Hypergraphic LP Relaxations for Steiner Trees

IPCO 2010

Deeparnab Chakrabarty,
Jochen Könemann,
David Pritchard

The Steiner Tree Problem

- Input: graph (R⊎N, E) with edge costs c_e,
  - *R: required vertices* or *terminals*
  - *N: optional vertices* or *Steiner nodes*
- Output: subgraph (R⊎N, F) connecting R
- Objective: min $\Sigma_{e \in F} c_e$
- NP-hard, APX-hard.
- Best approx alg has ratio 1.39 [BGRS10]

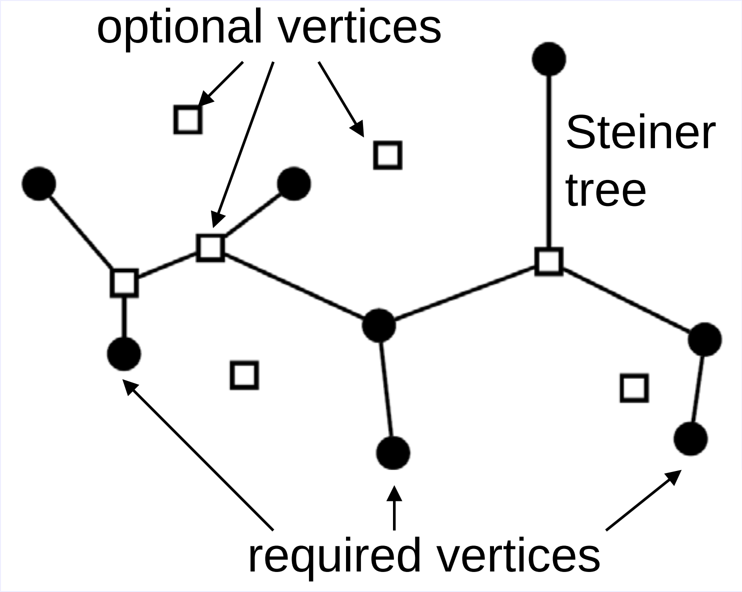

Summary

- Study LPs for the Steiner tree problem using *full components* as building blocks
- Give a new LP based on partitions
- It has same value as some previous LPs
  - Directed cut and subtour formulations
- Show all are equivalent to *bidirected cut formulation* on quasibipartite instances
- Integrality gap bound of $\sqrt{3} \sim 1.73$
  - Also 73/60 ~ 1.22 on "uniformly quasi-bipartite"

