Multi-commodity Flows and Cuts in *Polymatroidal* Networks

Chandra Chekuri

Univ. of Illinois, Urbana-Champaign

Joint work with

Sreeram Kannan, Adnan Raja, Pramod Viswanath

(UIUC ECE Department)

Paper available at http://arxiv.org/abs/1110.6832

Max-flow Min-cut Theorem

[Ford-Fulkerson, Menger]

G=(V,E) directed graph with non-negative edge-capacities

max s-t flow value equal to min s-t cut value

if capacities *integral* max flow can be chosen to be *integral*

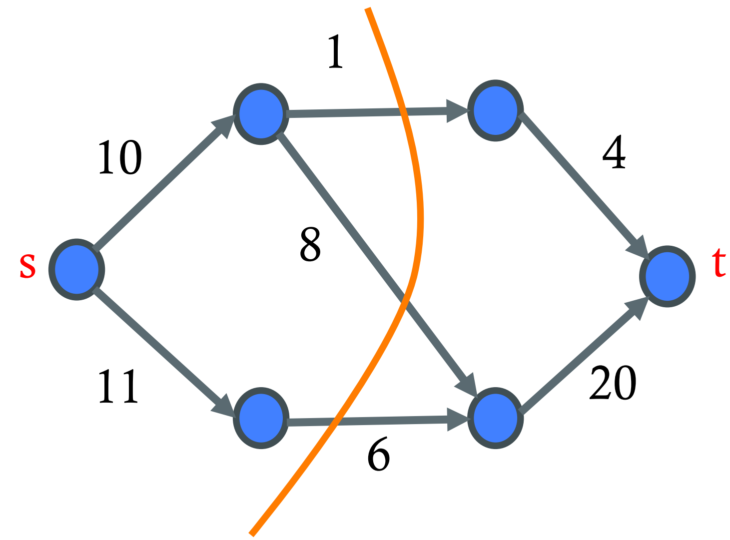

Multi-commodity Flows

Several pairs $(s_1,t_1),\ldots,(s_k,t_k)$ jointly use the network capacity to route their flow

$f_i(e)$: flow for pair i on edge e

$\sum_i f_i(e) \leq c(e)$ for all e

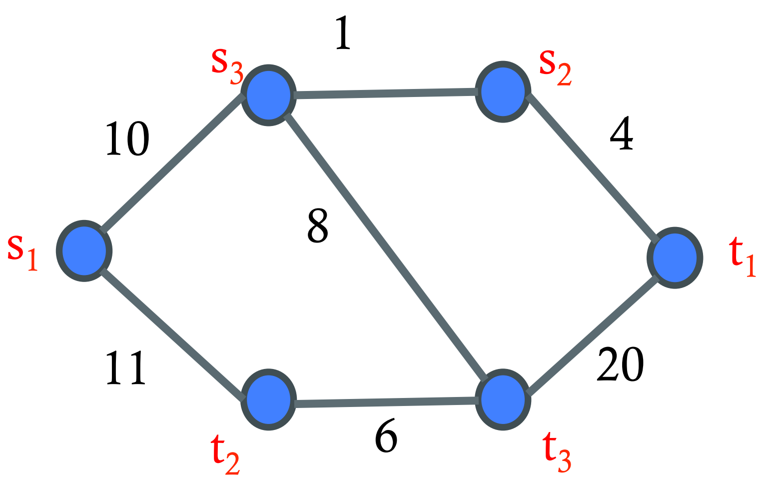

Max Throughput Flow and Min Multicut

$f_i(e)$: flow for pair i on edge e

$\sum_i f_i(e) \leq c(e)$ for all e

$\max \sum_i \mathrm{val}(f_i)$ (max throughput flow)

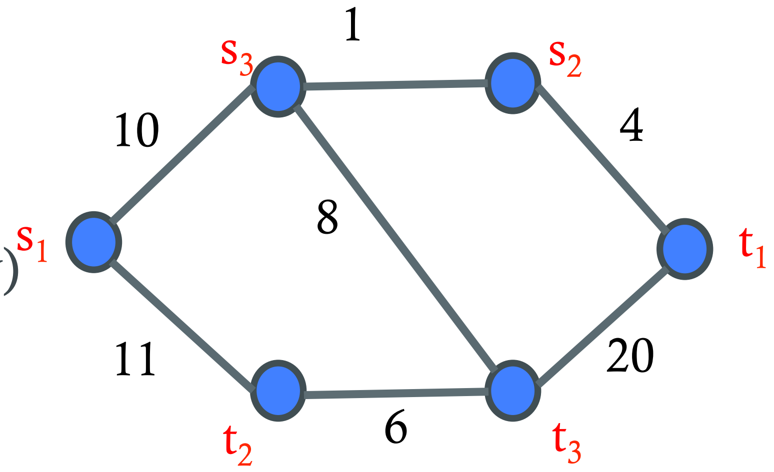

Max Throughput Flow and Min Multicut

$f_i(e)$: flow for pair i on edge e

$\sum_i f_i(e) \leq c(e)$ for all e

$\max \sum_i \mathrm{val}(f_i)$ (max throughput flow)

