A new Dynamic Programming Approach for Spanning Trees with Chain Constraints and Beyond

Martin Nägele   Rico Zenklusen

ETH Zürich

Introduction: Constrained Spanning Trees

Motivation, Applications, and Our Results

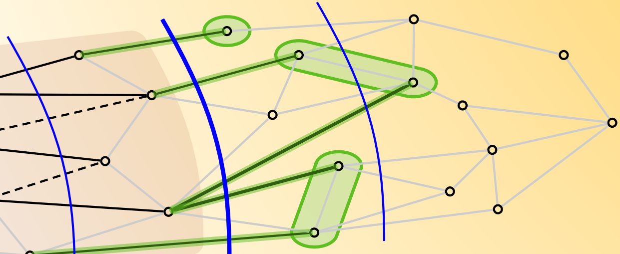

Problem Setting and Motivation

Constrained Spanning Tree Problem

Input: Graph $G = (V, E)$, edge costs $c\colon E \to \mathbb{R}$.

Goal: Find a minimum cost spanning tree $T \subseteq E$ satisfying a set of given constraints.

Constraint types:

- Degree constraints:
  $\deg_T(v) \leqslant b_v$ for $v \in V$.
- Cut constraints:
  $|T \cap \delta(S)| \leqslant b_S$ for $S \subseteq V$.
- Parity constraints:
  $|T \cap \delta(S)| \equiv 1 \pmod 2$ for $S \subseteq V$.
- ...

Motivation:

- Applications from Network Design:
  - ⇝ Bounded node capacities.
- Thin trees conjecture:
  - ⇝ Constraints on all cut sets.
- Parity-correction + uncrossing in Path TSP:
  - ⇝ Chain/laminar cut constraints.
  - ⇝ Parity constraints.

What's known?

Minimum Bounded Degree Spanning Trees (MBDST)

Find a minimum cost spanning tree such that

$$\forall v \in V: \quad \deg_T(v) \leqslant b_v \ .$$

Degree Constraints:

- Additive $+1$ violation.
  [Singh, Lau, 2007]
- Generalization: Constant violation if edges only in constantly many constraints.
  [Bansal, Kandekar, Nagarajan, 2009]

What's known?

Minimum Chain-/Laminarly-Constrained Spanning Trees (MCCST/MLCST)

Find a minimum cost spanning tree such that

$$\forall S \in \mathcal{F}: \quad |T \cap \delta(S)| \leqslant b_S \ ,$$

with $\mathcal{F} \subseteq 2^V$ a chain or laminar family, respectively.

Degree Constraints:

- Additive $+1$ violation.
  [Singh, Lau, 2007]
- Generalization: Constant violation if edges only in constantly many constraints.
  [Bansal, Kandekar, Nagarajan, 2009]

Chain or Laminar Constraints:

- Additive $\mathcal{O}(\log |V|)$ violation.
  [Bansal, Kandekar, Könemann, Nagarajan, Peis, 2013]
- Additive violation $\frac{c \cdot \log |V|}{\log \log |V|}$ impossible if P $\neq$ NP.
  [Olver, Zenklusen, 2013]
- Multiplicative guarantees: $(\frac{\lambda}{\lambda - 1}, 9\lambda)$-approximation for MCCST ($\lambda > 1$).
  [Olver, Zenklusen, 2013] [Linhares, Swamy, 2016]

What's known?

Minimum Chain-/Laminarly-Constrained Spanning Trees (MCCST/MLCST)

Find a minimum cost spanning tree such that

$$\forall S \in \mathcal{F}: \quad |T \cap \delta(S)| \leqslant b_S \ ,$$

with $\mathcal{F} \subseteq 2^V$ a chain or laminar family, respectively.

Degree Constraints:

- Additive $+1$ violation.
  [Singh, Lau, 2007]
- Generalization: Constant violation if edges only in constantly many constraints.
  [Bansal, Kandekar, Nagarajan, 2009]

Chain or Laminar Constraints:

- Additive $\mathcal{O}(\log |V|)$ violation.
  [Bansal, Kandekar, Könemann, Nagarajan, Peis, 2013]
- Additive violation $\frac{c \cdot \log |V|}{\log \log |V|}$ impossible if P $\neq$ NP.
  [Olver, Zenklusen, 2013]
- Multiplicative guarantees: $\left(\frac{\lambda}{\lambda-1}, 9\lambda\right)$-approximation for MCCST ($\lambda > 1$).
  [Olver, Zenklusen, 2013] [Linhares, Swamy, 2016]

$$c(T) \leqslant \frac{\lambda}{\lambda-1} \cdot c(\mathrm{OPT}),$$
$$|T \cap \delta(S)| \leqslant 9\lambda \cdot b_S.$$

Our Results

Chain Constraints:

Theorem 1: MCCST

Randomized $(1, 1+\varepsilon)$-approximation for MCCST with running time $|V|^{\mathcal{O}(\log |V|)/\varepsilon^2}$.

