

# Optimal Virtual Network Embeddings for Tree Topologies\*

Aleksander Figiel<sup>†</sup>  
Technische Universität Berlin  
Algorithmics and Computational  
Complexity  
Berlin, Germany

Matthias Rost  
SAP SE  
Potsdam, Germany  
Technische Universität Berlin  
Data Communications and  
Networking  
Berlin, Germany

Leon Kellerhals  
Technische Universität Berlin  
Algorithmics and Computational  
Complexity  
Berlin, Germany

Stefan Schmid<sup>‡</sup>  
University Vienna  
Faculty of Computer Science  
Vienna, Austria  
Technische Universität Berlin  
Data Communications and  
Networking  
Berlin, Germany

Rolf Niedermeier  
Technische Universität Berlin  
Algorithmics and Computational  
Complexity  
Berlin, Germany

Philipp Zschoche  
Technische Universität Berlin  
Algorithmics and Computational  
Complexity  
Berlin, Germany

## ABSTRACT

The performance of distributed and data-centric applications often critically depends on the interconnecting network. Applications are hence modeled as virtual networks, also accounting for resource demands on links. At the heart of provisioning such virtual networks lies the NP-hard Virtual Network Embedding Problem (VNEP): how to jointly map the virtual nodes and links onto a physical substrate network at minimum cost while obeying capacities.

This paper studies the VNEP in the light of parameterized complexity. We focus on tree topology substrates, a case often encountered in practice and for which the VNEP remains NP-hard. We provide the first fixed-parameter algorithm for the VNEP with running time  $O(3^r(s+r^2))$  for requests and substrates of  $r$  and  $s$  nodes, respectively. In a computational study our algorithm yields running time improvements in excess of 200× compared to state-of-the-art integer programming approaches. This makes it comparable in speed to the well-established ViNE heuristic while providing optimal solutions. We complement our algorithmic study with hardness results for the VNEP and related problems.

## CCS CONCEPTS

• **Networks** → **Network resources allocation**; • **Theory of computation** → **Fixed parameter tractability**;

\*A full version of this paper is available on arXiv: <https://arxiv.org/abs/2105.07006>.

<sup>†</sup>Supported by DFG, Project MaMu NI 369/19.

<sup>‡</sup>Funded by European Research Council (ERC), grant agreement 864228 (AdjustNet).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

SPAA '21, July 6–8, 2021, Virtual Event, USA

© 2021 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8070-6/21/07...\$15.00

<https://doi.org/10.1145/3409964.3461787>

## KEYWORDS

Virtual Network Embedding, Fixed-Parameter Tractability, Dynamic Programming, Computational Complexity

### ACM Reference Format:

Aleksander Figiel, Leon Kellerhals, Rolf Niedermeier, Matthias Rost, Stefan Schmid, and Philipp Zschoche. 2021. Optimal Virtual Network Embeddings for Tree Topologies. In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '21), July 6–8, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3409964.3461787>

## 1 INTRODUCTION

Data-centric and distributed applications, including batch processing, streaming, scale-out databases, or distributed machine learning, generate a significant amount of network traffic and their performance critically depends on the underlying network. As the network infrastructure is often shared and the bandwidth available can vary significantly over time, this can have a non-negligible impact on the application performance [18].

Network virtualization has emerged as a promising solution to ensure a predictable application performance over shared infrastructures, by providing a virtual network abstraction which comes with explicit bandwidth guarantees [5]. In a nutshell, a virtual network request is modeled as a directed graph  $G_R = (V_R, E_R)$  whose elements are attributed with resource demands. The nodes represent, e.g., containers or virtual machines, requesting, e.g., CPU cores and memory, while the edges represent communication channels of a certain bandwidth. Formally, the demands of a virtual network request  $G_R$  are a function  $d_R : G_R \rightarrow \mathbb{R}_{\geq 0}^\tau$ ,  $\tau \in \mathbb{N}$ , of every node and every edge onto a  $\tau$ -dimensional vector of nonnegative reals. To provision such a virtual network request in a (shared) physical substrate network, also modeled as a directed graph  $G_S = (V_S, E_S)$  with capacities  $d_S : G_S \rightarrow \mathbb{R}_{\geq 0}^\tau$ , we need to find an *embedding* that maps the request nodes onto the substrate nodes and the request edges onto paths in the substrate while respecting capacities.

The NP-hard VIRTUAL NETWORK EMBEDDING PROBLEM, asking to find such embeddings, poses the main challenge of provisioning

virtual networks and has been studied for various objectives [7]. In this paper, we study the following central cost-minimization variant (see Definition 1 for a formal definition):

MINIMUM-COST VIRTUAL NETWORK EMBEDDING (MIN-VNEP)

**Input:** A directed graph  $G_S = (V_S, E_S)$  on  $s$  nodes, called *substrate*, and a directed graph  $G_R = (V_R, E_R)$  on  $r$  nodes, called *request*, with  $\tau$ -dimensional demands  $d_R : G_R \rightarrow \mathbb{R}_{\geq 0}^\tau$ , capacities  $d_S : G_S \rightarrow \mathbb{R}_{\geq 0}^\tau$ , and costs  $c_S : G_S \rightarrow \mathbb{R}_{\geq 0}^\tau$ .

**Task:** Find mappings of the request onto substrate nodes and of the request edges onto paths in the substrate, such that

- (1) the node and edge capacities are respected
- (2) the cost of all nodes and edges used by the mapping is minimized.

We remark that several other variants of the VIRTUAL NETWORK EMBEDDING PROBLEM can be reduced to MIN-VNEP (see Section 1.2).

## 1.1 Contributions and Techniques

While the MIN-VNEP is known to be notoriously hard in general [21], real-world network optimization problems often exhibit a specific structure. In this work, we provide efficient, exact algorithms that exploit such a structural property. Our main theoretical contribution is a fixed-parameter algorithm for the MIN-VNEP onto tree substrates when parameterized by the number of nodes in the request, that is, we present an algorithm which performs very well for small request graphs:

**Theorem 1.** *MIN-VNEP can be solved in  $O(3^r(s+r)^2)$  time when the substrate  $G_S$  is a tree, where  $r = |V_R|$  and  $s = |V_S|$ .*

From a theoretical (worst-case) point of view there is almost no hope to obtain a substantially faster (exact) algorithm for tree substrates (see Section 2). A specific feature of the algorithm is its robustness: It can be easily modified to also support additional constraints such as mapping exclusions on a per-node or per-edge basis [24]. Furthermore, as a side result, we show that any instance of MIN-VNEP on tree substrates can be translated in linear time into an instance of MIN-VNEP in which the substrate is a binary tree and only its leaves have non-zero capacities. Hence, algorithms designed for such tree substrates, as, e.g., those by Ballani et al. [3] and Rost et al. [23], can also be applied on general tree substrates.

The algorithm of Theorem 1 also performs very well in practice. In an extensive computational study we compare our algorithm to the classical exact algorithm based on integer programming as well as to the well-established ViNE heuristic [5]. The results are clear: our algorithm outperforms the integer program on all instances, consistently yielding average speedups exceeding a factor of 100× and often even a factor of 200× for densely connected request graphs across small to medium-sized data center networks. The running time of ViNE lies in the same order of magnitude as the one of our algorithm, but produces feasible solutions only for a quarter of the instances for which our algorithm found an *optimal* solution. To ensure reproducibility and facilitate follow-up work, we will provide our implementation to the research community as open source code, together with all experimental artefacts.

As mentioned before, we complement our algorithm (Theorem 1) by showing that in theory there is little hope for improving its

running time substantially. This can be derived from a simple NP-hardness result for the decision version of MIN-VNEP, which we will call VNEP. Here, we are given an instance of MIN-VNEP together with an integer  $k$  and ask whether there is an embedding with costs at most  $k$ . We show the following.

**Theorem 2.** *VNEP is NP-hard, even if the substrate  $G_S$  consists of two nodes and the request  $G_R$  is edgeless, and  $k = 0$ .*

An intermediate question from Theorem 1 is whether we can find another graph parameter  $x$  of the request which is asymptotically smaller than  $r$  (number of vertices) but still admits an exact algorithm of running time  $f(x)(s+r)^{O(1)}$ , where  $f$  is a computable function. Assuming  $P \neq NP$ , such a running time cannot be achieved for any parameter  $x$  which is asymptotically smaller than the number of edges in the request. This is because the NP-hardness for the VNEP holds even if the request contains no edges. Also, Theorem 2 rules out the existence of any approximation algorithm for the MIN-VNEP, even if the degree of the polynomial may depend on the substrate’s number of nodes and the request’s number of edges.

Our last contribution is a conditional lower bound on the running time of the *Valid Mapping Problem* (VMP), a relaxation of the VNEP: Analogously to the VNEP, the question is whether there are node and edge mappings of the request onto the substrate such that the cost is below a given  $k \in \mathbb{R}_{\geq 0}$ , but we only enforce that the mapping of each individual virtual element does not exceed the capacities of the substrate (see Section 2 for a formal definition). This relaxation is used for instance by Rost et al. [24] to obtain an approximation algorithm for VNEP in the resource augmentation framework. Specifically, they present an algorithm for VMP running in  $\text{poly}(r) \cdot s^{O(\text{tw}(G_R))}$  time, where  $s$  and  $r$  are the number of nodes in the substrate and the request, respectively, and  $\text{tw}(G_R)$  is the treewidth of the request [8]. By proving a W[1]-hardness result, we show that there is presumably no fixed-parameter algorithm for VMP parameterized by the cost upper bound  $k$  combined with the number of nodes  $r$  in the request, and that the running time for VMP obtained by Rost et al. [24] is asymptotically optimal:

**Theorem 3.** *VMP parameterized by  $k + r$  is W[1]-hard and, unless the Exponential Time Hypothesis (ETH) fails, there is no algorithm for VMP running in  $f(r) \cdot s^{o(r)}$  time, where  $r$  and  $s$  are the number of nodes in the request and the substrate, respectively.*

## 1.2 Related Work and Novelty

The VIRTUAL NETWORK EMBEDDING PROBLEM has received tremendous attention by the networking community over the last 15 years: already by 2013 more than 80 algorithms had been published in the literature for its various flavors [7]. The particular MIN-VNEP objective, on which we focus in this paper, has received by far the most attention: there is extensive work on heuristics [5, 15, 17, 31] as well as exact algorithms based on mixed-integer programs [5, 13] for MIN-VNEP. Notably, however, there is no work so far on (non-trivial) combinatorial *exact* algorithms for MIN-VNEP.

