Pipage Rounding, Pessimistic Estimators and Matrix Concentration

Neil Olver

VU University Amsterdam

CWI Centrum Wiskunde & Informatica

BIRS, September 2018

Joint work with Nick Harvey

Thin trees

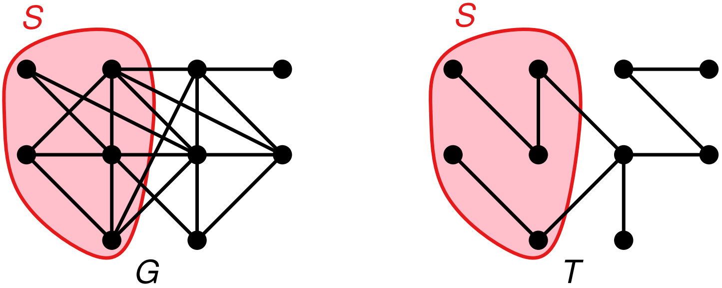

Spanning tree T of G is α-thin if

$$|\delta_T(S)| \leq \alpha|\delta_G(S)| \qquad \forall S.$$

Thin trees

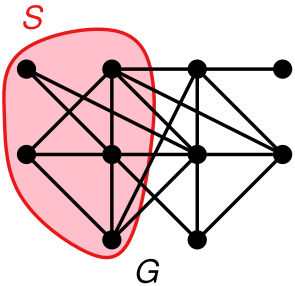

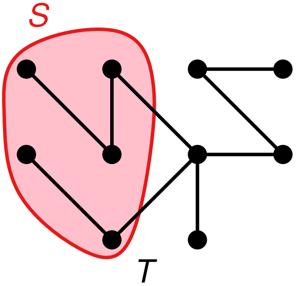

Spanning tree T of G is α-thin if

$$|\delta_T(S)| \leq \alpha |\delta_G(S)| \qquad \forall S.$$

edge connectivity

Goddyn's Conjecture

Every graph has a $O(1/K)$-thin tree, where $K = \min_e k_e$.

Would imply a (different) $O(1)$ approximation for asymmetric TSP. *Asadpour-Goemans-Madry-Oveis Gharan-Saberi '10*

An $O(\frac{1}{K}\log n)$-thin tree

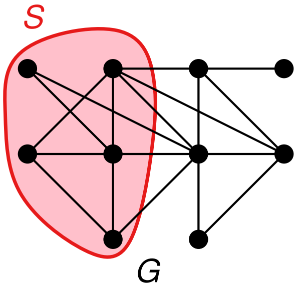

- H: include e independently with prob. $\min\{10\log n/K, 1\}$.
- For any S,

$$|\delta_H(S)| = \sum_{e\in\delta(S)} \mathbf{1}_{e\in H}$$

$$\mathbb{E}[|\delta_H(S)|] \leq \frac{10\log n}{K}|\delta_G(S)|.$$

Chernoff bounds $\Rightarrow 1 \leq |\delta_H(S)| \leq O(\frac{1}{K}\log n)|\delta_G(S)|$ w.h.p.

An $O(\frac{1}{K}\log n)$-thin tree

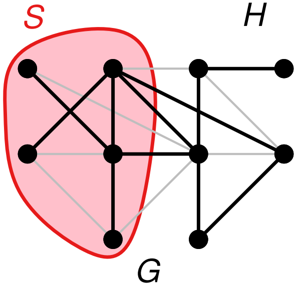

- H: include e independently with prob. $\min\{10\log n/K, 1\}$.
- For any S,

$$|\delta_H(S)| = \sum_{e\in\delta(S)} \mathbf{1}_{e\in H}$$

$$\mathbb{E}[|\delta_H(S)|] \leq \frac{10\log n}{K}|\delta_G(S)|.$$

Chernoff bounds $\Rightarrow 1 \leq |\delta_H(S)| \leq O(\frac{1}{K}\log n)|\delta_G(S)|$ w.h.p.

