lp}}$ and $\boldsymbol{A}x_{\epsilon,\delta} \geq b - \delta\vec{1}$

Notes

- One of many ways to formulate the problem.
- Key difficulty: how handle that constraint $\boldsymbol{A}x \geq b$?
- Recurring idea: penalty functions in the objective

Input: $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and $\epsilon, \delta, D \geq 0$
Problem: $\mathrm{OPT_{lp}} = \min_{x \in \mathbb{R}^n \mid \boldsymbol{A}x \geq b} c^\top x$
Promise: $\exists x_*^{\mathrm{lp}} \in \operatorname{argmin}_{x \in \mathbb{R}^n \mid \boldsymbol{A}x \geq b} c^\top x$ with $\left\| x_*^{\mathrm{lp}} \right\|_\infty \leq D$
Goal: find $x_{\epsilon,\delta}$ with $c^\top x \leq \mathrm{OPT_{lp}}$ and $\boldsymbol{A}x \geq b - \delta\vec{1}$

Linear Programming

Approach

- $p(x) \stackrel{\text{def}}{=} M \cdot \max\left\{0, \max_{i \in [m]} [b - \boldsymbol{A}x]_i\right\}$
- $\mathrm{OPT_p} = \min_{x \in \mathbb{R}^n \mid \|x\|_\infty \leq R} c^\top x + p(x)$

Claim: For $M = (\epsilon + 2\|c\|_1 R)\delta^{-1}$ any ϵ-approximate minimizer to OPT_p problem is (ϵ, δ)-approximate linear program solution.

Theorem: Can compute (ϵ, δ)-approximate linear program solution in

$$\tilde{O}\left(\mathrm{nnz}(\boldsymbol{A}) \cdot \frac{D\|\boldsymbol{A}\|_{op,\infty}}{\delta} \max\left\{1, \frac{D\|c\|_1}{\epsilon}\right\}\right)$$

Proof of Theorem from Claim

Can write penalized problem as box-simplex

- $\bar{x} = D^{-1}x$ and $\bar{c} = Dc$
- $\overline{\boldsymbol{A}} = \begin{pmatrix} -DM\boldsymbol{A} \\ \vec{0}_n^\top \end{pmatrix}$ and $b = \begin{pmatrix} -Mb \\ \vec{0}_n^\top \end{pmatrix}$

Penalized problem is the same as

$$\min_{\bar{x} \in B_\infty^n} \bar{c}^\top \bar{x} + \max_{i \in [m+1]} [\overline{\boldsymbol{A}}\bar{x} - b]_i$$

Note that $\|\overline{\boldsymbol{A}}\|_{op,\infty} = O\left(DM\|\boldsymbol{A}\|_{op,\infty}\right)$ and $M/\epsilon = O(\delta^{-1}\max\{1, D\|c\|_1\epsilon^{-1}\})$

Linear Programming

Input: $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and $\epsilon, \delta, D \geq 0$
Problem: $\mathrm{OPT_{lp}} = \min_{x \in \mathbb{R}^n \mid \boldsymbol{A}x \geq b} c^\top x$
Promise: $\exists x_*^{\mathrm{lp}} \in \operatorname{argmin}_{x \in \mathbb{R}^n \mid \boldsymbol{A}x \geq b} c^\top x$ with $\left\| x_*^{\mathrm{lp}} \right\|_\infty \leq D$
Goal: find $x_{\epsilon,\delta}$ with $c^\top x \leq \mathrm{OPT_{lp}}$ and $\boldsymbol{A}x \geq b - \delta\vec{1}$

Approach

- $p(x) \stackrel{\text{def}}{=} M \cdot \max\left\{0, \max_{i \in [m]} [b - \boldsymbol{A}x]_i\right\}$
- $\mathrm{OPT_p} = \min_{x \in \mathbb{R}^n \mid \|x\|_\infty \leq R} c^\top x + p(x)$

Claim: For $M = (\epsilon + 2\|c\|_1 R)\delta^{-1}$ any ϵ-approximate minimizer to OPT_p problem is (ϵ, δ)-approximate linear program solution.