*Closely Related Applications.* Various applications of the VIRTUAL NETWORK EMBEDDING PROBLEM have spawned independent research with dozens of proposed algorithms. Among the most prominent ones are the embedding problems pertaining to Virtual

Clusters (VCEP) [3], to Service Function Chains (SFCEP) [11], to Virtual Data Centers (VDCEP) [30], and to the Internet of Things [26]. In short, the VCEP studies the embedding of tree requests onto data center topologies, the SFCEP studies the embedding of sparse requests representing (virtualized) network functions, and the VDCEP focuses on the embedding of arbitrary requests across geographically distributed data centers in wide-area networks. While at times introducing additional constraints, the Virtual Network Embedding Problem lies at the heart of these problems as well.

*Applications of MIN-VNEP.* Various algorithms rely on solving the MIN-VNEP as a subroutine. The application domains include:

**Offline Objectives.** The offline setting of the VIRTUAL NETWORK EMBEDDING PROBLEM over several requests under cost objectives can be solved by MIN-VNEP by considering the union of the requests. Further, there are exponential-time (parameterized) approximations for the offline setting in the resource augmentation framework that use algorithms for the cost minimization variant of VMP or MIN-VNEP as a subroutine [19, 20, 24].

**Competitive online optimization.** Even et al. [6] showed how to construct competitive online algorithms for the profit variant of the VNEP from any exact algorithm for the MIN-VNEP.

**Congestion minimization.** Bansal et al. [4] studied the problem of minimizing the maximal load (while not enforcing capacities). They obtained competitive online and offline approximation algorithms that solve MIN-VNEP as a subroutine.

Given our fixed-parameter algorithm for the special case of tree substrates (see Theorem 1), novel parameterized algorithms for all of the above highlighted settings and objectives can be obtained.

*(Parameterized) Complexity.* Despite the popularity of the VNEP, until recently only little was known about its fine-grained computational complexity. Rost and Schmid [21] made the first step towards understanding the (parameterized) complexity of the VNEP, showing that any optimization variant of the VNEP (where MIN-VNEP is one of them) is inapproximable in polynomial time, unless  $P = NP$ , even when the request graph is planar and the substrate is acyclic. Rost et al. [24] gave the first approximation algorithm for the offline profit objective for requests of constant treewidth in the resource augmentation framework, also carrying over to the cost setting [19].

In contrast to the above works, we focus on efficient and exact fixed-parameter algorithms while restricting the substrate to be a tree. Tree substrates are most predominantly encountered in data centers, e.g., in the form of fat trees [1]. Fat trees or similar leaf-spine architectures are widely studied in the literature and used in industry [3, 10]. Additionally, by employing substrate transformations, such as computing Gomory-Hu trees [27], non-tree substrates may be transformed to trees, albeit optimality guarantees cannot be preserved. Bansal et al. [4] designed specific algorithms for tree substrates of bounded depth, where the objective is to minimize congestion. For the parameterization of the request size—the main focus of this paper—no results are known thus far.

*Small Request Graphs.* The application of our main result (cf. Theorem 1) yields algorithms of practical significance only when the number of request nodes is small and in our computational study we restrict our attention to request graphs on less than 12 nodes.

While this may be considered to be an unreasonably small number of nodes, many existing works on the VNEP [5, 7] and its applications in data centers [28, 29] consider requests of such size.

### 1.3 Preliminaries

For  $n \in \mathbb{N}$  let  $[n] := \{1, \dots, n\}$ . For two vectors  $a = (a_i)_{i=1}^\tau, b = (b_i)_{i=1}^\tau$  we write  $a \leq b$  if  $a_i \leq b_i$  for all  $i \in [\tau]$  and  $a \not\leq b$  otherwise.

Let  $G = (V, E)$  be a directed graph. For a node subset  $V' \subseteq V$ , we denote by  $G[V']$  the *subgraph of  $G$  induced by  $V'$* , and by  $V(G[V'])$  and  $E(G[V'])$  the node set and the edge set of  $G[V']$ , respectively. For a node  $v \in V$  we denote by  $N_G^+(v)$ , respectively  $N_G^-(v)$ , the set of nodes that are connected by an edge pointing away from, respectively towards  $v$ . By  $N_G(v) := N_G^+(v) \cup N_G^-(v)$  we denote the (combined) neighborhood of  $v$ . The degree  $\deg_G(v)$  of  $v$  is the number of nodes in the neighborhood of  $v$ . The *underlying undirected graph* of a directed graph  $G$  is the undirected graph without multiedges on the same node set and it has an edge  $\{u, v\}$  for every directed edge  $(u, v)$  in  $G$ . We say that a directed graph is a *tree* if its underlying undirected graph is a tree.

Given an instance of either MIN-VNEP, its decision variant VNEP or the VMP, we say that a pair of mappings  $(m_R^V, m_R^E)$  is a *valid mapping* if the edge mappings are valid, and capacities are respected per each individual virtual element, that is,

- (1) for every edge  $(u, v) \in E_r$ ,  $m_R^E(u, v)$  is a path from  $m_R^V(u)$  to  $m_R^V(v)$  in  $G_S$ ,
- (2)  $d_R(w) \leq d_S(m_R^V(w))$  for every  $w \in V_R$ , and
- (3)  $d_R(e) \leq d_S(e_S)$  for all virtual edges  $e \in E_R$  and their mappings  $e_S \in m_R^E(e)$ .

We call the mapping *feasible* if additionally all demands of the request nodes and edges can be fulfilled by the capacities of the substrate nodes and edges they are mapped onto, that is,

$$\sum_{w: m_R^V(w)=v} d_R(w) \leq d_S(v) \quad \text{for } v \in V_S, \quad \text{and}$$

$$\sum_{e_R: e_S \in m_R^E(e_R)} d_R(e_R) \leq d_S(e_S) \quad \text{for } e_S \in E_S.$$

The *cost* of a mapping  $(m_R^V, m_R^E)$  is defined as the sum of the cost of mapping all nodes plus the sum of the costs mapping all edges. Note that the latter consists of the cost of every single edge of the path onto which a request edge is mapped. Formally, the cost is

$$\sum_{v \in V(G_R)} d_R(v)^\top c_S(m_R^V(v)) + \sum_{e \in E(G_R)} \left( \sum_{e' \in m_R^E(e)} d_R(e)^\top c_S(e') \right).$$

We can now formally define MIN-VNEP:

**Definition 1** (Min. Cost Virtual Network Embedding (MIN-VNEP)).

**Input:** A directed graph  $G_S = (V_S, E_S)$  on  $s$  nodes, called *substrate*, and a directed graph  $G_R = (V_R, E_R)$  on  $r$  nodes, called *request*, with demands  $d_R : G_R \rightarrow \mathbb{R}_{\geq 0}^\tau$ , capacities  $d_S : G_S \rightarrow \mathbb{R}_{\geq 0}^\tau$ , and costs  $c_S : G_S \rightarrow \mathbb{R}_{\geq 0}^\tau$ .

**Task:** Find a feasible mapping of minimum cost.

In the decision variant, VNEP, we are additionally given a non-negative  $k \in \mathbb{R}_{\geq 0}$  with an instance of MIN-VNEP and decide whether there is a feasible mapping with cost at most  $k$ . Formally, it is defined as follows (note that in this definition we replace the  $\tau$ -dimensional vectors by scalars):

**Definition 2** (Virtual Network Embedding Problem (VNEP)).

**Input:** A directed graph  $G_S = (V_S, E_S)$  on  $s$  nodes, called *substrate*, and a directed graph  $G_R = (V_R, E_R)$  on  $r$  nodes, called *request*, with demands  $d_R : G_R \rightarrow \mathbb{R}_{\geq 0}$ , capacities  $d_S : G_S \rightarrow \mathbb{R}_{\geq 0}$ , costs  $c_S : G_S \rightarrow \mathbb{R}_{\geq 0}$ , and a cost upper bound  $k \in \mathbb{R}_{\geq 0}$ .

**Question:** Is there a feasible mapping of cost at most  $k$ ?

The Valid Mapping Problem (VMP) takes the same input as the VNEP and asks whether there is a *valid* (but not necessarily feasible) mapping with cost at most  $k$ .

We assume familiarity with standard notions regarding algorithms and complexity, but briefly review notions regarding parameterized complexity analysis. Let  $\Sigma$  denote a finite alphabet. A parameterized problem  $L \subseteq \{(x, k) \in \Sigma^* \times \mathbb{N}_0\}$  is a subset of all instances  $(x, k)$  from  $\Sigma^* \times \mathbb{N}_0$ , where  $k$  denotes the parameter. A parameterized problem  $L$  is *fixed-parameter tractable* (or contained in the class FPT) if there is an algorithm that decides every instance  $(x, k)$  for  $L$  in  $f(k) \cdot |x|^{O(1)}$  time, and it is *contained in the class XP* if there is an algorithm that decides every instance  $(x, k)$  for  $L$  in  $|x|^{f(k)}$  time, where  $f$  is any computable function only depending on the parameter and  $|x|$  is the size of  $x$ . For two parameterized problems  $L, L'$ , an instance  $(x, k) \in \Sigma^* \times \mathbb{N}_0$  of  $L$  is equivalent to an instance  $(x', k') \in \Sigma^* \times \mathbb{N}_0$  for  $L'$  if  $(x, k) \in L \iff (x', k') \in L'$ . A problem  $L$  is W[1]-hard if for every problem  $L' \in W[1]$  there is an algorithm that maps any instance  $(x, k)$  in  $f(k) \cdot |x|^{O(1)}$  time to an equivalent instance  $(x', k')$  with  $k' = g(k)$  for some computable functions  $f, g$ . It holds true that  $\text{FPT} \subseteq W[1] \subseteq \text{XP}$ . It is believed that  $\text{FPT} \neq W[1]$ , and that hence no W[1]-hard problem is believed to be fixed-parameter tractable. Another prominent assumption in the literature is the *Exponential Time Hypothesis* (ETH) which states that there is no  $2^{o(n)}$ -time algorithm for 3-SAT, where  $n$  is the number of variables [12].

