Charting the Complexity Landscape of Virtual Network Embeddings

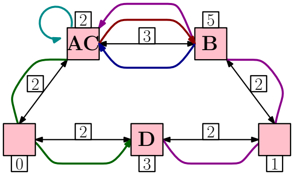

IFIP Networking 2018

Matthias Rost
Technische Universität Berlin, Internet Network Architectures

Stefan Schmid
Universität Wien, Communication Technologies

Introduction: Virtual Network Embeddings and Their Complexity

Virtualization: Resource Allocation Opportunities

'Classic' Cloud Computing

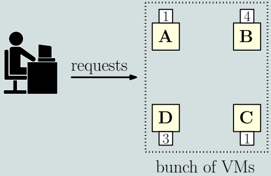

- User requests virtual machines
- No guarantee on network performance

Goal: Virtual Networks (since ≈ 2006)

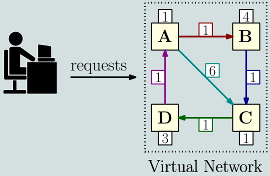

- Communication requirements given
- Network performance will be guaranteed

Introduction: Virtual Network Embeddings and Their Complexity

Virtualization: Resource Allocation Opportunities

'Classic' Cloud Computing

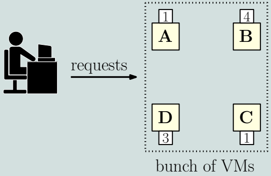

Goal: Virtual Networks (since ≈ 2006)

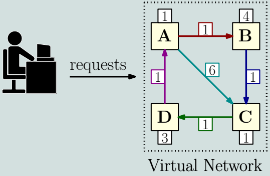

Novel Service Abstractions

- Virtual Networks – overlays (≈ 2006)
- Virtual Clusters – batch processing (≈ 2011)
- Service Chain – stitch functions (≈ 2013)

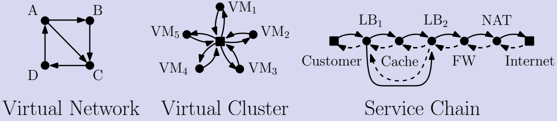

Introduction: Virtual Network Embeddings and Their Complexity

Virtualization: Resource Allocation Opportunities

'Classic' Cloud Computing

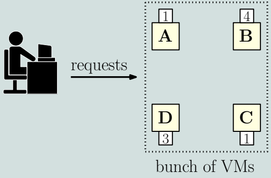

Goal: Virtual Networks (since ≈ 2006)

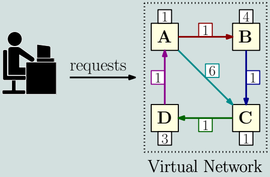

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

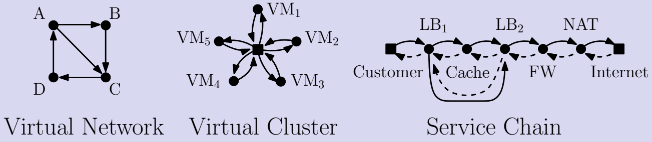

Introduction: Virtual Network Embeddings and Their Complexity

Virtualization: Resource Allocation Opportunities

'Classic' Cloud Computing

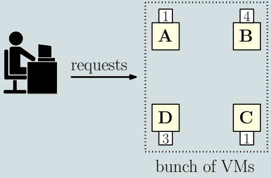

Goal: Virtual Networks (since ≈ 2006)

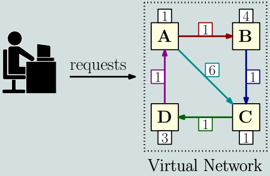

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

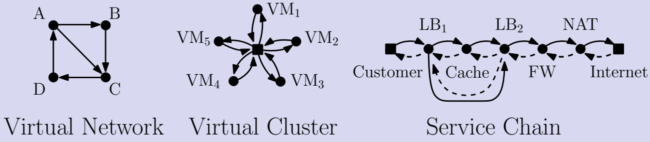

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

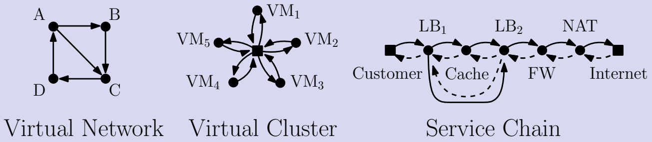

Virtual Network Embeddings at a Glance

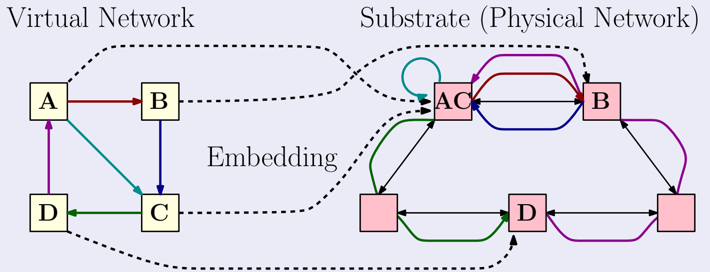

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

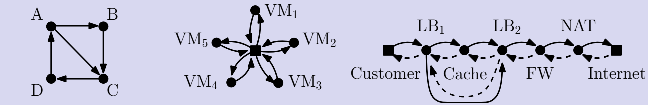

Virtual Network   Virtual Cluster   Service Chain

Virtual Network Embeddings at a Glance

Virtual Network   Substrate (Physical Network)

