On the complexity of the asymmetric VPN problem

Thomas Rothvoß & **Laura Sanità**

Institute of Mathematics
EPFL, Lausanne

ISMP'09

Concave Cost VPN

Given:

- Undirected graph $G = (V, E)$, costs $c : E \to \mathbb{Q}_+$
- Outgoing traffic bound $b_v^+ \in \mathbb{N}_0$, ingoing traffic bound $b_v^- \in \mathbb{N}_0$
- Concave non-decreasing function $f : \mathbb{Q}_+ \to \mathbb{Q}_+$

Find: Paths P_{uv}, capacities x_e s.t.

$$\sum_{e \in E} c(e) \cdot f(x_e) \to \text{minimized}$$

and every *valid* traffic matrix $(D_{u,v})_{u,v \in V}$ can be routed.
D is valid if v sends $\leq b_v^+$ and receives $\leq b_v^-$

Concave Cost VPN

Given:

- Undirected graph $G = (V, E)$, costs $c : E \to \mathbb{Q}_+$
- Outgoing traffic bound $b_v^+ \in \mathbb{N}_0$, ingoing traffic bound $b_v^- \in \mathbb{N}_0$
- Concave non-decreasing function $f : \mathbb{Q}_+ \to \mathbb{Q}_+$

Find: Paths P_{uv}, capacities x_e s.t.

$$\sum_{e \in E} c(e) \cdot f(x_e) \to \text{minimized}$$

and every *valid* traffic matrix $(D_{u,v})_{u,v \in V}$ can be routed.
D is valid if v sends $\leq b_v^+$ and receives $\leq b_v^-$

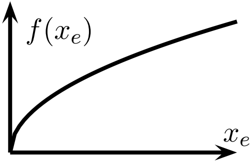

Concave Cost VPN

Given:

- Undirected graph $G = (V, E)$, costs $c : E \to \mathbb{Q}_+$
- Outgoing traffic bound $b_v^+ \in \mathbb{N}_0$, ingoing traffic bound $b_v^- \in \mathbb{N}_0$
- Concave non-decreasing function $f : \mathbb{Q}_+ \to \mathbb{Q}_+$

Find: Paths P_{uv}, capacities x_e s.t.

$$\sum_{e \in E} c(e) \cdot f(x_e) \to \text{minimized}$$

and every *valid* traffic matrix $(D_{u,v})_{u,v \in V}$ can be routed.
D is valid if v sends $\leq b_v^+$ and receives $\leq b_v^-$

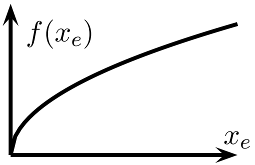

W.l.o.g.:

- senders $s \in S$: $b_s^+ = 1, b_s^- = 0$
- receivers $r \in R$: $b_r^+ = 0, b_r^- = 1$
- non-terminals v: $b_v^+ = b_v^- = 0$