An $O(\frac{1}{K}\log n)$-thin tree

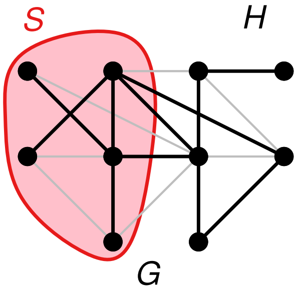

- H: include e independently with prob. $\min\{10\log n/K, 1\}$.
- For any S,

$$|\delta_H(S)| = \sum_{e\in\delta(S)} \mathbf{1}_{e\in H}$$

$$\mathbb{E}[|\delta_H(S)|] \le \frac{10\log n}{K}|\delta_G(S)|.$$

Chernoff bounds $\Rightarrow 1 \le |\delta_H(S)| \le O(\frac{1}{K}\log n)|\delta_G(S)|$ w.h.p.

- "Clever union bound" (Karger cut counting): holds for all S w.h.p. Hence any spanning tree of H is $O(\frac{1}{K}\cdot\log n)$-thin.

An $O(\frac{1}{K}\log n)$-thin tree

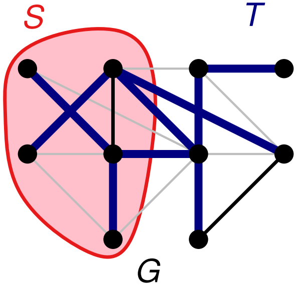

- H: include e independently with prob. $\min\{10 \log n/K, 1\}$.
- For any S,

$$|\delta_H(S)| = \sum_{e\in\delta(S)} \mathbf{1}_{e\in H}$$
$$\mathbb{E}[|\delta_H(S)|] \leq \frac{10\log n}{K}|\delta_G(S)|.$$

Chernoff bounds $\Rightarrow 1 \leq |\delta_H(S)| \leq O(\frac{1}{K}\log n)|\delta_G(S)|$ w.h.p.

- "Clever union bound" (Karger cut counting): holds for all S w.h.p.
  Hence any spanning tree of H is $O(\frac{1}{K}\cdot\log n)$-thin.

Can't improve the $O(\log n)$ using this approach, because we lose connectivity of H.

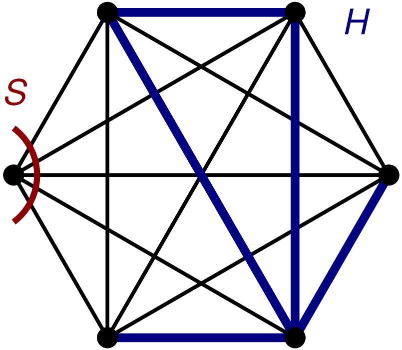

An $O(\frac{1}{K}\frac{\log n}{\log\log n})$-thin tree

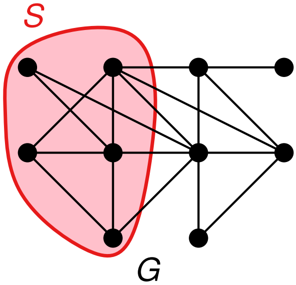

- Pick $z \in \mathbb{R}_+^E$ s.t. z is in the spanning tree polytope, and $z(\delta(S)) \leq \frac{2}{K}|\delta_G(S)| \quad \forall S$.

An $O(\frac{1}{K}\frac{\log n}{\log\log n})$-thin tree

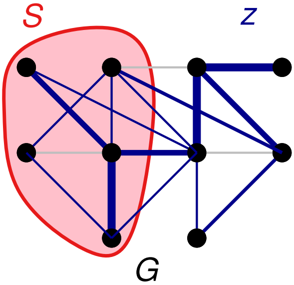

- Pick $z \in \mathbb{R}_+^E$ s.t. z is in the spanning tree polytope, and $z(\delta(S)) \leq \frac{2}{K}|\delta_G(S)| \quad \forall S$.

An $O(\frac{1}{K}\frac{\log n}{\log\log n})$-thin tree

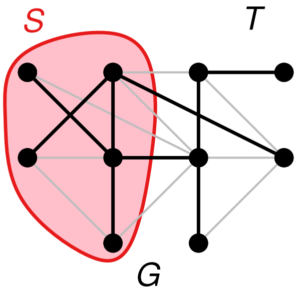

- Pick $z \in \mathbb{R}_+^E$ s.t. z is in the spanning tree polytope, and $z(\delta(S)) \le \frac{2}{K}|\delta_G(S)| \quad \forall S$.
- Randomly round to a tree T s.t. $\mathbb{P}[e \in T] = z_e$ for all e.