Theorem: Can compute (ϵ, δ)-approximate linear program solution in

$$\tilde{O}\left(\mathrm{nnz}(\boldsymbol{A}) \cdot \frac{D\|\boldsymbol{A}\|_{op,\infty}}{\delta} \max\left\{1, \frac{D\|c\|_1}{\epsilon}\right\}\right)$$

Proof of Claim

- Let x_ϵ be ϵ-approximate minimizer
- Since x_*^{lp} is feasible for penalized problem, $\mathrm{OPT_p} \leq \mathrm{OPT_{lp}}$
- $c^\top x_\epsilon + p(x_\epsilon) \leq \mathrm{OPT_p} + \epsilon \leq \mathrm{OPT_{lp}} + \epsilon$
- $p(x_\epsilon) \leq \epsilon + c^\top(x_*^{\mathrm{lp}} - x_\epsilon)$
- $c^\top\left(x_*^{\mathrm{lp}} - x_\epsilon\right) \leq \|c\|_1 \left\| x_*^{\mathrm{lp}} - x_\epsilon \right\|_\infty$
- $\left\| x_*^{\mathrm{lp}} - x_\epsilon \right\|_\infty \leq \left\| x_*^{\mathrm{lp}} \right\|_\infty + \|x_\epsilon\|_\infty$

Problem #3: Bipartite Matching

Maximum Cardinality (Bipartite) Matching (MCM)

- **Input**: undirected, bipartite graph $G = (V, E)$
- **Matching**: $M \subseteq E$ such that $e_1 \cap e_2 = \emptyset$ for all $e_1, e_2 \in M$ with $e_1 \neq e_2$
- **Problem**: compute matching M_* of maximum cardinality $|M_*|$
- **Goal**: find $(1 - \epsilon)$-approximate MCM, i.e. matching M_ϵ with $|M_\epsilon| \geq (1 - \epsilon)|M_*|$

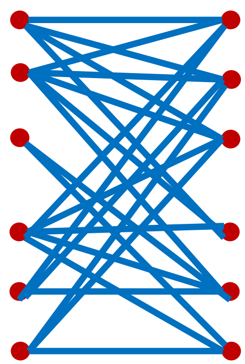

$G = (V, E)$

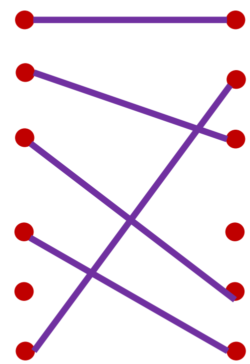

Matching M

MCM History *Fundamental, incredibly well-studied, notoriously difficult (to improve) problem.*

Year	Authors	Runtime $\tilde{O}(\cdot)$
1969-1973	Dinic, Karzanov, Hopcroft, Karp	$\|E\|\sqrt{\|V\|}$
1981	Ibarra, Moran	$\|V\|^{\omega}$
2013	Mądry	$\|E\|^{10/7}$
2020	Liu, **S**	$\|E\|^{11/8+o(1)}$
2020	Liu, Kathuria, **S**	$\|E\|^{4/3+o(1)}$
2020	Brand, Lee, Nanongkai, Peng, Saranurak, **S**, Song, Wang	$\|E\| + \|V\|^{1.5}$

Improvements since 1980s all use interior point methods which we may discuss on Friday.

Note: procedure will use very little graph structure.

- **Result**: can use box-simplex solver to compute $(1-\epsilon)$-approximate MCM in $\tilde{O}(|E|\epsilon^{-1})$ time and $\tilde{O}(\epsilon^{-1})$ depth
- Time matched by Dinic, Karzanov, Hopcroft, Karp and Allen-Zhu, Orecchia 2015
- Unaware of alternative method that gets this parallelism and this time.
- Alternative method either have large ϵ, $|E|$, or $|V|$ dependence
- Also, implementable semi-streaming (Assadi, Jambulapati, Jin, **S**, Tian 2021)

$w < 2.373$ is current fast matrix multiplication (FMM) constant [W13]

Approach

- **Input**: undirected, bipartite graph $G = (V, E)$
- **Matching**: $M \subseteq E$; $e_1 \cap e_2 = \emptyset$ for all $e_1, e_2 \in M$ with $e_1 \neq e_2$
- **Problem**: compute matching M_* maximizing $|M_*|$
- **Goal**: matching M_ϵ with $|M_\epsilon| \geq (1 - \epsilon)|M_*|$

$N(a) \stackrel{\text{def}}{=} \{b \in V \mid \{a, b\} \in E\}$ denotes the neighbors of A

Fractional Matching: in the MCM problem $f \in \mathbb{R}^E_{\geq 0}$ is a fractional matching if for all $a \in V$ it is the case that $\sum_{b \in N(a)} f_{\{a,b\}} \leq 1$.

Theorem [GPST91]: There is an algorithm which given any fractional matching $f \in \mathbb{R}^E_{\geq 0}$ can compute an integral matching of cardinality at least $\|f\|_1$ in time $\tilde{O}(|E|)$ and depth $\tilde{O}(1)$.