Steiner Trees as Hypergraphs

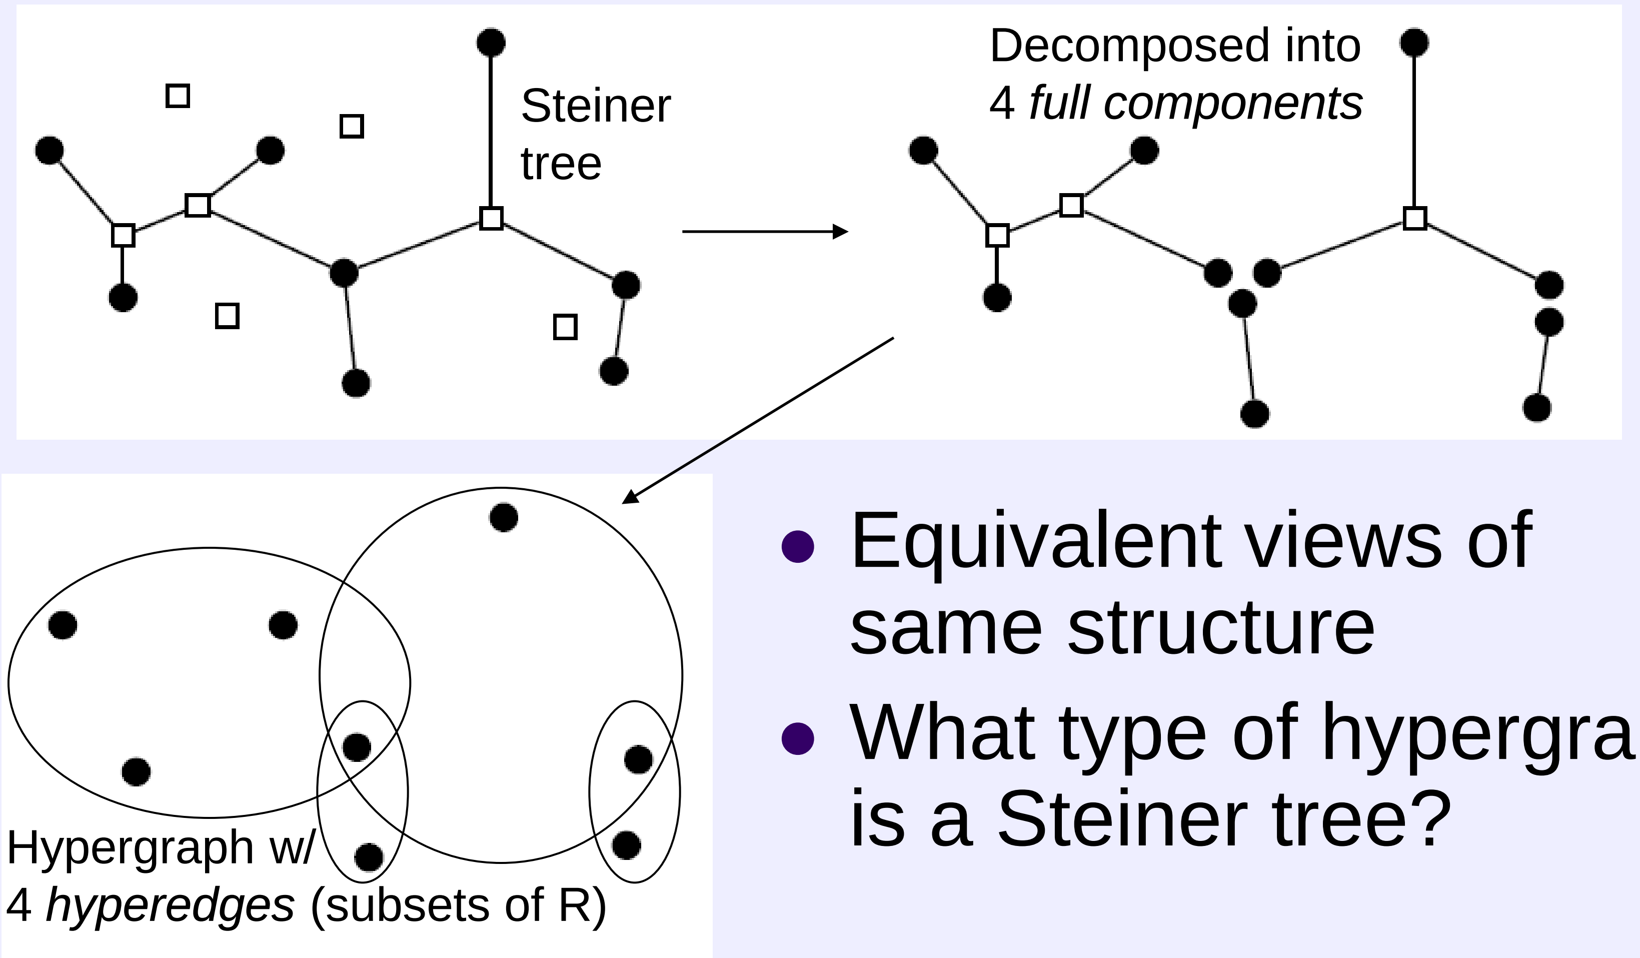

- Equivalent views of same structure
- What type of hypergraph is a Steiner tree?

Full Component ~ Hyperedge, Steiner Tree ~ Hyper-Spanning Tree

- Steiner trees are hyper-graphs that are
  - acyclic, connected, span R
  - called *hyper-spanning trees*

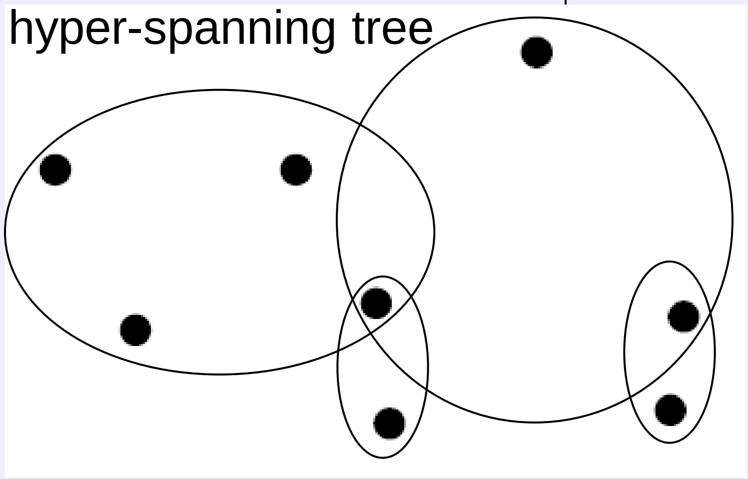

- Overall approach:
  - $\forall$ $K \subset R$ let C_K be cheapest full component on K
  - Then **Steiner tree problem** becomes special case of **Min Cost Hyper-Spanning Tree**
  - Motivates hypergraph-based LPs

