

Virtual Network Embedding Approximations: Leveraging Randomized Rounding

Matthias Rost¹ Stefan Schmid²

¹ Technische Universität Berlin, Germany ² Faculty of Computer Science, University of Vienna, Austria

Abstract—The Virtual Network Embedding Problem (VNEP) captures the essence of many resource allocation problems. In the VNEP, customers request resources in the form of Virtual Networks. An embedding of a virtual network on a shared physical infrastructure is the joint mapping of (virtual) nodes to physical servers together with the mapping of (virtual) edges onto paths in the physical network connecting the respective servers. This work initiates the study of approximation algorithms for the VNEP for general request graphs. Concretely, we study the offline setting with admission control: given multiple requests, the task is to embed the most profitable subset while not exceeding resource capacities. Our approximation is based on the randomized rounding of Linear Programming (LP) solutions. Interestingly, we uncover that the standard LP formulation for the VNEP exhibits an inherent structural deficit when considering general virtual network topologies: its solutions cannot be decomposed into valid embeddings. In turn, focusing on the class of cactus request graphs, we devise a novel LP formulation, whose solutions can be decomposed. Proving performance guarantees of our rounding scheme, we obtain the first approximation algorithm for the VNEP in the resource augmentation model. We propose different types of rounding heuristics and evaluate their performance in an extensive computational study. Our results indicate that good solutions can be achieved even without resource augmentations. Specifically, heuristical rounding achieves 77.2% of the baseline’s profit on average while respecting capacities.

I. INTRODUCTION

Cloud applications usually consist of multiple distributed components (e.g., virtual machines, containers), which results in substantial communication requirements. If the provider fails to ensure that these communication requirements are met, the performance can suffer dramatically [2]. Consequently, over the last decade, several proposals have been made to *jointly* provision the computational functionality *together* with appropriate network resources. The Virtual Network Embedding Problem (VNEP) captures the core of this problem: given a graph specifying computational requirements at the nodes and bandwidth requirements on the edges, an embedding of this *Virtual Network* in the physical network has to be found, such that both the computational and the network requirements are met. Figure 1 depicts an exemplary embedding.

While the VNEP originated in the context of testbed mappings and overlay networks, several further applications have been envisioned. Among the most studied ones are service chains [5] and virtual clusters [4]. Specifically, service chains are mostly studied in the context of stitching together (virtualized) network functions, while virtual clusters have

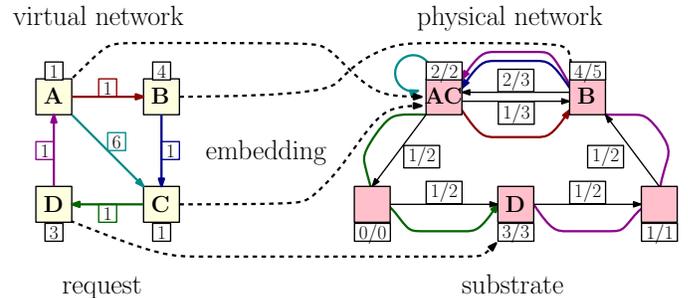


Fig. 1. Example of embedding a virtual network request (left) on the physical substrate network (right). The numeric labels of the virtual network elements denote the resource demands, while for the substrate network the resource usage and the total capacity are given. In the embedding, each request nodes is mapped to a single substrate node and each request edge is realized by a path in the substrate network. Notably, the model allows for the collocation of virtual nodes A and C on the same substrate node and, accordingly, the request edge (A, C) does not use any substrate edge resources.

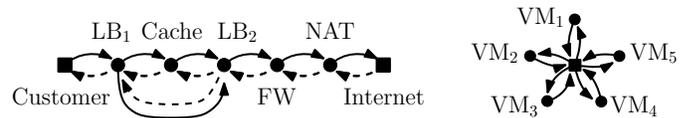


Fig. 2. Examples for virtual networks ‘in the wild’. The left graph shows a service chain for mobile operators [3]: load-balancers route (parts of the) traffic through a cache. Furthermore, a firewall and a network-address translation are used. The right graph depicts a Virtual Cluster abstraction for provisioning virtual machines (VMs) in data centers. The abstraction provides connectivity guarantees via a *logical switch* in the center [4].

been proposed in the context of batch-processing applications within data centers (see Figure 2 for examples).