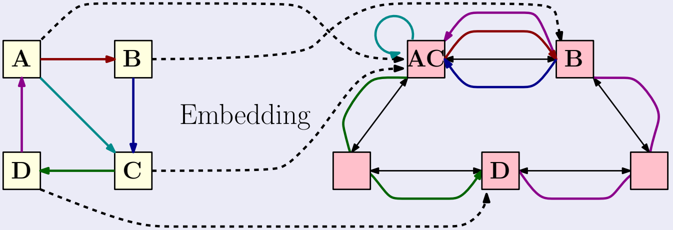

Embedding Restrictions

Capacity

- **V** Node
- **E** Edge

Additional

- **N** Node placement
- **R** Routing
- **L** Latencies

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

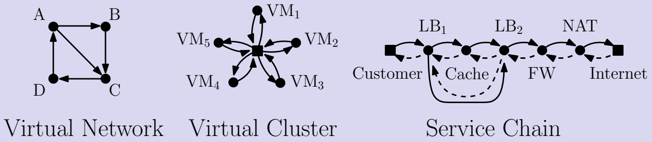

Virtual Network Embeddings at a Glance – Node Capacities

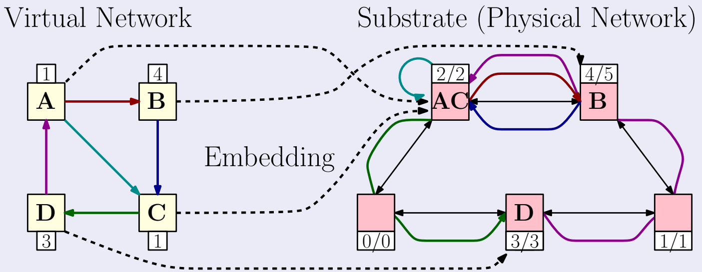

Embedding Restrictions

Capacity

- **V** Node
- **E** Edge

Additional

- **N** Node placement
- **R** Routing
- **L** Latencies

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

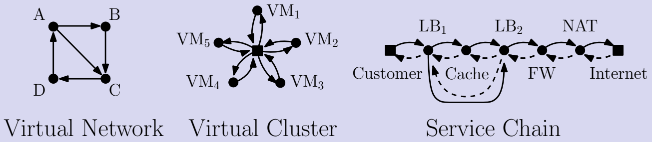

Virtual Network Embeddings at a Glance – Edge Capacities

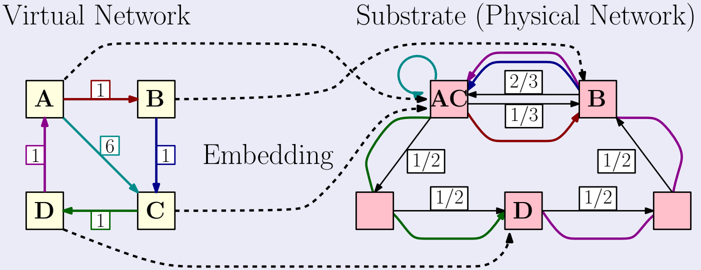

Embedding Restrictions

Capacity

- **V** Node
- **E** Edge

Additional

- **N** Node placement
- **R** Routing
- **L** Latencies

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

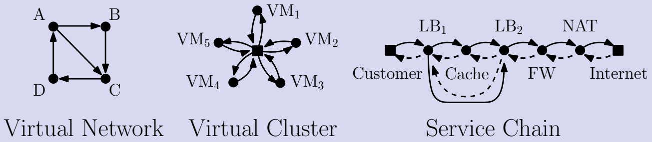

Virtual Network Embeddings at a Glance – Node Placement

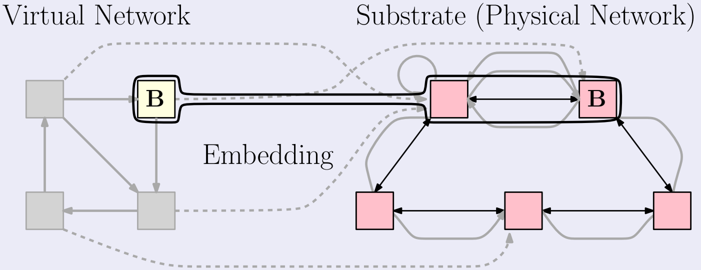

Embedding Restrictions

Capacity

V Node

E Edge

Additional

N Node placement

R Routing

L Latencies

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

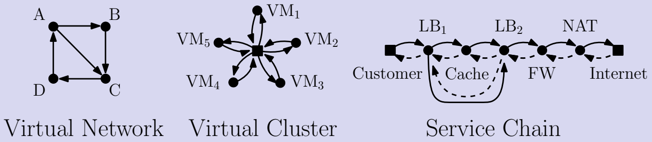

Virtual Network Embeddings at a Glance – Routing Restrictions

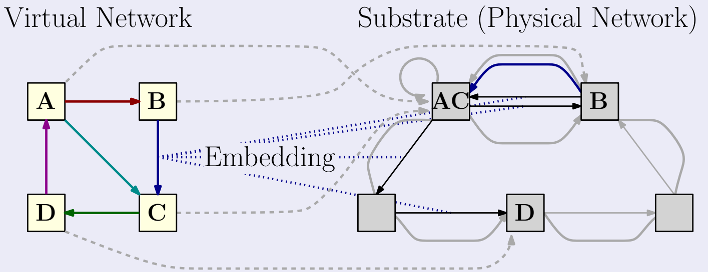

Embedding Restrictions

Capacity

V Node

E Edge

Additional

N Node placement

R Routing

L Latencies

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

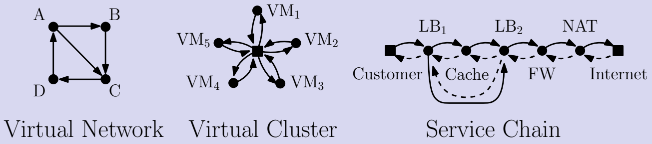

Virtual Network Embeddings at a Glance – Latencies

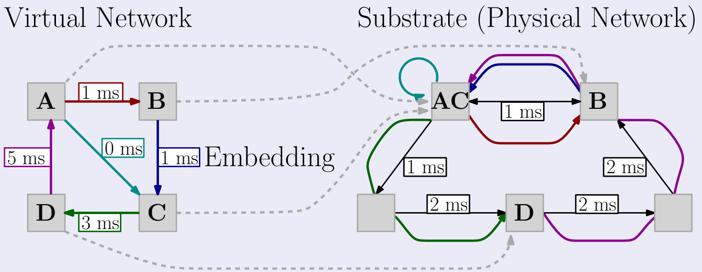

Embedding Restrictions

Capacity

V Node

E Edge

Additional

N Node placement

R Routing

L Latencies

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

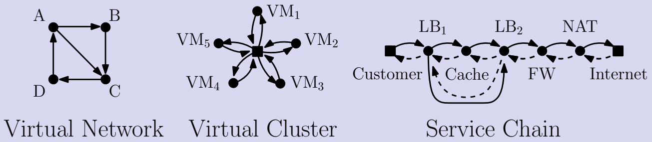

Taxonomy of VNEP Variants

VNEP $\langle\, \mathbf{C} \,|\, \mathbf{A} \,\rangle$: capacity restrictions $\mathbf{C}$ & additional restrictions $\mathbf{A}$