Laminar Constraints:

Theorem 2: MLCST

Randomized $(1, 1+\varepsilon)$-approximation for MLCST with running time $|V|^{\mathcal{O}(k \log |V|)/\varepsilon^2}$.

Minimum Chain-/Laminarly-Constrained Spanning Trees (MCCST/MLCST)

Find a minimum cost spanning tree such that

$$\forall S \in \mathcal{F}: \quad a_S \leqslant |T \cap \delta(S)| \leqslant b_S \ ,$$

with $\mathcal{F} \subseteq 2^V$ a chain or laminar family, respectively.

- Upper and lower bounds in constraints.
- Essentially best possible guarantees.
- Quasipolynomiality inherent to approach.

Our Results

Chain Constraints:

$$c(T) \leqslant c(\mathrm{OPT}),$$
$$\frac{a_S}{1+\varepsilon} \leqslant |T \cap \delta(S)| \leqslant (1+\varepsilon) b_S.$$

Theorem 1: MCCST

Randomized $(1, 1+\varepsilon)$-approximation for MCCST with running time $|V|^{\mathcal{O}(\log |V|)/\varepsilon^2}$.

Laminar Constraints:

Theorem 2: MLCST

Randomized $(1, 1+\varepsilon)$-approximation for MLCST with running time $|V|^{\mathcal{O}(k \log |V|)/\varepsilon^2}$.

Minimum Chain-/Laminarly-Constrained Spanning Trees (MCCST/MLCST)

Find a minimum cost spanning tree such that

$$\forall S \in \mathcal{F}: \quad a_S \leqslant |T \cap \delta(S)| \leqslant b_S \ ,$$

with $\mathcal{F} \subseteq 2^V$ a chain or laminar family, respectively.

- Upper and lower bounds in constraints.
- Essentially best possible guarantees.
- Quasipolynomiality inherent to approach.

Our Results

Chain Constraints:

$$c(T) \leqslant c(\text{OPT}),$$
$$\frac{a_S}{1+\varepsilon} \leqslant |T \cap \delta(S)| \leqslant (1+\varepsilon) b_S.$$

Theorem 1: MCCST

Randomized $(1, 1+\varepsilon)$-approximation for MCCST with running time $|V|^{\mathcal{O}(\log |V|)/\varepsilon^2}$.

Laminar Constraints:

Theorem 2: MLCST

Randomized $(1, 1+\varepsilon)$-approximation for MLCST with running time $|V|^{\mathcal{O}(k \log |V|)/\varepsilon^2}$.

$$k := \text{width}(\mathcal{F})$$

Minimum Chain-/Laminarly-Constrained Spanning Trees (MCCST/MLCST)

Find a minimum cost spanning tree such that

$$\forall S \in \mathcal{F}: \quad a_S \leqslant |T \cap \delta(S)| \leqslant b_S \ ,$$

with $\mathcal{F} \subseteq 2^V$ a chain or laminar family, respectively.

- Upper and lower bounds in constraints.
- Essentially best possible guarantees.
- Quasipolynomiality inherent to approach.

Our Results

$$c(T) \leqslant c(\mathrm{OPT}),$$
$$\frac{a_S}{1+\varepsilon} \leqslant |T \cap \delta(S)| \leqslant (1+\varepsilon) b_S.$$

Chain Constraints:

Theorem 1: MCCST

Randomized $(1, 1+\varepsilon)$-approximation for MCCST with running time $|V|^{\mathcal{O}(\log |V|)/\varepsilon^2}$.

Laminar Constraints:

Theorem 2: MLCST

Randomized $(1, 1+\varepsilon)$-approximation for MLCST with running time $|V|^{\mathcal{O}(k \log |V|)/\varepsilon^2}$.

$k := \mathrm{width}(\mathcal{F})$

Minimum Chain-/Laminarly-Constrained Spanning Trees (MCCST/MLCST)

Find a minimum cost spanning tree such that
$$\forall S \in \mathcal{F}: \quad a_S \leqslant |T \cap \delta(S)| \leqslant b_S ,$$
with $\mathcal{F} \subseteq 2^V$ a chain or laminar family, respectively.

- Upper and lower bounds in constraints.
- Essentially best possible guarantees.
- Quasipolynomiality inherent to approach.