We study the offline setting of the VNEP with admission control. Given are several requests, each being attributed with a profit that the provider of the physical infrastructure obtains when embedding it. The task is to decide which and how to embed requests while maximizing the obtained profit and respecting resource capacities. This work initiates the study of approximation algorithms for the VNEP for general request graphs and general physical network topologies. As the VNEP is generally \mathcal{NP} -hard and inapproximable [6], the resource augmentation model is considered, in which the approximation algorithms may exceed capacities by bounded factors. As resource augmentations imply a distinct degradation of service quality, we also derive and evaluate heuristics based on our approximation framework, which do respect capacities.

A. Formal Problem Statement

In the light of the recent interest in Service Chaining [5], we extend the VNEP’s general definition [7] by consid-

ering different *types* of computational nodes. We refer to the physical network as the *substrate network*. The substrate $G_S = (V_S, E_S)$ is offering a set \mathcal{T} of computational types. This set of types may contain, e.g., ‘FW’ (firewall), ‘x86 server’, etc. For a type $\tau \in \mathcal{T}$, the set $V_S^\tau \subseteq V_S$ denotes the substrate nodes that can host functionality of type τ . Denoting the node resources by $R_S^V = \{(\tau, u) \mid \tau \in \mathcal{T}, u \in V_S^\tau\}$ and all substrate resources by $R_S = R_S^V \cup E_S$, the capacity of nodes and edges is denoted by $d_S(x, y) > 0$ for $(x, y) \in R_S$.

For each request $r \in \mathcal{R}$, a directed graph $G_r = (V_r, E_r)$ together with a profit $b_r > 0$ is given. We refer to the respective nodes as virtual or request nodes and similarly refer to the respective edges as virtual or request edges. The types of virtual nodes are indicated by the function $\tau_r : V_r \rightarrow \mathcal{T}$.

Based on policies of the customer or the provider, the mapping of virtual node $i \in V_r$ is restricted to a set $V_S^{r,i} \subseteq V_S^{\tau_r(i)}$, while the mapping of virtual edge (i, j) is restricted to a subset of substrate edges $E_S^{r,i,j} \subseteq E_S$. Each virtual node $i \in V_r$ and each edge $(i, j) \in E_r$ is attributed with a resource demand $d_r(i) \geq 0$ and $d_r(i, j) \geq 0$, respectively. Virtual nodes and edges can only be mapped on substrate nodes and edges of sufficient capacity, i.e., $V_S^{r,i} \subseteq \{u \in V_S^{\tau_r(i)} \mid d_S(u) \geq d_r(i)\}$ and $E_S^{r,i,j} \subseteq \{(u, v) \in E_S \mid d_S(u, v) \geq d_r(i, j)\}$ holds.

We denote by $d_{\max}(r, x, y)$ the maximal demand that a request r may impose on a resource $(x, y) \in R_S$:

$$d_{\max}(r, \tau, u) = \max(\{0\} \cup \{d_r(i) \mid i \in V_r : \tau(i) = \tau \wedge u \in V_S^{r,i}\})$$

$$d_{\max}(r, u, v) = \max(\{0\} \cup \{d_r(i, j) \mid (i, j) \in E_r : (u, v) \in E_S^{r,i,j}\})$$

In the following the notions of valid mappings (respecting mapping constraints) and feasible embeddings (respecting resource capacities) are introduced to formalize the VNEP.

Definition 1 (Valid Mapping). *A valid mapping m_r of request $r \in \mathcal{R}$ is a tuple (m_r^V, m_r^E) of functions $m_r^V : V_r \rightarrow V_S$ and $m_r^E : E_r \rightarrow \mathcal{P}(E_S)$, such that the following holds:*

- *Virtual nodes are mapped to allowed substrate nodes: $m_r^V(i) \in V_S^{r,i}$ holds for all $i \in V_r$.*
- *The mapping $m_r^E(i, j)$ of virtual edge $(i, j) \in E_r$ is an edge-path connecting $m_r^V(i)$ to $m_r^V(j)$ only using allowed edges, i.e., $m_r^E(i, j) \subseteq \mathcal{P}(E_S^{r,i,j})$ holds.*

We denote by \mathcal{M}_r the set of valid mappings of request $r \in \mathcal{R}$.

Definition 2 (Allocations of Valid Mappings). *We denote by $A(m_r, x, y)$ the cumulative allocation induced by the valid mapping $m_r \in \mathcal{M}_r$ on resource $(x, y) \in R_S$:*

$$A(m_r, \tau, u) = \sum_{i \in V_r, \tau(i) = \tau, m_r^V(i) = u} d_r(i) \quad \forall (\tau, u) \in R_S^V$$

$$A(m_r, u, v) = \sum_{(i, j) \in E_r, (u, v) \in m_r^E(i, j)} d_r(i, j) \quad \forall (u, v) \in E_S$$

The maximal allocation that a valid mapping of request $r \in \mathcal{R}$ may impose on a substrate resource $(x, y) \in R_S$ is denoted by $A_{\max}(r, x, y) = \max_{m_r \in \mathcal{M}_r} A(m_r, x, y)$.