Corollary: The minimum ℓ_1-norm of a fractional matching is $|M_*|$ and it suffices to compute a fractional matching of ℓ_1-norm $\geq (1 - \epsilon)|M_*|$.

Linear Algebraic Representation

Unsigned (edge-vertex) Incidence Matrix: $|\boldsymbol{B}| \in \mathbb{R}^{E \times V}$ with

$$|\boldsymbol{B}|_{\{a,b\},c} = \begin{cases} 1 & c \in \{a,b\} \\ 0 & \text{otherwise} \end{cases} \text{ for all } \{a,b\} \in E \text{ and } c \in V$$

Lemma: $f \in \mathbb{R}^{E}_{\geq 0}$ is a fractional matching if and only if $|\boldsymbol{B}|^\top f \leq 1$.

Proof: $[|\boldsymbol{B}|^\top f]_a = \sum_{\{b,c\} \in E} f_{\{b,c\}} \, |\boldsymbol{B}|_{\{a,b\},c} = \sum_{b \in N(a)} f_{\{a,b\}}$

Upshot: it suffices to solve

$$\max_{f \in \mathbb{R}^{E}_{\geq 0} \,|\, |\boldsymbol{B}|^\top f \leq \vec{1}} \vec{1}^\top f \text{ or equivalently } \min_{f \in \mathbb{R}^{E}_{\geq 0} \,|\, |\boldsymbol{B}|^\top f \leq \vec{1}} \left(-\vec{1}\right)^\top f$$

Penalty and Rounding

In contrast to previous problem where we just solved approximately and bounded how infeasible, here we add a penalty term that allows us to reason more directly about obtaining a feasible solution.

Overflow (excess): $\mathrm{overflow}(f) \stackrel{\mathrm{def}}{=} \max\{\vec{0}, |\boldsymbol{B}|^\top f - \vec{1}\}$ entrywise

Note: $f \in \mathbb{R}^E$ is a fractional matching if and only if $\mathrm{overflow}(f) = \vec{0}$

Lemma: given $f \in \mathbb{R}^E_{\geq 0}$ let $\tilde{f} \in \mathbb{R}^E$ be defined for all $\{a, b\} \in E$ with $\tilde{f}_{\{a,b\}} = 0$ if $f_{\{a,b\}} = 0$ and otherwise

$$\tilde{f}_{\{a,b\}} = f_{\{a,b\}} \left(1 - \max\left\{\frac{[\mathrm{overflow}(f)]_a}{[|\boldsymbol{B}|^\top f]_a}, \frac{[\mathrm{overflow}(f)]_b}{[|\boldsymbol{B}|^\top f]_b}\right\}\right)$$

Then $0 \leq \tilde{f} \leq f$, $\tilde{f}$ is a fractional matching, and $\left\|f - \tilde{f}\right\|_1 \leq \|\mathrm{overflow}(f)\|_1$.

Proof: $f_{\{a,b\}} \cdot \frac{[\mathrm{overflow}(f)]_a}{[|\boldsymbol{B}|^\top f]_a}$ is the relative contribution of $f_{\{a,b\}}$ to overflow

Upshot: $-|M_*| = \min_{f \in \mathbb{R}^E_{\geq 0}} -\vec{1}^\top f + \sum_{a \in V} [\mathrm{overflow}(f)]_a$ and given any ϵ-additive minimizer can compute matching of size $\geq |M_*| - \epsilon$ in time $\tilde{O}(|E|)$.

The Result

- $\text{overflow}(f) \stackrel{\text{def}}{=} \max\{\vec{0}, |\boldsymbol{B}|^\top f - \vec{1}\}$
- $\epsilon|M_*|$ additive approximation to $\min_{f \in \mathbb{R}^E_{\geq 0}} -\vec{1}^\top f + \sum_{a \in V} [\text{overflow}(f)]_a$ suffices

Question #1: how to encode $\text{overflow}(f)$?

- Tool: $\max\{0, a\} = \frac{1}{2}[a + |a|]$
- Suffices to compute $\epsilon|M_*|$ additive approximation to

$$\min_{f \in \mathbb{R}^E_{\geq 0}} -\vec{1}^\top f + \frac{1}{2}\vec{1}^\top |\boldsymbol{B}|^\top f + \frac{1}{2}|V| + \left\| \frac{1}{2}|\boldsymbol{B}|^\top f - \vec{1} \right\|_1$$

What is $|\boldsymbol{B}|\vec{1}$? $=2 \cdot \vec{1}$!

Question #2: how to put f in simplex?