Max entropy distribution

Asadpour et al. '10

Pipage rounding

Chekuri-Vondrak-Zenklusen '10

An $O(\frac{1}{K}\frac{\log n}{\log\log n})$-thin tree

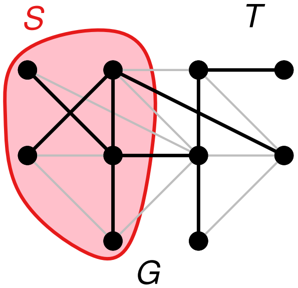

- Pick $z \in \mathbb{R}_+^E$ s.t. z is in the spanning tree polytope, and $z(\delta(S)) \leq \frac{2}{K}|\delta_G(S)| \quad \forall S$.
- Randomly round to a tree T s.t. $\mathbb{P}[e \in T] = z_e$ for all e.

Max entropy distribution

Asadpour et al. '10

Pipage rounding

Chekuri-Vondrak-Zenklusen '10

- Both yield negatively dependent distributions. Hence Chernoff bounds still hold for upper tail; yields $\log\log n$ improvement.

A difficulty with thin trees

$$|\delta_T(S)| \leq \alpha |\delta_G(S)| \qquad \forall S$$

How can one certify that a spanning tree T is α-thin?

Even approximately?

Laplacians and spectrally thin trees

Laplacian of a graph G with weights w

$$L_G = \sum_{e \in E} w_e L_e.$$

$$L_{\{i,j\}} = \begin{array}{c} \\ i \\ \\ j \end{array} \begin{array}{c} \begin{array}{ccc} i & \quad & j \end{array} \\ \begin{pmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{pmatrix} \end{array}$$

Löwner ordering: $A \preceq B$ if $B - A$ is PSD, i.e., $\boldsymbol{x}^T A \boldsymbol{x} \leq \boldsymbol{x}^T B \boldsymbol{x} \;\; \forall \boldsymbol{x}$.

Spanning tree T of G is α-spectrally thin if $L_T \preceq \alpha L_G$.

α-spectrally thin tree $\quad \Rightarrow \quad$ α-thin tree.

Laplacians and spectrally thin trees

Laplacian of a graph G with weights w

$$L_G = \sum_{e \in E} w_e L_e.$$

$$L_{\{i,j\}} = \begin{array}{cc} & \begin{array}{ccc} i & & j \end{array} \\ \begin{array}{c} i \\ \\ j \end{array} & \begin{pmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{pmatrix} \end{array}$$

Löwner ordering: $A \preceq B$ if $B - A$ is PSD, i.e., $\boldsymbol{x}^T A \boldsymbol{x} \leq \boldsymbol{x}^T B \boldsymbol{x} \ \ \forall \boldsymbol{x}$.

Spanning tree T of G is α-spectrally thin if $L_T \preceq \alpha L_G$.

α-spectrally thin tree $\quad \Rightarrow \quad$ α-thin tree.

Can be checked efficiently (compute $\lambda_{\max}(L_G^{\dagger/2} L_T L_G^{\dagger/2})$).

Spectrally thin trees

Goddyn's Conjecture

Every graph has a $O(1/K)$-thin tree.

Spectrally thin trees

Goddyn's Conjecture

Every graph has a $O(1/K)$-thin tree.

- Cannot hope for a $O(1/K)$-spectrally thin tree in general. Lowerbound is $O(\sqrt{n}/K)$. *Goemans; de Carli Silva et al.*
- Nonetheless, useful tool for providing certificates of thinness. *Anari & Oveis-Gharan – upcoming talks*

Spectrally thin trees

Goddyn's Conjecture

Every graph has a $O(1/K)$-thin tree.

$$K = \min_e k_e \qquad \longrightarrow \qquad C = \min_e c_e$$

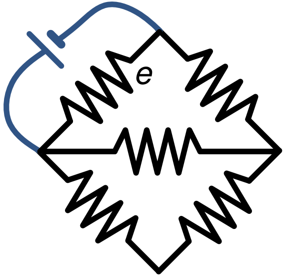

c_e = amount of current that flows if 1V battery attached to endpoints of e.

