


A Compact MIP for Aggregation and Multicast Trees under Flexible Routing and Function Placement

Matthias Rost, Technische Universität Berlin
joint work with Stefan Schmid, T-Labs & TU Berlin

July 17th, 2015
ISMP 2015, Pittsburgh

Communication Schemes: Multicast (same old! same old?)

processing = duplication + reroute



Communication Schemes: Multicast (same old! same old?)

processing = duplication + reroute





Figure: Hierarchy of processing nodes

Communication Schemes: Aggregation

processing = merge + reroute



Communication Schemes: Aggregation

processing = merge + reroute




Figure: Hierarchy of processing nodes

Problem Statement

Setting: Network Virtualization

- (Unsplittable) routes can be established arbitrarily (e.g. using Software-Defined Networks)
- Processing functionality can be placed on specific nodes (e.g. using Network Functions Virtualization)

Main Questions

- How to compute *virtual* aggregation / multicasting trees?
 - Where to place in-network processing functionality?
 - How to trade-off between traffic and processing?

Introductory Example

Aggregation scenario


- bi-directed grid graph



receiver



sender



Without in-network processing: Unicast

Solution Method

- minimal cost flow

Solution uses

- 41 edges
- 0 processing nodes



receiver



sender








Figure: Unicast solution

Solution Method

- ## Solution uses

-  receiver
  sender
-  processing node
  sender with processing






Figure: Aggregation tree

How to Trade-off?



vs.



What we aim for


Solution uses

- 26 edges
- 2 processing nodes

★ receiver

△ sender

□ processing node



Solution Structure




Figure: Virtual Arborescence




Figure: underlying routes

Input

Definition (Network $G = (V_G, E_G, c_E, u_E)$)

- integral capacities on the edges $u_E : E_G \rightarrow \mathbb{N}$
- positive edge costs $c_E : E_G \rightarrow \mathbb{R}^+$

Definition (Abstract Communication Request)



An abstract communication request on a graph G is defined as a 5-tuple $R_G = (r, S, T, u_r, c_S, u_S)$, where

- $T \subseteq V_G$ is the set of terminals,
- $r \in V_G \setminus T$ denotes the root with integral capacity $u_r \in \mathbb{N}$ and
- $S \subseteq V_G \setminus (\{r\} \cup T)$ denotes the set of possible *Steiner sites* with associated activation costs $c_S : S \rightarrow \mathbb{R}^+$ and integral capacities $u_S : S \rightarrow \mathbb{N}$.

Virtual Arborescence

Definition (Virtual Arborescence on G : $\mathcal{T}_G = (V_{\mathcal{T}}, E_{\mathcal{T}}, r, \pi)$)

- $\{r\} \subseteq V_{\mathcal{T}} \subseteq V_G$ and $E_{\mathcal{T}} \subseteq V_{\mathcal{T}} \times V_{\mathcal{T}}$
- $\pi : E_{\mathcal{T}} \rightarrow \mathcal{P}_G$ maps each edge of $E_{\mathcal{T}}$ on a (simple) path $P \in \mathcal{P}_G$, s.t.
 - (VA-1) $(V_{\mathcal{T}}, E_{\mathcal{T}}, r)$ is an rooted arborescence with edges either directed towards or away from r ,
 - (VA-2) for all $(u, v) \in E_{\mathcal{T}}$ the directed path $\pi(u, v)$ connects u to v in G .



Definition (Constrained Virtual Steiner Arborescence Problem)

Input: network $G = (V_G, E_G, c_E, u_E)$, request $R_G = (r, S, T, u_r, c_S, u_S)$.

Task: Find a minimal cost Virtual Arborescence $\mathcal{T}_G = (V_{\mathcal{T}}, E_{\mathcal{T}}, r, \pi)$ satisfying:

(CVSAP-1) $\{r\} \cup T \subseteq V_{\mathcal{T}}$ and $V_{\mathcal{T}} \subseteq \{r\} \cup S \cup T$,

(CVSAP-2) for all $t \in T$ holds $\delta_{E_{\mathcal{T}}}^+(t) + \delta_{E_{\mathcal{T}}}^-(t) = 1$,

(CVSAP-3) for the root $\delta_{E_{\mathcal{T}}}^+(r) + \delta_{E_{\mathcal{T}}}^-(r) \leq u_r$ holds,

(CVSAP-4) for all $s \in S \cap V_{\mathcal{T}}$ holds $\delta_{E_{\mathcal{T}}}^-(s) + \delta_{E_{\mathcal{T}}}^+(s) \leq u_S(s) + 1$ and

(CVSAP-5) for all $e \in E_G$ holds $|(\pi(E_{\mathcal{T}}))[e]| \leq u_E(e)$.

The cost of a Virtual Arborescence is defined to be

$$C_{\text{CVSAP}}(\mathcal{T}_G) = \sum_{e \in E_G} c_E(e) \cdot |(\pi(E_{\mathcal{T}}))[e]| + \sum_{s \in S \cap V_{\mathcal{T}}} c_S(s),$$

where $|(\pi(E_{\mathcal{T}}))[e]|$ denotes the number of times an edge is used.

Applications

Applications


	Network	Application	Technology, e.g.
multicast	ISP	service replication / cache placement [10, 11]	middleboxes / NFV + SDN
	backbone	optical multicast [6]	ROADM + SDH
	all	application-level multicast [16]	different
aggregation	sensor network	value & message aggregation [5, 8]	source routing
	ISP	network analytics: Gigascope [3]	middleboxes / NFV + SDN
	data center	big data / map-reduce: Camdoop [2]	SDN

edge capacities


processing node locations

processing node capacities

Service Replication




Service Replication




Service Replication

What if backend links are congested?




Service Replication




Service Replication

What if only '3' users can be handled?



Service Replication



Solution Approaches

Comprehensive algorithmic study

Computational Complexity
(Inapproximability)

Approximation
Algorithms

Exact Algorithms (MIPs)


LP-based Heuristics

- M.Sc. Thesis [13] *Matthias Rost (Advisors: Schmid, Bley, Feldmann)*
Optimal Virtualized In-Network Processing with Applications to
Aggregation and Multicast, TU Berlin '14
- Conference [15] *Matthias Rost and Stefan Schmid*
VirtuCast, Multicast and Aggregation with In-Network Processing,
OPODIS '13
- Tech. Report [14] *Matthias Rost and Stefan Schmid*
The Constrained Virtual Steiner Arborescence Problem: Formal
Definition, Single-Commodity Integer Programming Formulation
and Computational Evaluation, arXiv '13

Inapproximability

Inapproximability

Reduction via Set Cover: Does a set cover of size X exist?