Embedding Restrictions

Capacity

- **V** Node
- **E** Edge

Additional

- **N** Node placement
- **R** Routing
- **L** Latencies

Introduction: Virtual Network Embeddings and Their Complexity

Novel Service Abstractions

- Virtual Network Embedding Problem
- Virtual Clusters Embedding Problem
- Service Chain Embedding Problem

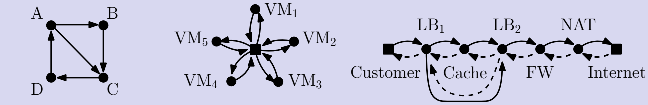

Virtual Network Virtual Cluster Service Chain

Taxonomy of VNEP Variants

VNEP $\langle \mathbf{C} \mid \mathbf{A} \rangle$: capacity restrictions $\mathbf{C}$ & additional restrictions $\mathbf{A}$

For example...

VNEP $\langle \mathbf{VE} \mid \text{-} \rangle$: node and edge capacities

VNEP $\langle \mathbf{V} \mid \mathbf{R} \rangle$: node capacities and routing restrictions

VNEP $\langle \text{-} \mid \mathbf{NL} \rangle$: node placement and latency restrictions

Embedding Restrictions

Capacity

- **V** Node
- **E** Edge

Additional

- **N** Node placement
- **R** Routing
- **L** Latencies

Introduction: Virtual Network Embeddings and Their Complexity

Related Work

Theoretical Results: **Few**

Andersen [2002] Considered $\langle\, \mathbf{VE} \,|\, \textbf{-} \,\rangle$ and argued for $\mathcal{NP}$-hardness

Amaldi et al. [2016] Considered $\langle\, \mathbf{VE} \,|\, \mathbf{N} \,\rangle$ under *profit objective*, proved $\mathcal{NP}$-hardness and derived inapproximability result.

Practical Results: **Many**

Generally More than 100 papers on VNEP alone, for example ...

Chowdhury et al. [2009] Developed algorithms for variant $\langle\, \mathbf{VE} \,|\, \mathbf{N} \,\rangle$ and hoped to obtain *approximations*.

VNEP is of **crucial** importance, yet is **hardly understood**!

Introduction: Virtual Network Embeddings and Their Complexity

VNEP is of **crucial** importance, yet is **hardly understood**!

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Practical Implications (unless $\mathcal{P} = \mathcal{NP}$)

There cannot exist a polynomial-time algorithm . . .

1. always yielding a solution to the VNEP under any of the above restrictions,
2. which does not violate capacities or latencies by less than some amount,
3. even when virtual networks are acyclic, planar, and degree-bounded.

Introduction: Virtual Network Embeddings and Their Complexity

VNEP is of **crucial** importance, yet is **hardly understood**!

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\text{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\text{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\text{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Practical Implications (unless $\mathcal{P} = \mathcal{NP}$)

There cannot exist a polynomial-time algorithm . . .

1. always yielding a solution to the VNEP under any of the above restrictions,
2. which does not violate capacities or latencies by less than some amount,
3. even when virtual networks are acyclic, planar, and degree-bounded.

Introduction: Virtual Network Embeddings and Their Complexity

VNEP is of **crucial** importance, yet is **hardly understood**!

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Practical Implications (unless $\mathcal{P} = \mathcal{NP}$)

There cannot exist a polynomial-time algorithm . . .

1. always yielding a solution to the VNEP under any of the above restrictions,
2. which does not violate capacities or latencies by less than some amount,
3. even when virtual networks are acyclic, planar, and degree-bounded.

Introduction: Virtual Network Embeddings and Their Complexity

VNEP is of **crucial** importance, yet is **hardly understood**!

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Practical Implications (unless $\mathcal{P} = \mathcal{NP}$)

There cannot exist a polynomial-time algorithm . . .

1. always yielding a solution to the VNEP under any of the above restrictions,
2. which does not violate capacities or latencies by less than some amount,
3. even when virtual networks are acyclic, planar, and degree-bounded.

Definition of the Virtual Network Embedding Problem

Definition of the Virtual Network Embedding Problem

Input

Substrate	$G_S = (V_S, E_S)$
Request	$G_r = (V_r, E_r)$
Restrictions	...

Feasible Embedding

A *feasible* embedding is a mapping of G_r to G_S respecting **all** restrictions.

Feasible embedding meets *all* requirements

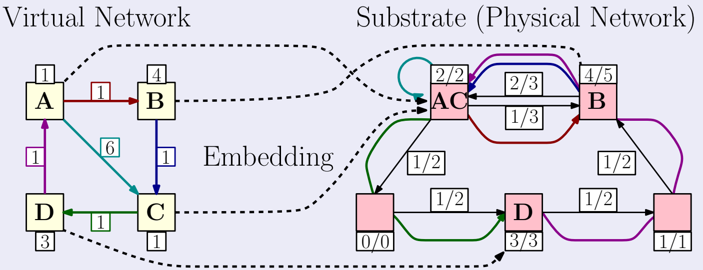

Definition of the Virtual Network Embedding Problem

Input

Substrate $G_S = (V_S, E_S)$
Request $G_r = (V_r, E_r)$
Restrictions ...

Feasible Embedding

A *feasible* embedding is a mapping of G_r to G_S respecting **all** restrictions.

Feasible embedding meets *all* requirements

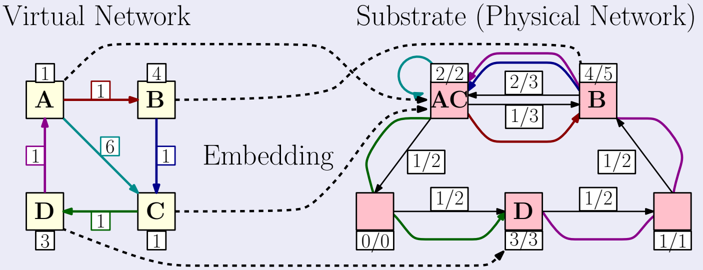

Virtual Network Embedding Problem (Decision Variant)

Decide whether a feasible embedding of request G_r on substrate G_S exists.
Output: Yes / No.