$c_e \leq k_e$

Spectrally thin trees

Goddyn's Conjecture

Every graph has a $O(1/K)$-thin tree.

$$K = \min_e k_e \qquad \longrightarrow \qquad C = \min_e c_e$$

Theorem *Marcus-Spielman-Srivastava '13*

Always exists a $O(1/C)$-spectrally thin tree.

- Implication of their solution to the Kadison-Singer Problem.
- Not constructive.

Spectrally thin trees

Goddyn's Conjecture

Every graph has a $O(1/K)$-thin tree.

$$K = \min_e k_e \qquad \longrightarrow \qquad C = \min_e c_e$$

Theorem *Marcus-Spielman-Srivastava '13*

Always exists a $O(1/C)$-spectrally thin tree.

- Implication of their solution to the Kadison-Singer Problem.
- Not constructive.

- What can we achieve constructively with simple randomized rounding?

Rounding for spectrally thin trees

- Let $z_e = 1/c_e$; then
$$\sum_{e \in E} z_e L_e \preceq \frac{1}{C} L_G.$$
- H: include e independently with prob. $\min\{10 \log n \cdot z_e, 1\}$.

Rounding for spectrally thin trees

- Let $z_e = 1/c_e$; then
$$\sum_{e \in E} z_e L_e \preceq \frac{1}{C} L_G.$$
- H: include e independently with prob. $\min\{10 \log n \cdot z_e, 1\}$.

$$L_H = \sum_{e \in E} \hat{X}_e L_e \qquad \text{where} \quad \hat{X}_e = \mathbf{1}_{e \in T}$$
$$\mathbb{E}[L_H] = 10 \log n \sum_{e \in E} z_e L_e \preceq \frac{10 \log n}{C} L_G.$$

Matrix Chernoff implies that whp,

$$H \text{ connected} \quad \text{and} \quad L_H \preceq O(\tfrac{1}{C} \cdot \log n) L_G.$$

Rounding for spectrally thin trees

Matrix Chernoff *Tropp '12*

Given $Y_1, \ldots, Y_m$, with $0 \preceq Y_i \preceq RI$. Let $S = \sum_i Y_i$, $\mu = \lambda_{\max}(\mathbb{E}S)$. Then

$$\mathbb{P}[\lambda_{\max}(S) > (1+\delta)\mu] \leq n \cdot \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^{\mu/R}.$$

$$L_H = \sum_{e \in E} \hat{X}_e L_e \qquad \text{where} \quad \hat{X}_e = \mathbf{1}_{e \in T}$$

$$\mathbb{E}[L_H] = 10 \log n \sum_{e \in E} z_e L_e \preceq \frac{10 \log n}{C} L_G.$$

Matrix Chernoff implies that whp,

$$H \text{ connected} \quad \text{and} \quad L_H \preceq O(\tfrac{1}{C} \cdot \log n) L_G.$$

Rounding for spectrally thin trees

- $z_e = 1/c_e$ is in the spanning tree polytope.

$$\sum_{e \in E} z_e L_e \preceq \frac{1}{C} L_G.$$

- Randomly round to a spanning tree T s.t. $\mathbb{P}[e \in T] = z_e$ for all e.

If we can obtain the same matrix concentration as independent rounding, T is $O(\frac{1}{C}\frac{\log n}{\log\log n})$-spectrally thin.

Rounding for spectrally thin trees

- $z_e = 1/c_e$ is in the spanning tree polytope.

$$\sum_{e \in E} z_e L_e \preceq \frac{1}{C} L_G.$$

- Randomly round to a spanning tree T s.t. $\mathbb{P}[e \in T] = z_e$ for all e.

If we can obtain the same matrix concentration as independent rounding, T is $O(\frac{1}{C}\frac{\log n}{\log\log n})$-spectrally thin.

Max entropy distribution **Pipage rounding**

Rounding for spectrally thin trees

- $z_e = 1/c_e$ is in the spanning tree polytope.

$$\sum_{e \in E} z_e L_e \preceq \frac{1}{C} L_G.$$

- Randomly round to a spanning tree T s.t. $\mathbb{P}[e \in T] = z_e$ for all e.