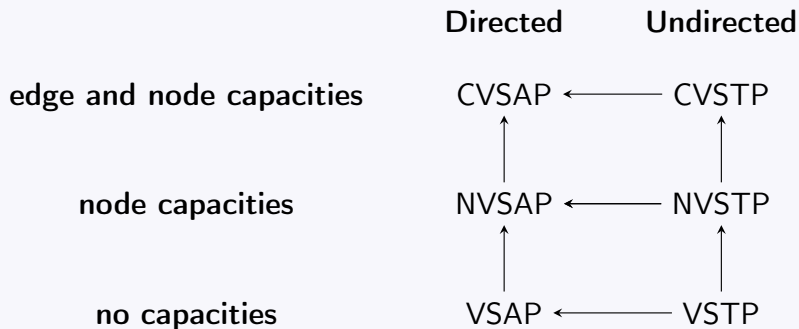


Theorem

Finding a *feasible* solution is already NP-complete.

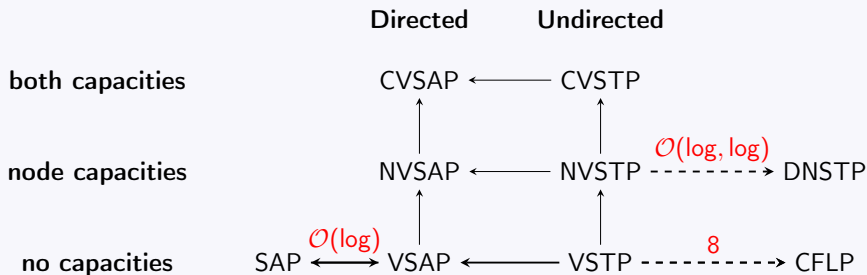
Approximation Algorithms for Variants

Variants



Approximation via related problems

Results




Exact Algorithms for CVSAP

Multi-Commodity Flow (MCF) Integer Program

First approach: MCF IP

- explicitly represent virtual arborescence
- necessitates independent construction of paths for all processing nodes




Multi-Commodity Flow (MCF) Integer Program

First approach: MCF IP

- explicitly represent virtual arborescence
- necessitates independent construction of paths for all processing nodes

Intuition: does not scale well

- number of binary variables:
 $\# \text{Steiner sites} \cdot \# \text{edges}$



Integer Program 1: A-CVSAP-MCF

$$\begin{aligned}
& \text{minimize} && C_{\text{MCF}} = \sum_{e \in E_G} c_e (f_e + \sum_{s \in S} f_{s,e}) && (\text{MCF-OBJ}) \\
& && + \sum_{s \in S} c_s \cdot x_s \\
& \text{subject to} && f^T(\delta_{\text{EMCF}}^+(v)) = f^T(\delta_{\text{EMCF}}^-(v)) + |\{v\} \cap T| && \forall v \in V_G && (\text{MCF-1}) \\
& && f^s(\delta_{\text{EMCF}}^+(v)) = f^s(\delta_{\text{EMCF}}^-(v)) + \delta_{s,v} \cdot x_s && \forall s \in S, v \in V_G && (\text{MCF-2}) \\
& && f_e^T + \sum_{s \in S} f_e^s \leq \begin{cases} u_s x_s, & e = (s, o^-), s \in S \\ u_r, & e = (r, o^-) \\ u_e, & e \in E_G \end{cases} && \forall e \in E_{\text{MCF}} && (\text{MCF-3}) \\
& && -|S|(1 - f_{\bar{s}, o^-}^s) \leq p_s - p_{\bar{s}} - 1 && \forall s, \bar{s} \in S && (\text{MCF-4}) \\
& && f_{(\bar{s}, o^-)}^s \leq x_{\bar{s}} && \forall s \in S, \bar{s} \in S - s && (\text{MCF-5}^*) \\
& && f_{s, o^-}^s = 0 && \forall s \in S && (\text{MCF-6}^*) \\
& && f_{\bar{s}, o^-}^s + f_{s, o^-}^{\bar{s}} \leq 1 && \forall s, \bar{s} \in S && (\text{MCF-7}^*) \\
& && x_s \in \{0, 1\} && \forall s \in S && (\text{MCF-8}) \\
& && f_e^T \in \mathbb{Z}_{\geq 0} && \forall e \in E_{\text{MCF}} && (\text{MCF-9}) \\
& && f_e^s \in \{0, 1\} && \forall s \in S, e \in E_{\text{MCF}} && (\text{MCF-10}) \\
& && p \in [0, |S| - 1] && \forall s \in S && (\text{MCF-11})
\end{aligned}$$

Integer Program 2: A-CVSAP-MCF

$$\begin{aligned}
& \text{minimize} && C_{\text{MCF}} = \sum_{e \in E_G} c_e (f_e + \sum_{s \in S} f_{s,e}) && (\text{MCF-OBJ}) \\
& && + \sum_{s \in S} c_s \cdot x_s \\
& \text{subject to} && f^T(\delta_{\text{EMCF}}^+(v)) = f^T(\delta_{\text{EMCF}}^-(v)) + |\{v\} \cap T| && \forall v \in V_G && (\text{MCF-1}) \\
& && f^s(\delta_{\text{EMCF}}^+(v)) = f^s(\delta_{\text{EMCF}}^-(v)) + \delta_{s,v} \cdot x_s && \forall s \in S, v \in V_G && (\text{MCF-2}) \\
& && f_e^T + \sum_{s \in S} f_e^s \leq \begin{cases} u_s x_s, & e = (s, o^-), s \in S \\ u_r, & e = (r, o^-) \\ u_e, & e \in E_G \end{cases} && \forall e \in E_{\text{MCF}} && (\text{MCF-3}) \\
& && -|S|(1 - f_{\bar{s}, o^-}^s) \leq p_s - p_{\bar{s}} - 1 && \forall s, \bar{s} \in S && (\text{MCF-4}) \\
& && f_{(\bar{s}, o^-)}^s \leq x_{\bar{s}} && \forall s \in S, \bar{s} \in S - s && (\text{MCF-5}^*) \\
& && f_{s, o^-}^s = 0 && \forall s \in S && (\text{MCF-6}^*) \\
& && f_{\bar{s}, o^-}^s + f_{s, o^-}^{\bar{s}} \leq 1 && \forall s, \bar{s} \in S && (\text{MCF-7}^*) \\
& && x_s \in \{0, 1\} && \forall s \in S && (\text{MCF-8}) \\
& && f_e^T \in \mathbb{Z}_{\geq 0} && \forall e \in E_{\text{MCF}} && (\text{MCF-9}) \\
& && f_e^s \in \{0, 1\} && \forall s \in S, e \in E_{\text{MCF}} && (\text{MCF-10}) \\
& && p \in [0, |S| - 1] && \forall s \in S && (\text{MCF-11})
\end{aligned}$$