Methodology

Reminder: 3-SAT and $\mathcal{NP}$-Completeness

3-SAT-Formula ϕ

$\phi = \bigwedge_{\mathcal{C}_i \in \mathcal{C}_\phi} \mathcal{C}_i$ with $\mathcal{C}_i \in \mathcal{C}_\phi$ being disjunctions of at most 3 (possible negated) literals.

Example 3-SAT formula ϕ over literals $\mathcal{L}_\phi = \{x_1, x_2, x_3, x_4\}$

$$\phi = \underbrace{(x_1 \vee x_2 \vee x_3)}_{\mathcal{C}_1} \wedge \underbrace{(\bar{x}_1 \vee x_2 \vee x_4)}_{\mathcal{C}_2} \wedge \underbrace{(x_2 \vee \bar{x}_3 \vee x_4)}_{\mathcal{C}_3}$$

Reminder: 3-SAT and $\mathcal{NP}$-Completeness

3-SAT-Formula ϕ

$\phi = \bigwedge_{\mathcal{C}_i \in \mathcal{C}_\phi} \mathcal{C}_i$ with $\mathcal{C}_i \in \mathcal{C}_\phi$ being disjunctions of at most 3 (possible negated) literals.

Example 3-SAT formula ϕ over literals $\mathcal{L}_\phi = \{x_1, x_2, x_3, x_4\}$

$$\phi = \underbrace{(x_1 \vee x_2 \vee x_3)}_{\mathcal{C}_1} \wedge \underbrace{(\bar{x}_1 \vee x_2 \vee x_4)}_{\mathcal{C}_2} \wedge \underbrace{(x_2 \vee \bar{x}_3 \vee x_4)}_{\mathcal{C}_3}$$

Definition of 3-SAT

Decide whether satisfying assignment $a : \mathcal{L}_\phi \to \{\mathrm{F}, \mathrm{T}\}$ exists for formula ϕ. Output: Yes/No.

Reminder: 3-SAT and $\mathcal{NP}$-Completeness

3-SAT-Formula ϕ

$\phi = \bigwedge_{\mathcal{C}_i \in \mathcal{C}_\phi} \mathcal{C}_i$ with $\mathcal{C}_i \in \mathcal{C}_\phi$ being disjunctions of at most 3 (possible negated) literals.

Example 3-SAT formula ϕ over literals $\mathcal{L}_\phi = \{x_1, x_2, x_3, x_4\}$

$$\phi = \underbrace{(x_1 \vee x_2 \vee x_3)}_{\mathcal{C}_1} \wedge \underbrace{(\bar{x}_1 \vee x_2 \vee x_4)}_{\mathcal{C}_2} \wedge \underbrace{(x_2 \vee \bar{x}_3 \vee x_4)}_{\mathcal{C}_3}$$

Definition of 3-SAT

Decide whether satisfying assignment $a : \mathcal{L}_\phi \to \{\mathrm{F}, \mathrm{T}\}$ exists for formula ϕ. Output: Yes/No.

Theorem: Karp [1972]

3-SAT is $\mathcal{NP}$-complete.

Reminder: 3-SAT and $\mathcal{NP}$-Completeness

Definition of 3-SAT

Decide whether satisfying assignment $a : \mathcal{L}_\phi \to \{\mathrm{F}, \mathrm{T}\}$ exists for formula ϕ. Output: Yes/No.

Theorem: Karp [1972]

3-SAT is $\mathcal{NP}$-complete.

A Decision Problem is $\mathcal{NP}$-complete if ...

... it lies in $\mathcal{NP}$ and *all* other decision problems in $\mathcal{NP}$ can be reduced to it.

Methodology: Proving $\mathcal{NP}$-completeness

Proving $\mathcal{NP}$-completeness of the VNEP

1. VNEP lies in $\mathcal{NP}$ (answer can be checked in polynomial time).
2. Reduction from 3-SAT to VNEP.

Methodology: Proving $\mathcal{NP}$-completeness

Proving $\mathcal{NP}$-completeness of the VNEP

1. VNEP lies in $\mathcal{NP}$ (answer can be checked in polynomial time).
2. Reduction from 3-SAT to VNEP.

Outline of Reduction Framework

3-SAT instance $\phi \longmapsto$ VNEP instance $(G_{r(\phi)}, G_{S(\phi)}, \text{mapping restrictions})$

ϕ satisfiable? $\Longleftrightarrow$ feasible embedding of $G_{r(\phi)}$ on $G_{S(\phi)}$ under restrictions?

Methodology: Proving $\mathcal{NP}$-completeness

Proving $\mathcal{NP}$-completeness of the VNEP

1. VNEP lies in $\mathcal{NP}$ (answer can be checked in polynomial time).
2. Reduction from 3-SAT to VNEP.

Outline of Reduction Framework

3-SAT instance $\phi \longmapsto$ VNEP instance $(G_{r(\phi)}, G_{S(\phi)}$, mapping restrictions$)$

ϕ satisfiable? $\Longleftrightarrow$ feasible embedding of $G_{r(\phi)}$ on $G_{S(\phi)}$ under restrictions?

Our Reduction Framework

Proving $\mathcal{NP}$-completeness of the VNEP

1. VNEP lies in $\mathcal{NP}$ (answer can be checked in polynomial time).
2. Reduction from 3-SAT to VNEP.

Outline of Reduction Framework

3-SAT instance $\phi \longmapsto$ VNEP instance $(G_{r(\phi)}, G_{S(\phi)},$ mapping restrictions$)$

ϕ satisfiable? $\rightleftarrows$ feasible embedding of $G_{r(\phi)}$ on $G_{S(\phi)}$ under restrictions?