## 2 HARDNESS

In this section, we show that there is no XP-algorithm to solve optimally, or approximate the costs of, MIN-VNEP for any combined parameter consisting of (i) any parameter of the substrate and (ii) the number of edges in the request, unless  $\text{P}=\text{NP}$ . In related work, we can find several special cases in which MIN-VNEP remains NP-hard [2, 4]. However, from the parameterized point of view the following simple polynomial-time many-one reduction from PARTITION to VNEP (the decision version of MIN-VNEP) excludes many potential parameters towards an FPT- or even an XP-algorithm.

**Theorem 2.** *VNEP is NP-hard, even if the substrate  $G_S$  consists of two nodes and the request  $G_R$  is edgeless, and  $k = 0$ .*

**PROOF.** We reduce from the NP-hard PARTITION problem, where we are given a multiset  $S$  of positive integers and ask whether there is a  $S' \subseteq S$  such that  $\sum_{x \in S'} x = \sum_{x \in S \setminus S'} x$  [14].

Let  $S$  be such a multiset of positive integers and assume without loss of generality that  $B := \sum_{x \in S} x$  is even. We construct an instance  $I = (G_S, G_R, d_R, d_S, c_S, k = 0)$  of VNEP such that  $G_S := (\{a, b\}, \{(a, b), (b, a)\})$ ,  $G_R := (S, \emptyset)$  and  $d_R(x) := x$  for all  $x \in S$ ,  $c_S(a) := c_S(b) := c_S(a, b) := c_S(b, a) := 0$ ,  $d_S(a, b) := d_S(b, a) := 0$ , and  $d_S(a) := d_S(b) := \frac{B}{2}$ . Clearly, this is doable in polynomial time.

We now show that there exists a solution  $S' \subseteq S$  if and only if there exists a feasible mapping  $(m_R^V, m_R^E)$  for  $I$  of cost 0.

( $\implies$ ): Let  $S' \subseteq S$  such that  $\sum_{x \in S'} x = \sum_{x \in S \setminus S'} x = \frac{B}{2}$ . Then, we set  $m_R^V(x) = a$ , for all  $x \in S'$ , and  $m_R^V(x) = b$ , for all  $x \in S \setminus S'$ . Observe that  $(m_R^V, m_R^E)$  is a feasible mapping of cost 0.

( $\impliedby$ ): Let  $(m_R^V, m_R^E)$  be a feasible mapping for  $I$  of cost 0. Let  $S' \subseteq S$  be the set of nodes of  $G_R$  which are mapped to  $a$ . Hence,  $\sum_{x \in S'} x \leq d_S(a) = \frac{B}{2}$  and  $\sum_{x \in S \setminus S'} x \leq d_S(b) = \frac{B}{2}$ . Since  $\sum_{x \in S} x = B$ , we have  $\sum_{x \in S'} x = \sum_{x \in S \setminus S'} x$ .  $\square$

Since VNEP is NP-hard even if the substrate is of constant size, we can conclude that there is no XP-algorithm for VNEP parameterized by any reasonable parameter of the substrate, unless  $\text{P}=\text{NP}$ . Otherwise, this would imply a polynomial-time algorithm for the NP-hard PARTITION problem. Furthermore, since VNEP is NP-hard even if the substrate graph is of constant size and the request is edgeless, we can exclude the existence of an XP-algorithm for VNEP parameterized by a combination of any ‘reasonable’ parameter for the substrate and the number of edges in the request. Note that this excludes among others the parameters *vertex cover number*, *feedback edge number*, *treewidth*, and *maximum degree* of the request, because these parameters are upper-bounded by the number of edges. Moreover, since  $k = 0$  in Theorem 2, any approximation algorithm<sup>1</sup> for MIN-VNEP would be able to solve PARTITION. Altogether, we have the following.

**Corollary 1.** *Let  $f : \mathcal{G} \rightarrow \mathbb{N}$  be a computable function, where  $\mathcal{G}$  is the set of directed graphs. Unless  $\text{P}=\text{NP}$ ,*

- (1) *there is no  $|I|^{h(f(G_S)+|E_R|)}$ -time algorithm for VNEP, and*
- (2) *there is no  $|I|^{h(f(G_S)+|E_R|)}$ -time approximation algorithm for MIN-VNEP,*

*where  $|I|$  is the instance size,  $G_S$  is the substrate,  $|E_R|$  is the number of edges in the request, and  $h : \mathbb{N} \rightarrow \mathbb{N}$  is a computable function.*

Given the hardness results of Corollary 1, we see two ways to develop efficient exact algorithms:

- (1) Restrict the input instances to special cases which are relevant in practice—this is what we do in Section 3.
- (2) Study a reasonable relaxation of the problem—such as the (NP-hard) VMP.

Towards (2), Rost et al. [24] studied and presented an algorithm for the VMP running in  $\text{poly}(r) \cdot s^{O(\text{tw}(G_R))}$  time, where  $\text{tw}(G_R)$  is the treewidth of the request. They then used this algorithm as a subroutine in an approximation algorithm for an offline variant of the Virtual Network Embedding Problem (see Section 1.2).

With Theorem 3, we show that the algorithm of Rost et al. [24] is asymptotically optimal, unless the Exponential Time Hypothesis fails. The proof of Theorem 3 can be found in the full version.

## 3 EFFICIENT VNEP ALGORITHM FOR SMALL REQUESTS ON TREES

We focus on the special case of VNEP where the substrate is a tree and show that it is fixed-parameter tractable when parameterized

<sup>1</sup>That is, an algorithm returning a feasible solution and giving provable guarantees on the distance of the returned solution to the optimal one.

by the number of nodes in the request. Thus, the main objective of this section is to show the following.

**Theorem 1.** *MIN-VNEP can be solved in  $O(3^r(s+r^2))$  time when the substrate  $G_S$  is a tree, where  $r = |V_R|$  and  $s = |V_S|$ .*

Recall that VNEP (and thus MIN-VNEP) on tree substrates is NP-hard (Theorem 2), even if the request contains no edges. Thus, we cannot improve on Theorem 1 by replacing the parameter *number of nodes in the request* with a smaller parameter like *vertex cover number*, *feedback edge number*, or *maximum degree*, unless P=NP.

Our algorithm for Theorem 1 works in three steps (see Algorithm 3.1 for a pseudocode illustration):

- (1) Introduce additional leaves to the substrate to ensure that all non-leaves have capacity zero (Lemma 1, method Leaf in the pseudocode).
- (2) Split nodes in the substrate with more than two children such that we obtain a binary tree (Lemma 2, method Split).
- (3) Use dynamic programming to solve MIN-VNEP with the substrate being restricted to such trees (method GetEntry).

We remark that the first two steps (Lemmata 1 and 2) can be used as a preprocessing for any algorithms that only work for binary tree substrates on which the capacity of all non-leaf nodes is zero [3, 23] to make them work for general tree substrates.

Throughout this section we assume without loss of generality that our substrate graph is bidirectional, that is, for every edge  $(u, v)$  in  $E_S$  we also have the edge  $(v, u)$ . Otherwise, we add the missing edge and set its capacity to zero. Further, we assume that our substrate graph  $G_S$  is a tree rooted at some vertex  $p$ .

We first show that we can assume that all non-leaf nodes of our substrate have capacity zero.

**Lemma 1.** *Given an instance  $I = (G_S, G_R, d_R, d_S, c_S)$  of MIN-VNEP, we can build in linear time an instance  $\tilde{I} = (\tilde{G}_S, G_R, d_R, \tilde{d}_S, \tilde{c}_S)$  of MIN-VNEP such that*

- (i) *each node  $v \in \tilde{V}_S$  of degree at least two fulfills  $\tilde{d}_S(v) = 0$ , and*
- (ii) *there is a solution for  $I$  of cost at most  $k$  if and only if there is a solution for  $\tilde{I}$  of cost at most  $k$ .*

**PROOF.** The idea is to add a fresh leaf for each non-leaf vertex with capacities above zero. Without loss of generality, we assume that each edge in  $G_S$  is bidirectional, otherwise we add the missing edge to which nothing can be mapped. We assume that  $G_S$  is rooted at some arbitrary node to avoid ambiguity in the following construction about whether a neighbor is a child or the parent. We construct  $\tilde{G}_S$  from  $G_S$  by adding a node  $v'$  and edges  $(v, v')$ ,  $(v', v)$  for each node  $v \in V_S$  which has children and set  $\tilde{d}_S(v) := 0$ ,  $\tilde{d}_S(v') := d_S(v)$ ,  $\tilde{c}_S(v') := c_S(v)$ ,  $\tilde{d}_S(v, v') := \tilde{d}_S(v', v) := \infty$ , and  $\tilde{c}_S(\{v, v'\}) := \tilde{c}_S(v, v) := 0$ . Note that we add at most  $O(|V_S|)$  nodes and edges to  $G_S$ . Hence,  $\tilde{I}$  can be constructed after linear time. We now show that  $I$  has a feasible mapping  $(m_R^V, m_R^E)$  of cost at most  $k$  if and only if  $\tilde{I}$  has a feasible mapping  $(\tilde{m}_R^V, \tilde{m}_R^E)$  of cost at most  $k$ .

( $\Rightarrow$ ): Let  $(m_R^V, m_R^E)$  be a solution for  $I$  of cost at most  $k$ . For all  $v \in V_R$ , we set  $\tilde{m}_R^V(v) := m_R^V(v)$  if  $m_R^V(v)$  is of degree at most one, otherwise we set  $\tilde{m}_R^V(v)$  to the new leaf  $m_R^V(v)'$  of  $m_R^V(v)$ . For all  $(u, v) \in E_R$ , we set  $\tilde{m}_R^E(u, v)$  to be the unique path from  $\tilde{m}_R^V(u)$  to  $\tilde{m}_R^V(v)$  in  $\tilde{G}_S$ . Then  $(\tilde{m}_R^V, \tilde{m}_R^E)$  has the same cost as  $(m_R^V, m_R^E)$ .

