Approximate Graph Embeddings in the Cloud

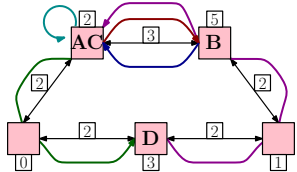

Highlights of Algorithms 2018

Matthias Rost
Technische Universität Berlin, Internet Network Architectures

Stefan Schmid
Universität Wien, Communication Technologies

Cloud Providers Offer Data Center Resources

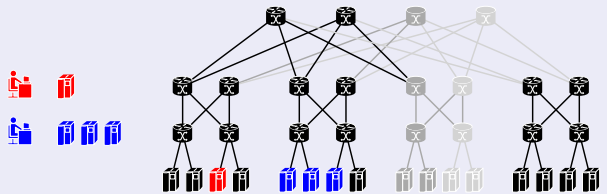

Customers Cloud Data Center (Amazon, Google, ...)

Cloud Providers Offer Data Center Resources

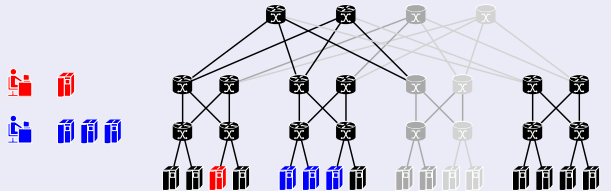

Customers Cloud Data Center (Amazon, Google, ...)

'Classic' Cloud Computing

- Customer specifies number and 'size' of Virtual Machines
- Communication between VMs not modeled

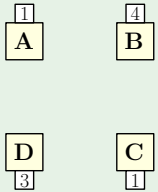

Cloud Providers Offer Data Center Resources

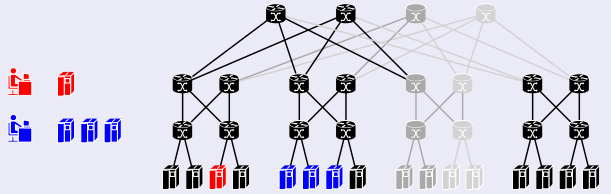

Customers Cloud Data Center (Amazon, Google, ...)

'Classic' Cloud Computing

- Customer specifies number and 'size' of Virtual Machines
- Communication between VMs not modeled

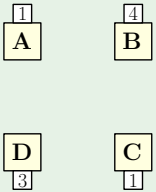

Goal: Virtual Networks (since ≈ 2006)

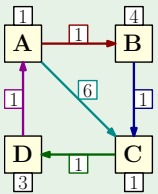

- Additionally: communication requirements given

‘Classic’ Cloud Computing

- Customer specifies number and ‘size’ of Virtual Machines
- Communication between VMs not modeled

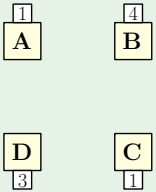

Goal: Virtual Networks (since ≈ 2006)

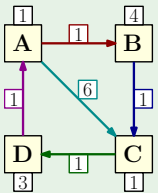

- Additionally: communication requirements given

The Virtual Network Embedding Problem (VNEP)

- Map virtual nodes to substrate nodes
- Map virtual edges to paths in the substrate
- Respecting capacities & mapping restrictions

Virtual Network

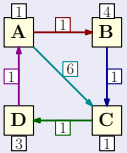

Substrate (Physical Network)

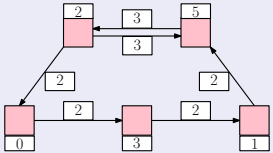

'Classic' Cloud Computing

- Customer specifies number and 'size' of Virtual Machines
- Communication between VMs not modeled

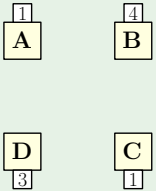

Goal: Virtual Networks (since ≈ 2006)

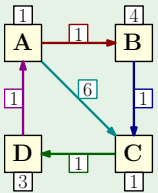

- Additionally: communication requirements given

The Virtual Network Embedding Problem (VNEP)

- Map virtual nodes to substrate nodes
- Map virtual edges to paths in the substrate
- Respecting capacities & mapping restrictions

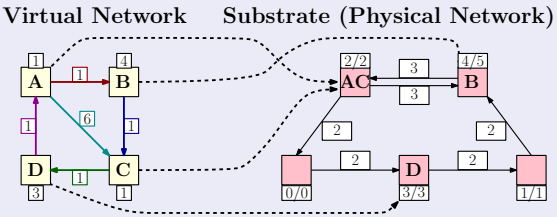

'Classic' Cloud Computing

- Customer specifies number and 'size' of Virtual Machines
- Communication between VMs not modeled

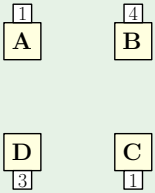

Goal: Virtual Networks (since ≈ 2006)