Definition 3 (Feasible Embedding). *A feasible embedding of a subset of requests $\mathcal{R}' \subseteq \mathcal{R}$ is a collection of valid mappings $\{m_r\}_{r \in \mathcal{R}'}$, such that the cumulative allocations on nodes and edges does not exceed the substrate capacities, i.e., $\sum_{r \in \mathcal{R}'} A(m_r, x, y) \leq d_S(x, y)$ holds for $(x, y) \in R_S$.*

Definition 4 (Offline Virtual Network Embedding Problem). *The VNEP asks for a feasible embedding $\{m_r\}_{r \in \mathcal{R}'}$ of a subset of requests $\mathcal{R}' \subseteq \mathcal{R}$ maximizing the profit $\sum_{r \in \mathcal{R}'} b_r$.*

B. Related Work

In the last decade, the VNEP has attracted much attention due to its many applications. A survey from 2013 already lists more than 80 different algorithms for its many variations [7]. The VNEP is known to be \mathcal{NP} -hard and inapproximable in general (unless $\mathcal{P} = \mathcal{NP}$) [6]. Based on the hardness of the VNEP, most works consider heuristics without any performance guarantee [7], [8]. Other works proposed exact methods as integer or constraint programming, coming at the cost of an exponential runtime [9], [10], [11].

A column generation approach was proposed by Jarray et al. in [10] to efficiently compute solutions to the VNEP by iteratively generating feasible mappings subject to a specific cost measure. In particular, Jarray et al. compute feasible mappings by relying on heuristics and, if no feasible heuristical solution was found, rely on Mixed-Integer Programming to compute a cost optimal feasible embedding in non-polynomial time. Our work can potentially bridge the gap between the heuristic and optimal generation of feasible mappings as our approach can be used to compute *approximate* mappings.

Acknowledging the hardness of the general VNEP and the diversity of applications, several subproblems of the VNEP have been studied recently by considering restricted graph classes for the virtual networks and the substrate graph. For example, virtual clusters with uniform demands are studied in [4], [12], line requests are studied in [13], [14], [15] and tree requests were studied in [14], [16].

Considering approximation algorithms, Even et al. employed randomized rounding in [14] to obtain a constant approximation for embedding line requests on arbitrary substrate graphs under strong assumptions on *both* the benefits and the capacities. In their interesting work, Bansal et al. [16] give an $n^{O(d)}$ time $O(d^2 \log(nd))$ -approximation algorithm for minimizing the load of embedding d -depth trees based on a n -node substrate. Their result is based on strong Linear Programming (LP) relaxations like the Sherali-Adams hierarchy.

To the best of our knowledge, this work is the first to give approximation algorithms for arbitrary substrate graphs and classes of virtual networks containing cyclic subgraphs. While our results are limited to the family of cactus request graphs, the approach presented in this work can be extended to arbitrary request graphs as shown in our technical report [17]. However, this generalization comes at the cost of polynomial-time solvability: for general graphs the size of the LP formulations may be exponential in the input. As finding LP formulations of minimal size is shown to be \mathcal{NP} -hard, this generalized approach may be hard to use in practice.

Besides the generalization of [17], another parametrized approach to compute decomposable solutions for general graphs has been recently presented in [18]. Specifically, in [18], column generation and dynamic programming are used to compute decomposable LP solutions in time exponential in the *treewidth* of the request graphs, while using the analysis and evaluation framework presented in this work.

C. Outline of Randomized Rounding for the VNEP

We shortly revisit the concept of randomized rounding [19]. Given an Integer Program for a certain problem, randomized rounding works by (i) computing a solution to its Linear Program relaxation, (ii) decomposing this solution into convex combinations of *elementary* solutions, and (iii) probabilistically selecting elementary solutions based on their weight.

Accordingly, for applying randomized rounding for the VNEP, a convex combination of valid mappings $\mathcal{D}_r = \{(f_r^k, m_r^k) | m_r^k \in \mathcal{M}_r, f_r^k > 0\}$ must be recovered from the Linear Programming solution for each request $r \in \mathcal{R}$, such that (i) the profit of these convex combinations equals the profit achieved by the Linear Program and (ii) the (fractional) cumulative allocations do not violate substrate capacities. *Rounding* a solution is then done as follows. For each request $r \in \mathcal{R}$, the mapping m_r^k is selected with probability f_r^k , rejecting r with probability $1 - \sum_k f_r^k$.