---

**Algorithm 3.1:** Algorithm for VNEP on tree substrates

---

```

1 Function Leaf( $v \in V_S$ ): // see Lemma 1
2   Add node  $v'$  to  $G_S$  as a child of  $v$ .
3    $d_S(v') \leftarrow d_S(v)$ ,  $c_S(v') \leftarrow c_S(v)$ .
4    $d_S(v) \leftarrow 0$ ,  $c_S(v) \leftarrow \infty$ .
5    $d_S(v, v')$ ,  $d_S(v', v) \leftarrow \infty$ ,  $c_S(v, v')$ ,  $c_S(v', v) \leftarrow 0$ .
6 Function Split( $v \in V_S$ ): // see Lemma 2
7   Let  $u_1, \dots, u_t$  be the children of  $v$ , let  $s = \lfloor t/2 \rfloor$ .
8   Add nodes  $v_\ell, v_r$  to  $G_S$ , with  $d_S(v_\ell)$ ,  $d_S(v_r) \leftarrow 0$ 
   and  $c_S(v_\ell)$ ,  $c_S(v_r) \leftarrow \infty$ .
9   Make  $v_\ell$  parent of  $u_1, \dots, u_s$  (keep capacities and costs).
10  Make  $v_r$  parent of  $u_{s+1}, \dots, u_t$  (keep capacities and costs).
11  Make  $v$  parent of  $v_\ell, v_r$ .
12   $d_S(v, v_r)$ ,  $d_S(v_r, v)$ ,  $d_S(v, v_\ell)$ ,  $d_S(v_\ell, v) \leftarrow \infty$ .
13   $c_S(v, v_r)$ ,  $c_S(v_r, v)$ ,  $c_S(v, v_\ell)$ ,  $c_S(v_\ell, v) \leftarrow 0$ .
14  if  $v_\ell$  has more than 2 children then call Split( $v_\ell$ )
15  if  $v_r$  has more than 2 children then call Split( $v_r$ )
16 Function GetEntry( $R \subseteq V_R, v \in V_S$ ):
   // returns the of entry in  $D$ , or computes it
17 if  $D[R, v]$  was already computed then return  $D[R, v]$ .
18 if  $v$  is a leaf then
19    $D[R, v] \leftarrow \begin{cases} \infty, & \text{if } \sum_{u \in R} d_R(u) \not\leq d_S(v), \\ \sum_{u \in R} d_R(u)^\top c_S(v), & \text{otherwise.} \end{cases}$ 
20 else if  $v$  has one child  $u$  then  $D[R, v] \leftarrow f(v, u, R)$ .
21 else if  $v$  has two children  $u$  and  $w$  then
22    $D[R, v] \leftarrow \min_{A \cup B = R} f(v, w, A) + f(v, u, B)$ .
   // Use  $f$  as defined in (3), but
   replace  $D[R, x]$  with GetEntry( $R, x$ ).
23 return  $D[R, v]$  (and mark it as computed).
24 Main Procedure ( $G_S, G_R, d_R, d_S, c_S$ ):
25   Let  $G_S$  be rooted at some node  $p$ .
26   for  $v \in V(G_S)$  do
27     if  $v$  is not a leaf and  $d_S(v) > 0$  then call Leaf( $v$ )
28   for  $v \in V(G_S)$  do
29     if  $v$  has more than two children then call Split( $v$ )
30   Initialize table  $D[R, v]$  for all  $R \subseteq V_R$  and  $v \in V_S$ .
31   return GetEntry( $V_R, p$ ).

```

---

( $\Leftarrow$ ): Let  $(\tilde{m}_R^V, \tilde{m}_R^E)$  be a solution for  $\tilde{I}$  of cost at most  $k$ . For all  $v \in V_R$ , we set  $m_R^V(v) := \tilde{m}_R^V(v)$  if  $\tilde{m}_R^V(v) \in V_S$ , otherwise  $\tilde{m}_R^V(v)$  is a leaf in  $\tilde{G}_S$  and we set  $m_R^V(v)$  to be the parent of  $\tilde{m}_R^V(v)$ . For all  $(u, v) \in E_R$  we set  $m_R^E(u, v)$  to be the unique path from  $u$  to  $v$  in  $G_S$ . Now the paths induced by  $m_R^E(u, v)$  and by  $\tilde{m}_R^E(u, v)$  may only differ in the leaves that were (possibly) added to the endpoints. Thus, by construction,  $(m_R^V, m_R^E)$  is a solution for  $I$  of cost at most  $k$ .  $\square$

Next, we show how to turn the substrate into a binary tree.

**Lemma 2.** *Given an instance  $I = (G_S, G_R, d_R, d_S, c_S)$  of MIN-VNEP with  $G_S$  being a tree, we can construct in linear time an instance  $\tilde{I} = (\tilde{G}_S, G_R, d_R, \tilde{d}_S, \tilde{c}_S)$  of MIN-VNEP such that  $\tilde{G}_S$  is a binary tree*

and there is a solution for  $I$  of cost at most  $k$  if and only if there is a solution for  $\tilde{I}$  of cost at most  $k$ .

**PROOF.** In a nutshell, we are going to replace a node with more than two children with a binary tree of sufficient size.

To construct  $\tilde{G}_S = (\tilde{V}_S, \tilde{E}_S)$  from  $G_S$ , as long as there is a node  $v$  with  $c > 2$  children, we replace it with a fresh rooted bidirectional binary tree  $T_v$  with root  $v'$  and  $c$  leaves. We add an edge between  $v'$  and the parent of  $v$ , and we add an edge between each child of  $v$  and one designated leaf of  $T_v$ . We set the capacity and cost of the root of  $T_v$  to  $d_S(v)$  and  $c_S(v)$ , respectively. All other nodes of  $T_v$  get capacity zero and cost  $k + 1$ . All edges of  $T_v$  get capacity  $\infty$  and cost zero. Let  $v_p$  be the parent of  $v$ , let  $u$  be a child of  $v$ , and let  $v_u$  be the leaf node in  $T_v$  which is adjacent to  $u$ . Then, we set  $\tilde{d}_S(v', v_p) := d_S(v, v_p)$ ,  $\tilde{d}_S(v_p, v') := d_S(v_p, v)$ ,  $\tilde{c}_S(v', v_p) := c_S(v, v_p)$ ,  $\tilde{c}_S(v_p, v') := c_S(v_p, v)$ ,  $\tilde{d}_S(v_u, u) := d_S(v, u)$ ,  $\tilde{d}_S(u, v_u) := d_S(u, v)$ ,  $\tilde{c}_S(v_u, u) := c_S(v, u)$ , and  $\tilde{c}_S(u, v_u) := c_S(u, v)$ . All other values of  $\tilde{d}_S$  and  $\tilde{c}_S$  are equal to  $d_S$  and  $c_S$ , respectively. Note that by the handshake lemma (the sum of degrees of is twice the number of edges in a graph),  $\tilde{G}_S$  is of size  $O(|V_S|)$ , because  $T_v$  is of size  $O(\deg_{G_S}(v))$ . Hence, we can construct  $\tilde{I}$  in linear time.

We show that  $I$  is feasible if and only if  $\tilde{I}$  is feasible.

( $\Rightarrow$ ): Let  $(m_R^V, m_R^E)$  be a feasible mapping for  $I$  of cost at most  $k$ . For all  $v \in V_R$ , we set  $\tilde{m}_R^V(v) := m_R^V(v)$  if  $m_R^V(v) \in \tilde{V}_S$ , otherwise we set  $\tilde{m}_R^V(v)$  to be the root of  $T_v$ . Hence, for all  $v \in \tilde{V}_S$  we have  $\sum_{w: \tilde{m}_R^V(w)=v} d_R(w) \leq \tilde{d}_S(v)$ . For all  $(u, v) \in E_R$ , set  $\tilde{m}_R^E(u, v)$  to be the unique path from  $\tilde{m}_R^V(u)$  to  $\tilde{m}_R^V(v)$  in  $\tilde{G}_S$  (recall that  $\tilde{G}_S$  is a tree). So for all  $e_S \in \tilde{E}_S$  we have  $\sum_{e_R: e_S \in \tilde{m}_R^E(e_R)} d_R(e_R) \leq \tilde{d}_S(e_S)$  and for all  $(u, v) \in E_R$ ,  $\tilde{m}_R^E(u, v)$  is a path from  $\tilde{m}_R^V(u)$  to  $\tilde{m}_R^V(v)$ . Moreover, by our construction, we get that

$$\begin{aligned} & \sum_{v \in V_R} d_R(v)^\top \tilde{c}_S(\tilde{m}_R^V(v)) + \sum_{e \in E_R} \left( \sum_{e' \in \tilde{m}_R^E(e)} d_R(e)^\top \tilde{c}_S(e') \right) \\ &= \sum_{v \in V_R} d_R(v)^\top c_S(m_R^V(v)) + \sum_{e \in E_R} \left( \sum_{e' \in m_R^E(e)} d_R(e)^\top c_S(e') \right) \leq k. \end{aligned}$$

Thus,  $(\tilde{m}_R^V, \tilde{m}_R^E)$  is a feasible mapping for  $\tilde{I}$  of cost at most  $k$ .

( $\Leftarrow$ ): Let  $(\tilde{m}_R^V, \tilde{m}_R^E)$  be a feasible mapping for  $\tilde{I}$  of cost at most  $k$ . Let  $v \in V_R$ . Note that if  $\tilde{m}_R^V(v) \notin V_S$ , then there must be a node  $w \in V_S$  such that  $\tilde{m}_R^V(v)$  is a node in  $T_w$ . Hence, we set  $m_R^V(v) := \tilde{m}_R^V(v)$  if  $\tilde{m}_R^V(v) \in V_S$ , otherwise we set  $m_R^V(v) := w$ , where  $w \in V_S$  is the node replaced by  $T_w$  and  $\tilde{m}_R^V(v)$  is a node of  $T_w$ . So, for all  $v \in V_S$ , we have  $\sum_{w: m_R^V(w)=v} d_R(w) \leq d_S(v)$ . For all  $(u, v) \in E_R$  we set  $m_R^E(u, v)$  to be the unique path in  $G_S$  from  $m_R^V(u)$  to  $m_R^V(v)$ . Note that the path induced by  $\tilde{m}_R^E(u, v)$  consists of a subset of edges of  $m_R^E(u, v)$ ; thus for all  $e_S \in E_S$  we have  $\sum_{e_R: e_S \in m_R^E(e_R)} d_R(e_R) \leq d_S(e_S)$ . Moreover, we have that

$$\begin{aligned} & \sum_{v \in V_R} d_R(v)^\top c_S(m_R^V(v)) + \sum_{e \in E_R} \left( \sum_{e' \in m_R^E(e)} d_R(e)^\top c_S(e') \right) \\ &= \sum_{v \in V_R} d_R(v)^\top \tilde{c}_S(\tilde{m}_R^V(v)) + \sum_{e \in E_R} \left( \sum_{e' \in E(\tilde{m}_R^E(e))} d_R(e)^\top \tilde{c}_S(e') \right) \leq k. \end{aligned}$$

Thus,  $(m_R^V, m_R^E)$  is a feasible mapping for  $I$  of cost at most  $k$ .  $\square$

Now that we have created an instance in which the substrate is a binary tree in which only the leaf nodes have nonzero capacity, we can formulate our dynamic program. Let  $p$  be the root of  $G_S$ . For each  $v \in V_S$ , let  $T_v$  be the induced subtree of  $G_S$  where  $v$  is the root. We assume that  $G_S$  is a full binary tree, i.e., each node is either a leaf or has two children (otherwise we add a fresh leaf to which nothing can be mapped).