- Suppose $\nu \geq |M_*|$, then suffices to work with $x = \left(\frac{1}{\nu} f, 1 - \frac{1}{\nu}\|f\|_1\right) \in \Delta^{|E|+1}$
- Let $b = (-\frac{\nu}{2}\vec{1}_{|E|}, 0)$, and let $\boldsymbol{A} = \frac{\nu}{2}|\boldsymbol{B}|^\top$ with 0 column added
- Suffices to compute $\epsilon|M_*|$ additive approximation to $\min_{x \in \Delta^{|E|+1}} b^\top x + \|\boldsymbol{A}x - \vec{1}\|_1$
- Suffices to compute $\epsilon|M_*|$ additive approximation to $\max_{x \in \Delta^{|E|+1}} -b^\top x - \|\boldsymbol{A}x - \vec{1}\|_1$
- Note that $\|\boldsymbol{A}\|_{\text{op},\infty} = \nu$ so can solve in $\tilde{O}\left(\frac{|E|\nu}{\epsilon|\boldsymbol{M}_*|}\right)$.
- Get result by picking ν as every power of 2 between 0 and $2|V|$!

Can also computing 2 approximation by greedy.

Improvable?

Theorem: Given any algorithm which compute an ϵ-approximate MCM for any input $\epsilon \in (0,1)$ in time $\tilde{O}(|E|\epsilon^{-\delta})$ for some fixed constant δ, there is an algorithm that computes exact MCM in time $\tilde{O}(|E| \cdot |V|^{\frac{\delta}{1+\delta}})$.

Proof

- Given any ϵ-approximate MCM, there are at most $\epsilon|M_*| \leq \epsilon|V|$ more edges that could be matched.
- Augmenting paths finds at least one more matched edge in time $O(|E|)$
- Total time: $\tilde{O}(|E|\epsilon^{-\delta} + \epsilon|E||V|)$ solving for δ yields result

Implication: $\tilde{O}(|E|\epsilon^{-1})$ time $(1-\epsilon)$-approximate MCM yields $\tilde{O}(|E|\sqrt{|V|})$ time exact MCM
Barrier to improving: only improvements known to date use interior point methods

Introduce more broadly

Problem #4: Flow Problems

Natural family of problems in combinatorial optimization.

- Graph $G = (V, E)$
- Vertices $s, t \in V$

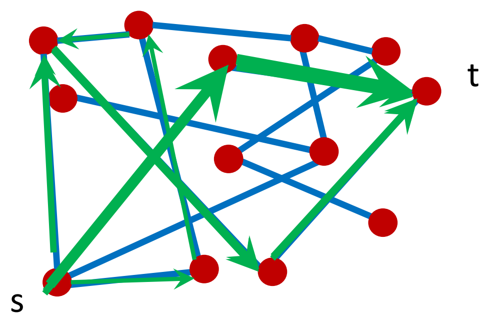

Goal
Send 1 unit of flow, $f \in \mathbb{R}^E$, between s and t in the "best" way possible.

If instead of sending 1 unit of flow from s to t, route arbitrary demand problem is called transshipment and is non-trivial

What should we minimize?

Shortest Path

$\tilde{O}(|E|)$

Length
$\sum_{e \in E} |f_e|$

$\|f\|_1$

Electric Flow
Laplacian System Solving

$\tilde{O}(|E|)$
[ST04]

Energy
$\sum_{e \in E} |f_e|^2$

$\|f\|_2$

See Rasmus Kyng's talks

Maximum Flow

$\tilde{O}(\min\{|E|^{3/2}, |E| \cdot |V|^{2/3}\})$
[K73,ET75,GR98]

No improvement until 2013, will discuss Friday.