If we can obtain the same matrix concentration as independent rounding, T is $O(\frac{1}{C}\frac{\log n}{\log\log n})$-spectrally thin.

Max entropy distribution **Pipage rounding**

Open!

Rounding for spectrally thin trees

- $z_e = 1/c_e$ is in the spanning tree polytope.

$$\sum_{e \in E} z_e L_e \preceq \frac{1}{C} L_G.$$

- Randomly round to a spanning tree T s.t. $\mathbb{P}[e \in T] = z_e$ for all e.

If we can obtain the same matrix concentration as independent rounding, T is $O(\frac{1}{C} \frac{\log n}{\log \log n})$-spectrally thin.

Max entropy distribution

Open!

Pipage rounding

Theorem *Harvey-O. '14*

Gives $O(\frac{1}{C} \cdot \frac{\log n}{\log \log n})$-spectrally thin tree.

Matrix Chernoff for pipage rounding

Theorem *Harvey-O. '14*

Given matroid base polytope $P \subseteq \mathbb{R}_+^m$, $x \in P$, and PSD matrices $L_1, \ldots, L_m$.

Let $\hat{X} \in \{0,1\}^m$ be the (random) outcome of "pipage rounding" starting from x.

Then $\sum_i \hat{X}_i L_i$ satisfies the same matrix Chernoff bounds as independent rounding from x.

Pipage rounding

Ageev-Sviridenko '04, Srinivasan '01, Calinescu et al. '07, Chekuri et al. '10

Let P be a matroid base polytope (e.g., spanning tree polytope).

- Swap directions:
  $\boldsymbol{e}_i - \boldsymbol{e}_j$ for $i \neq j$.

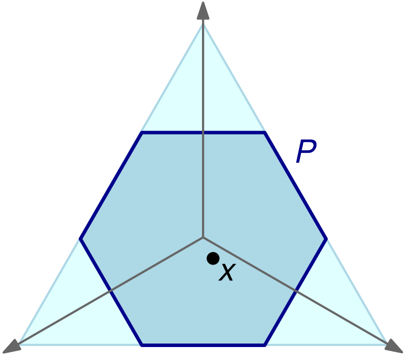

Pipage rounding

Ageev-Sviridenko ’04, Srinivasan ’01, Calinescu et al. ’07, Chekuri et al. ’10

Let P be a matroid base polytope (e.g., spanning tree polytope).

- Swap directions:
  $\boldsymbol{e}_i - \boldsymbol{e}_j$ for $i \neq j$.

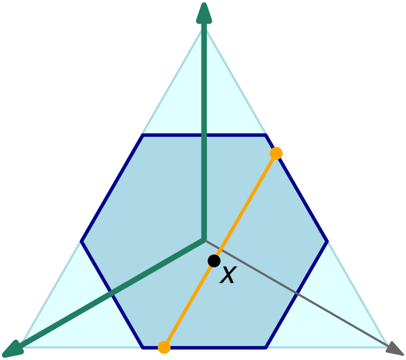

Pipage rounding

Ageev-Sviridenko '04, Srinivasan '01, Calinescu et al. '07, Chekuri et al. '10

Let P be a matroid base polytope (e.g., spanning tree polytope).

- Swap directions:
  $\boldsymbol{e}_i - \boldsymbol{e}_j$ for $i \neq j$.

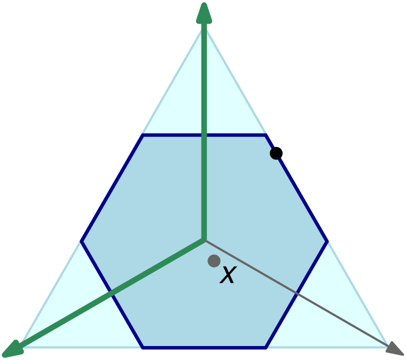

Pipage rounding

Ageev-Sviridenko '04, Srinivasan '01, Calinescu et al. '07, Chekuri et al. '10

Let P be a matroid base polytope (e.g., spanning tree polytope).