LPs Old and New

Hypergraphic Steiner tree relaxation	MST special case (size-2 full comps)
Subtour Relaxation [Warme 97] • equality constraint • every subset of R is fractionally acyclic	Matroid / Subtour LP [Edmonds 1970]
Directed Cut Relaxation [P-VD 03, B+10] • direct full components • every nontrivial cut fractionally spanned	Bidirected cut LP [Edmonds 1967]
Partition Relaxation [here, KPT09] • every partition π is fractionally spanned rank(π)-1 times, counting multiplicity	Partition LP [Fulkerson 1971]

Partitions

- *Partition of R*: family $\{V_1, V_2, \ldots, V_r\}$ of disjoint nonempty subsets of R (*parts*) with $\cup_i V_i = R$
- r: *rank* of π (number of parts)

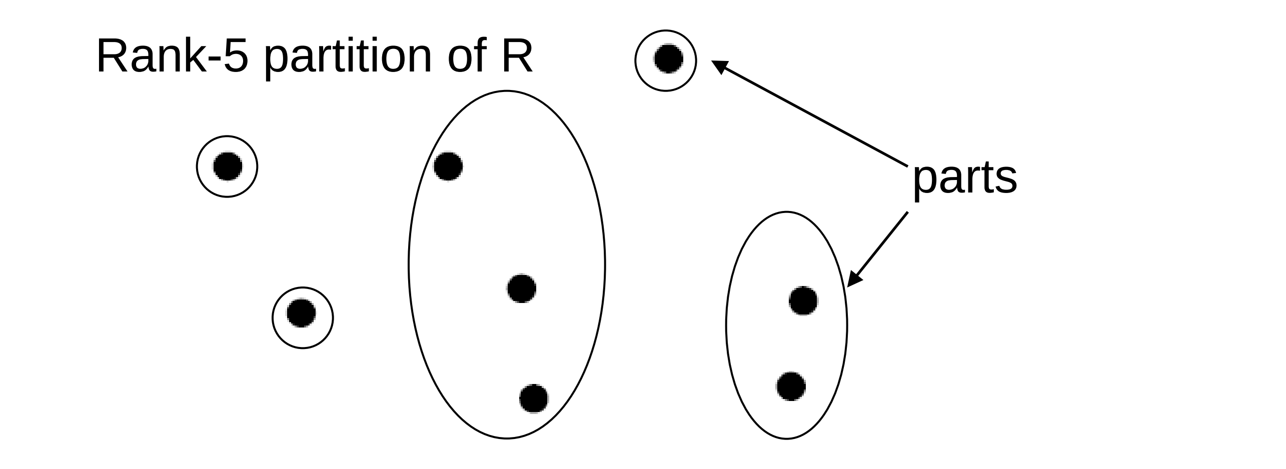

- How do partitions, hyper-spanning trees interact?

Rank Contribution

- For hyperedge K⊂V the *rank-contribution*
  $$\mathrm{rc}_K^\pi := |\{\text{parts of } \pi \text{ met by } K\}| - 1$$
  - (Equal to rank lost by π if we merge all its parts intersecting K)
- For any hyper-spanning tree T,
  $$\sum_{K \in T} \mathrm{rc}_K^\pi \geq r(\pi) - 1$$

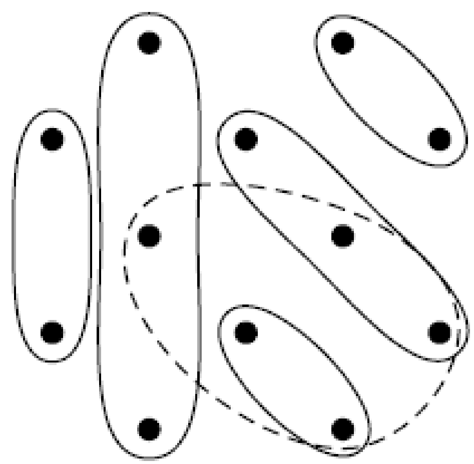

Solid partition π and dashed set K have $\mathrm{rc}_K^\pi = 2$

Hyperedge/Partition LP ($\mathcal{P}$)