Further application of new techniques:

- $(1.5+\varepsilon)$-approximations for Path TSP and shortest connected T-join problem.
  Running times: $|V|^{\mathcal{O}(1)/\varepsilon}$ and $|V|^{\mathcal{O}(|T|)/\varepsilon}$

Minimum Chain-Constrained Spanning Trees

An Overview of Our Techniques

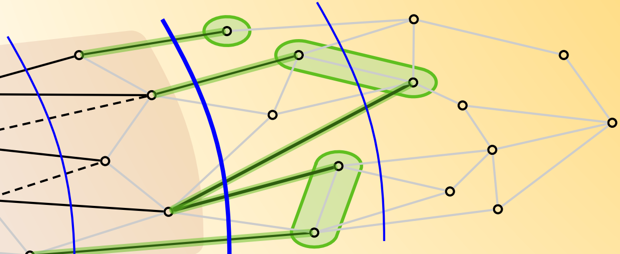

Techniques for MCCST

Minimum Chain-Constrained Spanning Trees (MCCST)

Find a minimum cost spanning tree such that

$$\forall i \in [k]: \quad a_i \leqslant |T \cap \delta(S_i)| \leqslant b_i \ ,$$

where $\emptyset \subsetneq S_1 \subsetneq S_2 \subsetneq \ldots \subsetneq S_k \subsetneq V$.

Techniques for MCCST

Minimum Chain-Constrained Spanning Trees (MCCST)

Find a minimum cost spanning tree such that

$$\forall i \in [k]: \quad a_i \leqslant |T \cap \delta(S_i)| \leqslant b_i \ ,$$

where $\emptyset \subsetneq S_1 \subsetneq S_2 \subsetneq \ldots \subsetneq S_k \subsetneq V$.

Three main steps:

1. Find good solution x of linear relaxation.
   ⇝ New DP approach, inspired by recent Path TSP approaches.
   [Traub, Vygen, 2018] [Zenklusen, 2019]

Techniques for MCCST

Minimum Chain-Constrained Spanning Trees (MCCST)

Find a minimum cost spanning tree such that

$$\forall i \in [k]: \quad a_i \leqslant |T \cap \delta(S_i)| \leqslant b_i \ ,$$

where $\emptyset \subsetneq S_1 \subsetneq S_2 \subsetneq \ldots \subsetneq S_k \subsetneq V$.

Three main steps:

1. Find good solution x of linear relaxation.
   - ⇝ New DP approach, inspired by recent Path TSP approaches.
     [Traub, Vygen, 2018] [Zenklusen, 2019]
2. Apply randomized rounding to obtain tree T from x.
   - ⇝ Marginal-preserving, negatively correlated rounding.
     [Asadpour, Goemans, Madry, Oveis Gharan, Saberi, 2010] [Chekuri, Vondrák, Zenklusen, 2010]
   - ⇝ Chernoff-type concentration bounds imply constraints up to $(1 \pm \varepsilon)$ with high probability.

Techniques for MCCST

Minimum Chain-Constrained Spanning Trees (MCCST)

Find a minimum cost spanning tree such that

$$\forall i \in [k]: \quad a_i \leqslant |T \cap \delta(S_i)| \leqslant b_i \ ,$$

where $\emptyset \subsetneq S_1 \subsetneq S_2 \subsetneq \ldots \subsetneq S_k \subsetneq V$.

Three main steps:

1. Find good solution x of linear relaxation.
   - ⇝ New DP approach, inspired by recent Path TSP approaches.
     [Traub, Vygen, 2018] [Zenklusen, 2019]
2. Apply randomized rounding to obtain tree T from x.
   - ⇝ Marginal-preserving, negatively correlated rounding.
     [Asadpour, Goemans, Madry, Oveis Gharan, Saberi, 2010] [Chekuri, Vondrák, Zenklusen, 2010]
   - ⇝ Chernoff-type concentration bounds imply constraints up to $(1 \pm \varepsilon)$ with high probability.
3. Perform local corrections to gain back potential loss in objective.
   - ⇝ In MCCST: One single edge swap.
   - ⇝ General procedure, applicable for similar rounding procedures in $\{0, 1\}$-polytopes.

What properties should x have?

Natural Relaxation:

$$Q = \underbrace{\left\{ x \in \mathbb{R}^E_{\geqslant 0} \;\middle|\; \begin{array}{c} x(E) = |V| - 1 \\ x(E[S]) \leqslant |S| - 1 \quad \forall\, S \subsetneq V,\ |S| \geqslant 2 \end{array} \right\}}_{\text{spanning tree polytope } P_{\mathrm{ST}}} \cap \underbrace{\left\{ x \in \mathbb{R}^E \;\middle|\; a_i \leqslant x(\delta(S_i)) \leqslant b_i \quad \forall\, i \in [k] \right\}}_{\text{chain constraints}}$$

What properties should x have?

Natural Relaxation:

$$Q = \underbrace{\left\{ x \in \mathbb{R}^E_{\geqslant 0} \;\middle|\; \begin{array}{c} x(E) = |V| - 1 \\ x(E[S]) \leqslant |S| - 1 \quad \forall\, S \subsetneq V,\ |S| \geqslant 2 \end{array} \right\}}_{\text{spanning tree polytope } P_{\mathrm{ST}}} \cap \underbrace{\left\{ x \in \mathbb{R}^E \;\middle|\; a_i \leqslant x(\delta(S_i)) \leqslant b_i \quad \forall\, i \in [k] \right\}}_{\text{chain constraints}}$$

- Lower bound for approximation with respect to Q: Factor 2.
  - ⇝ Hard limit for prior approaches.