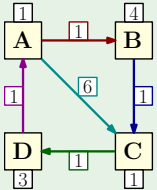

- Additionally: communication requirements given

The Virtual Network Embedding Problem (VNEP)

- Map virtual nodes to substrate nodes
- Map virtual edges to paths in the substrate
- Respecting capacities & mapping restrictions

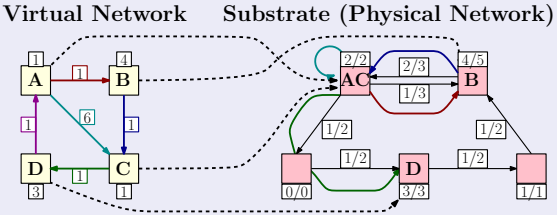

‘Classic’ Cloud Computing

- Customer specifies number and ‘size’ of Virtual Machines
- Communication between VMs not modeled

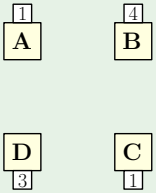

Goal: Virtual Networks (since ≈ 2006)

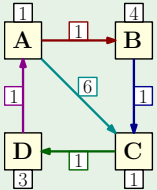

- Additionally: communication requirements given

The Virtual Network Embedding Problem (VNEP)

- Map virtual nodes to substrate nodes
- Map virtual edges to paths in the substrate
- Respecting capacities & mapping restrictions

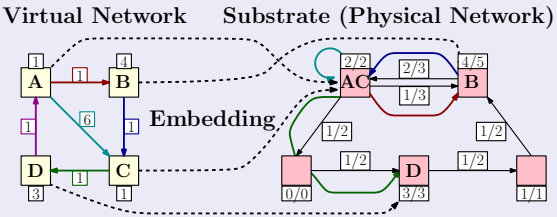

The Virtual Network Embedding Problem (VNEP)

- Map virtual nodes to substrate nodes
- Map virtual edges to paths in the substrate
- Respecting capacities & mapping restrictions

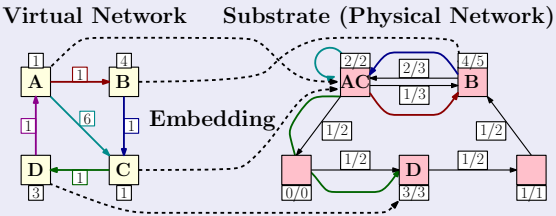

Related Work

- VNEP (and related problems) studied intensively in the networking community: > 100 papers.
- VNEP is related to classical problems as, e.g., subgraph isomorphism, **but different** ...
- No approximations known for arbitrary virtual networks graphs.

The Virtual Network Embedding Problem (VNEP)

- Map virtual nodes to substrate nodes
- Map virtual edges to paths in the substrate
- Respecting capacities & mapping restrictions

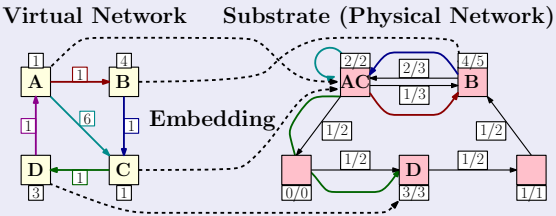

Related Work

- VNEP (and related problems) studied intensively in the networking community: > 100 papers.
- VNEP is related to classical problems as, e.g., subgraph isomorphism, **but different ...**
- No approximations known for arbitrary virtual networks graphs.

The Virtual Network Embedding Problem (VNEP)

- Map virtual nodes to substrate nodes
- Map virtual edges to paths in the substrate
- Respecting capacities & mapping restrictions

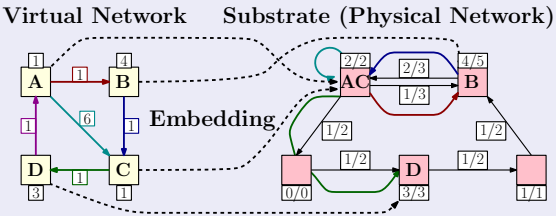

Related Work

- VNEP (and related problems) studied intensively in the networking community: > 100 papers.
- VNEP is related to classical problems as, e.g., subgraph isomorphism, **but different ...**
- **No approximations known for arbitrary virtual networks graphs.**

Related Work

- VNEP (and related problems) studied intensively in the networking community: > 100 papers.
- VNEP is related to classical problems as, e.g., subgraph isomorphism, **but different ...**
- **No approximations known for arbitrary virtual networks graphs.**

Focus: Offline Variant

Setting	Multiple Virtual Network requests are given
Objectives	Maximize profit (admission control) or minimize 'cost' *s.t. capacity constraints*.