Integer Program 3: A-CVSAP-MCF

$$\begin{aligned}
&\text{minimize} && C_{\text{MCF}} = \sum_{e \in E_G} c_e (f_e + \sum_{s \in S} f_{s,e}) && (\text{MCF-OBJ}) \\
&&& + \sum_{s \in S} c_s \cdot x_s \\
&\text{subject to} && f^T(\delta_{\text{EMCF}}^+(v)) = f^T(\delta_{\text{EMCF}}^-(v)) + |\{v\} \cap T| && \forall v \in V_G && (\text{MCF-1}) \\
&&& f^s(\delta_{\text{EMCF}}^+(v)) = f^s(\delta_{\text{EMCF}}^-(v)) + \delta_{s,v} \cdot x_s && \forall s \in S, v \in V_G && (\text{MCF-2}) \\
&&& f_e^T + \sum_{s \in S} f_e^s \leq \begin{cases} u_s x_s, & e = (s, o^-), s \in S \\ u_r, & e = (r, o^-) \\ u_e, & e \in E_G \end{cases} && \forall e \in E_{\text{MCF}} && (\text{MCF-3}) \\
&&& -|S|(1 - f_{\bar{s}, o^-}^s) \leq p_s - p_{\bar{s}} - 1 && \forall s, \bar{s} \in S && (\text{MCF-4}) \\
&&& f_{(\bar{s}, o^-)}^s \leq x_{\bar{s}} && \forall s \in S, \bar{s} \in S - s && (\text{MCF-5}^*) \\
&&& f_{s, o^-}^s = 0 && \forall s \in S && (\text{MCF-6}^*) \\
&&& f_{\bar{s}, o^-}^s + f_{s, o^-}^{\bar{s}} \leq 1 && \forall s, \bar{s} \in S && (\text{MCF-7}^*) \\
&&& x_s \in \{0, 1\} && \forall s \in S && (\text{MCF-8}) \\
&&& f_e^T \in \mathbb{Z}_{\geq 0} && \forall e \in E_{\text{MCF}} && (\text{MCF-9}) \\
&&& f_e^s \in \{0, 1\} && \forall s \in S, e \in E_{\text{MCF}} && (\text{MCF-10}) \\
&&& p \in [0, |S| - 1] && \forall s \in S && (\text{MCF-11})
\end{aligned}$$

Integer Program 4: A-CVSAP-MCF

$$\begin{aligned}
& \text{minimize} && C_{\text{MCF}} = \sum_{e \in E_G} c_e (f_e + \sum_{s \in S} f_{s,e}) && (\text{MCF-OBJ}) \\
& && + \sum_{s \in S} c_s \cdot x_s \\
& \text{subject to} && f^T(\delta_{E_{\text{MCF}}}^+(v)) = f^T(\delta_{E_{\text{MCF}}}^-(v)) + |\{v\} \cap T| && \forall v \in V_G && (\text{MCF-1}) \\
& && f^s(\delta_{E_{\text{MCF}}}^+(v)) = f^s(\delta_{E_{\text{MCF}}}^-(v)) + \delta_{s,v} \cdot x_s && \forall s \in S, v \in V_G && (\text{MCF-2}) \\
& && f_e^T + \sum_{s \in S} f_e^s \leq \begin{cases} u_s x_s, & e = (s, o^-), s \in S \\ u_r & , e = (r, o^-) \\ u_e & , e \in E_G \end{cases} && \forall e \in E_{\text{MCF}} && (\text{MCF-3}) \\
& && -|S|(1 - f_{\bar{s}, o^-}^s) \leq p_s - p_{\bar{s}} - 1 && \forall s, \bar{s} \in S && (\text{MCF-4}) \\
& && f_{(\bar{s}, o^-)}^s \leq x_{\bar{s}} && \forall s \in S, \bar{s} \in S - s && (\text{MCF-5}^*) \\
& && f_{s, o^-}^s = 0 && \forall s \in S && (\text{MCF-6}^*) \\
& && f_{\bar{s}, o^-}^s + f_{s, o^-}^{\bar{s}} \leq 1 && \forall s, \bar{s} \in S && (\text{MCF-7}^*) \\
& && x_s \in \{0, 1\} && \forall s \in S && (\text{MCF-8}) \\
& && f_e^T \in \mathbb{Z}_{\geq 0} && \forall e \in E_{\text{MCF}} && (\text{MCF-9}) \\
& && f_e^s \in \{0, 1\} && \forall s \in S, e \in E_{\text{MCF}} && (\text{MCF-10}) \\
& && p \in [0, |S| - 1] && \forall s \in S && (\text{MCF-11})
\end{aligned}$$

Single-Commodity Flow IP

Single-commodity flow formulation

- computes *aggregated* flow on edges independently of the origin
- does not represent virtual arborescence




Figure: Single-commodity

Multi- vs Single-Commodity

Example: 6000 edges and 200 Steiner sites

- Single-commodity: 6000 integer variables
- Multi-commodity: 1,200,000 binary variables




Figure: Single-commodity




Figure: Multi-commodity


VirtuCast Algorithm

VirtuCast Algorithm


Outline of VirtuCast

- 1 Solve single-commodity flow IP formulation.
- 2 Decompose IP solution into Virtual Arborescence.

How to
decompose?



(a) IP solution



(b) Virtual Arborescence



IP Formulation

Extended Graph

Additional nodes

- source o^+
- sinks o_r^- and o_s^-

Additional edges



IP Formulation I

$$\begin{array}{ll}
 \text{minimize} & C_{\text{IP}}(x, f) = \sum_{e \in E_G} c_e f_e + \sum_{s \in S} c_s x_s \\
 \text{subject to} & f(\delta_{E_{\text{ext}}}^+(v)) = f(\delta_{E_{\text{ext}}}^-(v)) \quad \forall v \in V_G \\
 & f(\delta_{E_{\text{ext}}}^+(W)) \geq x_s \quad \forall W \subseteq V_G, s \in W \cap S \neq \emptyset \\
 & f_e = 1 \quad \forall e = (o^+, t) \in E_{\text{ext}}^{T^+} \\
 & f_e = x_s \quad \forall e = (o^+, s) \in E_{\text{ext}}^{S^+} \\
 & x_s \in \{0, 1\} \quad \forall s \in S \\
 & f_e \in \mathbb{Z}_{\geq 0} \quad \forall e \in E_{\text{ext}}
 \end{array}$$

Connectivity Inequalities

STP Excursion [7]

$$\begin{array}{ll} \min & c^T x \\ \text{(uSP)} & \begin{array}{l} (i) \quad x(\delta(W)) \geq 1, \quad \text{for all } W \subset V, W \cap T \neq \emptyset, \\ & (V \setminus W) \cap T \neq \emptyset, \\ (ii) \quad 0 \leq x_e \leq 1, \quad \text{for all } e \in E, \\ (iii) \quad x \text{ integer,} \end{array} \end{array}$$

Connectivity Inequalities

STP Excursion [7]

$$\begin{aligned}
 & \min \quad c^T x \\
 (\text{uSP}) \quad & \begin{aligned}
 & (i) \quad x(\delta(W)) \geq 1, \quad \text{for all } W \subset V, W \cap T \neq \emptyset, \\
 & \quad \quad (V \setminus W) \cap T \neq \emptyset, \\
 & (ii) \quad 0 \leq x_e \leq 1, \quad \text{for all } e \in E, \\
 & (iii) \quad x \text{ integer,}
 \end{aligned}
 \end{aligned}$$

$$\forall W \subseteq V_G, s \in W \cap S \neq \emptyset. \quad f(\delta_{E_{\text{ext}}}^+(W)) \geq x_s$$

'From each activated Steiner site, there exists a path towards o_r^- .'