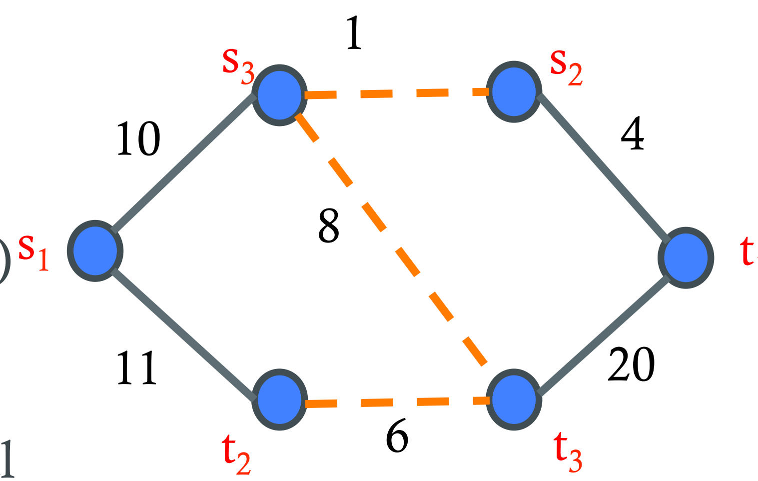

Multicut: set of edges whose removal disconnects all pairs

Max Throughput Flow $\leq$ Min Multicut Capacity

Max Concurrent Flow and Min Sparsest Cut

$f_i(e)$: flow for pair i on edge e

$\sum_i f_i(e) \leq c(e)$ for all e

$val(f_i) \geq \lambda D_i$ for all i

$\max \lambda$ (max concurrent flow)

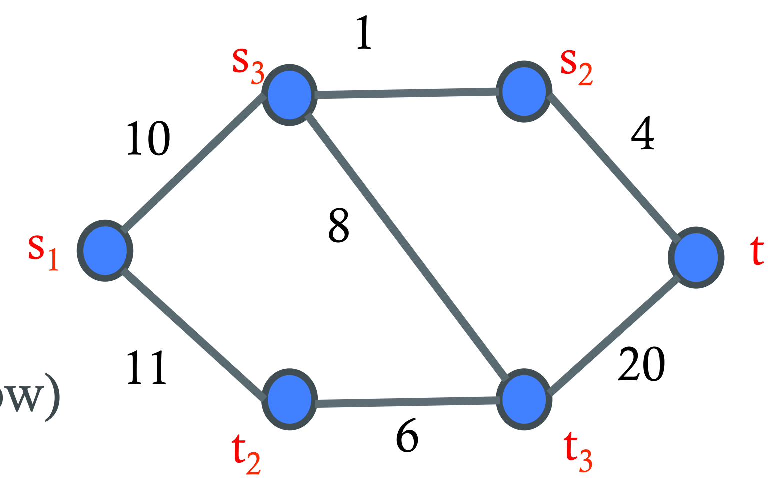

Max Concurrent Flow and Min Sparsest Cut

$f_i(e)$: flow for pair i on edge e

$\sum_i f_i(e) \leq c(e)$ for all e

$val(f_i) \geq \lambda D_i$ for all i

$\max \lambda$ (max concurrent flow)

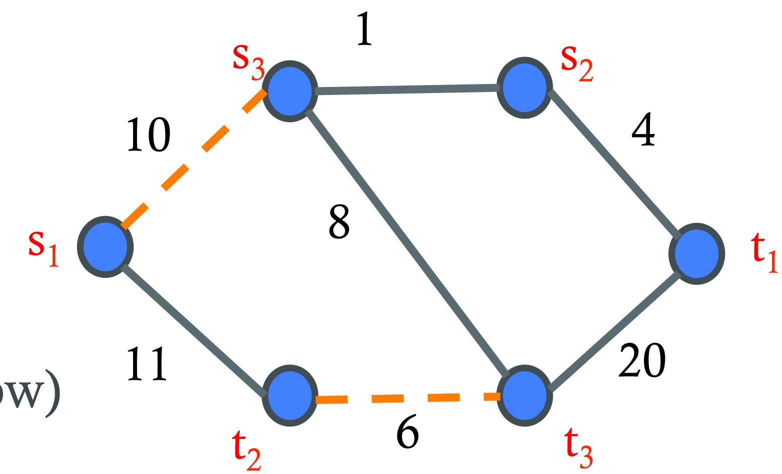

Sparsity of cut = capacity of cut / demand separated by cut

Max Concurrent Flow $\leq$ Min Sparsity

Flow-Cut Gap: Undir graphs

[Leighton-Rao'88] examples via expanders to show

Max Throughput Flow $\leq$ $O(1/\log k)$ Min Multicut

Max Concurrent Flow $\leq$ $O(1/\log k)$ Min Sparsity

$k = \Theta(n^2)$ in expander examples

Flow-Cut Gap: Undir graphs

[Leighton-Rao'88] for product multi-commodity flow

Max Concurrent Flow $\geq \Omega(1/\log k)$ Min Sparsity

[Garg-Vazirani-Yannakakis'93]