What properties should x have?

Natural Relaxation:

$$Q = \underbrace{\left\{ x \in \mathbb{R}^E_{\geqslant 0} \;\middle|\; \begin{array}{c} x(E) = |V| - 1 \\ x(E[S]) \leqslant |S| - 1 \quad \forall\, S \subsetneq V,\ |S| \geqslant 2 \end{array} \right\}}_{\text{spanning tree polytope } P_{\mathrm{ST}}} \cap \underbrace{\left\{ x \in \mathbb{R}^E \;\middle|\; a_i \leqslant x(\delta(S_i)) \leqslant b_i \quad \forall\, i \in [k] \right\}}_{\text{chain constraints}}$$

- Lower bound for approximation with respect to Q: Factor 2.
  - ⇝ Hard limit for prior approaches.
- Thought experiment: **What if $x(\delta(S_i)) \geqslant c \cdot \log k$ for all $i \in [k]$?**
  - ⇝ Chernoff Bounds:

$$\Pr\Big[\big|T \cap \delta(S_i)\big| \notin \big[(1-\varepsilon)x(\delta(S_i)), (1+\varepsilon)x(\delta(S_i))\big]\Big] \leqslant 2e^{-x(\delta(S_i))\cdot \varepsilon^2/3} = k^{-\Omega(1)} \ .$$

  - ⇝ Union bound is enough to conclude approximate chain bounds with high probability.

What properties should x have?

Conclusion: Cuts S_i with large value $x(\delta(S_i))$ are unproblematic!

⇝ What about small cuts?

What properties should x have?

Conclusion: Cuts S_i with large value $x(\delta(S_i))$ are unproblematic!

⇝ What about small cuts?

Definition: τ-integral point $x \in \mathbb{R}^E$

x is τ-integral wrt. $S_1, \ldots, S_k$ if for $i \in [k]$,

(i) $x(\delta(S_i)) \geqslant \tau + 1$, or

(ii) $x(\delta(S_i)) \leqslant \tau$ and x integral on $\delta(S_i)$.

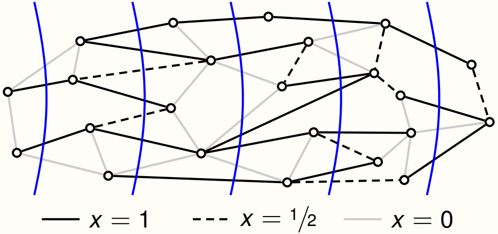

What properties should x have?

Conclusion: Cuts S_i with large value $x(\delta(S_i))$ are unproblematic!

⇝ What about small cuts?

Definition: τ-integral point $x \in \mathbb{R}^E$

x is τ-integral wrt. $S_1, \ldots, S_k$ if for $i \in [k]$,

(i) $x(\delta(S_i)) \geqslant \tau + 1$, or

(ii) $x(\delta(S_i)) \leqslant \tau$ and x integral on $\delta(S_i)$.

- Marginal-preserving rounding: $|T \cap \delta(S_i)| = x(\delta(S_i))$ for small cuts S_i.

⇝ Small cuts satisfy chain constraints exactly.

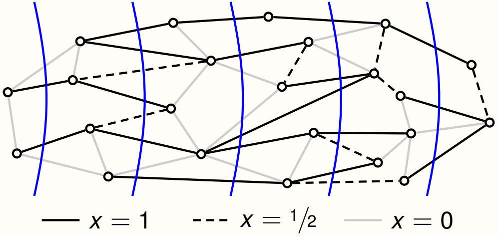

What properties should x have?

Conclusion: Cuts S_i with large value $x(\delta(S_i))$ are unproblematic!

⇝ What about small cuts?

Definition: τ-integral point $x \in \mathbb{R}^E$

x is τ-integral wrt. $S_1, \ldots, S_k$ if for $i \in [k]$,

(i) $x(\delta(S_i)) \geqslant \tau + 1$, or

(ii) $x(\delta(S_i)) \leqslant \tau$ and x integral on $\delta(S_i)$.

Theorem: Finding τ-integral points

A τ-integral point $x \in Q$ satisfying

$$c^\top x \leqslant c(\mathrm{OPT})$$

can be found in time $|V|^{\mathcal{O}(\tau)}$.

- Marginal-preserving rounding: $|T \cap \delta(S_i)| = x(\delta(S_i))$ for small cuts S_i.

⇝ Small cuts satisfy chain constraints exactly.

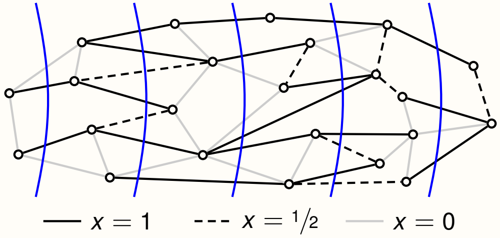

What properties should x have?