Exponentially many constraints, but ...

... can be separated in polynomial time.

Complete Formulation

$$\begin{array}{ll}
 \text{minimize} & C_{\text{IP}}(x, f) = \sum_{e \in E_G} c_e f_e + \sum_{s \in S} c_s x_s \\
 \text{subject to} & f(\delta_{E_{\text{ext}}}^+(v)) = f(\delta_{E_{\text{ext}}}^-(v)) \quad \forall v \in V_G \\
 & f(\delta_{E_{\text{ext}}}^+(W)) \geq x_s \quad \forall W \subseteq V_G, s \in W \cap S \neq \emptyset \\
 & f_e \leq u_s x_s \quad \forall e = (s, o_s^-) \in E_{\text{ext}}^{S^-} \\
 & f_{(r, o_r^-)} \leq u_r \\
 & f_e \leq u_e \quad \forall e \in E_G \\
 & f_e = 1 \quad \forall e \in E_{\text{ext}}^{T^+} \\
 & f_e = x_s \quad \forall e = (o^+, s) \in E_{\text{ext}}^{S^+} \\
 & x_s \in \{0, 1\} \quad \forall s \in S \\
 & f_e \in \mathbb{Z}_{\geq 0} \quad \forall e \in E_{\text{ext}}
 \end{array}$$


Decomposing flow is non-trivial!

Flow solution ...

- contains cycles and
- represents *arbitrary* hierarchies.

However, ...

- decomposition is *always* feasible
- constructive proof:
polynomial time algorithm



Outline of Decomposition Algorithm

Decomposition Approach

- 1 select a terminal $t \in T$
- 2 construct path P from t towards o_r^-
- 3 reduce flow along edges in P , *s.t. connectivity inequalities are valid*
- 4 connect t to the second last node of P and remove t

Outline of Decomposition Algorithm

Reduced problem must be feasible

Removing flow must not invalidate any connectivity inequalities.

Principle: Repair & Redirect


- decrease flow on path edge by edge
- if connectivity inequalities are violated
 - repair** increment flow on edge to regain feasibility
 - redirect** choose a different path from current node

Theorem

Given an optimal solution, the Decomposition Algorithm computes a Virtual Arborescence in time $\mathcal{O}(|V_G|^2 \cdot |E_G| \cdot (|V_G| + |E_G|))$.

Example

Scenario



sender




Steiner
site



receiver

Example

Extended Graph




sender

Steiner
site

receiver

Example

Solution




sender

activated
Steiner
site


receiver

Decomposition Example I




Decomposition Example I

$$P = \langle o^+, t_1, v, r, o_r^- \rangle$$




Decomposition Example I

$$P = \langle o^+, t_1, v, r, o_r^- \rangle$$




Decomposition Example I

$$P = \langle o^+, t_1, v, r, o_r^- \rangle$$




Decomposition Example I

$$P = \langle o^+, t_1, v, r, o_r^- \rangle$$




Redirecting Flow



Violation of Connectivity Inequality

$$f(\delta_{E_{\text{ext}}}^+(W)) \geq x_s \quad \forall W \subseteq V_G, s \in W \cap S \neq \emptyset$$


Redirecting Flow



Redirection towards o_s^- is possible!

There exists a path from v towards o_s^- in W .

Redirecting Flow




Redirection towards o_s^- is possible!

There exists a path from v towards o_s^- in W .

Reasoning


- ① Flow preservation holds within W .
- ② s could reach o_r^- via v before the reduction of flow.
- ③ v receives at least one unit of flow.
- ④ Flow leaving v must eventually terminate at o_s^- .

Decomposition Example II




Decomposition Example II


$$P = \langle o^+, t_1, v, \textcolor{red}{s}, o_S^- \rangle$$




Decomposition Example II




Solution




Decomposition Example II




Solution




Decomposition Example II




Solution



Decomposition Example II




Solution



Decomposition Example II

Final Solution



Linear Heuristics

Overview

Linear Relaxations

- The linear relaxation of an IP is obtained by relaxing the integrality constraints of the variables, thereby obtaining a Linear Program (LP).
- Solutions to linear relaxations are readily available when using branch-and-bound to solve an IP.
- May provide useful information to guide the construction of a solution.

Usage

- LP-based heuristics are employed within the *VirtuCast solver* to improve the bounding process.
- Yield polynomial time heuristics when used together with the root relaxation.

FlowDecoRound Heuristic

- computes a *flow* decomposition and connects nodes randomly according to the decomposition
- processing nodes are activated if another node connects to it, must be connected itself
- failsafe: shortest paths

Algorithm 1: FlowDecoRound

Input : Network $G = (V_G, E_G, c_E, u_E)$, Request $R_G = (r, s, T, u_r, c_s, u_s)$, LP relaxation solution $(\hat{r}, \hat{t}) \in \mathcal{F}_{LP}$ to ??