Related Work

- VNEP (and related problems) studied intensively in the networking community: > 100 papers.
- VNEP is related to classical problems as, e.g., subgraph isomorphism, **but different ...**
- **No approximations known for arbitrary virtual networks graphs.**

Focus: Offline Variant

Setting Multiple Virtual Network requests are given

Objectives Maximize profit (admission control) or minimize 'cost' *s.t. capacity constraints*.

Approach: Randomized Rounding à la Raghavan & Thompson

- Compute opt. 'convex combinations' of mappings: $\mathcal{D}_r = \{(\underbrace{f_r^k}_{\text{weight } \geq 0} , \underbrace{m_r^k}_{\text{mapping}})\}_k$ for request r
- Probabilistically select mapping m_r^k according to weight f_r^k for each request r
  - Yields: *approximate solutions of bounded resource augmentations with high probability*

Approach: Randomized Rounding à la Raghavan & Thompson

- Compute opt. 'convex combinations' of mappings: $\mathcal{D}_r = \{(\underbrace{f_r^k}_{\text{weight } \geq 0} , \underbrace{m_r^k}_{\text{mapping}})\}_k$ for request r
- Probabilistically select mapping m_r^k according to weight f_r^k for each request r
  - Yields: *approximate solutions of bounded resource augmentations with high probability*

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)

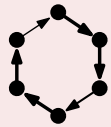

Substrate

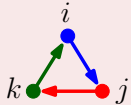

Request

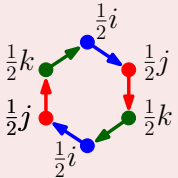

Classic LP Solution

- Observation: Need to fix *confluence targets* (here: node k) a priori.

Approach: Randomized Rounding à la Raghavan & Thompson

- Compute opt. 'convex combinations' of mappings: $\mathcal{D}_r = \{(\underbrace{f_r^k}_{\text{weight } \geq 0} , \underbrace{m_r^k}_{\text{mapping}})\}_k$ for request r
- Probabilistically select mapping m_r^k according to weight f_r^k for each request r
  - Yields: *approximate solutions of bounded resource augmentations with high probability*

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)

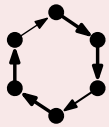

Substrate

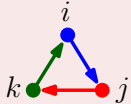

Request

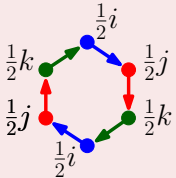

Classic LP Solution

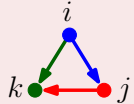

Extraction Order

- Observation: Need to fix *confluence targets* (here: node k) a priori.

Approach: Randomized Rounding à la Raghavan & Thompson

- Compute opt. 'convex combinations' of mappings: $\mathcal{D}_r = \{(\underbrace{f_r^k}_{\text{weight } \geq 0} , \underbrace{m_r^k}_{\text{mapping}})\}_k$ for request r
- Probabilistically select mapping m_r^k according to weight f_r^k for each request r
  - Yields: *approximate solutions of bounded resource augmentations with high probability*

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)

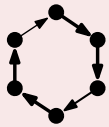

Substrate

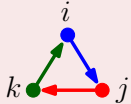

Request

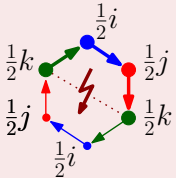

Classic LP Solution

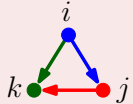

Extraction Order

- Observation: Need to fix *confluence targets* (here: node k) a priori.

Approach: Randomized Rounding à la Raghavan & Thompson

- Compute opt. 'convex combinations' of mappings: $\mathcal{D}_r = \{(\underbrace{f_r^k}_{\text{weight } \geq 0} , \underbrace{m_r^k}_{\text{mapping}})\}_k$ for request r
- Probabilistically select mapping m_r^k according to weight f_r^k for each request r
  - Yields: *approximate solutions of bounded resource augmentations with high probability*

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)

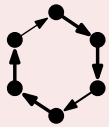

Substrate

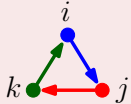

Request

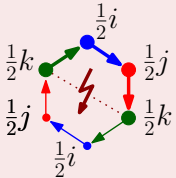

Classic LP Solution

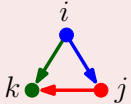

Extraction Order

- **Observation: Need to fix *confluence targets* (here: node k) a priori.**

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)
- **Observation: Need to fix *confluence targets* a priori.**

Main Contributions

- LP Formulations for **cactus request graphs** $\rightarrow$ **first approximation** algorithms[a]
- Derivation of **heuristics** & extensive **computational evaluation**[a]
- Extension to **arbitrary** virtual network **topologies** $\rightarrow$ **FPT-approximations**[b]
- **FPT required**: no poly.-time algorithms for computing valid mappings for general graphs[c]

---

[a] Matthias Rost and Stefan Schmid. "Virtual Network Embedding Approximations: Leveraging Randomized Rounding". In: *Proc. IFIP Networking*. 2018.