Max Throughput Flow $\geq \Omega(1/\log k)$ Min Multicut

[Linial-London-Rabinovich'95,Aumann-Rabani'95]

Max Concurrent Flow $\geq \Omega(1/\log k)$ Min Sparsity

Flow-Cut Gap: Undir graphs Node Capacities

[Feige-Hajiaghayi-Lee'05]

Max Concurrent Flow $\geq \Omega\,(1/\log k)$ Min Sparsity

[Garg-Vazirani-Yannakakis'93]

Max Throughput Flow $\geq \Omega(1/\log k)$ Min Multicut

Flow-Cut Gap: Dir graphs

[Saks-Samorodnitsky-Zosin'04]

Max Throughput Flow $\leq O(1/k)$ Min Multicut

[Chuzhoy-Khanna'07]

Max Throughput Flow $\leq O(1/n^{1/7})$ Min Multicut

[Agrawal-Alon-Charikar'07]

Max Throughput Flow $\geq \Omega(1/n^{11/23})$ Min Multicut

$\geq 1/k$ Min Multicut (trivial)

Flow-Cut Gap: Dir graphs

Symmetric demands: (s_i, t_i) and (t_i, s_i) for each pair and cut has to separate only one of the two

[Klein-Plotkin-Rao-Tardos'97]

Max Throughput Flow $\geq \Omega(1/\log^2 k)$ Min Multicut

Max Concurrent Flow $\geq \Omega(1/\log^3 k)$ Min Sparsity

[Even-Naor-Rao-Schieber'95]

Max Throu. Flow $\geq \Omega(1/\log n \log\log n)$ Min Multicut

Flow-Cut Gaps: Summary

k pairs in a graph G=(V,E)

- Θ(log k) for undir graphs
  - Throughput Flow vs Multicut
  - Concurrent Flow vs Sparsest Cut
  - Node-capacited flows [Feige-Hajiaghayi-Lee'05]
- O(polylog(k)) for dir graph with symmetric demands
- Polynomial-factor lower bounds for dir graphs

Polymatroidal Networks

Capacity of edges incident to v *jointly constrained* by a polymatroid (monotone non-neg submodular set func)

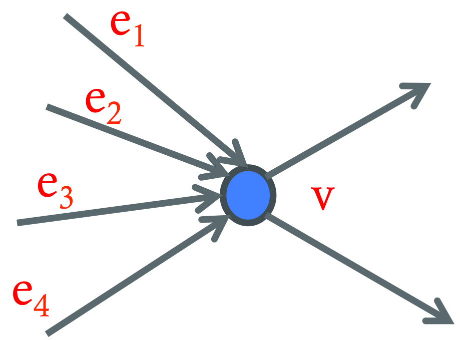

$\sum_{i \in S} c(e_i) \leq f(S)$ for every $S \subseteq \{1,2,3,4\}$

Detour: Network Information Theory

Question: What is the *information theoretic capacity* of a network?

Given $G=(V,E)$ and pairs $(s_1,t_1),\ldots,(s_k,t_k)$ and rates/demands $D_1,\ldots,D_k$: can the pairs use the network to successfully transmit information at these rates?

- Can use routing, (network) coding, and any other scheme …
- Network coding [Ahlswede-Cai-Li-Yeung'00]

Network Information Theory: Cut-Set Bound

Max Concurrent Rate $\leq$ Min Sparsity

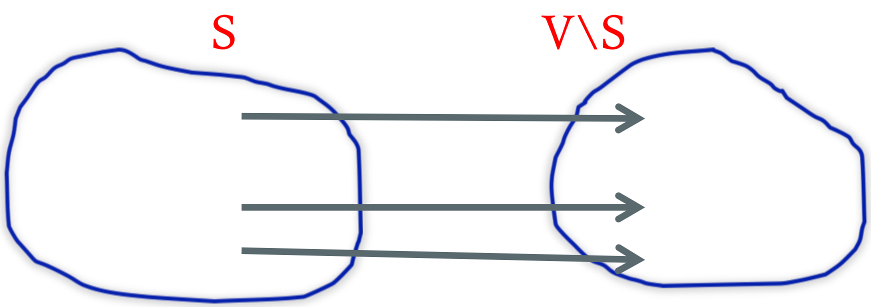

Network Information Theory

Max Concurrent Rate $\leq$ Min Sparsity

- In undirected graphs routing is near-optimal (within log factors). Follows from flow-cut gap upper bounds
- In directed graphs routing can be very far from optimal
- In directed graphs routing far from optimal even for multicast
- Capacity of networks poorly understood

Capacity of Wireless Networks

Capacity of wireless networks

Major issues to deal with:

- interference due to broadcast nature of medium
- noise

Capacity of wireless networks

Recent work: *understand/model/approximate wireless networks via wireline networks*

- Linear deterministic networks [Avestimehr-Diggavi-Tse'09]
  - *Unicast/multicast (single source)*. Connection to polylinking systems and submodular flows [Goemans-Iwata-Zenklusen'09]
- Polymatroidal networks [Kannan-Viswanath'11]
  - *Multiple unicast.*

Directed Polymatroidal Networks

[Lawler-Martel'82, Hassin'79]

Directed graph $G=(V,E)$

For each node v two polymatroids

- ρ_v^- with ground set $\delta^-(v)$
- ρ_v^+ with ground set $\delta^+(v)$

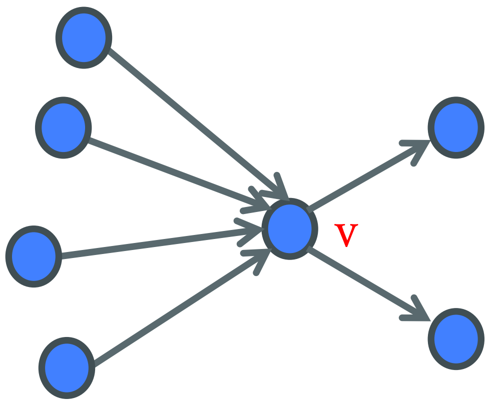

$\sum_{e \in S} f(e) \leq \rho_v^-(S)$ for all $S \subseteq \delta^-(v)$

$\sum_{e \in S} f(e) \leq \rho_v^+(S)$ for all $S \subseteq \delta^+(v)$

s-t flow

Flow from s to t: "standard flow" with polymatroidal capacity constraints

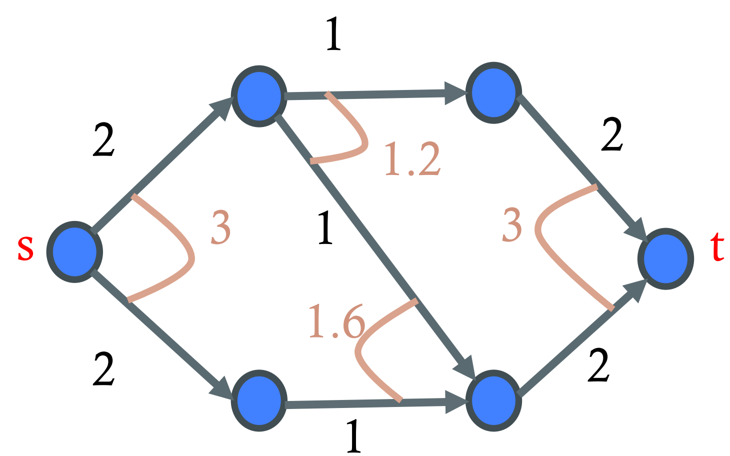

What is the cap. of a cut?

Assign each edge (a,b) of cut to either a or b

Value = sum of function values on assigned sets

Optimize over all assignments

min{1+1+1, 1.2+1, 1.6+1}

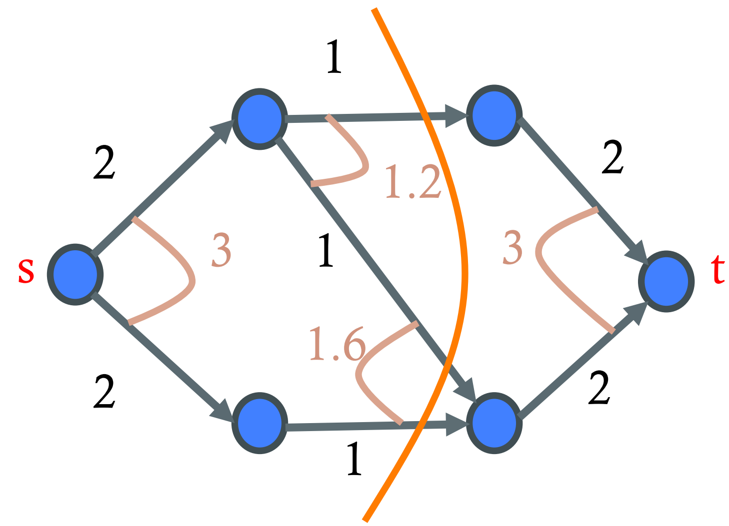

Maxflow-Mincut Theorem

[Lawler-Martel'82, Hassin'79]

Theorem: In a directed polymatroidal network the max s-t flow is equal to the min s-t cut value.