Example

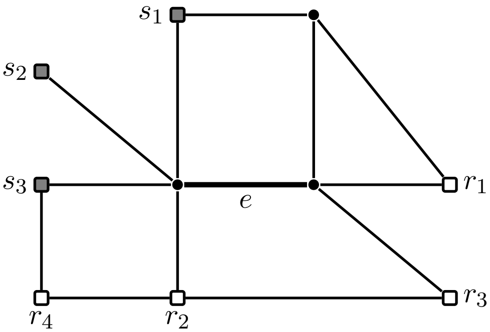

Example

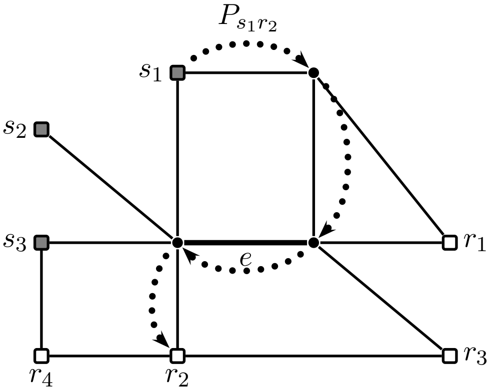

Example

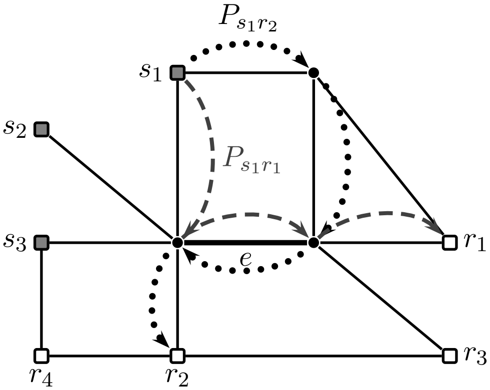

Example

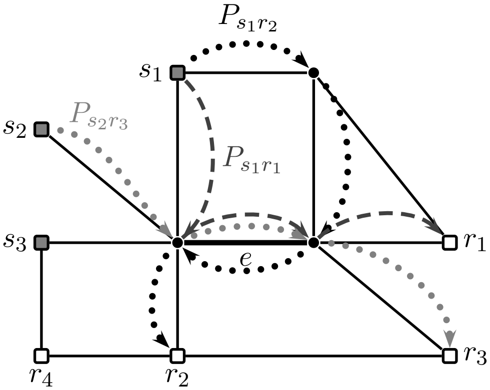

Example

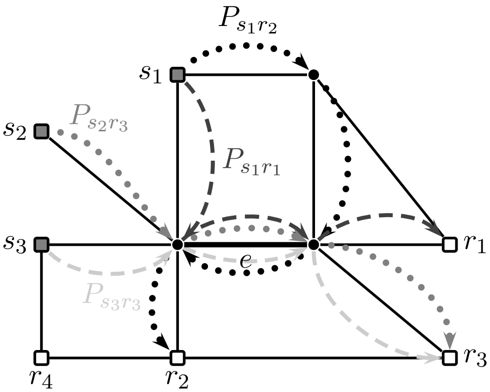

Example

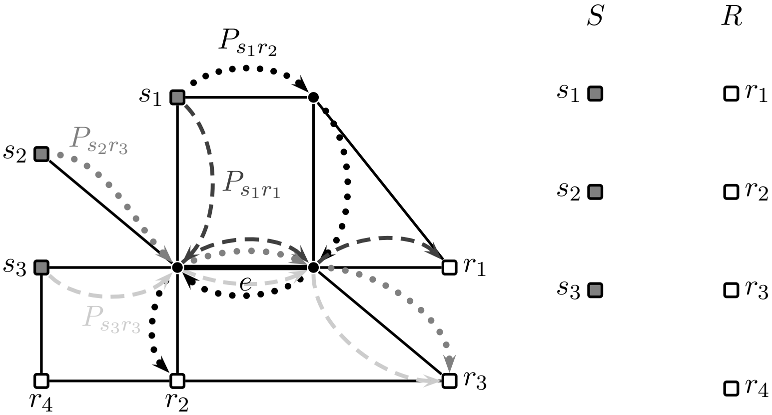
$P_{s_1r_2}$
s_1
s_2
$P_{s_2r_3}$
$P_{s_1r_1}$
s_3
e
r_1
$P_{s_3r_3}$
r_4
r_2
r_3
S
R
s_1
s_2
s_3
r_1
r_2
r_3
r_4

Example

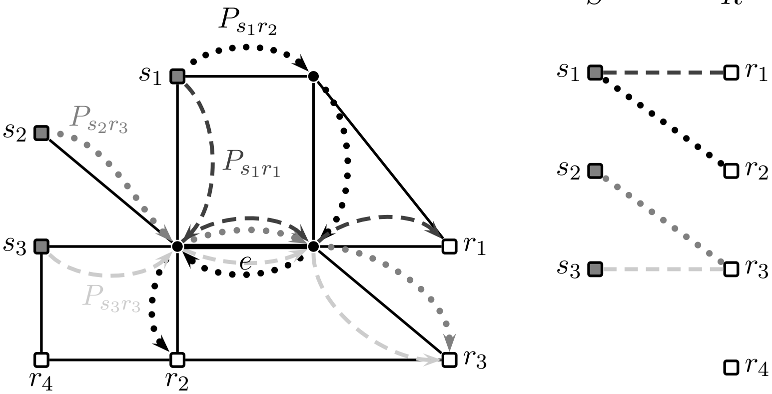

$P_{s_1r_2}$
s_1
s_2
$P_{s_2r_3}$
$P_{s_1r_1}$
s_3
e
r_1
$P_{s_3r_3}$
r_4
r_2
r_3
S
R
s_1
r_1
s_2
r_2
s_3
r_3
r_4

Example

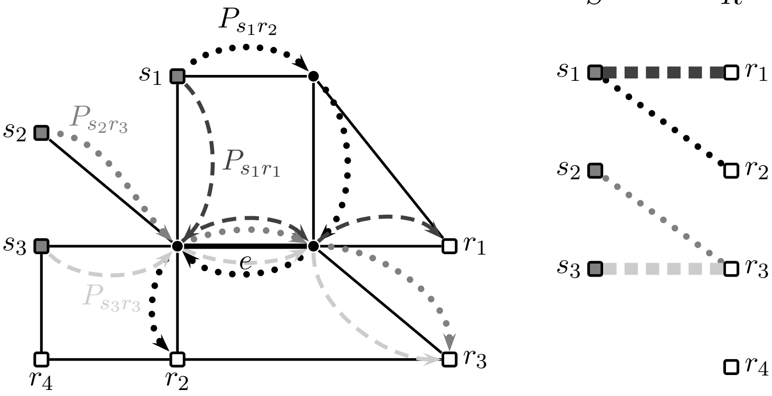

$$x_e = \text{maximal cardinality of a matching in } G_e = (S \cup R, E_e)$$
$$\text{with } (s,r) \in E_e \Leftrightarrow e \in P_{sr}$$