Our Reduction Framework

Input: 3-SAT formula $\phi = (x_1 \lor x_2 \lor x_3) \land (\bar{x}_1 \lor x_2 \lor x_4) \land (x_2 \lor \bar{x}_3 \lor x_4)$

Request $G_{r(\phi)}$

- $V_{r(\phi)} = \{v_i \mid \mathcal{C}_i \in \mathcal{C}_\phi\}$
- $E_{r(\phi)} = \{\ (v_i, v_j) \mid \mathcal{C}_i$ introduces literal used by $\mathcal{C}_j\ \}$

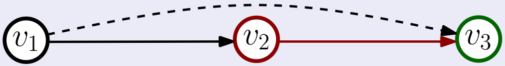

Substrate $G_{S(\phi)}$

- one node per clause and **per satisfying assignment**
- edges as for the requests, **if assignments do not contradict**

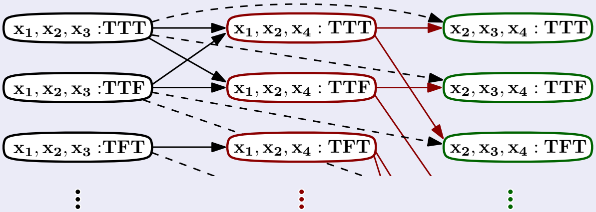

Complete Picture

ϕ: $(x_1 \vee x_2 \vee x_3) \quad \wedge \quad (\bar{x}_1 \vee x_2 \vee x_4) \quad \wedge \quad (x_2 \vee \bar{x}_3 \vee x_4)$

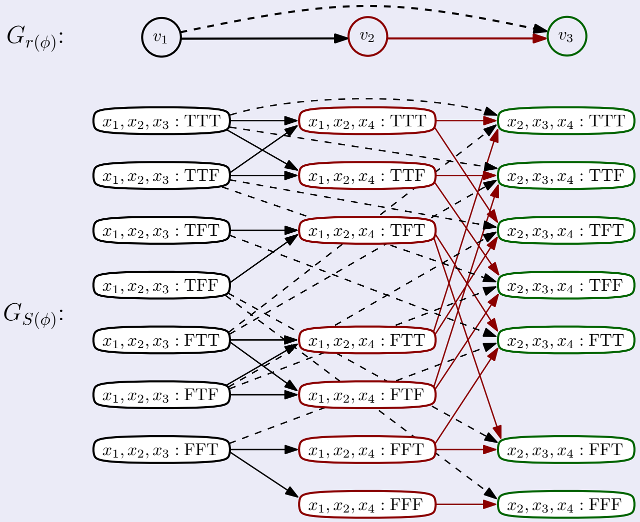

Our Reduction Framework

Outline of Reduction Framework

3-SAT instance $\phi \longmapsto$ VNEP instance $(G_{r(\phi)}, G_{S(\phi)}, \text{mapping restrictions})$

ϕ satisfiable? $\rightleftarrows$ feasible embedding of $G_{r(\phi)}$ on $G_{S(\phi)}$ under restrictions?

Our Reduction Framework

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

Visualization of Conditions (1) and (2)

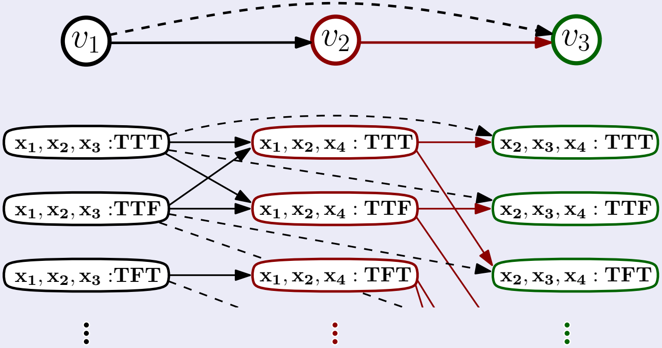

Our Reduction Framework

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

Visualization of Conditions (1) and (2)

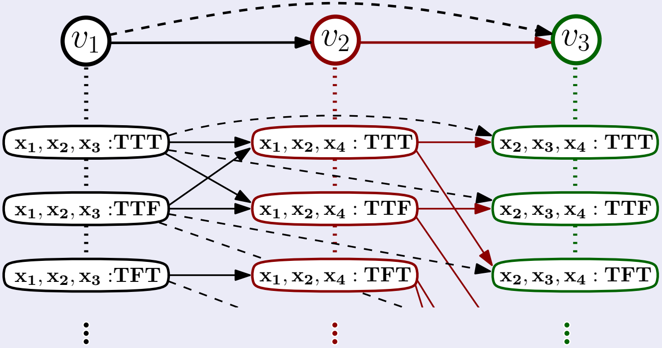

Our Reduction Framework

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

Visualization of Conditions (1) and (2)

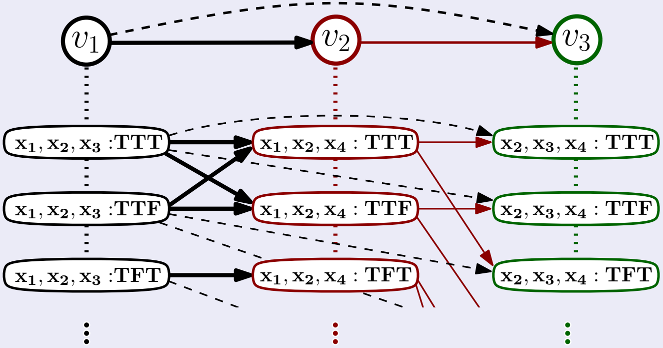

Our Reduction Framework

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

Visualization of Conditions (1) and (2)

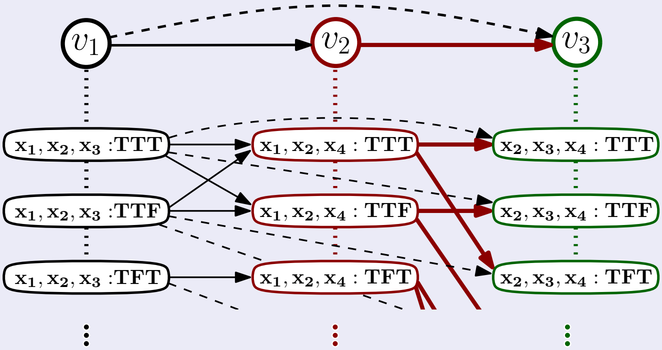

Our Reduction Framework

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

Visualization of Conditions (1) and (2)