D. Results and Organization

This paper initiates the study of approximation algorithms for the VNEP on general substrates *and* general virtual networks. Specifically, we employ randomized rounding to obtain the first approximation for the non-trivial class of cactus graph requests in the resource augmentation model.

Studying the classic multi-commodity flow (MCF) formulation for the VNEP in Section II, we show that its solutions can only be decomposed for tree requests: solutions for request graphs containing cycles can in general not be decomposed into valid mappings. This result has ramifications beyond the inability to apply randomized rounding: we prove that the MCF formulation exhibits an unbounded integrality gap. Investigating the root cause for this surprising result, we devise a novel *decomposable* Linear Programming formulation in Section III for the class of cactus graph requests. We then present our randomized rounding algorithm and prove its performance guarantees in Section IV, obtaining the first approximation algorithm for the VNEP. Section V presents a synthetic computational study, in which two types of rounding heuristics are evaluated. Our results indicate that high-quality solutions can be obtained even without resource augmentations: our heuristic achieves 77.2% of the baseline's profit on average, without exceeding resource capacities.

II. THE CLASSIC MULTI-COMMODITY FORMULATION FOR THE VNEP AND ITS LIMITATIONS

In this section, we study the relaxation of the standard multi-commodity flow (MCF) formulation for the VNEP (cf. [5], [8]). We first show the positive result that the formulation is sufficiently strong to decompose virtual networks being *trees* into convex combinations of valid mappings. Subsequently, we show that the formulation fails to allow for the decomposition of solutions for *cyclic requests*. This not only impacts its applicability for randomized rounding but renders the formulation useless for approximations in general: we show that the formulation's integrality gap is unbounded.

A. The Classic Multi-Commodity Formulation

The classic MCF formulation for the VNEP is presented as Formulation 1. We first describe its (mixed-)integer variant, which computes a single valid mapping for each request by using binary variables. The Linear Programming variant is obtained by relaxing the binary variables' domain to $[0, 1]$.

The variable $x_r \in \{0, 1\}$ indicates whether request $r \in \mathcal{R}$ is embedded or not. The variable $y_{r,i}^u \in \{0, 1\}$ indicates whether virtual node $i \in V_r$ was mapped on substrate node $u \in V_S$. Similarly, the flow variable $z_{r,i,j}^{u,v} \in \{0, 1\}$ indicates whether the substrate edge $(u, v) \in E_S$ is used to realize the virtual edge $(i, j) \in E_r$. The variable $a_r^{x,y} \geq 0$ denotes the cumulative allocations of request $r \in \mathcal{R}$ induced on resource $(x, y) \in R_S$.

By Constraint 2, the virtual node $i \in V_r$ of request $r \in \mathcal{R}$ must be placed on any of the suitable substrate nodes of $V_S^{r,i}$ iff. $x_r = 1$ holds and Constraint 3 forbids the mapping on nodes which may not host node i . Constraint 4 induces an unsplitable unit flow for each virtual edge $(i, j) \in E_r$ from the substrate location to which i was mapped to the substrate location to which j was mapped. By Constraint 5 virtual edges may only be mapped on *allowed* substrate edges. Constraints 6 and 7 compute the cumulative allocations and Constraint 8 guarantees that the substrate resource capacities are respected. The following lemma states the connectivity property enforced by Formulation 1.

Lemma 5 (Local Connectivity Property of Formulation 1).

For any virtual edge $(i, j) \in E_r$ and any substrate node $u \in V_S^{r,i}$ with $y_{r,i}^u > 0$, there exists a path $P_{r,i,j}^{u,v}$ in G_S from u to $v \in V_S^{r,j}$ with $y_{r,j}^v > 0$, such that the flow along any edge of $P_{r,i,j}^{u,v}$ with respect to the variables $z_{r,i,j}^{x,y}$ is greater 0.

The path $P_{r,i,j}^{u,v}$ can be computed in polynomial time.