[b] Matthias Rost and Stefan Schmid. *(FPT-)Approximation Algorithms for the Virtual Network Embedding Problem*. Tech. rep. Mar. 2018. URL: http://arxiv.org/abs/1803.04452.

[c] Matthias Rost and Stefan Schmid. "Charting the Complexity Landscape of Virtual Network Embeddings". In: *Proc. IFIP Networking*. 2018.

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)
- **Observation: Need to fix *confluence targets* a priori.**

Main Contributions

- LP Formulations for **cactus request graphs** $\rightarrow$ **first approximation** algorithms[a]
- Derivation of **heuristics** & extensive **computational evaluation**[a]
- Extension to **arbitrary** virtual network **topologies** $\rightarrow$ **FPT-approximations**[b]
- **FPT required**: no poly.-time algorithms for computing valid mappings for general graphs[c]

[a] Matthias Rost and Stefan Schmid. "Virtual Network Embedding Approximations: Leveraging Randomized Rounding". In: *Proc. IFIP Networking*. 2018.

[b] Matthias Rost and Stefan Schmid. *(FPT-)Approximation Algorithms for the Virtual Network Embedding Problem*. Tech. rep. Mar. 2018. URL: http://arxiv.org/abs/1803.04452.

[c] Matthias Rost and Stefan Schmid. "Charting the Complexity Landscape of Virtual Network Embeddings". In: *Proc. IFIP Networking*. 2018.

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)
- **Observation: Need to fix *confluence targets* a priori.**

Main Contributions

- LP Formulations for **cactus request graphs** $\rightarrow$ **first approximation** algorithms[a]
- Derivation of **heuristics** & extensive **computational evaluation**[a]
- Extension to **arbitrary** virtual network **topologies** $\rightarrow$ **FPT-approximations**[b]
- **FPT required**: no poly.-time algorithms for computing valid mappings for general graphs[c]

[a] Matthias Rost and Stefan Schmid. "Virtual Network Embedding Approximations: Leveraging Randomized Rounding". In: *Proc. IFIP Networking*. 2018.

[b] Matthias Rost and Stefan Schmid. *(FPT-)Approximation Algorithms for the Virtual Network Embedding Problem*. Tech. rep. Mar. 2018. URL: http://arxiv.org/abs/1803.04452.

[c] Matthias Rost and Stefan Schmid. "Charting the Complexity Landscape of Virtual Network Embeddings". In: *Proc. IFIP Networking*. 2018.

Main Challenge: Computing (Convex Combinations) of Valid Mappings

- Classic LP Formulation yields no meaningful solutions ($\rightarrow$ unbounded integrality gap)
- **Observation: Need to fix *confluence targets* a priori.**

Main Contributions

- LP Formulations for **cactus request graphs** $\rightarrow$ **first approximation** algorithms[a]
- Derivation of **heuristics** & extensive **computational evaluation**[a]
- Extension to **arbitrary** virtual network **topologies** $\rightarrow$ **FPT-approximations**[b]
- **FPT required**: no poly.-time algorithms for computing valid mappings for general graphs[c]

---

[a] Matthias Rost and Stefan Schmid. "Virtual Network Embedding Approximations: Leveraging Randomized Rounding". In: *Proc. IFIP Networking*. 2018.

[b] Matthias Rost and Stefan Schmid. *(FPT-)Approximation Algorithms for the Virtual Network Embedding Problem*. Tech. rep. Mar. 2018. URL: http://arxiv.org/abs/1803.04452.

[c] Matthias Rost and Stefan Schmid. "Charting the Complexity Landscape of Virtual Network Embeddings". In: *Proc. IFIP Networking*. 2018.

Main Contributions

- LP Formulations for **cactus request graphs** → **first approximation** algorithms[a]
- Derivation of **heuristics** & extensive **computational evaluation**[a]
- Extension to **arbitrary** virtual network **topologies** → **FPT-approximations**[b]
- **FPT required**: no poly.-time algorithms for computing valid mappings for general graphs[c]

---

[a] Matthias Rost and Stefan Schmid. "Virtual Network Embedding Approximations: Leveraging Randomized Rounding". In: *Proc. IFIP Networking*. 2018.

[b] Matthias Rost and Stefan Schmid. *(FPT-)Approximation Algorithms for the Virtual Network Embedding Problem*. Tech. rep. Mar. 2018. URL: http://arxiv.org/abs/1803.04452.

[c] Matthias Rost and Stefan Schmid. "Charting the Complexity Landscape of Virtual Network Embeddings". In: *Proc. IFIP Networking*. 2018.

Thanks for your attention!