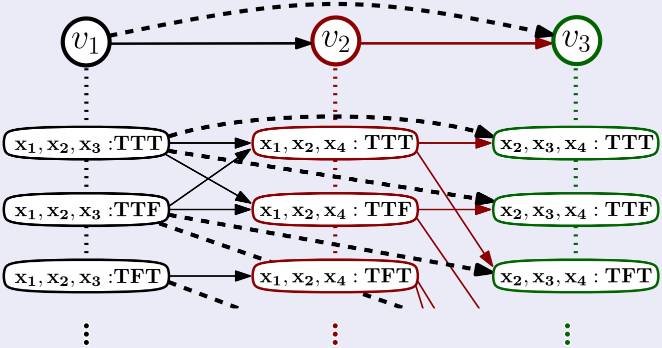

Our Reduction Framework

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a ‘satisfying assignment node’ of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

Example

$\phi = (x_1 \vee x_2 \vee x_3) \wedge (\bar{x}_1 \vee x_2 \vee x_4) \wedge (x_2 \vee \bar{x}_3 \vee x_4)$

Assignment for ϕ

$x_1 \mapsto \mathrm{T}$

$x_2 \mapsto \mathrm{T}$

$x_3 \mapsto \mathrm{F}$

$x_4 \mapsto \mathrm{F}$

Request

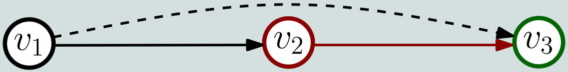

Embedding satisfying conditions (1) and (2)

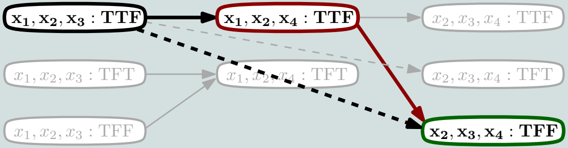

Our Reduction Framework

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

Application of Base Lemma for VNEP Variant $\langle \mathbf{X} \,|\, \mathbf{Y} \rangle$

VNEP $\langle \mathbf{X} \,|\, \mathbf{Y} \rangle$ is $\mathcal{NP}$-complete if we can enforce all *feasible* embeddings to satisfy (1) and (2).

3-SAT instance $\phi \longmapsto$ VNEP instance $(G_{r(\phi)}, G_{S(\phi)},$ under mapping restrictions$)$

ϕ satisfiable? $\Longleftrightarrow$ feasible embedding of $G_{r(\phi)}$ on $G_{S(\phi)}$ under restrictions?

Results

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

VNEP $\langle$ - | **NR** $\rangle$ is $\mathcal{NP}$-complete.

Node placement enforce (1)

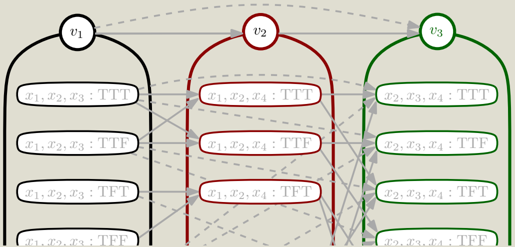

Routing restrictions enforce (2)

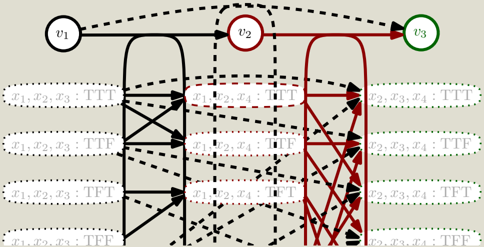

Results

Base Lemma

Formula ϕ is satisfiable **if and only if** there exists a mapping of $G_{r(\phi)}$ on $G_{S(\phi)}$, s.t.

(1) each virtual node v_i is mapped to a 'satisfying assignment node' of the i-th clause, and

(2) all virtual edges are mapped on exactly one substrate edge.

VNEP $\langle$ - | NL $\rangle$ is $\mathcal{NP}$-complete.

Node placement enforce (1)

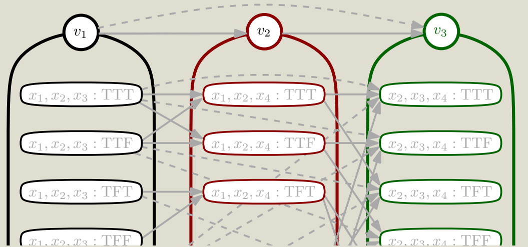

Placement and latency restrictions enforce (2)

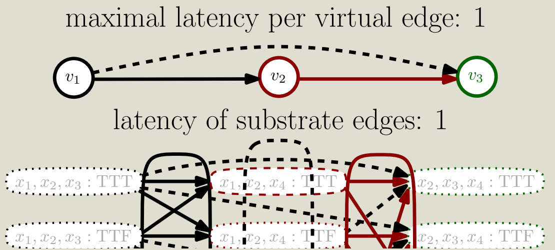

Results

$\mathcal{NP}$-Completeness shown for $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$ and $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$

In the paper: $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$.

Results

$\mathcal{NP}$-Completeness shown for $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$ and $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$

In the paper: $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$.

Implications of $\mathcal{NP}$-Completeness

- Finding a feasible embedding for the VNEP is $\mathcal{NP}$-complete.
- Finding an optimal feasible embedding **subject to any objective** is $\mathcal{NP}$-hard.
- There cannot exist polynomial-time approximation algorithms (unless $\mathcal{P} = \mathcal{NP}$).

$\mathcal{NP}$-Completeness of Computing Approximate Embeddings

$\mathcal{NP}$-Completeness of Computing Approximate Embeddings

Insight: If the problem is too hard, relax the model.

How hard are the VNEP variants when we allow for capacity violations or latency violations?

Allowing for Node Capacity Violations

Relaxation: We allow for substrate node capacity violations by a factor $\alpha < 2$.

Result: $\langle \mathbf{VE} \,|\, \textbf{-} \rangle$ and $\langle \mathbf{V} \,|\, \mathbf{R} \rangle$ stay $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Latency Violations

Relaxation: We allow for latency violations by a factor $\gamma < 2$.