Output: A Feasible Virtual Arborescence \hat{T}_G or null

```

1 set  $\hat{S} \triangleq \emptyset$  and  $\hat{T} \triangleq \emptyset$  and  $U = T$ 
2 set  $\hat{V}_T \triangleq \{r\}$ ,  $\hat{E}_T \triangleq \emptyset$  and  $\hat{\pi} : \hat{E}_T \rightarrow \mathcal{P}_G$ 
3 set  $u(e) \triangleq \begin{cases} u_E(e) & , \text{ if } e \in E_G \\ u_r(r) & , \text{ if } e = (r, o_r^-) \\ u_s(s) & , \text{ if } e = (s, o_s^-) \in E_{\text{ext}}^{S^-} \\ 1 & , \text{ else} \end{cases} \text{ for all } e \in E_{\text{ext}}$ 
4 while  $U \neq \emptyset$  do
5   choose  $t \in U$  uniformly at random and set  $U \leftarrow U - t$ 
6   set  $\Gamma_t \triangleq \text{MinCostFlow}(G_{\text{ext}}, \hat{r}, \hat{t}(o^+, t), t, \{o_s^-, o_r^-\})$ 
7   set  $\hat{t} \leftarrow \hat{t} - \sum_{(P,f) \in \Gamma_t, e \in P} f$ 
8   set  $\Gamma_t \leftarrow \Gamma_t \setminus \{(P, f) \in \Gamma_t \mid \exists e \in P, u(e) = 0\}$ 
9   set  $\Gamma_t \leftarrow \Gamma_t \setminus \{(P, f) \in \Gamma_t \mid (\hat{V}_T + t, \hat{E}_T + (t, P_{|P|-1})) \text{ is not acyclic}\}$ 
10  if  $\Gamma_t \neq \emptyset$  then
11    choose  $(P, f) \in \Gamma_t$  with probability  $f / \left( \sum_{(P_j, f_j) \in \Gamma_t} f_j \right)$ 
12    if  $P_{|P|-1} \notin \hat{V}_T$  then
13      set  $U \leftarrow U + P_{|P|-1}$  and  $\hat{V}_T \leftarrow \hat{V}_T + P_{|P|-1}$ 
14      set  $\hat{V}_T \leftarrow \hat{V}_T + t$  and  $\hat{E}_T \leftarrow \hat{E}_T + (t, P_{|P|-1})$ 
15      and  $\hat{\pi}(t, P_{|P|-1}) \triangleq P$ 
16      set  $u(e) \leftarrow u(e) - 1$  for all  $e \in P$ 
17  set  $u(e) \leftarrow 0$  for all  $e = (s, o_s^-) \in E_{\text{ext}}^{S^-}$  with  $s \in S \wedge s \notin \hat{V}_T$ 
18  set  $\hat{T} \triangleq (\hat{T} \setminus \hat{V}_T) \cup (\{s \in S \cap \hat{V}_T \mid \delta_{\hat{E}_T}^+(s) = 0\})$ 
19  for  $t \in \hat{T}$  do
20    choose  $P \leftarrow \text{ShortestPath}(G_{\text{ext}}^{u, c_E, t}, \{o_s^-, o_r^-\})$ 
21    such that  $(\hat{V}_T + t, \hat{E}_T + (t, P_{|P|-1}))$  is acyclic
22    if  $P = \emptyset$  then
23      return null
24    set  $\hat{V}_T \leftarrow \hat{V}_T + t$  and  $\hat{E}_T \leftarrow \hat{E}_T + (t, P_{|P|-1})$  and  $\hat{\pi}(t, P_{|P|-1}) \triangleq P$ 
25    set  $u(e) \leftarrow u(e) - 1$  for all  $e \in P$ 
26  for  $e \in \hat{E}_T$  do
27    set  $P \triangleq \hat{\pi}(e)$ 
28    set  $\hat{\pi}(e) \leftarrow (P_1, \dots, P_{|P|-1})$ 
29  set  $\hat{T}_G \triangleq \text{Virtual Arborescence}(\hat{V}_T, \hat{E}_T, r, \hat{\pi})$ 
30  return PruneSteinerNodes( $\hat{T}_G$ )

```

Intermezzo: VCPriMConnect

Important Observation

If all placed processing nodes are already connected, all senders can be assigned *optimally* using a minimum cost flow.

Outline

- 1 connect all opened processing nodes in tree via a adaption of Prim's MST algorithm
- 2 assign all sending nodes using min-cost flow

Algorithm 2: VCPriMConnect

Input : Network $G = (V_G, E_G, c_E, u_E)$, Request

$R_G = (r, S, T, u_r, c_S, u_S)$,

Partial Virtual Arborescence $\mathcal{T}_G^P = (V_T^P, E_T^P, r, \pi^P)$

Output: Feasible Virtual Arborescence $\mathcal{T}_G = (V_T, E_T, r, \pi)$ or null

```

1 set  $U \triangleq \{v | v \in V_T^P \setminus \{r\}, \delta_{E_T^P}^+(v) = 0\}$ 
2 set  $\tilde{S} \triangleq U \cap S$ 
3 set  $V_T \triangleq V_T^P$ ,  $E_T \triangleq E_T^P$  and  $\pi(u, v) = \pi^P(u, v)$  for all  $(u, v) \in E_T$ 
4 set  $u(e) \triangleq u_E(e) - |\pi(E_T)[e]|$  for all  $e \in E_G$ 
5 while  $\tilde{S} \neq \emptyset$  do
6   compute  $R \leftarrow \{r' | r' \in \{r\} \cup (V_T \cap S), r' \text{ reaches } r \text{ in } \mathcal{T}_G, \delta_{E_T}^-(r') < u_{r,S}(r')\}$ 
7   compute  $P = \text{MinAllShortestPath}(G^u, c_E, \tilde{S}, R)$ 
8   if  $P = \text{null}$  then
9     return null
10  end
11  set  $\tilde{S} \leftarrow \tilde{S} - P_1$ 
12  set  $E_T \leftarrow E_T + (P_1, P_{|P|})$  and  $\pi(P_1, P_{|P|}) \triangleq P$ 
13  set  $u(e) \leftarrow u(e) - 1$  for all  $e \in P$ 
14 end
15 set  $\tilde{T} \triangleq U \cap T$ 
16 set  $u_v(r') \triangleq u_{r,S}(r') - \delta_{E_T}^-(r')$  for all  $r' \in \{r\} \cup (V_T \cap S)$ 
17 compute  $\Gamma = \{P^t\} \leftarrow \text{MinCostAssignment}(G, c_E, u, u_v, \tilde{T}, \{r\} \cup V_T \cap S)$ 
18 if  $\Gamma = \emptyset$  then
19   return null
20 end
21 set  $E_T \leftarrow E_T + (t, P_{|P^t|}^t)$  and  $\pi(t, P_{|P^t|}^t) \triangleq P^t$  for all  $P^t \in \Gamma$ 
22 return  $\mathcal{T}_G \triangleq (V_T, E_T, r, \pi)$ 

```


MultipleShots

- treats node variables as probabilities and iteratively places processing functionality accordingly
- tries to generate a feasible solution in each round via VCPriMConnect

Algorithm 3: MultipleShots

Input : Network $G = (V_G, E_G, c_E, u_E)$, Request

$R_G = (r, S, T, u_r, c_S, u_S)$,

LP relaxation solution $(\hat{x}, \hat{r}) \in \mathcal{F}_{LP}$ to ??