Removing the edges  $(v, u), (u, v) \in E_S$  splits the tree  $G_S$  into two rooted trees. Without loss of generality assume that  $v$  is the parent of  $u$  in  $G_S$ . Hence, one of the trees is  $T_u$  and the other one is  $T' := G_S[V_S \setminus V(T_u)]$ . Note that for a given solution  $(m_R^V, m_R^E)$  of  $I$ , the cut  $\{(v, u), (u, v)\}$  also splits the mapping of  $G_R$  into two parts  $B := \{w \in V_R \mid m_R^V(w) \in V(T_u)\}$  and  $A := V_R \setminus B$ . Further, for each edge  $e \in E_R$  we have that  $(v, u) \in m_R^E(e)$  if and only if  $e \in \text{cut}_{G_R}(A) := \{(x, y) \in E_R \mid x \in A, y \notin A\}$ , and moreover  $(u, v) \in m_R^E(e)$  if and only if  $e \in \text{cut}_{G_R}^-(A) := \text{cut}_{G_R}(V_R \setminus A)$ , since every path from  $T'$  to  $T_u$  must contain  $(v, u)$  and every path from  $T_u$  to  $T'$  must contain  $(u, v)$ . We use this observation to describe a dynamic program in which each entry  $D[R, v]$  contains the minimum cost for a feasible mapping of  $G_R[R]$  into  $T_v$  plus the induced cost of  $\text{cut}_{G_R}(A) \cup \text{cut}_{G_R}^-(A)$  on edges in  $T_v$ .

Let  $v \in V_S$  and  $R \subseteq V_R$ . If  $v$  is a leaf, then

$$D[R, v] := \begin{cases} \infty, & \text{if } \sum_{u \in R} d_R(u) \not\leq d_S(v) \\ \sum_{u \in R} d_R(u)^\top c_S(v), & \text{otherwise.} \end{cases} \quad (1)$$

If  $v$  is not a leaf, then

$$D[R, v] := \min_{A \uplus B = R} f(v, w, A) + f(v, u, B), \quad (2)$$

where  $w$  and  $u$  are the neighbors of  $v$  in  $T_v$  and for  $x \in \{w, u\}$  the function  $f$  is defined as  $f(v, x, R)$

$$:= \begin{cases} \infty, & \text{if } \sum_{e \in \text{cut}_{G_R}(R)} d_R(e) \not\leq d_S(v, x), \\ \infty, & \text{if } \sum_{e \in \text{cut}_{G_R}^-(R)} d_R(e) \not\leq d_S(v, x), \\ D[R, x] + \sum_{e \in \text{cut}_{G_R}^-(R) \cup \text{cut}_{G_R}(R)} d_R(e)^\top c_S(v, x), & \text{otherwise.} \end{cases} \quad (3)$$

To show the correctness of the dynamic program (defined in (1) and (2)), we introduce the following notations and definitions. For  $v \in V_S$ , for  $(x, y) \in E(T_v)$ , for  $X \subseteq V_R$ , and  $m_R^V: X \rightarrow V(T_v)$ , let  $\mathcal{P}_{(x,y)}^v(X)$  be the set of paths  $P$  within  $T_v$  between  $v$  and a node  $m_R^V(u^*)$  such that  $(x, y)$  is in  $P$ , and if  $v$  is the start node of  $P$ ,  $u^* \in X$  is the sink of an edge in  $\text{cut}_{G_R}^-(X)$ , otherwise  $u^* \in X$  is the source of an edge in  $\text{cut}_{G_R}(X)$ . Furthermore, let  $E_{(x,y)}^v(X) :=$

$$\begin{aligned} & \{(u^*, w^*) \in \text{cut}_{G_R}(X) \mid (x, y) \text{ is on the } m_R^V(u^*)\text{-}v\text{-path in } T_v\} \\ & \cup \{(w^*, u^*) \in \text{cut}_{G_R}^-(X) \mid (x, y) \text{ is on the } v\text{-}m_R^V(u^*)\text{-path in } T_v\}. \end{aligned}$$

**Definition 3.** For a node  $v \in V_S$  and a subset  $X \subseteq V_R$ . We call a feasible mapping  $(m_R^V, m_R^E)$  of  $G_R[X]$  to  $T_v$  *desirable* if for every

edge  $e_S \in E(T_v)$  we have

$$\sum_{e_R: e_S \in m_R^E(e_R)} d_R(e_R) \leq d_S(e_S) - \sum_{e \in E_{e_S}^v(X)} d_R(e). \quad (4)$$

Furthermore, we say that the *induced cost* of  $(m_R^V, m_R^E)$  is

$$\begin{aligned} & \sum_{w \in X} d_R(w)^\top c_S(m_R^V(w)) + \sum_{e \in E(G_R[X])} \left( \sum_{e' \in m_R^E(e)} d_R(e)^\top c_S(e') \right) + \\ & \sum_{e \in \text{cut}_{G_R}^+(X)} \left( \sum_{e' \in P_e} d_R(e)^\top c_S(e') \right) + \sum_{e \in \text{cut}_{G_R}^-(X)} \left( \sum_{e' \in P_e^-} d_R(e)^\top c_S(e') \right). \end{aligned} \quad (5)$$

Here  $P_e$  is the set of edges of the path from the source of  $e$  to  $v$  in  $T_v$  and  $P_e^-$  is the set of edges of the path from  $v$  to the target of  $e$  in  $T_v$ .

Later, our algorithm will report that the minimum cost for a solution is  $D[V_R, p]$ . We show that indeed there is such a solution.

**Lemma 3.** *Let  $v \in V_S$  and  $X \subseteq V_R$ . If  $D[X, v] < \infty$ , then there is a desirable mapping  $(m_R^V, m_R^E)$  of  $G_R[X]$  onto  $T_v$  where the induced cost is at most  $D[X, v]$ .*

**PROOF.** We show this by induction over the tree  $G_S$ . By the definition in (1), every mapping of  $G_R[X]$ ,  $X \subseteq V_R$ , onto a leaf  $v \in V_S$  is desirable and has induced costs of  $D[X, v]$ .

For the induction step, let  $v \in V_S$  be a non-leaf, let  $X \subseteq V_R$ , and assume that for all  $u \in V(T_v) \setminus \{v\}$  we have that if  $D[Y, u] < \infty$ . Then there is a desirable mapping of  $G_R[Y]$  onto  $T_u$  with induced cost of at most  $D[Y, u]$ . Assume further that  $D[X, v] < \infty$ , and let  $a$  and  $b$  be the children of  $v$ . Then by the definition in (2) there is a partition  $A \uplus B = X$  such that  $D[X, v] = D[A, a] + D[B, b] +$

$$\begin{aligned} & \sum_{e \in \text{cut}_{G_R}^+(A)} d_R(e)^\top c_S(v, a) + \sum_{e \in \text{cut}_{G_R}^-(A)} d_R(e)^\top c_S(a, v) \\ & + \sum_{e \in \text{cut}_{G_R}^+(B)} d_R(e)^\top c_S(v, b) + \sum_{e \in \text{cut}_{G_R}^-(B)} d_R(e)^\top c_S(b, v). \end{aligned} \quad (6)$$

Thus,  $D[A, a] < \infty$  and  $D[B, b] < \infty$ , and we get by assumption that there are desirable mappings  $(m_R^V, m_R^E)$  and  $(m_R^V, m_R^E)$  of  $G_R[A]$  onto  $T_a$  and of  $G_R[B]$  onto  $T_b$ , respectively.

We create a mapping  $(m_R^V, m_R^E)$  of  $G_R[X]$  onto  $T_v$  with

$$\begin{aligned} m_R^V(x) &:= \begin{cases} m_R^V(x), & x \in A, \\ m_R^V(x), & x \in B, \end{cases} & \text{and} \\ m_R^E(x, y) &:= \begin{cases} m_R^E(x, y), & x, y \in A, \\ m_R^E(x, y), & x, y \in B, \\ \text{path from } m_R^V(x) \text{ to } m_R^V(y) \text{ in } T_v, & x \in A, y \in B, \\ \text{path from } m_R^V(x) \text{ to } m_R^V(y) \text{ in } T_v, & x \in B, y \in A. \end{cases} \end{aligned} \quad (7)$$

Observe that  $(m_R^V, m_R^E)$  is a feasible mapping of  $G_R[X]$  onto  $T_v$ : Let  $(x, y)$  be an edge in  $E(G_R[X])$  such that one endpoint is in  $A$  and the other endpoint is in  $B$ . Then every edge in  $T_v$  that is on a path from  $m_R^V(x)$  to  $m_R^V(y)$  has sufficient capacity to map all edges of  $m_R^E(x, y)$  as defined in (7). Hence, (4) for  $(m_R^V, m_R^E)$  and  $(m_R^V, m_R^E)$  implies that for every edge  $e_S \in E(T_v)$  we have

$$\sum_{e_R: e_S \in m_R^E(e_R)} d_R(e_R) \leq d_S(e_S) - \sum_{e \in E_{e_S}^v(X)} d_R(e).$$

Moreover, for all  $c \in \{a, b\}$ , a path from a node in  $V(T_c)$  to  $v$  contains the edge  $(v, c)$  and a path from  $v$  to some node in  $V(T_c)$  contains the edge  $(c, v)$ . Hence, the induced cost of  $(m_R^V, m_R^E)$  is the sum of the induced cost of  $(m_R^V, m_R^E)$  and  $(m_R^V, m_R^E)$  and

$$\begin{aligned} & \sum_{e \in \text{cut}_{G_R}^+(A)} d_R(e)^\top c_S(v, a) + \sum_{e \in \text{cut}_{G_R}^-(A)} d_R(e)^\top c_S(a, v) + \\ & \sum_{e \in \text{cut}_{G_R}^+(B)} d_R(e)^\top c_S(v, b) + \sum_{e \in \text{cut}_{G_R}^-(B)} d_R(e)^\top c_S(b, v). \end{aligned}$$

Thus, by (6) the induced cost of  $(m_R^V, m_R^E)$  is at most  $D[X, v]$ , because the induced cost of  $(m_R^V, m_R^E)$  is at most  $D[A, a]$  and the induced cost of  $(m_R^V, m_R^E)$  is at most  $D[B, b]$ .