Result: $\langle \textbf{-} \,|\, \mathbf{NL} \rangle$ stays $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Edge Capacity Violations (proven in our technical report [Rost and Schmid, 2018])

Relaxation: We allow for substrate edge capacity violations by a factor $\beta < 2$.

Result: $\langle \mathbf{VE} \,|\, \textbf{-} \rangle$ and $\langle \mathbf{E} \,|\, \mathbf{N} \rangle$ stay inapproximable for $\beta \in \mathcal{O}(\log n / \log \log n)$, $n = |V_S|$, unless $\mathcal{NP} \subseteq \mathcal{BP}\text{-}TIME$[a] $(\bigcup_{d \geq 1} n^{d \log \log n})$.

[a] $\mathcal{BP}\text{-}TIME$: Bounded-Error Probabilistic Polynomial-Time

$\mathcal{NP}$-Completeness of Computing Approximate Embeddings

Insight: If the problem is too hard, relax the model.

How hard are the VNEP variants when we allow for capacity violations or latency violations?

Allowing for Node Capacity Violations

Relaxation: We allow for substrate node capacity violations by a factor $\alpha < 2$.

Result: $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$ and $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$ stay $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Latency Violations

Relaxation: We allow for latency violations by a factor $\gamma < 2$.

Result: $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$ stays $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Edge Capacity Violations (proven in our technical report [Rost and Schmid, 2018])

Relaxation: We allow for substrate edge capacity violations by a factor $\beta < 2$.

Result: $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$ and $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$ stay inapproximable for $\beta \in \mathcal{O}(\log n / \log\log n)$, $n = |V_S|$, unless $\mathcal{NP} \subseteq \mathcal{BP}\text{-}\mathit{TIME}^a(\bigcup_{d\geq 1} n^{d \log\log n})$.

[a] $\mathcal{BP}\text{-}\mathit{TIME}$: Bounded-Error Probabilistic Polynomial-Time

$\mathcal{NP}$-Completeness of Computing Approximate Embeddings

Insight: If the problem is too hard, relax the model.

How hard are the VNEP variants when we allow for capacity violations or latency violations?

Allowing for Node Capacity Violations

Relaxation: We allow for substrate node capacity violations by a factor $\alpha < 2$.

Result: $\langle\, \mathbf{VE} \,|\, \textbf{-} \,\rangle$ and $\langle\, \mathbf{V} \,|\, \mathbf{R} \,\rangle$ stay $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Latency Violations

Relaxation: We allow for latency violations by a factor $\gamma < 2$.

Result: $\langle\, \textbf{-} \,|\, \mathbf{NL} \,\rangle$ stays $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Edge Capacity Violations (proven in our technical report [Rost and Schmid, 2018])

Relaxation: We allow for substrate edge capacity violations by a factor $\beta < 2$.

Result: $\langle\, \mathbf{VE} \,|\, \textbf{-} \,\rangle$ and $\langle\, \mathbf{E} \,|\, \mathbf{N} \,\rangle$ stay inapproximable for $\beta \in \mathcal{O}(\log n / \log \log n)$, $n = |V_S|$, unless $\mathcal{NP} \subseteq \mathcal{BP}\text{-}\mathit{TIME}$[a]$(\bigcup_{d \geq 1} n^{d \log \log n})$.

[a] $\mathcal{BP}\text{-}\mathit{TIME}$: Bounded-Error Probabilistic Polynomial-Time

$\mathcal{NP}$-Completeness of Computing Approximate Embeddings

Insight: If the problem is too hard, relax the model.

How hard are the VNEP variants when we allow for capacity violations or latency violations?

Allowing for Node Capacity Violations

Relaxation: We allow for substrate node capacity violations by a factor $\alpha < 2$.

Result: $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$ and $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$ stay $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Latency Violations

Relaxation: We allow for latency violations by a factor $\gamma < 2$.

Result: $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$ stays $\mathcal{NP}$-complete and inapproximable (unless $\mathcal{P} = \mathcal{NP}$).

Allowing for Edge Capacity Violations (proven in our technical report [Rost and Schmid, 2018])

Relaxation: We allow for substrate edge capacity violations by a factor $\beta < 2$.

Result: $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$ and $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$ stay inapproximable for $\beta \in \mathcal{O}(\log n / \log\log n)$, $n = |V_S|$, unless $\mathcal{NP} \subseteq \mathcal{BP}\text{-}\mathit{TIME}^a(\bigcup_{d \geq 1} n^{d \log\log n})$.

[a] $\mathcal{BP}\text{-}\mathit{TIME}$: Bounded-Error Probabilistic Polynomial-Time

$\mathcal{NP}$-Completeness when Restricting Graph Classes

$\mathcal{NP}$-Completeness for Restricted Graph Classes

Insight: If the problem is too hard, restrict the model inputs.

How hard are the VNEP variants when we restrict the graph classes for the substrate and the requests?

Restriction on Directed Acyclic Graphs

By construction, the graphs $G_{r(\phi)}$ and $G_{S(\phi)}$ are directed acyclic graphs (DAGs). Accordingly, the hardness results pertain when restricting the input graphs to be DAGs.

Restriction of Requests to Planar Degree-Bounded Graphs

Restriction: The request graph must be a planar and degree-bounded.

Result: All previous results pertain based on a reduction from a **special planar** 3-SAT variant.

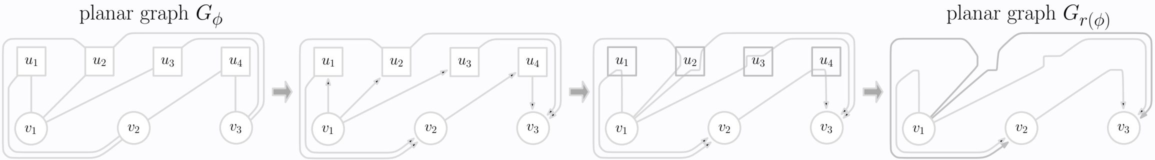

$\mathcal{NP}$-Completeness for Restricted Graph Classes

Insight: If the problem is too hard, restrict the model inputs.