Congestion
$\max_{e \in E} |f_e|$

$\|f\|_\infty$

Focus for today

Introduce maximum flow problem more formally

The Maximum Flow Problem

Graph $G = (V, E)$

- n vertices V
- m edges E

Capacities

- $u \in \{1, \ldots, U\}^E$

Terminals

- Source $s \in V$
- Sink $t \in V$

Value of Flow
total flow leaving s or entering t

$s \to t$ Flow
flow in = flow out
for all $v \notin \{s, t\}$

Capacity Constraints

- Directed: $f_e \in [0, u_e]$
- Undirected: $f_e \in [-u_e, u_e]$

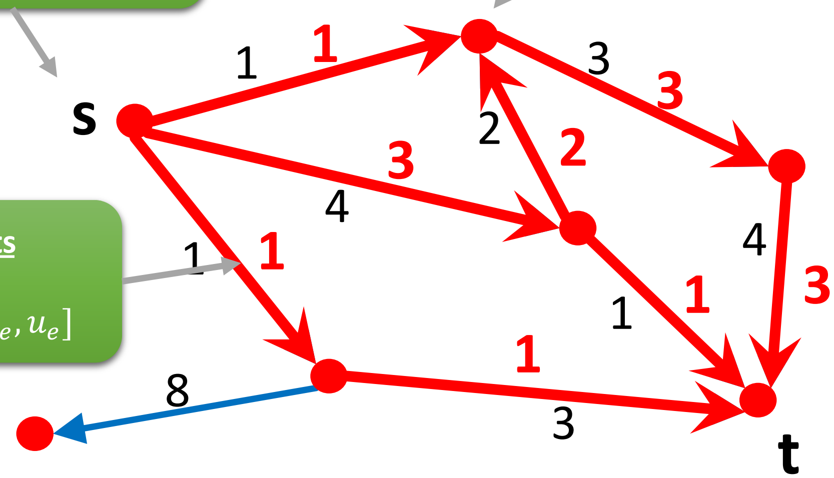

Goal
compute maximum $s \to t$ flow

Flow
$f \in \mathbb{R}^E$ where $f_e =$ amount of flow on edge e

Note: there are additional improvements with $\log(\epsilon^{-1})$*.*
Such results give exact directed flow algorithms.

Undirected Maxflow

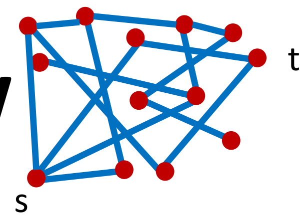

$\underline{(1-\epsilon)\text{-Approximate Flow}}$
feasible $s \to t$ flow of value $\geq (1-\epsilon)OPT$

	Authors	Time for ϵ-Approximate Undirected Flow	Capacitated (U $\neq$ 1)
ℓ_1-*ish*	⋮	⋮	⋮
	[Kar98]	$\tilde{O}(m\sqrt{n}\epsilon^{-1})$	Yes
ℓ_2	[CKMST11]	$\tilde{O}(mn^{1/3}\epsilon^{-11/3})$	Yes
	[LRS13]	$\tilde{O}(mn^{1/3}\epsilon^{-2/3}$	No
ℓ_∞	[S13,KLO**S**14]	$O(m^{1+o(1)}\epsilon^{-2})$	Yes
	[P16]	$\tilde{O}(m\epsilon^{-2})$	Yes
	[S17]	$\tilde{O}(m\epsilon^{-1})$	Yes
	[**S**T18]	$\tilde{O}(m+\sqrt{mn}\epsilon^{-1})$	Yes

How?
Work more directly in ℓ_∞. Reduction to and methods for box-simplex-like games.

Step 1 *(Combinatorial Advance)*
Build coarse ℓ_∞-approximator (e.g. oblivious routing or congestion approximator) to change representation.

Step 2 (*Optimization Advance*)
Apply iterative method to boost accuracy (e.g. gradient descent, area-convex dual extrapolation, mirror prox, coordinate descent)

Note (*Further Implications*)
Parallel optimal transport [J**S**T19], streaming matching [J**S**T20], optimization methods [C**S**T21]

Talk Plan

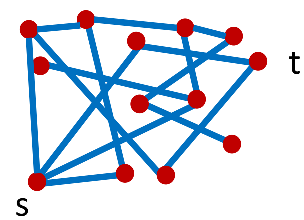

$(1-\epsilon)$-Approximate Flow
feasible $s \to t$ flow of value $\geq (1-\epsilon)OPT$

	Authors	Time for ϵ-Approximate Undirected Flow	Capacitated ($U \neq 1$)
ℓ_1-ish	⋮	⋮	⋮
	[Kar98]	$\tilde{O}(m\sqrt{n}\epsilon^{-1})$	Yes
ℓ_2	[CKMST11]	$\tilde{O}(mn^{1/3}\epsilon^{-11/3})$	Yes
	[LRS13]	$\tilde{O}(mn^{1/3}\epsilon^{-2/3}$	No
ℓ_∞	[S13,KLO**S**14]	$O(m^{1+o(1)}\epsilon^{-2})$	Yes
	[P16]	$\tilde{O}(m\epsilon^{-2})$	Yes
	[S17]	$\tilde{O}(m\epsilon^{-1})$	Yes
	[**S**T18]	$\tilde{O}(m+\sqrt{mn}\epsilon^{-1})$	Yes

- Talk 1 & 2: Focus on $\tilde{O}(m\epsilon^{-1})$ runtime.
- Talk 3: Discuss state-of-the art small ϵ results

Thank you

Questions?

Aaron Sidford

Contact Info:

- email: *sidford@stanford.edu*
- website: *www.aaronsidford.com*