Example

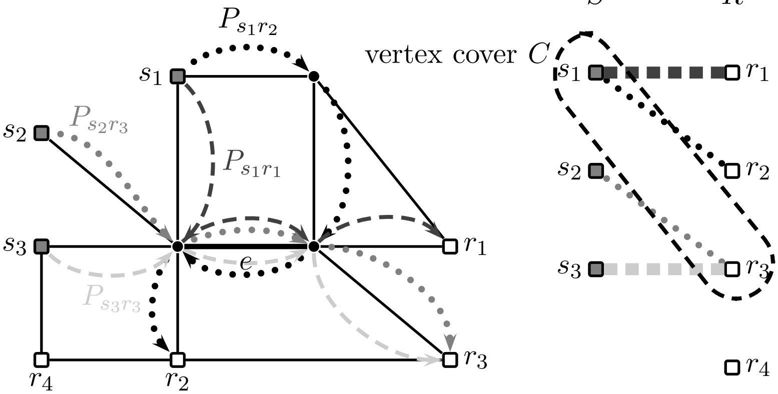

$$x_e = \text{maximal cardinality of a matching in } G_e = (S \cup R, E_e)$$
$$\text{with } (s, r) \in E_e \Leftrightarrow e \in P_{sr}$$

Known results

Linear costs:

- **APX**-hard
- 5.55-apx [Gupta, Kumar, Roughgarden '03]
- 4.74-apx [Eisenbrand, Grandoni '05]
- 3.55-apx [Eisenbrand, Grandoni, Oriolo, Skutella '07]

Known results

Linear costs:

- **APX**-hard
- 5.55-apx [Gupta, Kumar, Roughgarden ’03]
- 4.74-apx [Eisenbrand, Grandoni ’05]
- 3.55-apx [Eisenbrand, Grandoni, Oriolo, Skutella ’07]

Linear costs/symmetric ($b_v^+ = b_v^-$):

- Opt. solution is a tree [Goyal, Olver, Shepherd ’08]
- Opt. tree solution = best shortest path tree [Fingerhut et al. ’97; Gupta et al. ’01]

Known results

Linear costs:

- **APX**-hard
- 5.55-apx [Gupta, Kumar, Roughgarden '03]
- 4.74-apx [Eisenbrand, Grandoni '05]
- 3.55-apx [Eisenbrand, Grandoni, Oriolo, Skutella '07]

Linear costs/symmetric ($b_v^+ = b_v^-$):

- Opt. solution is a tree [Goyal, Olver, Shepherd '08]
- Opt. tree solution = best shortest path tree [Fingerhut et al. '97; Gupta et al. '01]

Linear costs/balanced ($|R| = |S|$):

- Opt. tree solution = best shortest path tree [Italiano, Leonardi, Oriolo '06]

Known results

Linear costs:

- **APX**-hard
- 5.55-apx [Gupta, Kumar, Roughgarden '03]
- 4.74-apx [Eisenbrand, Grandoni '05]
- 3.55-apx [Eisenbrand, Grandoni, Oriolo, Skutella '07]

Linear costs/symmetric ($b_v^+ = b_v^-$):

- Opt. solution is a tree [Goyal, Olver, Shepherd '08]
- Opt. tree solution = best shortest path tree [Fingerhut et al. '97; Gupta et al. '01]

Linear costs/balanced ($|R| = |S|$):

- Opt. tree solution = best shortest path tree [Italiano, Leonardi, Oriolo '06]

Theorem

There is a polytime 50-approximation for Concave Cost VPN that also gives a tree solution.

Single Sink Buy-at-Bulk

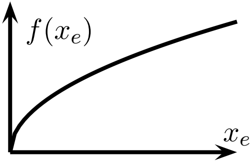

Given:

- $G = (V, E)$, costs $c : E \to \mathbb{Q}_+$
- *clients* $D \subseteq V$, *root* r
- Concave non-decreasing function $f : \mathbb{Q}_+ \to \mathbb{Q}_+$

Find: Capacities x_e s.t.

$$\sum_{e \in E} c(e) \cdot f(x_e) \to \text{minimize}$$

and each client can send a flow of 1 to r (simultaneously).

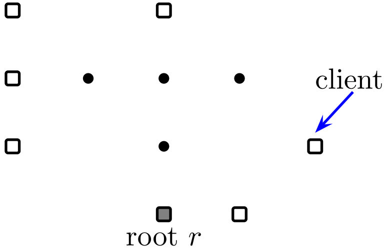

Single Sink Buy-at-Bulk

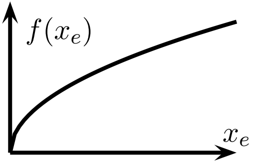

Given:

- $G = (V, E)$, costs $c : E \to \mathbb{Q}_+$
- *clients* $D \subseteq V$, *root* r
- Concave non-decreasing function $f : \mathbb{Q}_+ \to \mathbb{Q}_+$

Find: Capacities x_e s.t.

$$\sum_{e \in E} c(e) \cdot f(x_e) \to \text{minimize}$$

and each client can send a flow of 1 to r (simultaneously).