- Swap directions:
  $\boldsymbol{e}_i - \boldsymbol{e}_j$ for $i \neq j$.

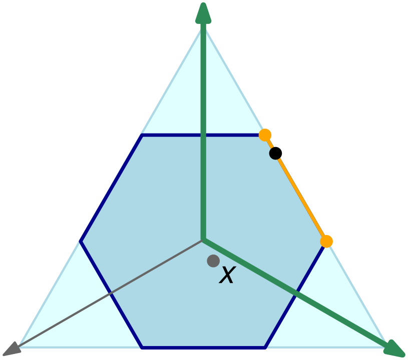

Pipage rounding

Ageev-Sviridenko '04, Srinivasan '01, Calinescu et al. '07, Chekuri et al. '10

Let P be a matroid base polytope (e.g., spanning tree polytope).

- Swap directions: $\boldsymbol{e}_i - \boldsymbol{e}_j$ for $i \neq j$.
- Martingale: $\mathbb{E}[\hat{X}] = x$.
- $\hat{X}$ satisfies negative cylinder dependence.

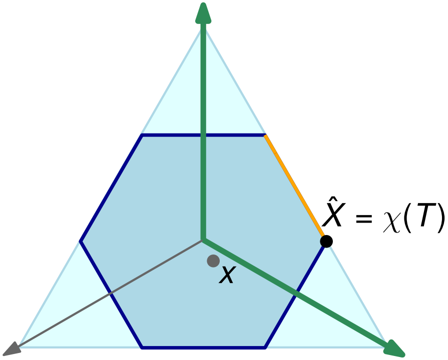

Pipage rounding

Ageev-Sviridenko '04, Srinivasan '01, Calinescu et al. '07, Chekuri et al. '10

Let P be a matroid base polytope (e.g., spanning tree polytope).

- Swap directions: $\boldsymbol{e}_i - \boldsymbol{e}_j$ for $i \neq j$.
- Martingale: $\mathbb{E}[\hat{X}] = x$.
- $\hat{X}$ satisfies negative cylinder dependence.

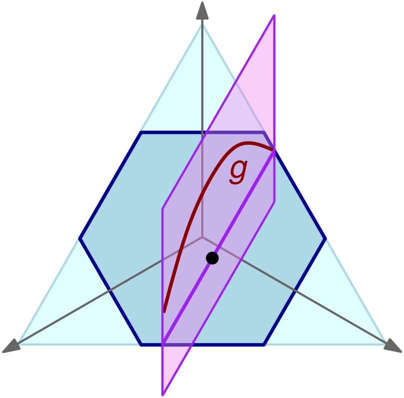

- If $g : \mathbb{R}^m \to \mathbb{R}$ is concave under swaps, then $\mathbb{E}[g(\hat{X})] \leq g(x)$.

$z \to g(x + z(\boldsymbol{e}_i - \boldsymbol{e}_j))$ concave for any $x \in P$, $i \neq j$

Warmup: Chernoff bounds for pipage rounding

Let $D(x)$ be the product distribution on $\{0,1\}^m$ where $\mathbb{P}[X_i = 1] = x_i$.

Usual Chernoff proof: for all $\theta > 0$,

$$\begin{aligned}\mathbb{P}_{X\sim D(x)}[\sum_i X_i > t] &= \mathbb{P}_{X\sim D(x)}[e^{\theta \sum_i X_i} > e^{\theta t}] \\ &\leq e^{-\theta t}\mathbb{E}_{X\sim D(x)}[e^{\theta \sum_i X_i}] \\ &= e^{-\theta t}\prod \mathbb{E}_{X\sim D(x)}[e^{\theta X_i}] =: g_{t,\theta}(x).\end{aligned}$$

Then show that

$$\inf_{\theta>0} g_{(1+\delta)\mu,\theta}(x) \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu}.$$

Warmup: Chernoff bounds for pipage rounding

Let $D(x)$ be the product distribution on $\{0,1\}^m$ where $\mathbb{P}[X_i = 1] = x_i$.