Conclusion: Cuts S_i with large value $x(\delta(S_i))$ are unproblematic!

⇝ What about small cuts?

Definition: τ-integral point $x \in \mathbb{R}^E$

x is τ-integral wrt. $S_1, \dots, S_k$ if for $i \in [k]$,

(i) $x(\delta(S_i)) \geqslant \tau + 1$, or

(ii) $x(\delta(S_i)) \leqslant \tau$ and x integral on $\delta(S_i)$.

Theorem: Finding τ-integral points

A τ-integral point $x \in Q$ satisfying

$$c^\top x \leqslant c(\mathrm{OPT})$$

can be found in time $|V|^{\mathcal{O}(\tau)}$.

- Marginal-preserving rounding: $|T \cap \delta(S_i)| = x(\delta(S_i))$ for small cuts S_i.

⇝ Small cuts satisfy chain constraints exactly.

Corollary

$(1+\varepsilon, 1+\varepsilon)$-approximation for MCCST with running time $|V|^{\mathcal{O}(\log k)/\varepsilon^2}$.

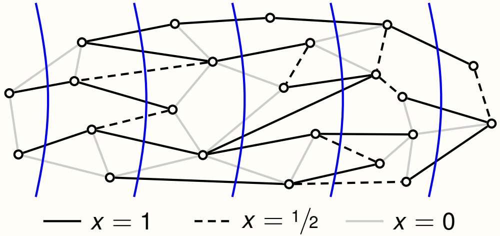

The Dynamic Program

Finding Cheap τ-Integral Points

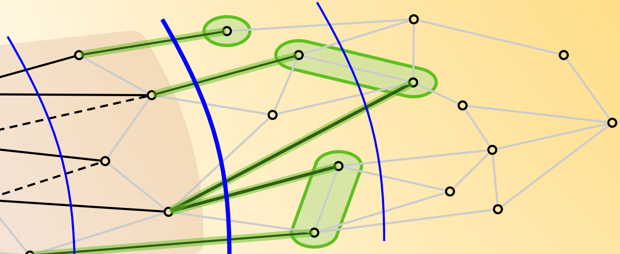

Finding τ-integral points using a DP

Definition: τ-integral point $x \in P_{ST}$

$x \in P_{ST}$ is τ-integral wrt. $S_1, \ldots, S_k$ if for $i \in [k]$,

(i) $x(\delta(S_i)) \geqslant \tau + 1$, or

(ii) $x(\delta(S_i)) \leqslant \tau$ and x integral on $\delta(S_i)$.

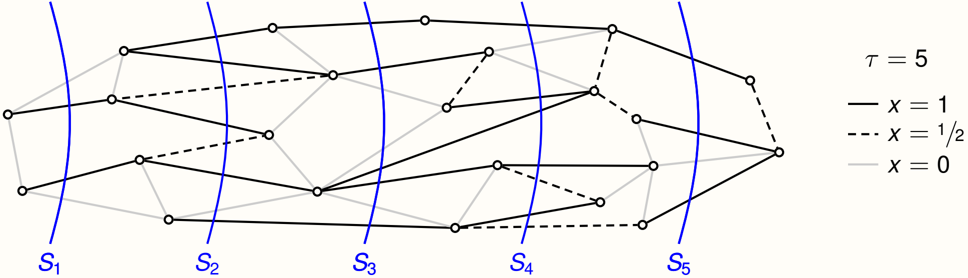

Finding τ-integral points using a DP

Definition: τ-integral point $x \in P_{ST}$

$x \in P_{ST}$ is τ-integral wrt. $S_1, \ldots, S_k$ if for $i \in [k]$,

(i) $x(\delta(S_i)) \geqslant \tau + 1$, or

(ii) $x(\delta(S_i)) \leqslant \tau$ and x integral on $\delta(S_i)$.

large cut → (i)

small cut → (ii)

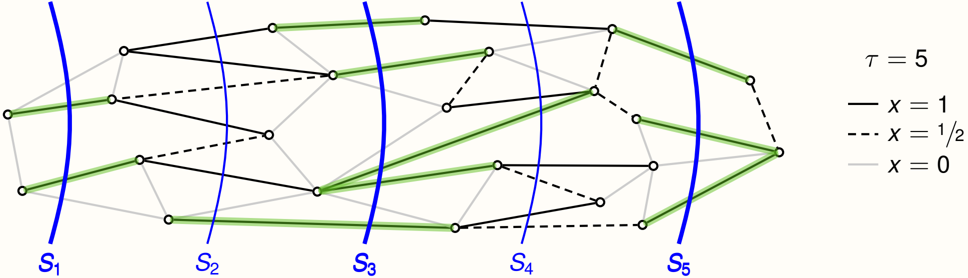

Finding τ-integral points using a DP

Definition: τ-integral point $x \in P_{ST}$

$x \in P_{ST}$ is τ-integral wrt. $S_1, \ldots, S_k$ if for $i \in [k]$,

large cut → (i) $x(\delta(S_i)) \geqslant \tau + 1$, or

small cut → (ii) $x(\delta(S_i)) \leqslant \tau$ and x integral on $\delta(S_i)$.