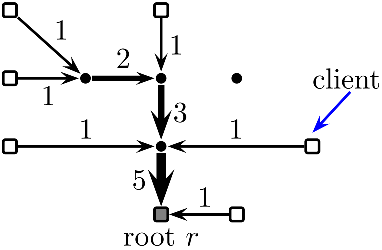

Single Sink Buy-at-Bulk

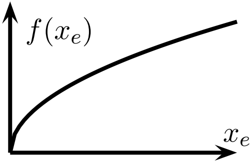

Given:

- $G = (V, E)$, costs $c : E \to \mathbb{Q}_+$
- *clients* $D \subseteq V$, *root* r
- Concave non-decreasing function $f : \mathbb{Q}_+ \to \mathbb{Q}_+$

Find: Capacities x_e s.t.

$$\sum_{e \in E} c(e) \cdot f(x_e) \to \text{minimize}$$

and each client can send a flow of 1 to r (simultaneously).

Known results:

- **APX**-hard
- Opt. solution is tree [Karger, Minkoff '00]
- For cable-based formulation:
  - 76.8-apx [Gupta, Kumar, Roughgarden '03]
  - improved to 25-apx [Grandoni, Italiano '06]

The algorithm

Algorithm:

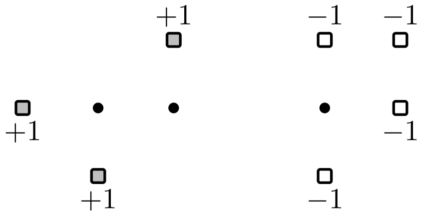

The algorithm

Algorithm:

1. Choose a sender $s^* \in S$ uniformly at random

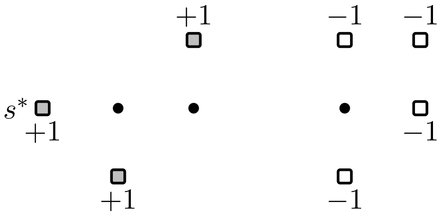

The algorithm

Algorithm:

1. Choose a sender $s^* \in S$ uniformly at random
2. Define *central hub* instance with single sender s^* (but $b^+_{s^*} = |S|$), receivers $S \cup R$

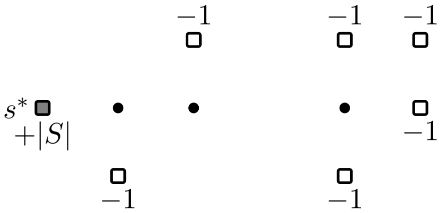

The algorithm

Algorithm:

1. Choose a sender $s^* \in S$ uniformly at random
2. Define *central hub* instance with single sender s^* (but $b^+_{s^*} = |S|$), receivers $S \cup R$
3. Compute 25-apx solution x'_e using a SSBB algo for the central hub VPN instance

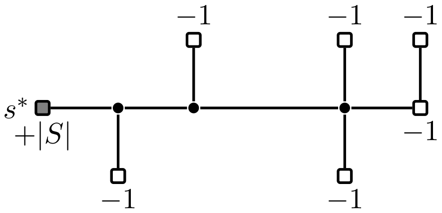

The algorithm

Algorithm:

1. Choose a sender $s^* \in S$ uniformly at random
2. Define *central hub* instance with single sender s^* (but $b^+_{s^*} = |S|$), receivers $S \cup R$
3. Compute 25-apx solution x'_e using a SSBB algo for the central hub VPN instance

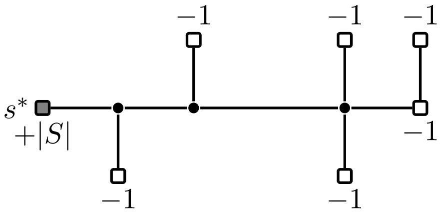

Claim: Capacities x'_e suffice for orig. instance

The algorithm

Algorithm:

1. Choose a sender $s^* \in S$ uniformly at random
2. Define *central hub* instance with single sender s^* (but $b^+_{s^*} = |S|$), receivers $S \cup R$
3. Compute 25-apx solution x'_e using a SSBB algo for the central hub VPN instance

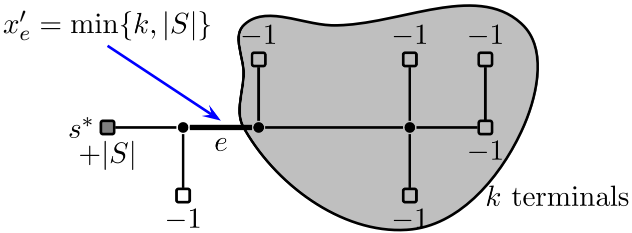

Claim: Capacities x'_e suffice for orig. instance

SSBB vs. VPN with central hub

VPN:

- sender s^* ($b^+_{s^*} = |S|$), receivers $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(x_e)$

SSBB:

- root s^*, clients $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(\min\{x_e, |S|\})$