Output: A Feasible Virtual Arborescence \hat{T}_G or null

```

1 set  $|S| \triangleq \{s \in S | \hat{x}_s \leq 0.01\}$  and  $|S| \triangleq \{s \in S | \hat{x}_s \geq 0.99\}$ 
2 addConstraintsLocally( $\{x_s = 0 | s \in |S|\} \cup \{x_s = 1 | s \in |S|\}$ )
3 set  $\hat{S}_0 \triangleq |S| \cup$  and  $\hat{S}_1 \triangleq |S|$ 
4 disableGlobalPrimalBound()
5 repeat
6    $(\hat{x}, \hat{r}) \leftarrow \text{solveSeparateSolve}()$ 
7   if infeasibleLP() return null
8   set  $|S| \triangleq \{s \in S | \hat{x}_s \leq 0.01\}$  and  $|S| \triangleq \{s \in S | \hat{x}_s \geq 0.99\}$ 
9   addConstraintsLocally( $\{x_s = 0 | s \in |S|\} \cup \{x_s = 1 | s \in |S|\}$ )
10  set  $\hat{S}_0 \leftarrow \hat{S}_0 \cup |S|$  and  $\hat{S}_1 \leftarrow \hat{S}_1 \cup |S|$ 
11  set  $\hat{S} \triangleq \hat{S}_0 \setminus (\hat{S}_0 \cup \hat{S}_1)$ 
12  if  $\hat{S} \neq \emptyset$  then
13    repeat
14      set  $S_1 \triangleq \hat{S}$ 
15      remove  $s$  from  $S_1$  with probability  $1 - \hat{x}_s$  for all  $s \in S_1$ 
16      if  $S_1 = \emptyset$  and  $|S \setminus (\hat{S}_0 \cup \hat{S}_1)| < 10$  then
17        set  $S_1 \leftarrow S \setminus (\hat{S}_0 \cup \hat{S}_1)$ 
18    until  $S_1 \neq \emptyset$ 
19    addConstraintsLocally( $\{x_s = 1 | s \in S_1\}$ )
20    set  $\hat{S}_1 \leftarrow \hat{S}_1 \cup S_1$ 
21   $\hat{T}_G^P \triangleq (\hat{V}_T^P, \hat{E}_T^P, r, \emptyset)$  where  $\hat{V}_T^P \triangleq \{r\} \cup T \cup \hat{S}_1$  and  $\hat{E}_T \triangleq \emptyset$ 
22  set  $\hat{T}_G \triangleq \text{VCPriMConnect}(G, R_G, \hat{T}_G^P)$ 
23  if  $\hat{T}_G \neq \text{null}$  then
24    return PruneSteinerNodes( $\hat{T}_G$ )
25 until  $\hat{S}_0 \cup \hat{S}_1 = S$ 
26 return null

```

GreedyDiving

- aims at generating a feasible *IP* solution
- iteratively bounds at least a single variable from below, first fixing node variables
- complex failsafe:
PartialDecompose + VCPPrimConnect

Algorithm 4: GreedyDiving

Input : Network $G = (V_G, E_G, c_E, u_E)$, Request $R_G = (r, S, T, u_r, c_S, u_S)$, LP relaxation solution $(\hat{x}, \hat{r}) \in \mathcal{F}_{LP}$ to ??

Output: A Feasible Virtual Arborescence \tilde{T}_G or null


```

1 set  $|S| \triangleq \{s \in S | \hat{x}_s \leq 0.01\}$  and  $|S| \triangleq \{s \in S | \hat{x}_s \geq 0.99\}$ 
2 addConstraintsLocally( $\{x_s = 0 | s \in |S|\} \cup \{x_s = 1 | s \in |S|\}$ )
3 set  $\hat{S} \triangleq |S| \cup |S|$  and  $\hat{E} \triangleq \emptyset$ 
4 do
5    $(\hat{x}', \hat{r}') \leftarrow \text{solveSeparateSolve}()$ 
6   if  $\text{infeasibleLP}()$  and  $\hat{S} = S$  then
7     break
8   else if  $\text{infeasibleLP}()$  or  $\text{objectiveLimit}()$  then
9     return null
10  set  $(\hat{x}, \hat{r}) \leftarrow (\hat{x}', \hat{r}')$ 
11  if  $\hat{S} \neq S$  then
12    set  $|S| \triangleq \{s \in S | \hat{x}_s \leq 0.01\}$  and  $|S| \triangleq \{s \in S | \hat{x}_s \geq 0.99\}$ 
13    addConstraintsLocally( $\{x_s = 0 | s \in |S|\} \cup \{x_s = 1 | s \in |S|\}$ )
14    set  $\hat{S} \leftarrow \hat{S} \cup |S| \cup |S|$ 
15    set  $\hat{S} \triangleq \hat{S}$ 
16    if  $\hat{S} \neq \emptyset$  then
17      choose  $\hat{s} \in \hat{S}$  with  $c_S(\hat{s})/\hat{x}_{\hat{s}}$  minimal
18      addConstraintsLocally( $\{x_{\hat{s}} = 1\}$ )
19      set  $\hat{S} \leftarrow \hat{S} + \hat{s}$ 
20  else if  $\hat{E} \neq E_{\text{ext}}$  then
21    set  $|E| \triangleq \{e \in E_{\text{ext}} | |\hat{r}_e - \lceil \hat{r}_e \rceil| \leq 0.001\}$ 
22     $|E| \triangleq \{e \in E_{\text{ext}} | \lceil \hat{r}_e - \lceil \hat{r}_e \rceil \rceil \leq 0.001\}$ 
23    addConstraintsLocally( $\{\lceil \hat{r}_e \rceil | e \in |E|\} \cup \{\lceil \hat{r}_e \rceil | e \in |E|\}$ )
24    set  $\hat{E} \leftarrow \hat{E} \cup |E| \cup |E|$ 
25    set  $\hat{E} \triangleq E_{\text{ext}} \setminus \hat{E}$ 
26    if  $\hat{E} \neq \emptyset$  then
27      choose  $\hat{e} \in \hat{E}$  with  $\lceil \hat{r}_{\hat{e}} \rceil - \hat{r}_{\hat{e}}$  minimal
28      addConstraintsLocally( $\{\lceil \hat{r}_{\hat{e}} \rceil\}$ )
29      set  $\hat{E} \leftarrow \hat{E} + \hat{e}$ 
29  else
30    break
31  set  $\hat{r}_e \leftarrow \lceil \hat{r}_e \rceil$  for all  $e \in E_{\text{ext}} \setminus \hat{E}$ 
32  set  $\tilde{T}_G^D \leftarrow \text{PartialDecompose}(G, R_G, (\hat{x}, \hat{r}))$ 
33  return VCPPrimConnect( $G, R_G, \tilde{T}_G^D$ )


```

Computational Evaluation


Topologies



3D torus



Fat tree



An ISP topology generated by IGen with 2400 nodes.

Instances

Generation Parameters

- five graph sizes I-V
- 15 instances per graph size: different Steiner costs, different edge capacities

	Nodes	Edges	Processing Locations	Senders
Fat tree	1584	14680	720	864
3D torus	1728	10368	432	864
IGen	4000	16924	401	800