Usual Chernoff proof: for all $\theta > 0$,

$$\begin{aligned}
\mathbb{P}_{X\sim D(x)}[\sum_i X_i > t] &= \mathbb{P}_{X\sim D(x)}[e^{\theta \sum_i X_i} > e^{\theta t}] \\
&\leq e^{-\theta t}\mathbb{E}_{X\sim D(x)}[e^{\theta \sum_i X_i}] \\
&= e^{-\theta t} \prod \mathbb{E}_{X\sim D(x)}[e^{\theta X_i}] =: g_{t,\theta}(x).
\end{aligned}$$

Then show that

$$\inf_{\theta>0} g_{(1+\delta)\mu,\theta}(x) \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu}.$$

Claim

$g_{t,\theta}$ is concave under swaps for any $t \in \mathbb{R}, \theta > 0$.

Warmup: Chernoff bounds for pipage rounding

$$\mathbb{P}_{X \sim D(x)}\left[\sum X_i > t\right] \leq g_{t,\theta}(x)$$

$$x \xrightarrow{\text{pipage rounding}} \hat{X}.$$

$$\mathbb{P}[\sum_i \hat{X}_i > t] \leq \mathbb{E}[g_{t,\theta}(\hat{X})]$$
$$\leq g_{t,\theta}(x).$$

Hence get precisely the same tail bounds for pipage rounding as for independent rounding.

Noncommutative difficulties

Usual Chernoff proof: for all $\theta > 0$,

$$\begin{aligned}\mathbb{P}[\sum_i X_i > t] &= \mathbb{P}[e^{\theta \sum_i X_i} > e^{\theta t}] \\ &\leq e^{-\theta t}\mathbb{E}[e^{\theta \sum_i X_i}] \\ &= e^{-\theta t}\mathbb{E}[\prod e^{\theta X_i}] \\ &= e^{-\theta t}\prod \mathbb{E}[e^{\theta X_i}].\end{aligned}$$

Noncommutative difficulties

Usual Chernoff proof: for all $\theta > 0$,

$$\begin{aligned}
\mathbb{P}[\sum_i X_i > t] &= \mathbb{P}[e^{\theta \sum_i X_i} > e^{\theta t}] \\
&\leq e^{-\theta t}\mathbb{E}[e^{\theta \sum_i X_i}] \\
&= e^{-\theta t}\mathbb{E}[\prod e^{\theta X_i}] \\
&= e^{-\theta t}\prod \mathbb{E}[e^{\theta X_i}].
\end{aligned}$$

For matrices, $e^{A+B} \neq e^A \cdot e^B$!

Noncommutative difficulties

Usual Chernoff proof: for all $\theta > 0$,

$$\begin{aligned}
\mathbb{P}[\sum_i X_i > t] &= \mathbb{P}[e^{\theta \sum_i X_i} > e^{\theta t}] \\
&\leq e^{-\theta t}\mathbb{E}[e^{\theta \sum_i X_i}] \\
&= e^{-\theta t}\mathbb{E}[\prod e^{\theta X_i}] \\
&= e^{-\theta t} \prod \mathbb{E}[e^{\theta X_i}].
\end{aligned}$$

For matrices, $e^{A+B} \neq e^A \cdot e^B$!

Golden-Thompson: $\operatorname{tr} e^{A+B} \leq \operatorname{tr}(e^A \cdot e^B)$.

Lieb's Theorem.

Matrix Chernoff bounds for pipage rounding

Let $A_1, \ldots, A_m$ be $n \times n$ symmetric matrices with $0 \preceq A_i \preceq I$.

Tropp '12:

$$\mathbb{P}_{X \sim D(x)}[\lambda_{\max}(\textstyle\sum_i X_i A_i) > t] \leq \underbrace{e^{-\theta t} \operatorname{tr} \exp(\textstyle\sum_i \log \mathbb{E}_{X \sim D(x)}[e^{\theta X_i A_i}])}_{g_{t,\theta}(x)}$$

$$\inf_{\theta > 0} g_{(1+\delta)\mu,\theta}(x) \leq n \cdot \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu}.$$

Matrix Chernoff bounds for pipage rounding

Let $A_1, \ldots, A_m$ be $n \times n$ symmetric matrices with $0 \preceq A_i \preceq I$.