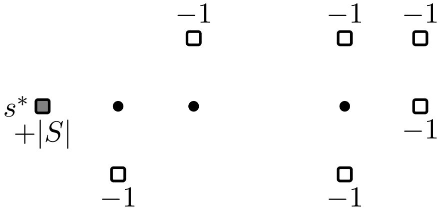

SSBB vs. VPN with central hub

VPN:

- sender s^* ($b^+_{s^*} = |S|$), receivers $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(x_e)$

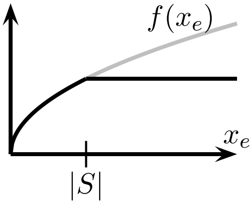

SSBB:

- root s^*, clients $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(\min\{x_e, |S|\}) \rightarrow$ concave

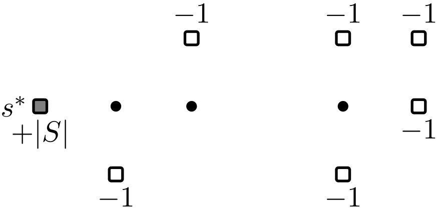

SSBB vs. VPN with central hub

VPN:

- sender s^* ($b^+_{s^*} = |S|$), receivers $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(x_e)$

SSBB:

- root s^*, clients $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(\min\{x_e, |S|\}) \rightarrow$ concave

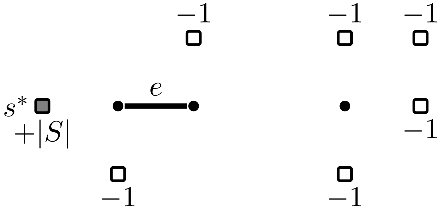

SSBB vs. VPN with central hub

VPN:

- sender s^* ($b^+_{s^*} = |S|$), receivers $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(x_e)$

SSBB:

- root s^*, clients $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(\min\{x_e, |S|\}) \to$ concave

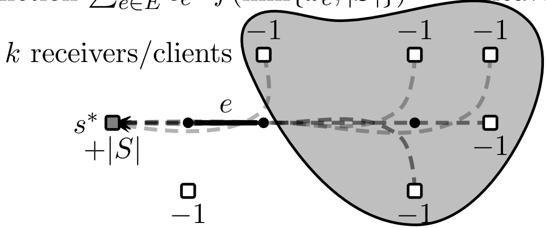

	capacity on e	cost for e
VPN		
SSBB		

SSBB vs. VPN with central hub

VPN:

- sender s^* ($b^+_{s^*} = |S|$), receivers $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(x_e)$

SSBB:

- root s^*, clients $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(\min\{x_e, |S|\}) \to$ concave

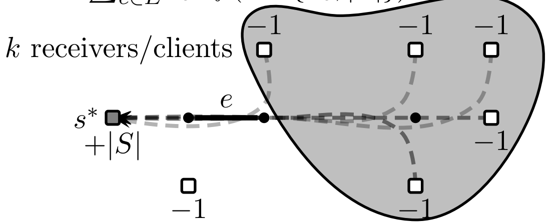

	capacity on e	cost for e
VPN	$\min\{k, \|S\|\}$	$c(e) \cdot f(\min\{k, \|S\|\})$
SSBB		

SSBB vs. VPN with central hub

VPN:

- sender s^* ($b^+_{s^*} = |S|$), receivers $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(x_e)$

SSBB:

- root s^*, clients $S \cup R$
- cost function $\sum_{e \in E} c_e \cdot f(\min\{x_e, |S|\}) \rightarrow$ concave

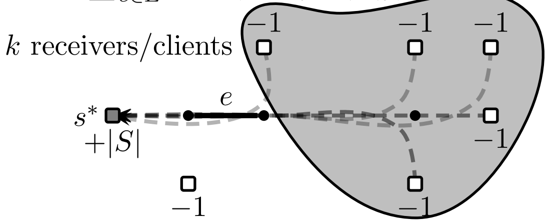

	capacity on e	cost for e
VPN	$\min\{k, \|S\|\}$	$c(e) \cdot f(\min\{k, \|S\|\})$
SSBB	k	$c(e) \cdot f(\min\{k, \|S\|\})$

Analysis

Theorem

There is a central hub solution of expected cost $\leq 2 \cdot OPT$.

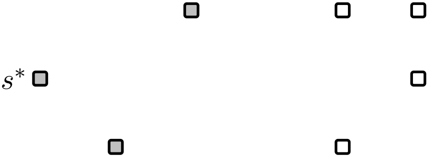

Analysis

Theorem

There is a central hub solution of expected cost $\leq 2 \cdot OPT$.

Proof:

- Let (x_e, P_{sr}) be optimal solution for orig. instance

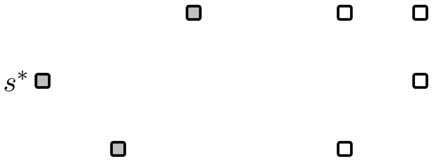

Analysis

Theorem

There is a central hub solution of expected cost $\leq 2 \cdot OPT$.