Table: Largest graph sizes

Computational Setup

Implementation

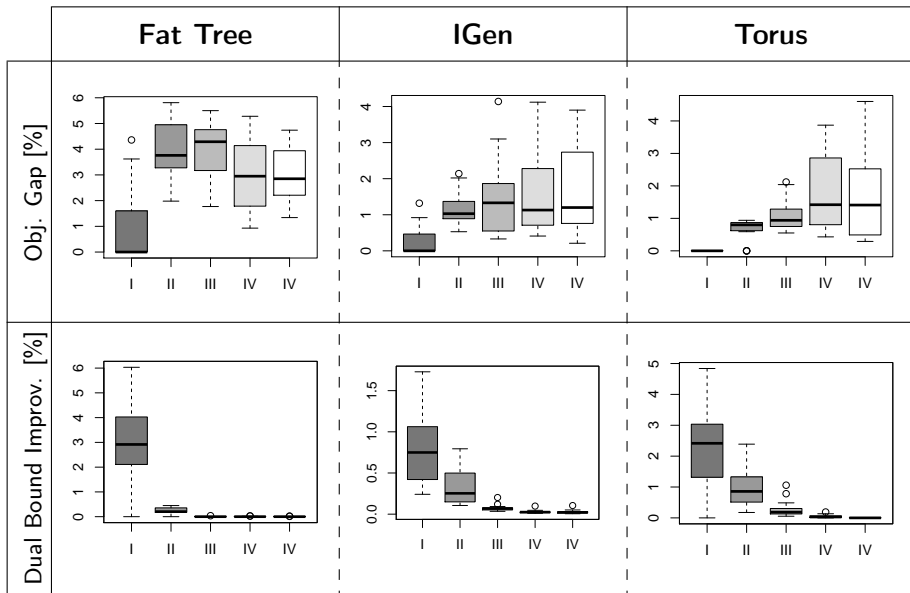
- *all* algorithms (except MCF-IP) are implemented in C/C++
- VirtuCast uses SCIP [1], many different parameters to consider
 - separation
 - branching
 - heuristics
 - **separation procedure: nested cuts, creep flow, cyclic generation...**
- MCF-IP is implemented using GMPL + CPLEX

Objective

Solve instances within reasonable time: 1 hour runtime limit

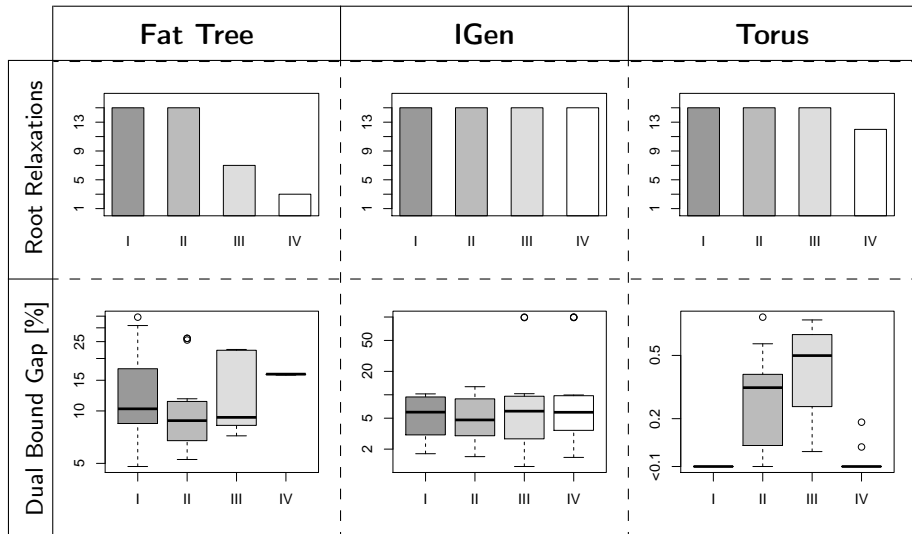
VirtuCast + LP-based Heuristics

VirtuCast + LP-based Heuristics




MCF-IP

MCF-IP: Performance

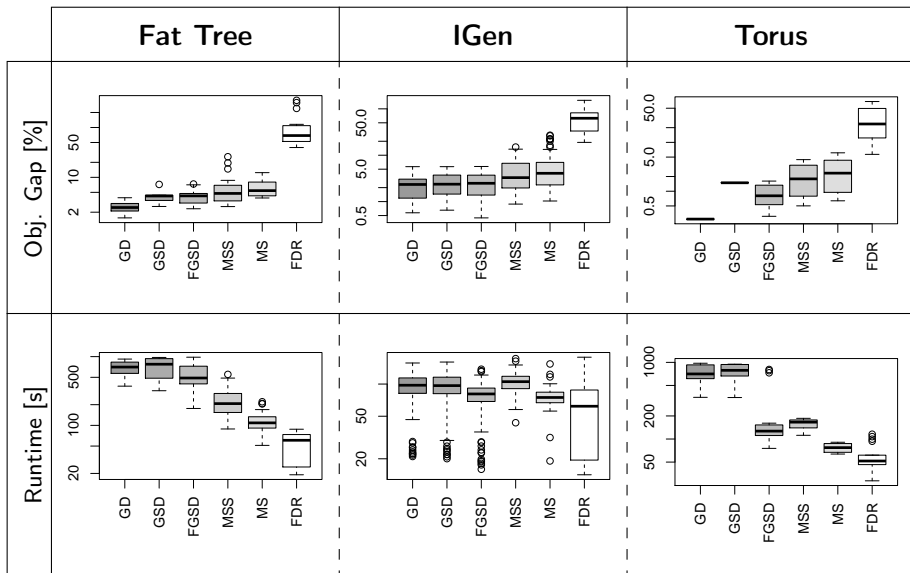


LP-based Heuristics

LP-based Heuristics: Efficacy



LP-based Heuristics: Performance on graph size V



Conclusion

Publications

Matthias Rost, Stefan Schmid: OPODIS 2013 & arXiv [15, 14]
Matthias Rost (Adv. Stefan Schmid): M.Sc. Thesis [13]

Concise definition of CVSAP

Inapproximability

Approximations

- NVSTP
- VSTP
- VSAP

Exact Algorithms

- multi-commodity flow
- single-commodity flow
→ VirtuCast

Heuristics

- FlowDecoRound
- MultipleShots
- GreedyDiving

Extensive explorative Computational Evaluation

Related Work

Molnar: Constrained Spanning Tree Problems [9]

- Shows that optimal solution is a 'spanning hierarchy' and not a DAG.

Oliveira et. al: Flow Streaming Cache Placement Problem [11]

- Consider a weaker variant of multicasting CVSAP without bandwidth
- Use a (faulty) MIP to define the problem
- Give weak approximation algorithm

Shi: Scalability in Overlay Multicasting [16]

- Provided heuristic and showed improvement in scalability.

Future Work

Model Extensions


- prize-collecting variants
- concurrent multicast / aggregation sessions
- 'extend' MIP formulation for weaker variants

Speeding-up Separation / Public Service Announcement

- Koch et al. [7] stated that using Hao-Orlin the computation could be sped up.
- Cronholm et al. show that this is not really the case, but derive an adaptation [4]:

For single node, all separations can be computed in $\mathcal{O}(nm \log(n^2/m))$

Thanks



References I

- [1] T. Achterberg.
SCIP: solving constraint integer programs.
Mathematical Programming Computation, 1(1):1–41, 2009.
- [2] P. Costa, A. Donnelly, A. Rowstron, and G. O. Shea.
Camdoop: Exploiting In-network Aggregation for Big Data Applications.
In Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2012.
- [3] C. Cranor, T. Johnson, O. Spataschek, and V. Shkapenyuk.
Gigascop: A Stream Database for Network Applications.
In Proc. ACM SIGMOD International Conference on Management of Data, pages 647–651, 2003.
- [4] W. Cronholm, F. Ajili, and S. Panagiotidi.
On the minimal steiner tree subproblem and its application in branch-and-price.
In Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, pages 125–139. Springer, 2005.
- [5] M. Ding, X. Cheng, and G. Xue.
Aggregation tree construction in sensor networks.
In Vehicular Technology Conference, 2003. VTC 2003-Fall. 2003 IEEE 58th, volume 4, pages 2168–2172. IEEE, 2003.