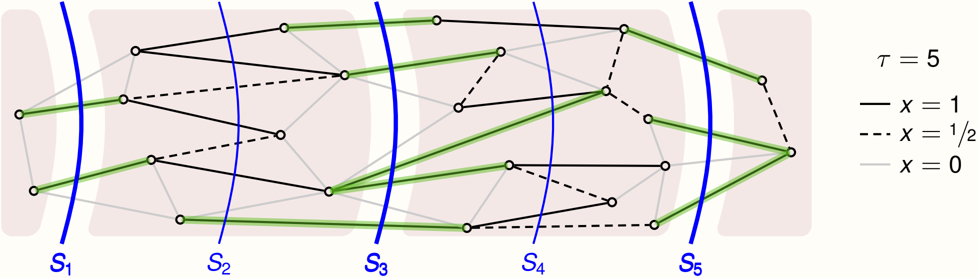

DP idea: Extend solution from one small cut to another.

A special case: $\tau = 1$

- Small cuts separate instance into independent subproblems.
- LP for optimizing subproblems.
  - ⇝ Enforcing large cuts: Linear constraints.
- Standard DP finds cheapest 1-integral point.

$$\begin{aligned} \min \quad & c^\top x \\ & x \in P_{\mathrm{ST}}(S_j \setminus S_i) \\ x(\delta(S_\ell)) \geqslant 2 \quad & \forall i < \ell < j \end{aligned}$$

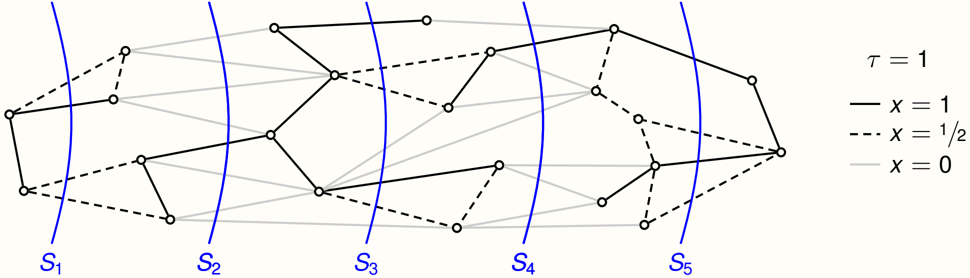

A special case: $\tau = 1$

- Small cuts separate instance into independent subproblems.
- LP for optimizing subproblems.
  - ⇝ Enforcing large cuts: Linear constraints.
- Standard DP finds cheapest 1-integral point.

$$\begin{aligned} \min \quad & c^\top x \\ & x \in P_{\mathrm{ST}}(S_j \setminus S_i) \\ x(\delta(S_\ell)) & \geqslant 2 \quad \forall i < \ell < j \end{aligned}$$

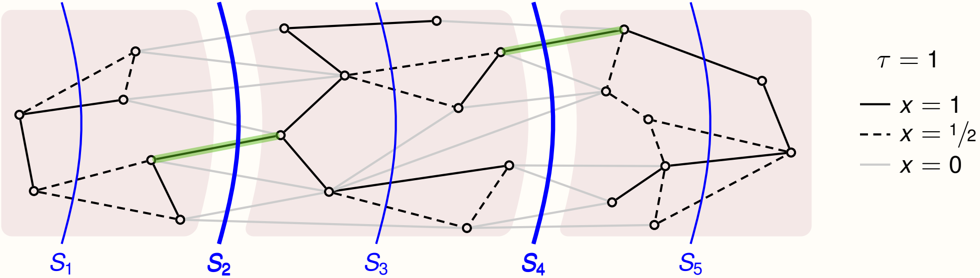

The general case

Problem: Small cuts no longer separate into independent subproblems.

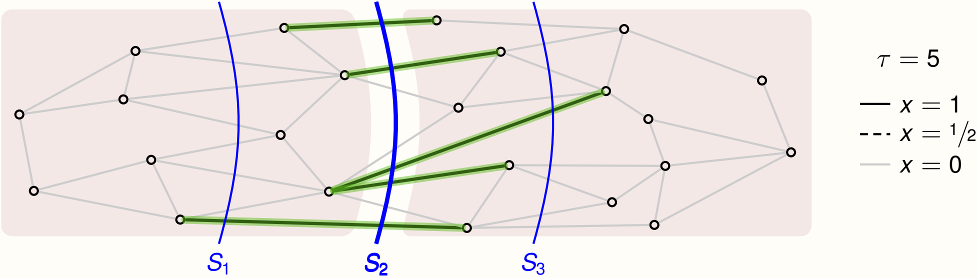

The general case

Problem: Small cuts no longer separate into independent subproblems.