Model equivalent to submodular-flow model of[Edmonds-Giles'77] that can derive as special cases

- polymatroid intersection theorem
- maxflow-mincut in standard network flows
- Lucchesi-Younger theorem

Undirected Polymatroidal Networks

"New" model:

Undirected graph $G=(V,E)$

For each node v single polymatroids

- ρ_v with ground set $\delta(v)$

$\sum_{e \in S} f(e) \leq \rho_v(S)$ for all $S \subseteq \delta(v)$

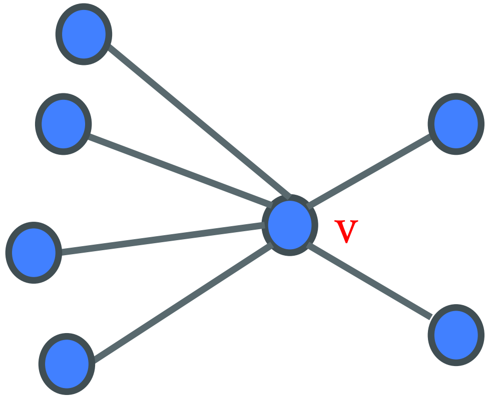

Note: maxflow-mincut does not hold, only within factor of 2!

Why Undirected Polymatroidal Networks?

- captures node-capacitated flows in undirected graphs
- within factor of 2 approximates bi-directed polymatroidal networks relevant to wireless networks which have reciprocity
- ability to use metric methods, large flow-cut gaps for multicommodity flows in directed networks

Multi-commodity Flows

Polymatroidal network $G=(V,E)$

k pairs $(s_1,t_1),\ldots,(s_k,t_k)$

Multi-commodity flow:

- f_i is s_i-t_i flow
- $f(e) = \sum_i f_i(e)$ is total flow on e
- flows on edges constrained by polymatroid constraints at nodes

Multi-commodity Cuts

Polymatroidal network $G=(V,E)$

k pairs $(s_1,t_1),\ldots,(s_k,t_k)$

Multicut: set of edges that separates all pairs

Sparsity of cut: cost of cut/demand separated by cut

Cost of cut: as defined earlier via optimization

Main Results

- $\Theta(\log k)$ flow-cut gap for undir polymatroidal networks
  - throughput flow vs multicut
  - concurrent flow vs sparsest cut
- $O(\sqrt{\log k})$-approximation in undir polymatroidal networks for *separators* (via tool from [Arora-Rao-Vazirani'04])
- Directed graphs and symmetric demands
  - $O(\log^2 k)$ flow-cut gap for throughput flow vs multicut
  - $O(\log^3 k)$ flow-cut gap for concurrent flow vs sparsest cut

Flow-cut gap results match the known bounds for standard networks

Other Results

See paper …

Remark: Two “new” proofs of maxflow-mincut theorem for s-t flow in polymatroidal networks

Implications for network information theory

[Kannan-Viswanath'11] + these results imply

capacity of a class of wireless networks understood to within $O(\log k)$ factor for k-unicast

Local vs Global Polymatroid Constraints

A more general model:

$G=(V,E)$ graph

$f: 2^E \rightarrow R$ is a polymatroid on the set of *edges*

$f(S)$ is the total capacity of the set of edges S

Function is global but problems become intractable

[Jegelka-Bilmes'10,Svitkina-Fleischer'09]

Technical Ideas

- Directed polymatroidal networks: a *reduction via uncrossing* in the dual to standard edge-capacitated directed networks
- Undirected polymatroidal networks: *dual via Lovasz-extension*
  - **sparsest cut**: round via line embeddings inspired by [Feige-Hajiaghayi-Lee'05] on undir node-capacitated graphs
  - **multicut**: line embedding idea plus region growing [Leighton-Rao'88,Garg-Vazirani-Yannakakis'93]

Rest of talk

$O(\log k)$ upper bound on gap between max concurrent flow and min sparsity in undir polymatroidal networks

Relaxation for Sparsest Cut

Want to find edge set $E' \subseteq E$ to

minimize cost(E')/dem-sep(E')

Variables:

x(e) whether e is cut or not

y(i) whether pair $s_i t_i$ is separated or not

Relaxation for Sparsest Cut

Relaxation for standard networks:

$\min \sum_e c(e)\, x(e)$

$\sum_i D_i\, y(i) = 1$

$\mathrm{dist}_x(s_i,t_i) \geq y(i)$ for all pairs i

$\mathbf{x}, \mathbf{y} \geq 0$

Dual of LP for max concurrent flow

Relaxation for Sparsest Cut

Relaxation for polymatroidal networks:

min *cost of cut*

$\sum_i D_i \, y(i) = 1$

$\mathrm{dist}_x(s_i,t_i) \geq y(i)$ for all pairs i

$\mathbf{x}, \mathbf{y} \geq 0$

Modeling cost of cut

- Each cut edge uv has to be assigned to u or v
  - Introduce variables x(e,u) and x(e,v) for each edge uv
  - Add constraint $x(e,u) + x(e,v) = x(e)$
- For a node v if $S \subseteq \delta(v)$ are cut edges assigned to v then cost at v is $\rho_v(S)$