Proof:

- Let (x_e, P_{sr}) be optimal solution for orig. instance
- Choose a receiver r^* randomly

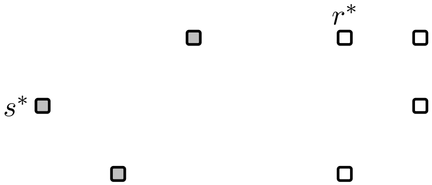

Analysis

Theorem

There is a central hub solution of expected cost $\leq 2 \cdot OPT$.

Proof:

- Let (x_e, P_{sr}) be optimal solution for orig. instance
- Choose a receiver r^* randomly
- Install $|S|$ units of capacity on $P_{s^*r^*}$

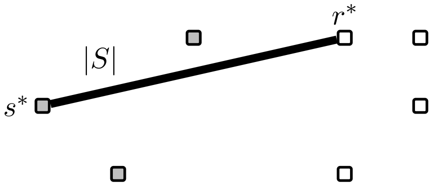

Analysis

Theorem

There is a central hub solution of expected cost $\leq 2 \cdot OPT$.

Proof:

- Let (x_e, P_{sr}) be optimal solution for orig. instance
- Choose a receiver r^* randomly
- Install $|S|$ units of capacity on $P_{s^*r^*}$
- Install (cumulatively) 1 unit on each P_{sr^*} and P_{s^*r} (in total never more then $|S|$)

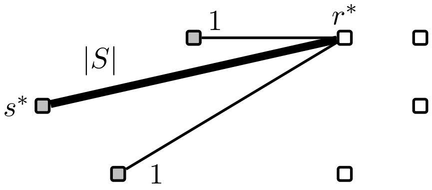

Analysis

Theorem

There is a central hub solution of expected cost $\leq 2 \cdot OPT$.

Proof:

- Let (x_e, P_{sr}) be optimal solution for orig. instance
- Choose a receiver r^* randomly
- Install $|S|$ units of capacity on $P_{s^*r^*}$
- Install (cumulatively) 1 unit on each P_{sr^*} and P_{s^*r} (in total never more then $|S|$)

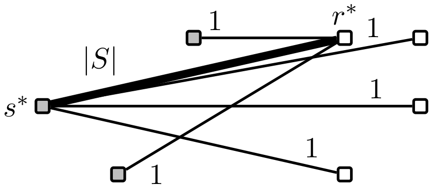

Analysis

Theorem

There is a central hub solution of expected cost $\leq 2 \cdot OPT$.

Proof:

- Let (x_e, P_{sr}) be optimal solution for orig. instance
- Choose a receiver r^* randomly
- Install $|S|$ units of capacity on $P_{s^*r^*}$
- Install (cumulatively) 1 unit on each P_{sr^*} and P_{s^*r} (in total never more then $|S|$)
- Claim: $E[\text{capacity on } e] \leq 2 \cdot x_e$

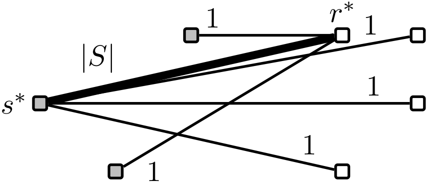

Analysis (2)

Lemma

$E[capacity\ on\ e] \leq 2 \cdot x_e.$

Analysis (2)

Lemma

$E[capacity\ on\ e] \leq 2 \cdot x_e$.

- Consider $G_e = (S \cup R, E_e)$ with edges $(s, r) \in E_e \Leftrightarrow e \in P_{s,r}$

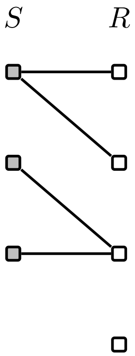

Analysis (2)

Lemma

$E[\mathit{capacity\ on\ e}] \leq 2 \cdot x_e$.

- Consider $G_e = (S \cup R, E_e)$ with edges $(s, r) \in E_e \Leftrightarrow e \in P_{s,r}$
- Let C be vertex cover with $|C| = x_e$

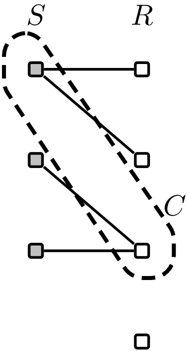

Analysis (2)

Lemma

$E[\textit{capacity on } e] \leq 2 \cdot x_e.$

- Consider $G_e = (S \cup R, E_e)$ with edges $(s, r) \in E_e \Leftrightarrow e \in P_{s,r}$
- Let C be vertex cover with $|C| = x_e$
- Case: s^* or r^* are in C
  - Prob: $\leq \frac{|S \cap C|}{|S|} + \frac{|R \cap C|}{|R|} \leq \frac{|C|}{|S|}$
  - Capacity: $\leq |S|$
  - Contribution: $\leq \frac{|C|}{|S|} \cdot |S| = |C| = x_e$

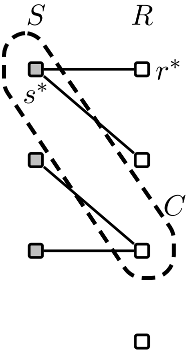

Analysis (2)

Lemma

$E[\text{capacity on } e] \leq 2 \cdot x_e$.