References II

- [6] C. Hermsmeyer, E. Hernandez-Valencia, D. Stoll, and O. Tamm.
Ethernet aggregation and core network models for efficient and reliable iptv services.
Bell Labs Technical Journal, 12(1):57–76, 2007.
- [7] T. Koch and A. Martin.
Solving steiner tree problems in graphs to optimality.
Networks, 32(3):207–232, 1998.
- [8] B. Krishnamachari, D. Estrin, and S. Wicker.
Modelling data-centric routing in wireless sensor networks.
In *IEEE infocom*, volume 2, pages 39–44, 2002.
- [9] M. Molnár.
Hierarchies to Solve Constrained Connected Spanning Problems.
Technical Report Irimm-00619806, University Montpellier 2, LIRMM, 2011.
- [10] S. Narayana, W. Jiang, J. Rexford, and M. Chiang.
Joint Server Selection and Routing for Geo-Replicated Services.
In *Proc. Workshop on Distributed Cloud Computing (DCC)*, 2013.

References III

- [11] C. Oliveira and P. Pardalos.
Streaming cache placement.
In *Mathematical Aspects of Network Routing Optimization*, Springer Optimization and Its Applications, pages 117–133. Springer New York, 2011.
- [12] R. Ravi, M. V. Marathe, S. Ravi, D. J. Rosenkrantz, and H. B. Hunt III.
Approximation algorithms for degree-constrained minimum-cost network-design problems.
Algorithmica, 31(1):58–78, 2001.
- [13] M. Rost.
Optimal Virtualized In-Network Processing with Applications to Aggregation and Multicast, 2014.
- [14] M. Rost and S. Schmid.
The Constrained Virtual Steiner Arborescence Problem: Formal Definition, Single-Commodity Integer Programming Formulation and Computational Evaluation.
Technical report, arXiv, 2013.
- [15] M. Rost and S. Schmid.
Virtucast: Multicast and aggregation with in-network processing.
In R. Baldoni, N. Nisse, and M. Steen, editors, *Principles of Distributed Systems*, volume 8304 of *Lecture Notes in Computer Science*, pages 221–235. Springer International Publishing, 2013.

References IV

[16] S. Shi.

A proposal for a scalable internet multicast architecture.


In *Washington Universtiy*, 2001.

Approximation of NVSTP via DNSTP

Approximation of NVSTP via DNSTP

NVSTP

- undirected version
- no edge capacities
- Steiner nodes have capacities
- connect terminals using Steiner nodes to root



Approximation of NVSTP via DNSTP: [12]

Definition (Degree-Constrained Node Weighted Steiner Tree Problem [12])

Given: Undirected network $G = (V_G, E_G, c_E, c_V, u_V)$ with edge costs $c_E : E_G \rightarrow \mathbb{R}_{\geq 0}$ ^a, node costs $c_V : V_G \rightarrow \mathbb{R}_{\geq 0}$, and a degree bound function $u_V : V_G \rightarrow \mathbb{N}_{\geq 2}$ and set of terminals $T \subset V_G$.

Task: Find a Steiner tree $\mathcal{T} \subseteq E_G$ connecting all terminals T , such that for each node v that is contained in \mathcal{T} the degree bound is not violated, i.e. that $\delta_{\mathcal{T}}(v) \leq u_V(v)$ holds, minimizing the cost $C_{\text{DNSTP}}(\mathcal{T}) = \sum_{e \in \mathcal{T}} c_E(e) + \sum_{v \in \mathcal{T}} c_V(v)$.

^aThe original definition and the corresponding theorem only considers the node weighted case.

Approximation of NVSTP via DNSTP

Theorem (Logarithmic bi-criteria approximation for DNSTP [12])


There exists a polynomial-time algorithm that returns a solution where node capacities are (individually) violated at most by a factor $\mathcal{O}(\log |T|)$ and of cost within a factor of $\mathcal{O}(\log |T|)$ the optimum solution.

Differences of NVSTP w.r.t. DNSTP

- NVSTP constructs a tree, i.e. terminals have degree 1.
- NVSTP may use arbitrary paths to connect nodes.
- Not all nodes may be used as Steiner nodes.

NVSTP via DNSTP: Construction

- bipartite mesh connecting any terminal to any Steiner node
- clique between all Steiner nodes and the root
- all edges have cost of respective shortest path




Checklist

- NVSTP constructs a tree, i.e. terminals have degree 1.
- ~~NVSTP may use arbitrary paths to connect nodes.~~
- ~~Not all nodes may be used as Steiner nodes.~~

NVSTP via DNSTP: Construction

- bipartite mesh connecting any terminal to any Steiner node
- clique between all Steiner nodes and the root
- all edges have cost of respective shortest path




Checklist

- NVSTP constructs a tree, i.e. terminals have degree 1.
- ~~NVSTP may use arbitrary paths to connect nodes.~~
- ~~Not all nodes may be used as Steiner nodes.~~

Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above



Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution



Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution
- 3 consider bipartite subgraph of terminals with degree > 1



Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution
- 3 consider bipartite subgraph of terminals with degree > 1
- 4 compute maximum matching with size = number of terminals



Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution
- 3 consider bipartite subgraph of terminals with degree > 1
- 4 compute maximum matching with size = number of terminals
- 5 perform 'leafify' operation on terminals



Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution
- 3 consider bipartite subgraph of terminals with degree > 1
- 4 compute maximum matching with size = number of terminals
- 5 perform 'leafify' operation on terminals



Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution
- 3 consider bipartite subgraph of terminals with degree > 1
- 4 compute maximum matching with size = number of terminals
- 5 perform 'leafify' operation on terminals



Outline of Bicriteria Approximation for NVSTP

Algorithm


- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution
- 3 consider bipartite subgraph of terminals with degree > 1
- 4 compute maximum matching with size = number of terminals
- 5 perform 'leafify' operation on terminals



Outline of Bicriteria Approximation for NVSTP

Algorithm

- 1 construct graph as described above
- 2 use Approximation by Ravi et al. to obtain DNSTP solution
- 3 consider bipartite subgraph of terminals with degree > 1
- 4 compute maximum matching with size = number of terminals
- 5 perform 'leafify' operation on terminals



Theorem

- 1 Cost of introduced edges is bounded by triangle equation
- 2 Degree of non-terminals in matching is increased by 1
- 3 $\mathcal{O}(\log |T|, \log |T|)$ for DNSTP $\Rightarrow \mathcal{O}(\log |T|, \log |T|)$ for NVSTP