Variable x_K in [0,1] for each hyperedge $K \subset R$

Inequality for each partition π of V:

$$\sum_K x_K \mathrm{rc}_K^\pi \geq r(\pi) - 1$$

Objective: minimize $\Sigma C_K x_K$

- Main results:
  - integrality gaps by *dual fitting* and *MST-exactness*
  - equivalence theorems by *partition uncrossing*

Partition Uncrossing

- **Prop**. If constraints for partitions π, σ hold with equality, same holds for their *meet* & *join*

Meet: intersect parts of π with parts of σ in all possible ways

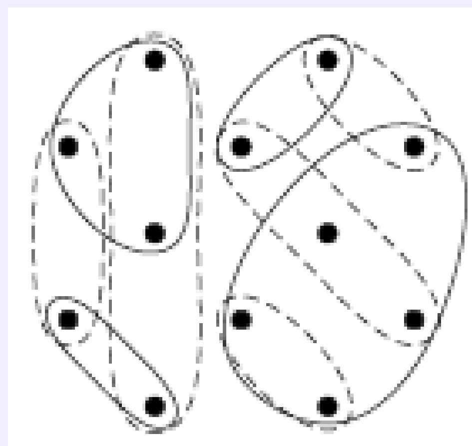

Join: transitively join parts of π with parts of σ

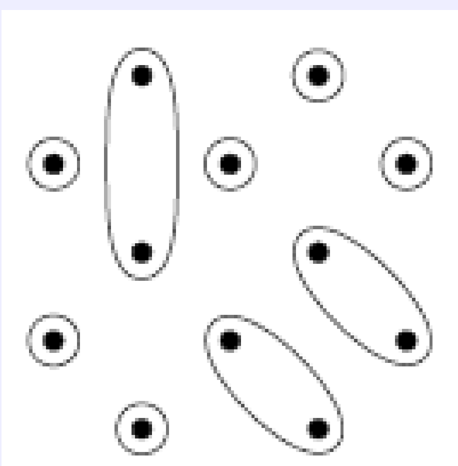

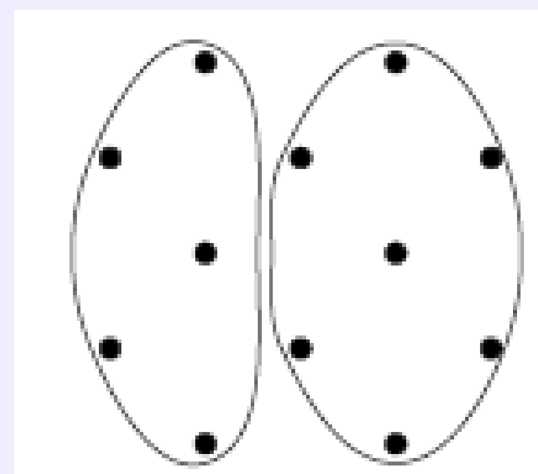

Partition Uncrossing

- Proving **Prop** looks easy in that $(\mathcal{P})$ resembles a *lattice polyhedron…* but typical uncrossing approach fails on small examples
- We use a new partition uncrossing technique; it shows extreme duals are supported by a chain (non-crossing set)
- Implies extreme primal solutions have at most |R|-1 nonzeroes
  - Spanning trees show |R|-1 is tight