- Consider $G_e = (S \cup R, E_e)$ with edges $(s, r) \in E_e \Leftrightarrow e \in P_{s,r}$
- Let C be vertex cover with $|C| = x_e$
- Case: s^* or r^* are in C
  - Prob: $\leq \frac{|S \cap C|}{|S|} + \frac{|R \cap C|}{|R|} \leq \frac{|C|}{|S|}$
  - Capacity: $\leq |S|$
  - Contribution: $\leq \frac{|C|}{|S|} \cdot |S| = |C| = x_e$
- Case: Neither s^* nor r^* are in C
  - Prob: ≤ 1
  - $(s^*, r^*) \notin E_e \Rightarrow e \notin P_{s^* r^*}$
  - Capacity: $\leq \deg(s^*) + \deg(r^*) \leq |R \cap C| + |S \cap C| = |C|$
  - Contribution: $\leq |C| = x_e$

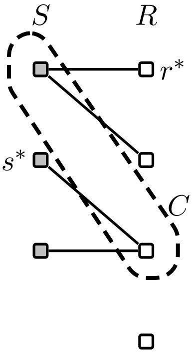

Balanced VPN is NP-hard

Theorem

VPN with linear costs and $|S| = |R|$ is still $\mathbf{NP}$*-hard.*

- $G = (V, E)$ Steiner tree instance with terminals $v_1, \ldots, v_{k+1}$
- $C > \sum_{e \in E} c(e)$, $M \gg (k+1)C$

Balanced VPN is NP-hard

Theorem

VPN with linear costs and $|S| = |R|$ is still $\mathbf{NP}$*-hard.*

- $G = (V, E)$ Steiner tree instance with terminals $v_1, \ldots, v_{k+1}$
- $C > \sum_{e \in E} c(e)$, $M \gg (k+1)C$

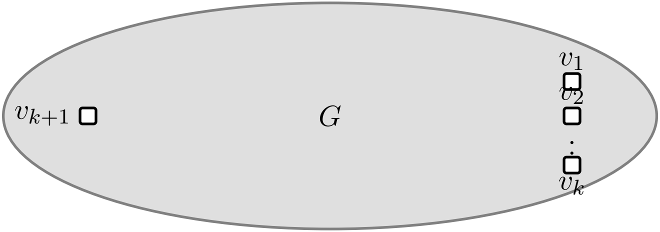

Balanced VPN is NP-hard

Theorem

VPN with linear costs and $|S| = |R|$ is still **NP***-hard.*

- $G = (V, E)$ Steiner tree instance with terminals $v_1, \ldots, v_{k+1}$
- $C > \sum_{e \in E} c(e)$, $M \gg (k+1)C$

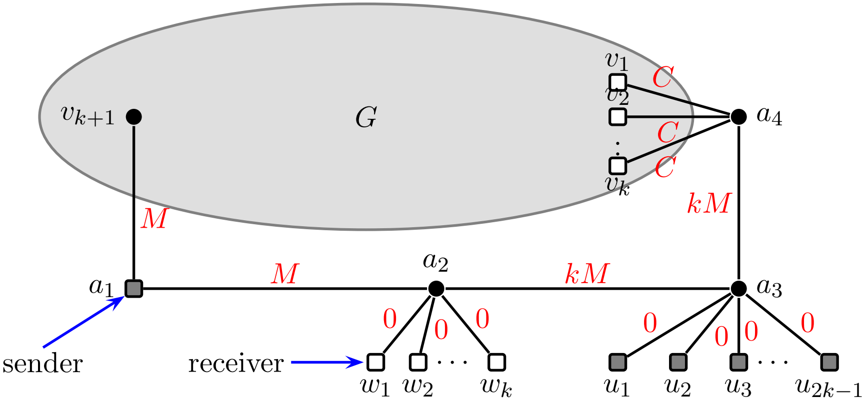

Balanced VPN is NP-hard

Theorem

VPN with linear costs and $|S| = |R|$ is still $\mathbf{NP}$*-hard.*

- $G = (V, E)$ Steiner tree instance with terminals $v_1, \ldots, v_{k+1}$
- $C > \sum_{e \in E} c(e)$, $M \gg (k+1)C$