Relaxation for Sparsest Cut

Relaxation for polymatroidal networks:

min *cost of cut*

$\sum_i D_i \, y(i) = 1$

$x(e,u) + x(e,v) = x(e)$ for each edge uv

$\mathrm{dist}_x(s_i,t_i) \geq y(i)$ for all pairs i

$\mathbf{x}, \mathbf{y} \geq 0$

Modeling cost of cut

- Each cut edge uv has to be assigned to u or v
  - Introduce variables x(e,u) and x(e,v) for each edge uv
  - Add constraint $x(e,u) + x(e,v) = x(e)$
- For a node v if $S \subseteq \delta(v)$ are cut edges assigned to v then cost at v is $\rho_v(S)$
  - $\mathbf{x}_v$ is the vector $(x(e_1,v),x(e_2,v),\dots,x(e_h,v))$ where $e_1,e_2,\dots,e_h$ are edges in $\delta(v)$
  - Use continuous extension $\rho^*_v(\mathbf{x}_v)$ to model $\rho_v(S)$

Relaxation for Sparsest Cut

Relaxation for polymatroidal networks:

$\min \sum_v \rho^*_v(\mathbf{x}_v)$

$\sum_i D_i\, y(i) = 1$

$x(e,u) + x(e,v) = x(e)$ for each edge uv

$\mathrm{dist}_x(s_i,t_i) \geq y(i)$ for all pairs i

$\mathbf{x}, \mathbf{y} \geq 0$

Lovasz-extension of f

$f^*(\mathbf{x}) = \mathbf{E}_{\theta \in [0,1]}[\ f(\mathbf{x}^\theta)\ ] = \int_0^1 f(\mathbf{x}^\theta)\, d\theta$

where $\mathbf{x}^\theta = \{\ i \mid x_i \geq \theta\ \}$

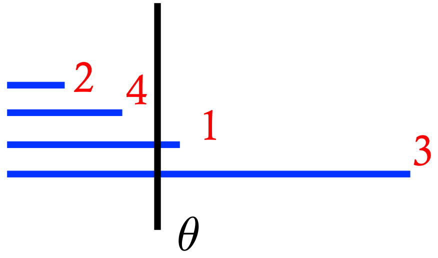

Example: $\mathbf{x} = (0.3, 0.1, 0.7, 0.2)$

$\mathbf{x}^\theta = \{1,3\}$ for $\theta = 0.21$ and $\mathbf{x}^\theta = \{3\}$ for $\theta = 0.6$

$$f^*(\mathbf{x}) = (1-0.7)\, f(\emptyset) + (0.7-0.3) f(\{3\}) + (0.3-0.2)\, f(\{1,3\}) + (0.2-0.1)\, f(\{1,3,4\}) + (0.1-0)\, f(\{1,2,3,4\})$$

Properties of f^*

- f^* is convex iff f is submodular
- Easy to evaluate f^*
- $f^*(\mathbf{x}) = f^-(\mathbf{x})$ for all $\mathbf{x}$ when f is submodular
- If f is monotone and $\mathbf{x} \leq \mathbf{y}$ then $f^*(\mathbf{x}) \leq f^*(\mathbf{y})$

Relaxation for Sparsest Cut

Relaxation for polymatroidal networks:

$\min \sum_v \rho^*_v(\mathbf{x}_v)$

$\sum_i D_i \, y(i) = 1$

$x(e,u) + x(e,v) = x(e)$ for each edge uv

$\mathrm{dist}_x(s_i,t_i) \geq y(i)$ for all pairs i

$\mathbf{x}, \mathbf{y} \geq 0$

Lemma: Dual to LP for maximum concurrent flow

Rounding of Relaxation

Standard undirected networks:

- Edge capacities: round via *l_1 embedding* [Linial-London-Rabinovich'95,Aumanna-Rabani'95]
- Node-capacities: round via *line embedding* [Feige-Hajiaghayi-Lee'05]

Line Embeddings

[Matousek-Rabinovich'01]

(V,d) metric space w(uv) non-neg weight for each uv

$g : V \to R$ is a line embedding with average weighted distortion $\alpha \geq 1$ if

- $|g(u) - g(v)| \leq d(u,v)$ for all u,v (contraction)
- $\sum_{uv} w(uv)\ |g(u)-g(v)| \geq \sum_{uv} w(uv)\ d(uv)/\alpha$

Line Embeddings

[Matousek-Rabinovich'01]

(V,d) metric space w(uv) non-neg weight for each uv

$g : V \rightarrow R$ is a line embedding with average weighted distortion α if

- $|g(u) - g(v)| \leq d(u,v)$ for all u,v (contraction)
- $\sum_{uv} w(uv)\ |g(u)-g(v)| \geq \sum_{uv} w(uv)\ d(uv)/\alpha$

Theorem [Bourgain]: Any metric space on n nodes admits line embedding with O(log n) average weighted distortion.