Proof. First, note that $\sum_{u \in V_S^{r,i}} y_{r,i}^u = \sum_{v \in V_S^{r,j}} y_{r,j}^v$ holds by

Formulation 1: Classic MCF Formulation for the VNEP

$$\max \sum_{r \in \mathcal{R}} b_r x_r \quad (1)$$

$$\sum_{u \in V_S^{r,i}} y_{r,i}^u = x_r \quad \forall r \in \mathcal{R}, i \in V_r \quad (2)$$

$$\sum_{u \in V_S \setminus V_S^{r,i}} y_{r,i}^u = 0 \quad \forall r \in \mathcal{R}, i \in V_r \quad (3)$$

$$\left[\begin{array}{c} \sum_{(u,v) \in \delta^+(u)} z_{r,i,j}^{u,v} \\ - \sum_{(v,u) \in \delta^-(u)} z_{r,i,j}^{v,u} \end{array} \right] = \left[\begin{array}{c} y_{r,i}^u \\ -y_{r,j}^u \end{array} \right] \quad \forall \left[\begin{array}{c} r \in \mathcal{R}, (i, j) \in E_r, \\ u \in V_S \end{array} \right] \quad (4)$$

$$z_{r,i,j}^{u,v} = 0 \quad \forall \left[\begin{array}{c} r \in \mathcal{R}, (i, j) \in E_r, \\ (u, v) \in E_S \setminus E_S^{r,i,j} \end{array} \right] \quad (5)$$

$$\sum_{i \in V_r, \tau_r(i) = \tau} d_r(i) \cdot y_{r,i}^u = a_r^{\tau,u} \quad \forall r \in \mathcal{R}, (\tau, u) \in R_S^V \quad (6)$$

$$\sum_{(i,j) \in E_r} d_r(i, j) \cdot z_{r,i,j}^{u,v} = a_r^{u,v} \quad \forall r \in \mathcal{R}, (u, v) \in E_S \quad (7)$$

$$\sum_{r \in \mathcal{R}} a_r^{x,y} \leq d_S(x, y) \quad \forall (x, y) \in R_S \quad (8)$$

Constraint 2 and hence both virtual nodes i and j must be mapped to an equal extent on suitable substrate nodes. Fix any substrate node $u \in V_S^{r,i}$ with $y_{r,i}^u > 0$. If j is also partially mapped on u , i.e., if $y_{r,j}^u > 0$ holds, then the result follows directly, as u connects to u using (and allowing) the empty path $P_{r,i,j}^{u,u} = \langle \rangle$. If, on the other hand, $y_{r,j}^u = 0$ holds, then Constraint 4 induces a flow of value $y_{r,i}^u$ at u in the commodity $z_{r,i,j}^{\cdot}$. As the right hand side of Constraint 4 may only attain negative values at nodes $v \in V_S^{r,j}$ for which $y_{r,j}^v > 0$ holds, the flow emitted at node u must eventually reach a node $v \in V_S^{r,j}$ with $y_{r,j}^v > 0$ and hence the result follows. The path $P_{r,i,j}^{u,v}$ can be constructed in polynomial-time by a breadth-first search only considering edges $(u', v') \in E_S$ with $z_{r,i,j}^{u',v'} > 0$. \square

B. Decomposing Solutions for Tree Requests

Given Lemma 5, we now present Algorithm 1 to decompose solutions to the LP Formulation 1 into convex combinations of valid mappings $\mathcal{D}_r = \{(f_r^k, m_r^k) | m_r^k \in \mathcal{M}_r, f_r^k > 0\}$ (cf. Section I-C), if the request's underlying undirected graph is a tree. Recall that in the LP formulation the binary variables are relaxed to take any value in the interval $[0, 1]$.

Given a request $r \in \mathcal{R}$, the algorithm processes all virtual edges according to an arbitrary acyclic representation $G_r^A = (V_r, E_r^A, r_r)$ of the undirected interpretation of G_r being rooted at $r_r \in V_r$. Concretely, the edge set E_r^A is obtained from E_r by reorienting (some of the) edges, such that any node $i \in V_r$ can be reached from r_r . Considering tree requests for now, G_r^A is an arborescence and can be computed by a simple graph search of the underlying undirected graph starting at an arbitrary root node $r_r \in V_r$. We denote by $\overleftarrow{E}_r^A = E_r \setminus E_r^A$ the edges whose orientations were reversed in the process of computing G_r^A .

The algorithm extracts mappings m_r^k of value f_r^k iteratively, as long as $x_r > 0$ holds. Initially, in the k -th iteration, none of the virtual nodes and edges are mapped. As $x_r > 0$ holds, there must exist a node $u \in V_S^{r,r_r}$ with $y_{r,r_r}^u > 0$ by Constraint 2 and the algorithm accordingly sets $m_r^V(r_r) = u$. Given this initial fixing, the algorithm iteratively extracts nodes from the queue \mathcal{Q} which have been already mapped and considers all outgoing virtual edges $(i, j) \in E_r^A$. If an outgoing edge (i, j) is contained in E_r , Lemma 5 can be readily applied to obtain a joint mapping of the edge (i, j) and its head j . If the edge's orientation was reversed, i.e. if $(i, j) \in \overleftarrow{E}_r^A$ holds, Lemma 5 is applied *while reversing* the flow's direction (see Lines 13-16).