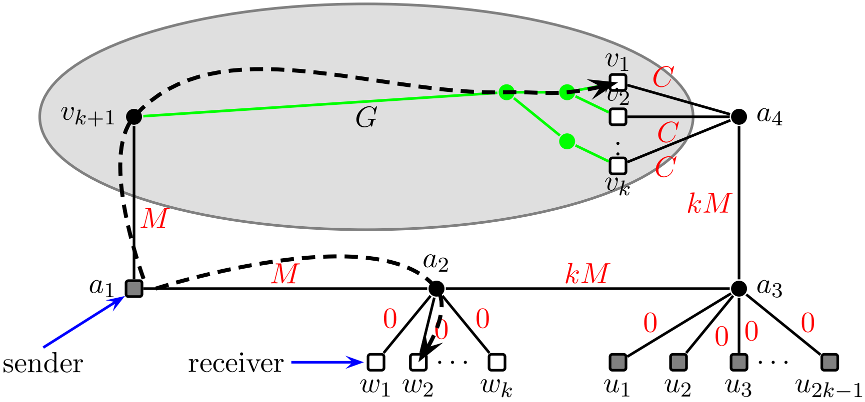

Balanced VPN is NP-hard

Theorem

VPN with linear costs and $|S| = |R|$ is still **NP***-hard.*

- $G = (V, E)$ Steiner tree instance with terminals $v_1, \ldots, v_{k+1}$
- $C > \sum_{e \in E} c(e)$, $M \gg (k+1)C$

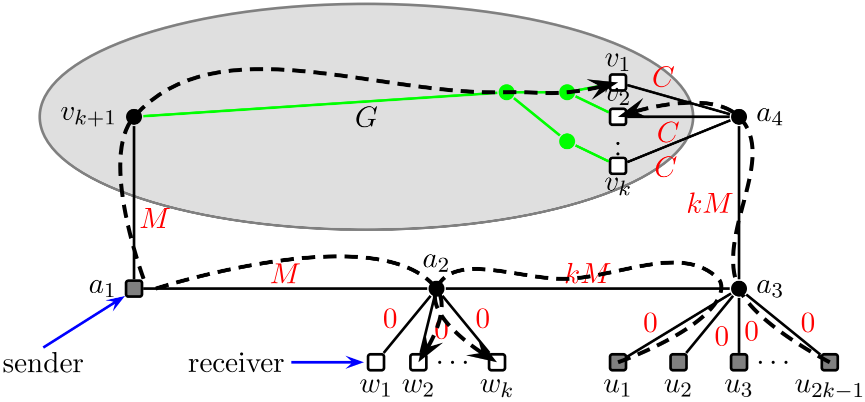

Balanced VPN is NP-hard

Theorem

VPN with linear costs and $|S| = |R|$ is still $\mathbf{NP}$*-hard.*

- $G = (V, E)$ Steiner tree instance with terminals $v_1, \ldots, v_{k+1}$
- $C > \sum_{e \in E} c(e)$, $M \gg (k+1)C$

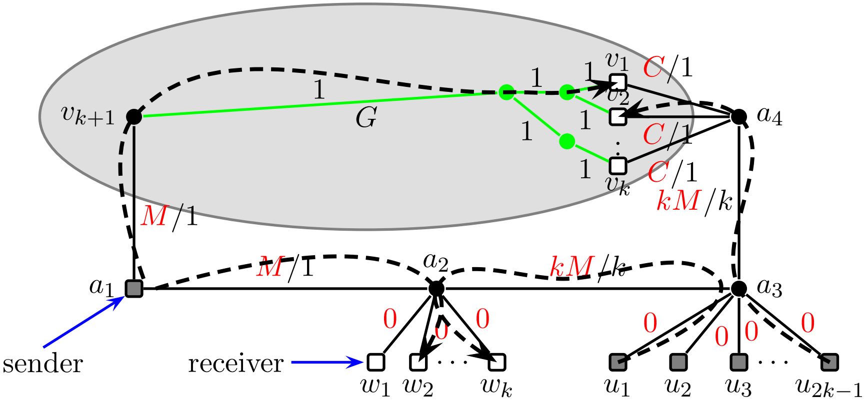

Conclusions

Conclusions

Theorem

There is a 50*-apx for Concave Cost VPN*

Conclusions

Theorem

There is a 50*-apx for Concave Cost VPN*

Corollary

ρ*-apx for SSBB* $\Rightarrow 2\rho$*-apx for Concave Cost VPN*

Conclusions

Theorem

There is a 50-apx for Concave Cost VPN

Corollary

ρ*-apx for SSBB* $\Rightarrow 2\rho$*-apx for Concave Cost VPN*

Corollary

$OPT_{tree} \leq 2 \cdot OPT$ *for Concave Cost VPN*

Conclusions

Theorem

There is a 50-apx for Concave Cost VPN

Corollary

ρ*-apx for SSBB* $\Rightarrow 2\rho$*-apx for Concave Cost VPN*

Corollary

$OPT_{tree} \leq 2 \cdot OPT$ *for Concave Cost VPN*

Theorem

VPN with linear costs and $|S| = |R|$ *is still* **NP***-hard.*

Conclusions

Theorem

There is a 50-apx for Concave Cost VPN

Corollary

ρ*-apx for SSBB* $\Rightarrow 2\rho$*-apx for Concave Cost VPN*

Corollary

$OPT_{tree} \leq 2 \cdot OPT$ *for Concave Cost VPN*

Theorem

VPN with linear costs and $|S| = |R|$ *is still* **NP***-hard.*

Thanks for your attention