Finding left-compatible solution

- Guess connectivity pattern on the right.
- Solve LP for guessed pattern.

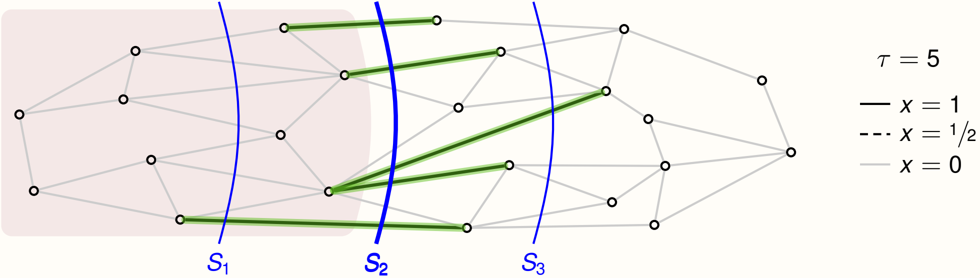

The general case

Problem: Small cuts no longer separate into independent subproblems.

Finding left-compatible solution

- Guess connectivity pattern on the right.
- Solve LP for guessed pattern.

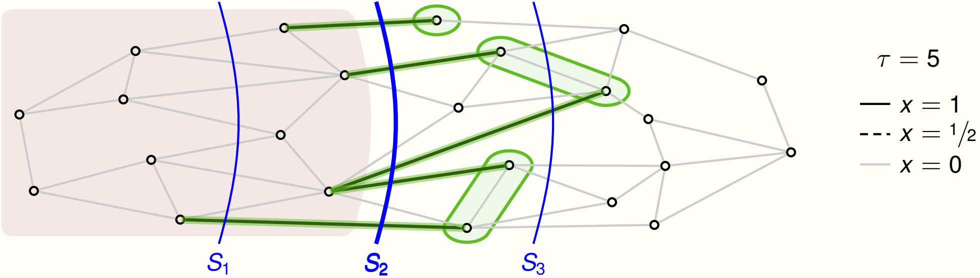

The general case

Problem: Small cuts no longer separate into independent subproblems.

Finding left-compatible solution

- Guess connectivity pattern on the right.
- Solve LP for guessed pattern.

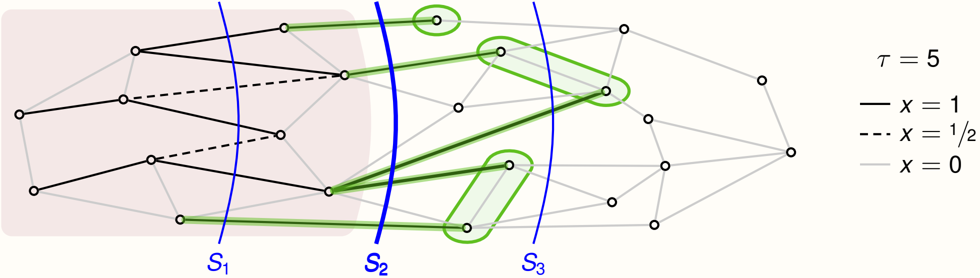

The general case

Problem: Small cuts no longer separate into independent subproblems.

Finding left-compatible solution

- Guess connectivity pattern on the right.
- Solve LP for guessed pattern.

Extending to right-compatible solution

- Completing connectivity patterns is hard.
- Solution: relax connectivity requirement, extend explicit subsolution.

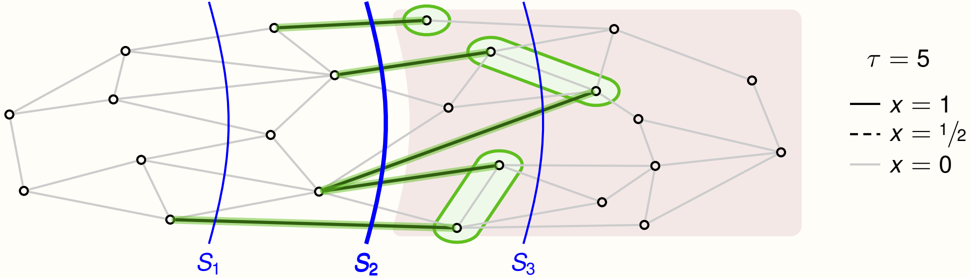

The general case

Problem: Small cuts no longer separate into independent subproblems.

Finding left-compatible solution

- Guess connectivity pattern on the right.
- Solve LP for guessed pattern.

Extending to right-compatible solution

- Completing connectivity patterns is hard.
- Solution: relax connectivity requirement, extend explicit subsolution.