Rounding Algorithm

- Solve Lovasz-extension based convex relaxation
- $x(e)$ values induce metric on V
- Embed metric into line with $O(\log n)$ average distortion w.r.t to weights $w(uv) = D(uv)$
- Pick the best cut S_θ among all cuts on the line

Rounding Algorithm

- Solve Lovasz-extension based convex relaxation
- $x(e)$ values induce metric on V
- Embed metric into line with $O(\log n)$ average distortion w.r.t to weights $w(uv) = D(uv)$
- Pick the best cut S_θ among all cuts on the line
- **Remark:** Clean algorithm that generalizes edge/node/polymatroid cases since cut is defined on edges though cost is more complex

Rounding Algorithm

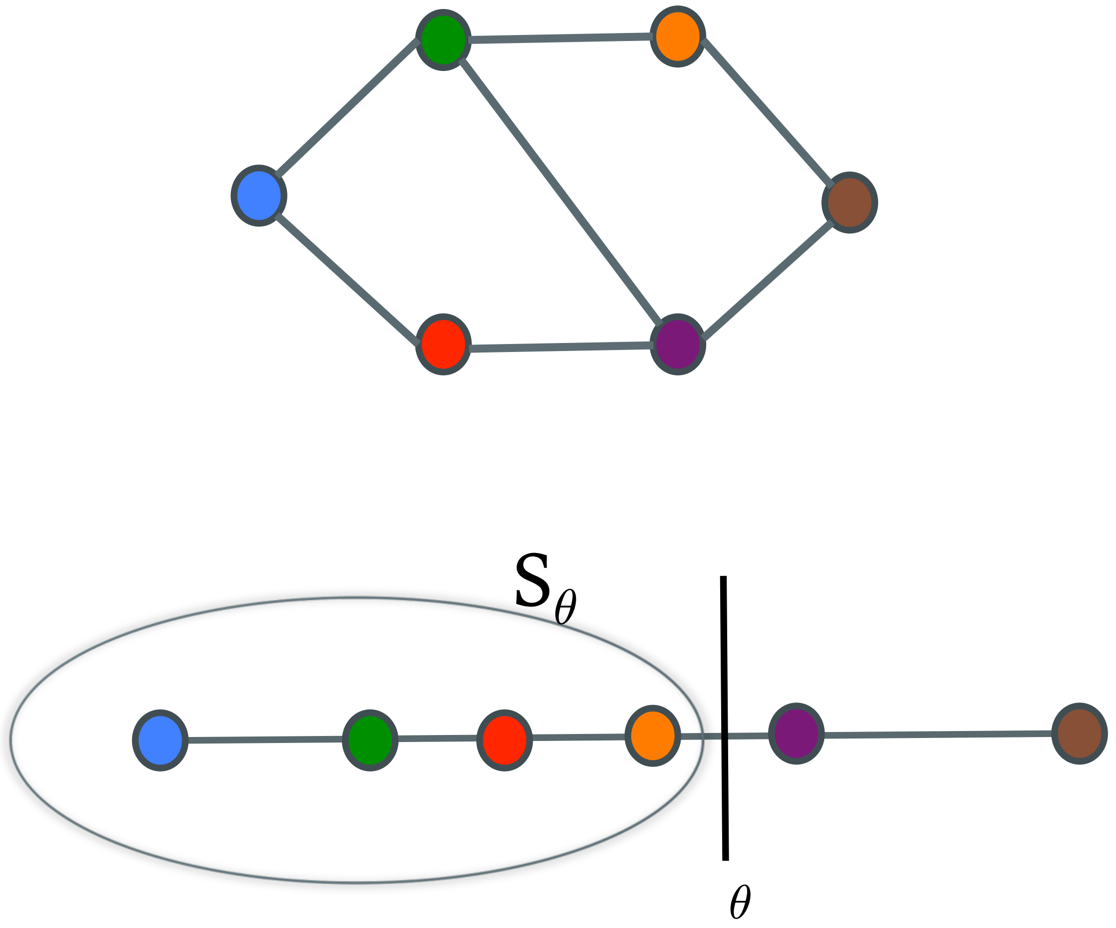

Analysis

$\nu(\delta(S_\theta))$: cost of cut at θ

Lemma: $\int \nu(\delta(S_\theta))\, d\theta \leq 2 \sum_v \rho^*_v(\mathbf{x}_v) = 2\ \mathrm{OPT}_{\mathrm{frac}}$

$D(\delta(S_\theta))$: demand separated by θ cut

Lemma: $\int D(\delta(S_\theta))\, d\theta \geq \sum_i D_i\, \mathrm{dist}_x(s_i t_i)/\log n$

Therefore:

$\int \nu(\delta(S_\theta))\, d\theta \ / \ \int D(\delta(S_\theta))\, d\theta \leq O(\log n)\ \mathrm{OPT}_{\mathrm{frac}}$

Proof of lemma

Lemma: $\int \nu(\delta(S_\theta))\, d\theta \leq 2 \sum_v \rho^*_v(\mathbf{x}_v)$