Finally, since  $D[X, v] < \infty$  we get by (2) that (4) holds for  $(m_R^V, m_R^E)$  as well. Thus,  $(m_R^V, m_R^E)$  is a desirable mapping of  $G_R[X]$  onto  $T_v$  of induced cost at most  $D[X, v]$ , and we are done.  $\square$

Moreover, we also need to show that if there is feasible mapping for  $I$  of cost at most  $k$ , then  $D[V_R, p] \leq k$ . More formally, we show:

**Lemma 4.** *Let  $v \in V_S$  and  $(m_R^V, m_R^E)$  be a feasible mapping for  $I$  of cost at most  $k$ . Then,*

$$\begin{aligned} D[X, v] &\leq \sum_{w \in X} d_R(w)^\top c_S(m_R^V(w)) \\ &+ \sum_{e \in \text{cut}_{G_R}^+(X) \cup \text{cut}_{G_R}^-(X) \cup E(G_R[X])} \left( \sum_{e' \in m_R^E(e) \cap E(T_u)} d_R(e)^\top c_S(e') \right), \end{aligned}$$

where  $X := \{w \in V_R \mid m_R^V(w) \in V(T_v)\}$ .

**PROOF.** We show the statement of the lemma by structural induction over the tree  $G_S$ . By the definition in (1), this is true for all leaves  $v \in G_S$  as  $E(T_v) = \emptyset$ . For the induction step let  $v \in V_S$  be a non-leaf node, let  $X := \{w \in V_R \mid m_R^V(w) \in V(T_v)\}$ , and assume that for all  $u \in V(T_v) \setminus \{v\}$  we have

$$\begin{aligned} D[Y, u] &\leq \sum_{w \in Y} d_R(w)^\top c_S(m_R^V(w)) + \\ & \sum_{e \in \text{cut}_{G_R}^+(Y) \cup \text{cut}_{G_R}^-(Y) \cup E(G_R[Y])} \left( \sum_{e' \in m_R^E(e) \cap E(T_u)} d_R(e)^\top c_S(e') \right), \end{aligned} \quad (8)$$

where  $Y := \{w \in V_R \mid m_R^V(w) \in V(T_u)\}$ . Let  $a$  and  $b$  be the children of  $v$ , and let  $A := \{w \in V_R \mid m_R^V(w) \in V(T_a)\}$  and  $B := \{w \in V_R \mid m_R^V(w) \in V(T_b)\}$ . Node  $v$  is not a leaf; thus  $d_S(v) = 0$ , so no node of  $G_R$  is mapped onto  $v$ . By (2) we obtain  $D[X, v] \leq$

$$\begin{aligned} & D[A, a] + \sum_{e \in \text{cut}_{G_R}^+(A)} d_R(e)^\top c_S(v, a) + \sum_{e \in \text{cut}_{G_R}^-(A)} d_R(e)^\top c_S(a, v) + \\ & D[B, b] + \sum_{e \in \text{cut}_{G_R}^+(B)} d_R(e)^\top c_S(v, b) + \sum_{e \in \text{cut}_{G_R}^-(B)} d_R(e)^\top c_S(b, v). \end{aligned}$$

By assumption, (8) holds for  $D[A, a]$  and  $D[B, b]$ ; so

$$\begin{aligned} D[X, v] &\leq \sum_{w \in A \cup B} d_R(w)^\top c_S(m_R^V(w)) + \\ & \sum_{e \in \text{cut}_{G_R}^+(A) \cup \text{cut}_{G_R}^-(A) \cup E(G_R[A])} \left( \sum_{e' \in m_R^E(e) \cap E(T_a)} d_R(e)^\top c_S(e') \right) + \\ & \sum_{e \in \text{cut}_{G_R}^+(B) \cup \text{cut}_{G_R}^-(B) \cup E(G_R[B])} \left( \sum_{e' \in m_R^E(e) \cap E(T_b)} d_R(e)^\top c_S(e') \right) + \end{aligned}$$

$$\sum_{e \in \text{cut}_{G_R}^-(A)} d_R(e)^\top c_S(v, a) + \sum_{e \in \text{cut}_{G_R}(A)} d_R(e)^\top c_S(a, v) + \sum_{e \in \text{cut}_{G_R}^-(B)} d_R(e)^\top c_S(v, b) + \sum_{e \in \text{cut}_{G_R}(B)} d_R(e)^\top c_S(b, v).$$

Note that every path from a node in  $T_a$  ( $T_b$ ) to a node in  $T_b$  ( $T_a$ ) contains the edges  $(a, v)$ ,  $(v, b)$  ( $(b, v)$ ,  $(v, a)$ ). Moreover, for  $c \in \{a, b\}$  every path from  $T_c$  to some node in  $G_S - V(T_v)$  contains the edge  $(c, v)$  and every path from  $G_S - V(T_v)$  to some node in  $T_c$  contains the edge  $(v, c)$ . Hence, we obtain

$$D[X, v] \leq \sum_{w \in A \cup B} d_R(w)^\top c_S(m_R^V(w)) + \sum_{e \in E(G_R[X])} \left( \sum_{e' \in m_R^E(e) \cap E(T_v)} d_R(e)^\top c_S(e') \right) + \sum_{e \in \text{cut}_{G_R}(X) \cup \text{cut}_{G_R}^-(X)} \left( \sum_{e' \in m_R^E(e) \cap E(T_v)} d_R(e)^\top c_S(e') \right). \quad \square$$

Now we have everything at hand to prove Theorem 1.

**PROOF OF THEOREM 1.** Let  $I = (G_S, G_R, d_R, d_S, c_S)$  be some instance of MIN-VNEP. By Lemmata 1 and 2 we can assume that  $G_S$  is a binary tree rooted at some arbitrary node  $p$  and each node  $v \in V_S$  with degree at least two fulfills  $d_S(v) = 0$ . We apply the dynamic program stated in (1) and (2). Since  $G_S = T_p$  and  $\text{cut}_{G_R}(V_R) = \emptyset$ , Lemmata 3 and 4 imply that  $D[V_R, p]$  contains the minimum cost for a feasible mapping for  $I$ , where  $D[V_R, p] = \infty$  if and only if there is no feasible mapping for  $I$ .

Let  $r := |V_R|$ . It remains to be shown that  $D[V_R, p]$  can be computed in  $O(3^r(|V_S| + r^2))$  time. We first compute for every  $A \subseteq V_R$  the demand of the cut  $\text{cut}_{G_R}(A)$ . There are  $2^r$  subsets  $A$ , for each of which we need to iterate over the  $O(r^2)$  edges; thus this step takes  $O(2^r \cdot r^2)$  time. With this at hand we can compute  $D[X, v]$  in constant time for each leaf  $v \in V_S$  and for each subset  $X \subseteq V_R$ . For a non-leaf node  $v$ , computing the entries  $D[X, v]$  for each  $X \subseteq V_R$  can be done in  $O(3^r)$  operations: For a partition  $X = A \uplus B$  we require constant time. Observe that there are  $3^r$  partitions of  $V_R$  into three parts  $A \uplus B \uplus C$ . Thus, choosing  $X = V_R \setminus C$  gives us all partitions of all subsets  $X \subseteq V_R$  into two parts  $A$  and  $B$ . Thus, for all non-leaf nodes  $v$  and all subsets  $X \subseteq V_R$  combined we require  $O(3^r \cdot |V_S|)$  time. Altogether, this yields the claimed running time of  $O(3^r \cdot (|V_S| + r^2))$ .  $\square$

As a final note, we highlight that our dynamic program is rather simple to implement and robust in the sense that it also works if one has further natural constraints or other objectives.

## 4 EVALUATION

We evaluate the performance of our exact dynamic programming algorithm for tree substrates (presented in Section 3 and henceforth abbreviated with DP) on common fat tree topologies as they are widely deployed, e.g., in data centers [1]. Specifically, we compare the performance of our algorithm with two well-established approaches for solving the VNEP. The first is the standard integer programming formulation (IP) which gives exact results. The second is the ViNE heuristic by Chowdhury et al. [5], which takes the relaxation of an IP formulation and then applies randomized rounding to fix node mappings and realizes edges via shortest paths. In our comparisons the focus is on the running time and the solution quality of the three approaches. Since the running time of the IP

may take hours for medium-sized instances, we set a time limit on the IP, and we also report on the quality of the sub-optimal solutions obtained by the IP when the imposed time limit was reached. Recall that the solution obtained by our DP is *always* optimal.

*Testing Methodology.* For our evaluation, we employ fat trees [1] as our substrate network topology. Fat trees are common topologies, e.g., in data centers built using commodity switches, where each switch has the same number  $f \geq 4$  of ports. Fat trees are highly structured: servers are located at the bottom and are connected by a three-layer hierarchy of switches (see Figure 1). A fat tree constructed of  $f$ -port switches connects up to  $f^3/4$  servers. While the actual physical infrastructure is not a tree, the *forwarding abstraction* provided by fat trees is a tree. Specifically, based on link aggregation techniques [25], switches and their interconnections are logically aggregated from an application-level perspective. Hence, embeddings can and must be computed on this tree forwarding abstraction. Note that MIN-VNEP is clearly NP-hard on such trees (see Theorem 2).

We consider seven different fat tree forwarding abstractions for  $f \in \{4, 6, \dots, 16\}$ , hosting between 16 and 1024 servers and using between 5 and 145 switches. Considering a single node resource type, we set the computational capacities on servers to 1 and on switches to 0. For edges of the bottom layer, i.e., connecting to servers we set a bandwidth of one. Due to the aggregation of edges, the edge bandwidth of the above layers is set accordingly to  $f/2$  and  $(f/2)^2$ . To simulate heterogeneous usage patterns within the data center, we perturb node and edge capacities by random factors drawn from the interval  $[1, 10]$  and draw costs from  $[1, 10]$ .