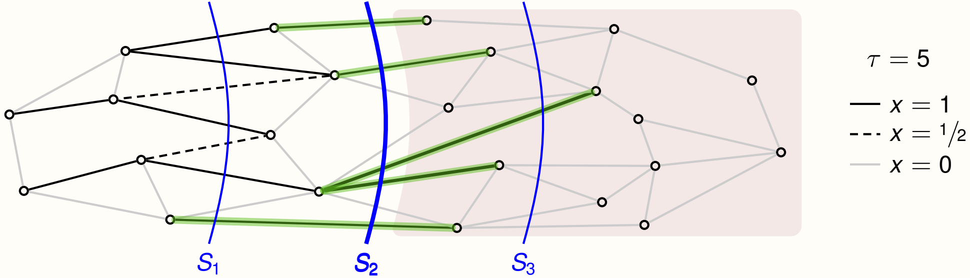

The general case

Problem: Small cuts no longer separate into independent subproblems.

Finding left-compatible solution

- Guess connectivity pattern on the right.
- Solve LP for guessed pattern.

Extending to right-compatible solution

- Completing connectivity patterns is hard.
- Solution: relax connectivity requirement, extend explicit subsolution.

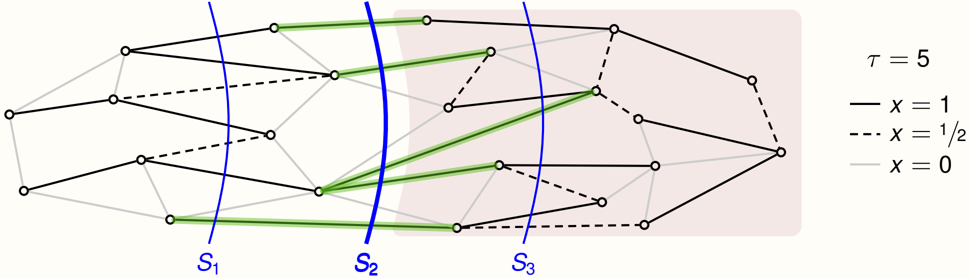

The general case

Problem: Small cuts no longer separate into independent subproblems.

Finding left-compatible solution

- Guess connectivity pattern on the right.
- Solve LP for guessed pattern.

Extending to right-compatible solution

- Completing connectivity patterns is hard.
- Solution: relax connectivity requirement, extend explicit subsolution.

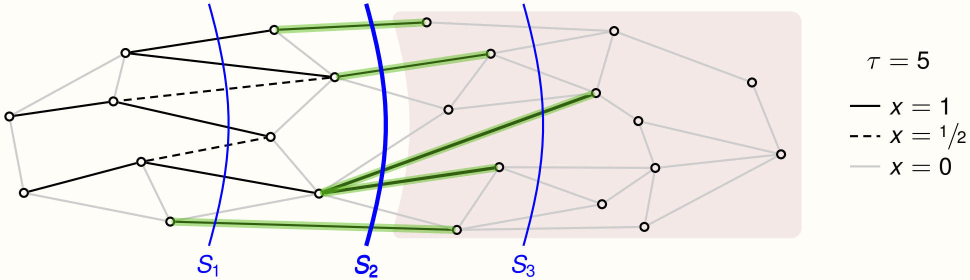

Consequence: Cannot find cheapest τ-integral point, but remain better than OPT.

Conclusions

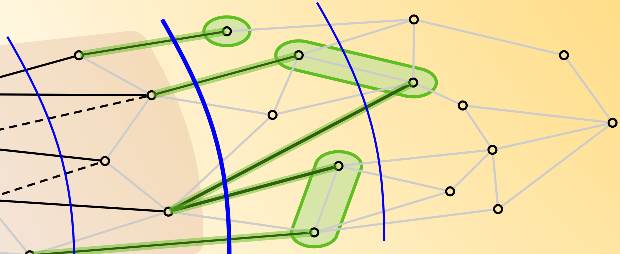

Conclusions

Randomized $(1, 1+\varepsilon)$-approximation algorithm for MCCST

1. Use the DP to find a τ-integral point x for $\tau = \left\lfloor {}^{96\log(2|V|)}/_{\varepsilon^2} \right\rfloor$.
2. Obtain spanning tree T from x by marginal-preserving, negatively correlated rounding.
3. Return cheapest tree obtained from T by one edge swap.

Conclusions

Randomized $(1, 1+\varepsilon)$-approximation algorithm for MCCST

1. Use the DP to find a τ-integral point x for $\tau = \left\lfloor {}^{96\log(2|V|)}/_{\varepsilon^2} \right\rfloor$.
2. Obtain spanning tree T from x by marginal-preserving, negatively correlated rounding.
3. Return cheapest tree obtained from T by one edge swap.

Open questions:

- Polynomial-time algorithm for MCCST?
- Reducing exponential dependence on $\mathtt{width}(\mathcal{L})$ in running time for MLCST?
- Connected T-join problem: Efficient algorithms for arbitrary T?