2 partitions that do not cross

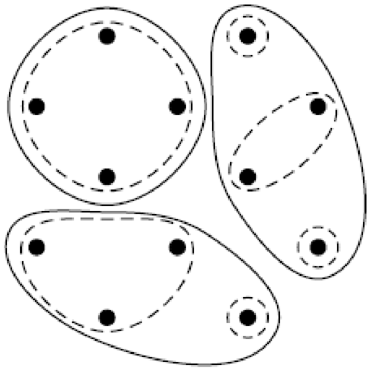

LP Equivalences

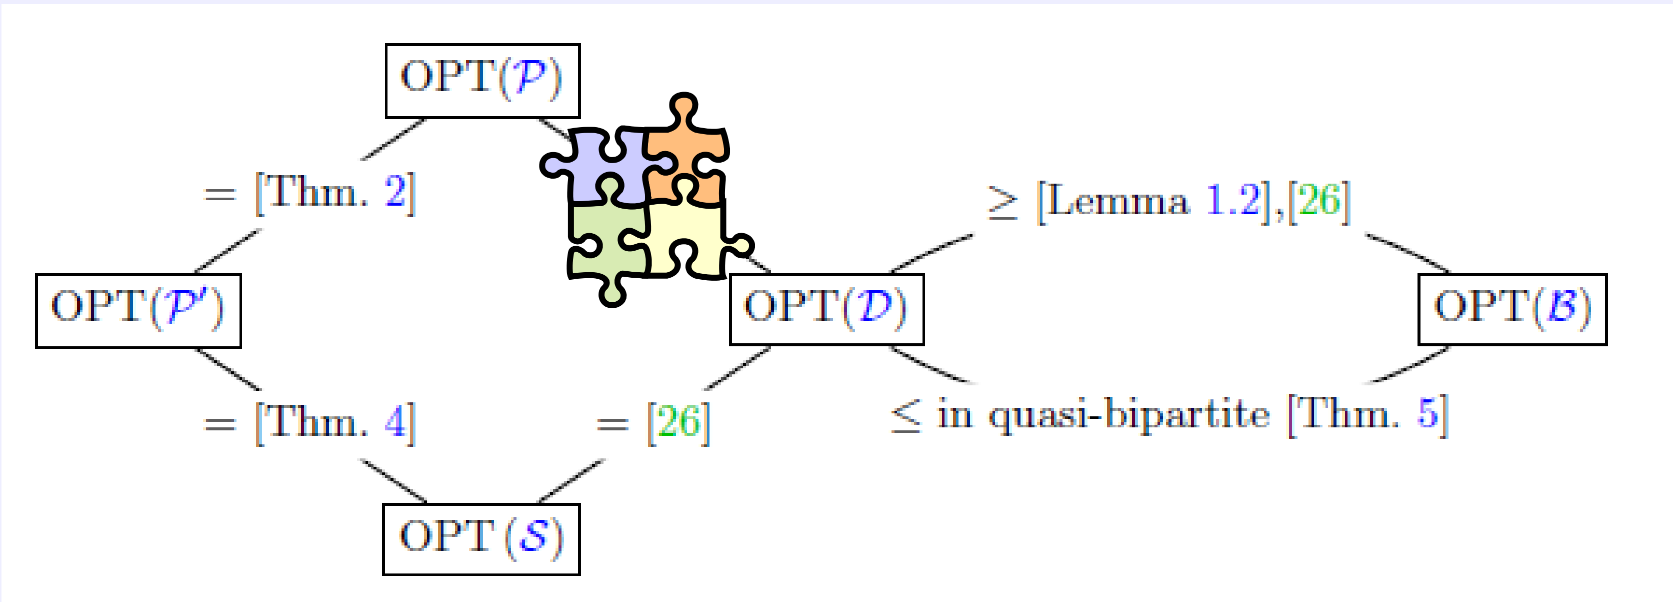

- 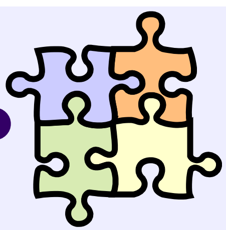 immediate from hypergraph orientation results by Frank, Kiraly, Kiraly 2003

Quasibipartite Equivalence

- One LP equivalence still surprises me a lot
- In a *quasibipartite instance,* there are no edges connecting two Steiner nodes
- The *bidirected cut relaxation* for Steiner tree:
    - introduce a variable x_a for each arc *a* (two per undirected edge in the input)
    - pick any terminal as root (doesn’t matter which)
    - require all cuts from root to any other Steiner node to be crossed by x-value of ≥ 1

Quasibipartite Equivalence

- Previously studied hypergraphic LPs known to strengthen bidirected cut [P-VD 03]
- **Thm**. In quasi-bipartite instances, both LPs have the same value
- We found two proofs:
  - implicit, using theorem about total unimodularity
  - explicit algorithmic proof
- Both “lift” duals. Is there a more direct proof?