$\nu(\delta(S_\theta))$ is difficult to estimate exactly

Recall: $uv \in \delta(S_\theta)$ has to be assigned to u or v

Assign according to x(e,u) and x(e,v) *proportionally*

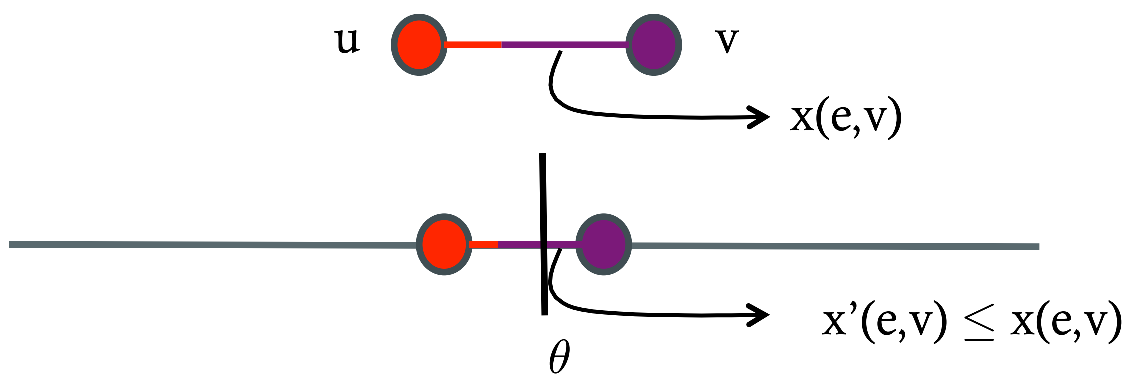

Proof of lemma

Lemma: $\int \nu(\delta(S_\theta))\, d\theta \leq 2 \sum_v \rho^*_v(\mathbf{x}_v)$

$\nu(\delta(S_\theta))$ is difficult to estimate exactly

Recall: $uv \in \delta(S_\theta)$ has to be assigned to u or v

Assign according to $x(e,u)$ and $x(e,v)$ *proportionally*

With assignment defined, estimate $\int \nu(\delta(S_\theta))\, d\theta$ by summing over nodes

Proof of lemma

Lemma: $\int \nu(\delta(S_\theta))\, d\theta \leq 2 \sum_v \rho^*_v(\mathbf{x}_v)$

With assignment defined, estimate $\int \nu(\delta(S_\theta))\, d\theta$ by summing over nodes

$\int \nu(\delta(S_\theta))\, d\theta \leq 2 \sum_v \rho^*_v(\mathbf{x'}_v) \leq 2 \sum_v \rho^*_v(\mathbf{x}_v)$

$\mathbf{x'}_v = (x'(e_1,v),\ldots,x'(e_h,v))$ where $\delta(v)=\{e_1,\ldots,e_h\}$

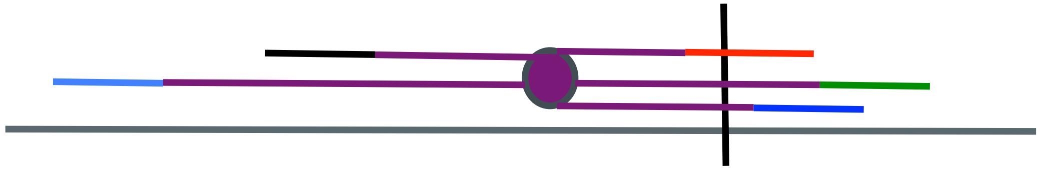

Concluding Remarks

- Flow-cut gaps for polymatroidal networks match those for standard networks

Questions:

- L_1 embeddings characterize flow-cut gap in undirected edge-capacitated networks. What characterizes flow-cut gaps of node-capacitated and polymatroidal networks?
- What are flow-cut gaps for say planar graphs? Okamura-Seymour instances?

Thanks!

Continuous extensions of f

For $f : 2^N \rightarrow \mathbb{R}^+$ define $g : [0,1]^N \rightarrow \mathbb{R}^+$ s.t

- for any $S \subseteq N$ want $f(S) = g(\mathbf{1}_S)$
- given $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in [0,1]^N$ want polynomial time algorithm to evaluate $g(\mathbf{x})$
- for *minimization* want g to be *convex* and for *maximization* want g to be *concave*

Canonical extension

$\mathbf{x} = (x_1, x_2, \ldots, x_n) \in [0,1]^N$

$$\min/\max \sum_S \alpha_S f(S)$$

$$\sum_S \alpha_S = 1$$

$$\sum_S \alpha_S = x_i \quad \text{for all } i$$

$$\alpha_S \geq 0 \quad \text{for all } S$$

$f^-(x)$ for minimization and $f^+(x)$ for maximization: convex and concave closure of f

Submodular f

- For minimization $f^{-}(x)$ can be evaluated in poly-time via submodular function minimization
  - Equivalent to the *Lovasz-extension*
- For maximization $f^{+}(x)$ is NP-Hard to evaluate even when f is monotone submodular