Tropp '12:

$$\mathbb{P}_{X \sim D(x)}[\lambda_{\max}(\textstyle\sum_i X_i A_i) > t] \leq \underbrace{e^{-\theta t} \operatorname{tr} \exp(\textstyle\sum_i \log \mathbb{E}_{X \sim D(x)}[e^{\theta X_i A_i}])}_{g_{t,\theta}(x)}$$

$$\inf_{\theta > 0} g_{(1+\delta)\mu,\theta}(x) \leq n \cdot \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu}.$$

Key Theorem

$g_{t,\theta}$ is concave under swaps for any $t \in \mathbb{R}, \theta > 0$.

$$x \xrightarrow{\text{pipage rounding}} \hat{X}.$$

$$\mathbb{P}\left[\lambda_{\max}(\textstyle\sum_i \hat{X}_i A_i) > t\right] \leq \mathbb{E}[g_{t,\theta}(\hat{X})] \leq g_{t,\theta}(x).$$

Lieb's theorem and a variant

Lieb '73 (used by Tropp '12):
If A, B symmetric and C PSD, then

$$z \to \operatorname{tr}\exp\big(A + \log(C + zB)\big)$$

is concave.

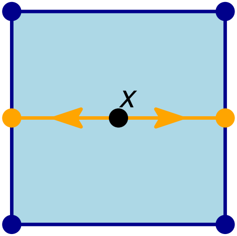

$z \to g_{t,\theta}(x + z\boldsymbol{e}_i)$

Harvey-O. '14:
If A symmetric, B_1, B_2 PSD and C_1, C_2 PD, then

$$z \to \operatorname{tr}\exp\big(A + \log(C_1 + zB_1) + \log(C_2 - zB_2)\big)$$

is concave.

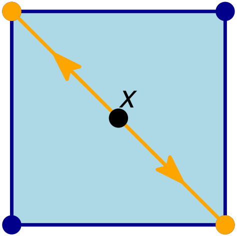

$z \to g_{t,\theta}(x + z(\boldsymbol{e}_i - \boldsymbol{e}_j))$

Conclusion

- MSS implies $O(1/C)$-spectral thin trees exist.
  Polynomial time algorithm?

- Do $O(1/K)$-thin trees exist?

  Anari & Oveis Gharan '15: $O(\text{polyloglog}\, n/K)$-thin trees exist

- Concentration bounds for negatively dependent sums of matrices?

Conclusion

- MSS implies $O(1/C)$-spectral thin trees exist.
  Polynomial time algorithm?

- Do $O(1/K)$-thin trees exist?

  Anari & Oveis Gharan '15: $O(\text{polyloglog}\, n/K)$-thin trees exist

- Concentration bounds for negatively dependent sums of matrices?

  Kyng-Song '18: A Chernoff-type bound for strongly Rayleigh measures.

  Implies that a max-entropy spanning tree satisfies
  $L_T \preceq O(\log n) L_G$ (but nothing better).

Conclusion

- MSS implies $O(1/C)$-spectral thin trees exist.
  Polynomial time algorithm?

- Do $O(1/K)$-thin trees exist?

  Anari & Oveis G exist

Thank you!

- Concentration bounds for negatively dependent sums of matrices?

  Kyng-Song '18: A Chernoff-type bound for strongly Rayleigh measures.

  Implies that a max-entropy spanning tree satisfies
  $L_T \preceq O(\log n) L_G$ (but nothing better).

Kyng-Song '18

Suppose

- $(X_1, \ldots, X_m) \in \{0,1\}^m$ is strongly Rayleigh, with $\sum_i X_i = k$ surely
- $A_1, \ldots, A_m$ are PSD, $A_i \preceq I$

Let $S = \sum_i X_i A_i$, $\mu = \|S\|$. Then for some universal $C > 0$

$$\mathbb{P}[\|S\| > (1+\delta)\mu] \leq n \cdot \exp\left(-C\frac{\mu\delta^2}{\log k + \delta}\right)^{\mu}$$

- Implies that $O(\log^2 n/\epsilon^2)$ random spanning trees (from max entropy distribution), with edge weights correctly chosen, is a $(1 \pm \epsilon)$-spectral sparsifier.