For generating requests, we follow the standard approach of sampling Erdős-Rényi-topologies of various sizes [5, 7]. In this model, for a specific number of nodes, edges between pairs of nodes are created probabilistically using a connection probability  $p$ . This approach is attractive, as it does not impose assumptions on the applications modeled by the requests albeit allowing to easily vary the interconnection density. Again, following the standard evaluation methodology [5, 7], node and edge demands are also sampled uniformly at random. Specifically, node demands are drawn from the interval  $[1, 5]$ . For edge demands, we proceed as follows. For each node, we draw the total cumulative outgoing bandwidth from  $[1, 5]$  and then distribute the bandwidth randomly across the actual edges. By this construction, the expected total bandwidth (per request size) is independent from the connection probability  $p$ .

For our evaluation we focus on requests of 5 to 12 nodes and consider ten different connection probabilities  $p \in \{0.1, 0.2, \dots, 1.0\}$  (disconnected graphs are discarded and resampled). For each combination of graph size and connection probability, we sample ten instances. Together with the 7 different fat tree topologies, our computational study encompasses 5.6k instances.

*Computational Setup.* We first discuss the implementation of our dynamic program (DP), the integer programming (IP), and ViNE.

We have implemented the dynamic program presented in Section 3 in C++ using only the standard library. While implemented for single node and edge resources, our implementation can be easily extended to an arbitrary number of resources. Our implementation is tweaked to skip computations that involve table entries

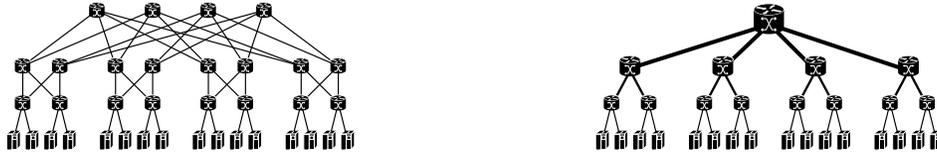


Figure 1: Fat tree topology [1] constructed using  $f = 4$ -port switches (left) and corresponding forwarding abstraction (right).

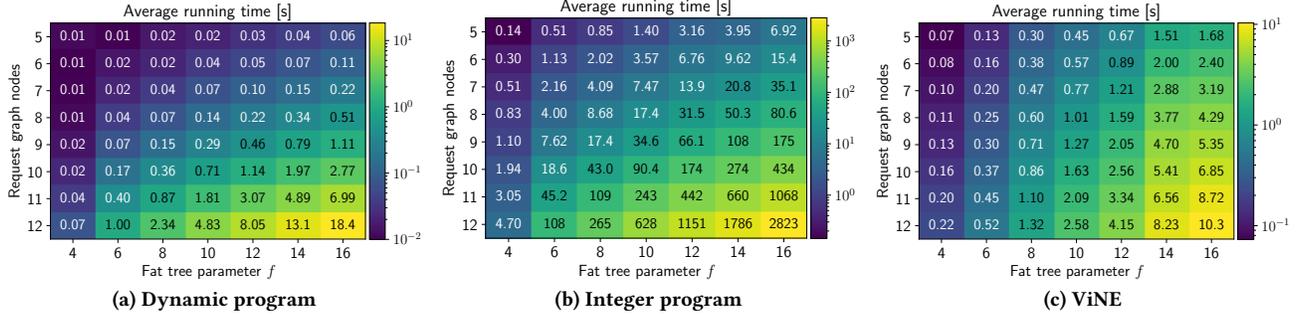


Figure 2: Running time statistics in seconds. Each heatmap cell averages 100 instances of different Erdős-Rényi request graphs, 10 for each connection probability  $p \in \{0.1, \dots, 1.0\}$ . Recall that for the integer program the time limit is set to  $200 \times$  the dynamic program's running time. Note the different (logarithmic) z-axes.

containing  $\infty$ , as these cannot lead to a feasible solution. Furthermore we do not store table entries that contain  $\infty$ . To facilitate this, we store the table entries for a node  $v$  as a set-trie, rather than a simple array, to allow for fast subset and superset queries. During our experiments we discovered that on instances with 12-node requests this tweak resulted in a decrease of 90% in table size, and a corresponding drop in the running time is to be expected. The source of our implementation is available online.<sup>2</sup>

Existing exact algorithms for the VNEP in the literature are essentially all based on integer programming [7]. Especially one integer programming formulation, based on multi-commodity flows, has been studied extensively [5, 13, 20, 22]. More details on the Integer Program can be found in the full version of this paper.

To construct the integer program, we employ a simple GMPL model and translate it into an LP-file using GLPSOL. We then solve the integer program using the commercial solver Gurobi 8.1.1. We set the thread limit of Gurobi to 1, to allow for a fair comparison with the single-threaded dynamic program. As the running time of the IP drastically exceeds the running time of the DP, for each instance we employ a time limit of 200 times the running time of the DP. Notably, the time to construct the LP-files using the unoptimized GLPSOL command is not counted towards the running time of the IP, as it often exceeded it by a factor of 3 even on smaller instances.

For the ViNE baseline, we use the Python 2 implementation of Rost et al. [22] with Gurobi 8.1.1 to solve the LP relaxation. Given a solution for the LP relaxation, we try 25 times to obtain a feasible solution by randomized rounding. For more details on the ViNE heuristic, we refer to Chowdhury et al. [5] and Rost et al. [22].

**Results.** We compared the implementations on servers equipped with an Intel Xeon W-2125 4-core, 8-thread CPU clocked at 4.0 GHz and 256GB of RAM running Ubuntu 18.04. In Figure 2, the running times of our dynamic program (DP) as well as the integer program (IP) and the ViNE heuristic are depicted. The running time of the

DP increases on average by a factor of 2 to 3 with the number of nodes of the request graph. Notably, this factor lies beneath the proven factor of 3 (see Section 3), as our implementation of the DP skips some redundant computations. The running time of the IP increases exponentially as well, however due to the enforced time limit, specific growth values could not be gathered. The running time of the IP exceeds the one of the DP by at least  $10 \times$  for more than 98.5% of the instances and by at least  $100 \times$  for more than 61.4% of the instances. The DP is faster than ViNE in 85% of the instances; the running time of ViNE is better than the one of the DP whenever both the request and the substrate graphs become large.

In Figure 3a we further analyze the speedup of the DP over the IP and how it relates to the parameters that control the size of the substrate and density of the requests. It can be seen that the speedup of the DP increases for larger values of  $p$  and  $f$ . This is likely due to the fact that the number of variables in the IP is  $O(|V_S| \cdot (|V_R| + |E_R|))$ , while the running time of the DP has exponential dependence only on  $|V_R|$ . The average speedup on instances with large  $f$  and  $p$  is close to 200, meaning that almost always the  $200 \times$  time limit was reached. To better understand the impact of this premature termination, we also report on the (empirical) approximation ratio achieved by the integer program in Figure 3b: For instances that the IP could not solve exactly within the time limit, there is a substantial gap in the embedding cost. Moreover, there were 152 instances (2.7%) for which the IP could not produce an initial feasible solution within the time limit; note that the DP produced the optimal solution while being  $200 \times$  quicker. Figure 3c gives insights into the instances for which this case was encountered. One can see that the IP struggles to construct solutions for requests with high connectivity  $p$ . The peak number of instances for which the IP did not produce a solution was observed for requests of graph size 9. We believe the reason for this to be that the IP spent more time on initialization efforts, such as computing the root linear programming relaxation.

Next, we compare our DP to the ViNE heuristic in terms of approximation quality (see Figure 4). One can observe that, as opposed

<sup>2</sup><https://git.tu-berlin.de/akt-public/vnep-for-trees>

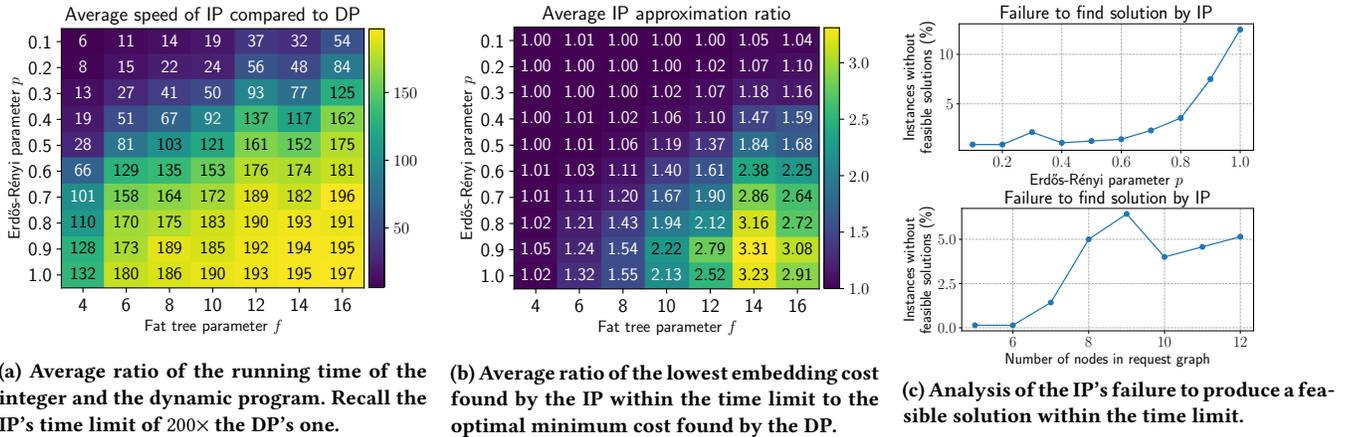


Figure 3: Comparison to the IP in terms of running time ratio, approximation ratio, and feasible solutions. Each heatmap cell averages 80 instances of Erdős-Rényi request graphs of sizes 5–12.