How hard are the VNEP variants when we restrict the graph classes for the substrate and the requests?

Restriction on Directed Acyclic Graphs

By construction, the graphs $G_{r(\phi)}$ and $G_{S(\phi)}$ are directed acyclic graphs (DAGs). Accordingly, the hardness results pertain when restricting the input graphs to be DAGs.

Restriction of Requests to Planar Degree-Bounded Graphs

Restriction: The request graph must be a planar and degree-bounded.

Result: All previous results pertain based on a reduction from a **special planar** 3-SAT variant.

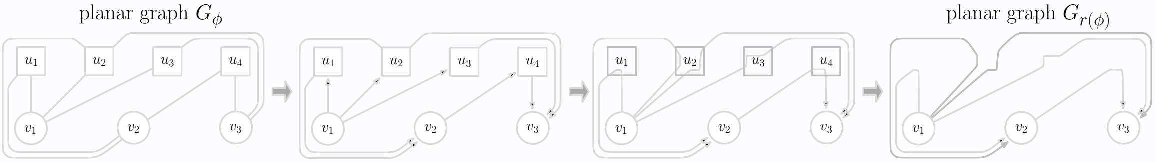

$\mathcal{NP}$-Completeness for Restricted Graph Classes

Insight: If the problem is too hard, restrict the model inputs.

How hard are the VNEP variants when we restrict the graph classes for the substrate and the requests?

Restriction on Directed Acyclic Graphs

By construction, the graphs $G_{r(\phi)}$ and $G_{S(\phi)}$ are directed acyclic graphs (DAGs). Accordingly, the hardness results pertain when restricting the input graphs to be DAGs.

Restriction of Requests to Planar Degree-Bounded Graphs

Restriction: The request graph must be a planar and degree-bounded.

Result: All previous results pertain based on a reduction from a **special planar** 3-SAT variant.

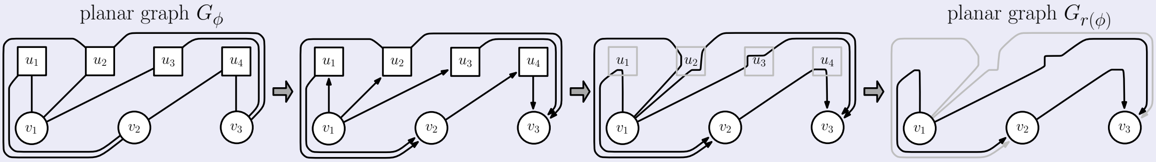

Conclusion

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Outlook

Further Results: Proof $\mathcal{NP}$-completeness for $\langle\,\mathbf{V}\,|\,\mathbf{RL}\,\rangle$, consider *uniform* capacities, ...

Improvements: Improvement of lower bounds for approximate embeddings.

Gained Insights: The VNEP is really hard.

- Justifies using heuristics and exponential-time algorithms.
- Approximations require model relaxations/restrictions.

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Outlook

Further Results Proof $\mathcal{NP}$-completeness for $\langle\,\mathbf{V}\,|\,\mathbf{RL}\,\rangle$, consider *uniform* capacities, ...

Improvements Improvement of lower bounds for approximate embeddings.

Gained Insights The VNEP is really hard.

- Justifies using heuristics and exponential-time algorithms.
- Approximations require model relaxations/restrictions.

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Outlook

Further Results	Proof $\mathcal{NP}$-completeness for $\langle\,\mathbf{V}\,|\,\mathbf{RL}\,\rangle$, consider *uniform* capacities, ...
Improvements	Improvement of lower bounds for approximate embeddings.
Gained Insights	The VNEP is really hard.

- Justifies using heuristics and exponential-time algorithms.
- Approximations require model relaxations/restrictions.

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\textbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\textbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Outlook

Further Results	Proof $\mathcal{NP}$-completeness for $\langle\,\mathbf{V}\,|\,\mathbf{RL}\,\rangle$, consider *uniform* capacities, ...
Improvements	Improvement of lower bounds for approximate embeddings.
Gained Insights	The VNEP is really hard.

- Justifies using heuristics and exponential-time algorithms.
- Approximations require model relaxations/restrictions.
  More about VNEP *approximations in my talk tomorrow.*

Our Contributions

1. $\mathcal{NP}$**-completeness** under restrictions $\langle\,\mathbf{VE}\,|\,\mathbf{-}\,\rangle$, $\langle\,\mathbf{E}\,|\,\mathbf{N}\,\rangle$, $\langle\,\mathbf{V}\,|\,\mathbf{R}\,\rangle$, $\langle\,\mathbf{-}\,|\,\mathbf{NR}\,\rangle$, $\langle\,\mathbf{-}\,|\,\mathbf{NL}\,\rangle$.
2. **Relaxed model:** $\mathcal{NP}$**-completeness** of computing **approximate embeddings**.
3. **Restricted input:** $\mathcal{NP}$**-completeness** pertains when restricting request topologies.

Outlook

Further Results Proof $\mathcal{NP}$-completeness for $\langle\,\mathbf{V}\,|\,\mathbf{RL}\,\rangle$, consider *uniform* capacities, ...

Improvements Improvement of lower bounds for approximate embeddings.

Gained Insights The VNEP is really hard.

- Justifies using heuristics and exponential-time algorithms.
- Approximations require model relaxations/restrictions.

More about VNEP *approximations in my talk tomorrow.*

Thank you! Questions?

References I

Edoardo Amaldi, Stefano Coniglio, Arie M.C.A. Koster, and Martin Tieves. On the computational complexity of the virtual network embedding problem. *Electronic Notes in Discrete Mathematics*, 52:213 – 220, 2016.

David G. Andersen. Theoretical approaches to node assignment. [Online]. Available: `http://repository.cmu.edu/compsci/86/`, December 2002.

NM Mosharaf Kabir Chowdhury, Muntasir Raihan Rahman, and Raouf Boutaba. Virtual network embedding with coordinated node and link mapping. In *IEEE INFOCOM*, 2009.

Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.

Matthias Rost and Stefan Schmid. NP-Completeness and Inapproximability of the Virtual Network Embedding and Its Variants. *CoRR*, abs/1801.03162, 2018. URL `http://arxiv.org/abs/1801.03162`.