First, note that by the repeated application of Lemma 5, the mapping of virtual nodes and edges is valid. As G_r^A is an arborescence, each edge and each node of G_r^A will eventually be mapped and hence m_r^k is a valid mapping. The mapping value f_r^k is computed as the minimum of the mapping variables \mathcal{V}_k used for constructing m_r^k . Reducing the values of the mapping variables together with the allocation variables \vec{a}_r (Lines 20-21), the Constraints 2-7 continue to hold.

As the decomposition process continues as long as $x_r > 0$ holds and in the k -th iteration at least one variable's value is set to 0, the algorithm terminates with a complete decomposition for which $\sum_k f_r^k = x_r$ holds. Furthermore, the algorithm has polynomial runtime, as in each iteration at least one variable is

Algorithm 1: Decompositioning MCF solutions for Tree Requests

Input : Tree request $r \in \mathcal{R}$, solution $(x_r, \vec{y}_r, \vec{z}_r, \vec{a}_r)$ to LP Formulation 1, acyclic reorientation $G_r^A = (V_r, E_r^A, r_r)$
Output: Convex combination $\mathcal{D}_r = \{(f_r^k, m_r^k)\}_k$

- 1 **set** $\mathcal{D}_r \leftarrow \emptyset$ **and** $k \leftarrow 1$
- 2 **while** $x_r > 0$ **do**
- 3 **set** $m_r^k = (m_r^V, m_r^E) \leftarrow (\emptyset, \emptyset)$
- 4 **set** $\mathcal{Q} = \{r_r\}$
- 5 **choose** $u \in V_S^{r,r_r}$ **with** $y_{r,r_r}^u > 0$ **and set** $m_r^V(r_r) \leftarrow u$
- 6 **while** $|\mathcal{Q}| > 0$ **do**
- 7 **choose** $i \in \mathcal{Q}$ **and set** $\mathcal{Q} \leftarrow \mathcal{Q} \setminus \{i\}$
- 8 **foreach** $(i, j) \in E_r^A$ **do**
- 9 **if** $(i, j) \in E_r$ **then**
- 10 **compute** $\vec{P}_{r,i,j}^{u,v}$ connecting $m_r^V(i) = u$ to $v \in V_S^{r,j}$
- 11 according to Lemma 5
- 12 **set** $m_r^V(j) = v$ **and** $m_r^E(i, j) = \vec{P}_{r,i,j}^{u,v}$
- 13 **else**
- 14 **let** $\overleftarrow{z}_{r,i,j}^{v',u'} \triangleq z_{r,i,j}^{u',v'}$ **for all** $(u', v') \in E_S$
- 15 **compute** $\overleftarrow{P}_{r,i,j}^{v,u}$ connecting $m_r^V(i) = v$ to $u \in V_S^{r,j}$
- 16 according to Lemma 5
- 17 **set** $\overleftarrow{P}_{r,i,j}^{u,v} = \text{reverse}(\overleftarrow{P}_{r,i,j}^{v,u})$
- 18 **set** $m_r^V(i) = u$ **and** $m_r^E(j, i) = \overleftarrow{P}_{r,i,j}^{u,v}$
- 19 **set** $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{j\}$
- 20 **set** $\mathcal{V}_k \leftarrow \left(\begin{array}{l} \{x_r\} \cup \{y_{r,i}^{m_r^V(i)} | i \in V_r\} \\ \cup \{z_{r,i,j}^{u,v} | (i, j) \in E_r, (u, v) \in m_r^E(i, j)\} \end{array} \right)$
- 21 **set** $f_r^k \leftarrow \min \mathcal{V}_k$
- 22 **set** $v \leftarrow v - f_r^k$ **for all** $v \in \mathcal{V}_k$
- 23 **set** $a_{r,x,y}^{x,y} \leftarrow a_{r,x,y}^{x,y} - f_r^k \cdot A(m_r^k, x, y)$ **for all** $(x, y) \in R_S$
- 24 **add** $D_r^k = (f_r^k, m_r^k)$ to \mathcal{D}_r **and set** $k \leftarrow k + 1$
- 25 **return** \mathcal{D}_r .

set to 0 and the number of variables for request r is bounded by $\mathcal{O}(|G_r| \cdot |G_S|)$. Hence, we obtain the following:

Lemma 6. *Given a virtual network request $r \in \mathcal{R}$, whose underlying undirected graph is a tree, Algorithm 1 decomposes a solution $(x_r, \vec{y}_r, \vec{z}_r, \vec{a}_r)$ to LP Formulation 1 into valid a convex combination of valid mappings $\mathcal{D}_r = \{(m_r^k, f_r^k) | m_r^k \in \mathcal{M}_r, f_r^k > 0\}_k$, such that:*

- *The decomposition is complete, i.e., $x_r = \sum_k f_r^k$ holds.*
- *The decomposition's allocations are bounded by \vec{a}_r , i.e., $a_{r,x,y}^{x,y} \geq \sum_k f_r^k \cdot A(m_r^k, x, y)$ holds for $(x, y) \in R_S$.*

C. Limitations of the Classic MCF Formulation

Above it was shown that LP solutions to the classic MCF formulation can be decomposed into convex combinations of valid mappings if the underlying graph is a tree. However, this does not hold anymore when considering cyclic virtual networks:

Theorem 7. *Solutions to the LP Formulation 1 can in general not be decomposed into convex combinations of valid mappings if the virtual networks contain cycles.*

Proof. In Figure 3 we visually depict an example of a solution to the LP Formulation 1 from which *not a single* valid mapping

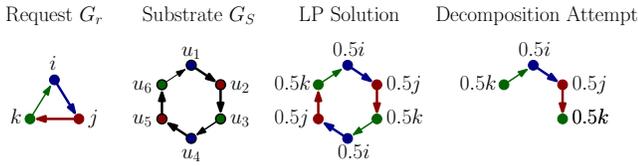


Fig. 3. Example showing that solutions to the LP Formulation 1 can in general not be decomposed into convex combinations of valid mappings. Request r is a simple cyclic graph which shall be mapped on the substrate graph G_S . We assume following node mapping restrictions $V_S^{r,i} = \{u_1, u_4\}$, $V_S^{r,j} = \{u_2, u_5\}$, $V_S^{r,k} = \{u_3, u_6\}$. The LP solution with $x_r = 1$ is depicted as follows. Substrate nodes are annotated with the mapping of virtual nodes. Hence, $0.5i$ at node u_1 indicates $y_{r,i}^{u_1} = 1/2$, i.e., that virtual node i is mapped with 0.5 on substrate node u_1 . Substrate edges are colored according to the color of virtual links mapped onto it. Virtual links are all mapped using flow values $1/2$. Accordingly, e.g., $z_{r,i,j}^{u_1, u_2} = 1/2$ holds.

can be extracted. The validity of the depicted solution follows from the fact that the virtual node mappings sum to 1 and each virtual node connects to its neighboring node with half a unit of flow. Assume for the sake of contradiction that the depicted solution can be decomposed. As virtual node $i \in V_r$ is mapped onto substrate node $u_1 \in V_S$, and $u_2 \in V_S$ is the only neighboring node with respect to variables $z_{r,i,j}^{u_1, u_2}$ that hosts $j \in V_r$, there must exist a mapping (m_r^V, m_r^E) with $m_r^V(i) = u_1$ and $m_r^V(j) = u_2$. Similarly, $m_r^V(k) = u_3$ must hold. However, for $m_r^V(i) = u_1$, the virtual node k must be mapped to u_6 , as otherwise the embedding of (k, i) cannot lead to substrate node u_1 . Hence the virtual node $k \in V_r$ must be mapped both on u_6 and u_3 . As this is not possible and the same argument holds when considering mapping i onto u_4 , no valid mapping can be extracted from the LP solution. \square

The non-decomposability of the solutions to the MCF formulation shows that it cannot be used for applying randomized rounding, as no valid mappings can be extracted in the first place. Even more, the non-decomposability does not only forbid the application of randomized rounding but cannot be at all used for obtaining approximations. Specifically, given the non-decomposability it is easy to construct instances in which the MCF formulation attains the maximal profit while any optimal solution cannot embed any request. Hence, the integrality gap of the MCF formulation is unbounded and can – for graphs containing cycles – never lead to approximation algorithms. At this point, we only state the respective theorems and refer the reader to the appendix for the proofs.