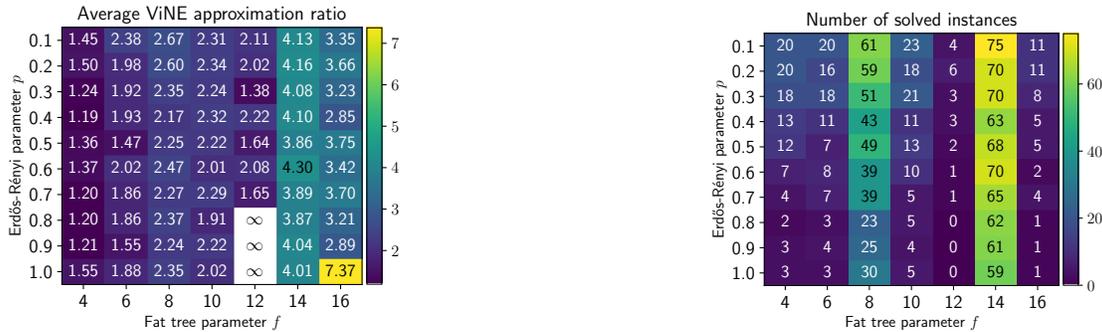


Figure 4: Comparison to ViNE in terms of approximation quality and number of solved instances. Each heatmap cell averages 80 instances of Erdős-Rényi request graphs of sizes 5–12.

to the IP, the approximation ratio of ViNE slightly improves with growing connection probability  $p$  (see Figure 4a). But with growing fat tree parameter  $f$ , the solution quality decreases, with ViNE returning a feasible solution only for very few instances (see Figure 4b). Notably, except for fat tree sizes  $f = 8$  and  $f = 14$ , ViNE finds feasible solutions for only 26% of all instances. Considering running time and approximation ratio combined we observed that there are 839 instances (15%) for which ViNE was faster than the DP. In 130 of those, ViNE found feasible solutions with an average approximation ratio of 3.64 and a speedup factor of 2.64.

*Discussion.* The above results have shown that our dynamic programming algorithm (DP) consistently outperforms the classical integer programming formulation (IP) for MIN-VNEP as well as the well-established ViNE heuristic. While the formulations of the IP and of ViNE may be improved, e.g., by exploiting the tree structure of the substrate, we believe it to be highly unlikely to be possible to close the tremendous performance gap. Accordingly, we consider the DP a valuable alternative to integer programming based algorithms as well as heuristics based on linear programming relaxations, for request graphs of small or medium size. For requests on dozens of nodes, a direct application of our DP seems prohibitive, however. Here, an interesting approach would be to reduce the size

of requests to speed up the algorithm heuristically by using clustering techniques. As already shown by Fuerst et al. [9], heuristic and optimal (pre-)clustering schemes to reduce the request size can be beneficial. Also Mano et al. [16] discuss request graph reductions and showed that the cost of embedding reduced request graphs only increases linearly while reducing the running times by exponential factors. We hence consider this an interesting avenue for developing heuristics based on the dynamic program presented in this work; in this way, one may scale beyond medium-sized requests.

## 5 CONCLUSION

We initiated the study of a parameterized algorithmic approach for the fundamental Virtual Network Embedding Problem which lies at the heart of emerging innovative network architectures that can be tailored to the application needs. In particular, we have shown that despite the general hardness of the problem, efficient and exact algorithms do exist for practically relevant scenarios. We understand our work as a first step and believe that it opens several interesting avenues for future research. In particular, it would be interesting to further investigate the power of polynomial-time data reduction through a parameterized lens, also known as kernelization in parameterized algorithmics.

## REFERENCES

- [1] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. A scalable, commodity data center network architecture. In *Proceedings of the ACM SIGCOMM Conference on Data Communication*, pages 63–74, 2008. doi: 10.1145/1402958.1402967.
- [2] 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.
- [3] Hitesh Ballani, Paolo Costa, Thomas Karagiannis, and Ant Rowstron. Towards predictable datacenter networks. In *Proceedings ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 242–253, 2011. doi: 10.1145/2018436.2018465.
- [4] Nikhil Bansal, Kang-Won Lee, Viswanath Nagarajan, and Murtaza Zafer. Minimum congestion mapping in a cloud. *SIAM Journal on Computing*, 44(3):819–843, 2015. doi: 10.1137/110845239.
- [5] Mosharaf Chowdhury, Muntasir Raihan Rahman, and Raouf Boutaba. ViNEYard: Virtual network embedding algorithms with coordinated node and link mapping. *IEEE/ACM Transactions on Networking*, 20(1):206–219, 2012. doi: 10.1109/TNET.2011.2159308.
- [6] Guy Even, Moti Medina, Gregor Schaffrath, and Stefan Schmid. Competitive and deterministic embeddings of virtual networks. *Theoretical Computer Science*, 496: 184–194, 2013. doi: 10.1016/j.tcs.2012.10.036.
- [7] Andreas Fischer, Juan F. Botero, Michael T. Beck, Hermann de Meer, and Xavier Hesselbach. Virtual network embedding: A survey. *IEEE Communications Surveys & Tutorials*, 15(4):1888–1906, 2013. doi: 10.1109/SURV.2013.013013.00155.
- [8] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Springer, 2006. doi: 10.1007/3-540-29953-X.
- [9] Carlo Fuerst, Stefan Schmid, and Anja Feldmann. Virtual network embedding with collocation: Benefits and limitations of pre-clustering. In *Proceedings of the IEEE International Conference on Cloud Networking*, pages 91–98, 2013. doi: 10.1109/CloudNet.2013.6710562.
- [10] Albert G. Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, Parveen Patel, and Sudipta Sengupta. VL2: a scalable and flexible data center network. *Communications of the ACM*, 54(3):95–104, 2011. doi: 10.1145/1897852.1897877.
- [11] Juliver Gil Herrera and Juan Felipe Botero. Resource allocation in NFV: A comprehensive survey. *IEEE Transactions on Network and Service Management*, 13(3): 518–532, 2016. doi: 10.1109/TNSM.2016.2598420.
- [12] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001. doi: 10.1006/jcss.2001.1774.
- [13] Johannes Inführ and Günther R. Raidl. Introducing the virtual network mapping problem with delay, routing and location constraints. In *Proceedings of the International Conference on Network Optimization*, pages 105–117, 2011. doi: 10.1007/978-3-642-21527-8\_14.
- [14] Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972. doi: 10.1007/978-1-4684-2001-2\_9.
- [15] Jens Lischka and Holger Karl. A virtual network mapping algorithm based on subgraph isomorphism detection. In *Proceedings of the ACM Workshop on Virtualized Infrastructure Systems and Architectures*, pages 81–88, 2009. doi: 10.1145/1592648.1592662.
- [16] Toru Mano, Takeru Inoue, Kimihiro Mizutani, and Osamu Akashi. Reducing dense virtual networks for fast embedding. *IEICE Transactions on Communications*, 103-B(4):347–362, 2020.
- [17] Xiaoqiao Meng, Vasileios Pappas, and Li Zhang. Improving the scalability of data center networks with traffic-aware virtual machine placement. In *Proceedings of the IEEE International Conference on Computer Communications*, pages 1154–1162, 2010. doi: 10.1109/INFCOM.2010.5461930.
- [18] Jeffrey C. Mogul and Lucian Popa. What we talk about when we talk about cloud network performance. *ACM SIGCOMM Computer Communication Review*, 42(5): 44–48, 2012. doi: 10.1145/2378956.2378964.
- [19] Balázs Németh, Yvonne Anne Pignolet, Matthias Rost, Stefan Schmid, and Balázs Vass. Cost-efficient embedding of virtual networks with and without routing flexibility. In *Proceedings of the IFIP Networking Conference*, pages 476–484, 2020.
- [20] Matthias Rost and Stefan Schmid. Virtual network embedding approximations: Leveraging randomized rounding. *IEEE/ACM Transactions on Networking*, 27(5): 2071–2084, 2019. doi: 10.1109/TNET.2019.2939950.
- [21] Matthias Rost and Stefan Schmid. On the hardness and inapproximability of virtual network embeddings. *IEEE/ACM Transactions on Networking*, 28(2):791–803, 2020. doi: 10.1109/TNET.2020.2975646.
- [22] Matthias Rost, Stefan Schmid, and Anja Feldmann. It’s about time: On optimal virtual network embeddings under temporal flexibilities. In *Proceedings of the International Parallel and Distributed Processing Symposium*, pages 17–26, 2014. doi: 10.1109/IPDPS.2014.14.
- [23] Matthias Rost, Carlo Fuerst, and Stefan Schmid. Beyond the stars: Revisiting virtual cluster embeddings. *ACM SIGCOMM Computer Communication Review*, 45(3):12–18, 2015. doi: 10.1145/2805789.2805792.
- [24] Matthias Rost, Elias Döhne, and Stefan Schmid. Parametrized complexity of virtual network embeddings: dynamic & linear programming approximations. *ACM SIGCOMM Computer Communication Review*, 49(1):3–10, 2019. doi: 10.1145/3314212.3314214.
- [25] C. J. Sher Decusatis, A. Carranza, and C. M. Decusatis. Communication within clouds: open standards and proprietary protocols for data center networking. *IEEE Communications Magazine*, 50(9):26–33, 2012. doi: 10.1109/MCOM.2012.6295708.
- [26] Balázs Sonkoly, Dávid Haja, Balázs Németh, Márk Szalay, János Czentye, Róbert Szabó, Rehmat Ullah, Byung-Seo Kim, and László Toka. Scalable edge cloud platforms for iot services. *Journal of Network and Computer Applications*, 170: 102785, 2020. doi: 10.1016/j.jnca.2020.102785.
- [27] Oussama Soualah, Ilhem Fajjari, Makhlof Hadji, Nadjib Aitsaadi, and Djamel Zeghlache. A novel virtual network embedding scheme based on Gomory-Hu tree within cloud’s backbone. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium*, pages 536–542, 2016. doi: 10.1109/NOMS.2016.7502855.
- [28] Gang Sun, Zhenrong Chen, Hongfang Yu, Xiaojiang Du, and Mohsen Guizani. Online parallelized service function chain orchestration in data center networks. *IEEE Access*, 7:100147–100161, 2019. doi: 10.1109/ACCESS.2019.2930295.
- [29] Ying Yuan, Cong Wang, Sancheng Peng, and Keshav Sood. Topology-oriented virtual network embedding approach for data centers. *IEEE Access*, 7:2429–2438, 2018. doi: 10.1109/ACCESS.2018.2886270.
- [30] Qi Zhang, Mohamed Faten Zhani, Maissa Jabri, and Raouf Boutaba. Venice: Reliable virtual data center embedding in clouds. In *Proceedings of the IEEE Conference on Computer Communications*, pages 289–297, 2014. doi: 10.1109/INFCOM.2014.6847950.
- [31] Yong Zhu and Mostafa H. Ammar. Algorithms for assigning substrate network resources to virtual network components. In *Proceedings of the IEEE International Conference on Computer Communications*, 2006. doi: 10.1109/INFCOM.2006.322.