Integrality Gap ($\mathcal{P}$) ≤ √3 ~ 1.73

- In other words, there is always a Steiner tree T with cost(T) ≤ √3 · OPT($\mathcal{P}$)
- Integrality gap of 2 is trivial but until recently no better bound was known for *any* LP
  - [BGRS10] got a 1.55 bound first via RZ algorithm
  - We noticed different techniques give an “online” √3 bound. Give 2√2 – 1 ~ 1.82 in talk.
- Techniques: *cost reduction* [CDV08] and *MST-exactness*

MST-Exactness

- Suppose that for some Steiner tree instance, the minimum spanning tree MST(G[R]) of the terminal-induced subgraph is an optimal Steiner tree. (An *MST-exact* instance.)
- **Thm**. This tree is optimal for the LP $(\mathcal{P})$; i.e. OPT$(\mathcal{P})$ is the optimal Steiner cost.
- Algorithmic leverage: reduce some costs to get an MST-exact instance; gives lower bound on LP value of original instance.

$2\sqrt{2} - 1$ integrality gap algorithm

1. Divide all terminal-terminal costs by $\sqrt{2}$
2. Calculate initial MST(G[R])
3. For each full component K, in any order
   - Contract terminal subset K to single pseudonode and pay C_K, **if** MST cost would drop by > C_K

- **Analysis idea**: contracted instance at end of algorithm is MST-exact
- Also use fact MST ≤ 2 · OPT($\mathcal{P}$) in final contracted instance

Better Bound in a Special Case

- A *uniformly* quasi-bipartite instance is one in which for every Steiner node, all incident edges have the same cost

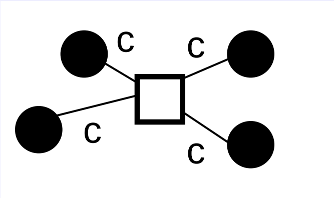

- For this special case the best approx algo known has ratio 73/60 [Gröpl et al. ‘00]
- We get a 73/60 integrality gap bound
  - Somewhat simpler proof of a stronger result
  - Only class of instances where best known integrality gap matches best known approx ratio

Proof of the 73/60 bound

Cost-per-connection algorithm of Gröpl et al:

1. For each K in increasing order of $C_K/(|K|-1)$,
   A. If K forms no hypercycle w/ previous purchases,
      Purchase K.

- Analysis: Define a natural dual solution with cost equal to the algorithm's output
- Dual is not feasible, but can prove it becomes feasible if scaled down by 73/60

Illustration of Proof Idea

- Algorithm selects bcde (has min cost per connection, 20)
- Then it selects ab (c.p.c. 28)
- Dual solution in proof assigns value 20 to {a, b, c, d, e} and value 8 to {a, bcde}

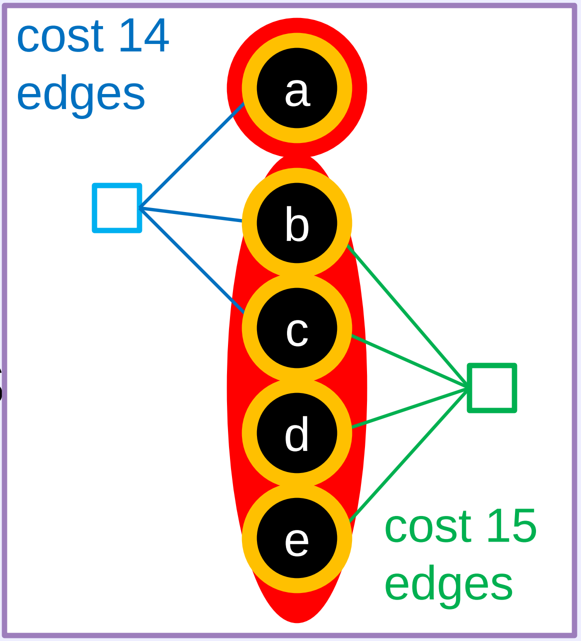

- Dual load on full component K := abc is

$$\sum_{\pi} y_\pi \mathrm{rc}_K^\pi = 20 \cdot 2 + 8 \cdot 1 = 48 < \frac{73}{60} C_K = \frac{73}{60} 42$$

Future Work

- Apply LP technology for approx algorithms?
  - Degree-bounded Steiner tree, Steiner forest, prize-collecting Steiner tree, k-Steiner tree, …
- Funny technical point in quasi-bipartite case:
  - Forget usual "*r*-restricted full component" trick
  - Can compute OPT of bidir and hence OPT $(\mathcal{P})$
  - Can we compute explicit primal opt of $(\mathcal{P})$?
    - Would save ε in quasibipartite approx ratio of [BGRS10]
- Better direct understanding of bidirected cut?