Theorem 8. *The integrality gap of the MCF formulation is unbounded. This even holds under infinite substrate capacities.*

Theorem 9. *The integrality gap of the MCF formulation lies in $\Omega(|V_S|)$, when only considering node mapping restrictions.*

III. NOVEL DECOMPOSABLE LP FORMULATION

In this section, we present a novel LP formulation and its accompanying decomposition algorithm for the class of cactus request graphs, i.e., graphs for which cycles intersect in at most a single node (in its undirected interpretation). Accordingly, these graphs can be uniquely decomposed into cycles and a single forest (cf. Lemma 10 below).

Before delving into the details of our novel LP formulation, we discuss our main insight on how to overcome the limitations of the MCF formulation and accordingly how to derive decomposable formulations. To this end, it is instructive, to revisit the non-decomposable example of Figure 3 by applying the decomposition Algorithm 1 on the depicted LP solution. Concretely, we consider the acyclic reorientation $G_r^A = (V_r, E_r^A, r_r)$ with $E_r^A = \{(i, k), (i, j), (j, k)\}$, such that i is the root, $r_r = i$. Assuming that i is initially mapped on node u_1 , Algorithm 1 will map edges (i, k) and (i, j) first, setting $m_r^V(k) = u_6$ and $m_r^V(j) = u_2$. However, when the edge (j, k) is processed, k must be mapped on substrate node $u_3 \neq m_r^V(k)$ and the algorithm hence fails to produce a valid mapping. Accordingly, to avoid such *diverging* node mappings, our key idea is to decide the mapping location of nodes with more than one incoming edge in the request's acyclic reorientation G_r^A a priori.

By considering only cactus request graphs, the above can be implemented rather easily as exactly one node of each cycle has more than one incoming edge: one only needs to ensure *compatibility* of node mappings for this node. To resolve potential conflicts for the mapping of this unique *cycle target*, our formulation employs *multiple* copies of the MCF formulation for the respective cyclic subgraph. Specifically, considering a cycle with virtual *target* node k , we instantiate one MCF formulation per substrate node $w \in V_S^{r,k}$ onto which k can be mapped. Accordingly, this yields at most $|V_S|$ many copies and for each of these copies k is *fixed* to one specific (substrate) mapping location. Accordingly, as the mapping location of k is fixed to a specific node, valid mappings for the respective cycles can always be extracted from such a MCF copy: the mappings of k cannot possibly diverge.

A. Cactus Request Graph Decomposition and Notation

Based on the assumed cactus graph nature of request graphs, we consider the following decomposition of request graphs (see the appendix for the proof).

Lemma 10. *Consider a cactus request graph G_r and an acyclic reorientation G_r^A of G_r . The graph G_r^A can be uniquely partitioned into subgraphs $\{G_r^{A,C_1}, \dots, G_r^{A,C_n}\} \sqcup G_r^{A,\mathcal{F}}$, such that the following holds:*

- 1) *The subgraphs $\{G_r^{A,C_1}, \dots, G_r^{A,C_n}\}$ correspond to the (undirected) cycles of G_r and $G_r^{A,\mathcal{F}}$ is the forest remaining after removing the cyclic subgraphs. We denote the index set of the cycles by $\mathcal{C}_r = \{C_1, \dots, C_n\}$.*
- 2) *The subgraphs partition the edges of E_r^A : an edge $(i, j) \in E_r^A$ is contained in exactly one of the subgraphs.*
- 3) *The edge set E_r^{A,C_k} of each cycle $C_k \in \mathcal{C}_r$ can itself be partitioned into two branches $\mathcal{B}_1^{C_k}$ and $\mathcal{B}_2^{C_k}$, such that both lead from $s_r^{C_k} \in V_r^{A,C_k}$ to $t_r^{C_k} \in V_r^{A,C_k}$.*

Besides the above introduced notation, we denote by $G_r^{C_k}$ and $G_r^{\mathcal{F}}$ the subgraphs that agree with G_r on the edge orientations and use $V_{S,t}^{C_k} = V_S^{r,t_r^{C_k}}$ to denote the substrate nodes on which $t_r^{C_k}$ can be mapped.

B. Novel LP Formulation for Cactus Request Graphs

Our novel Formulation 2 uses the a priori partition of G_r^A into cycles G_r^{A,C_k} and the forest $G_r^{A,\mathcal{F}}$ to construct MCF formulations for the respective subgraphs: for the subgraph $G_r^{\mathcal{F}}$ a single copy is used (cf. Constraint 10) while for the cyclic subgraphs a single MCF formulation is employed *per* potential target location contained in the set $V_{S,t}^{C_k}$ (cf. Constraint 11). We index the variables of these sub-LPs by employing square brackets.

To bind together these (at first) independent MCF formulations, we reuse the variables \vec{x} , \vec{y} , and \vec{a} introduced already for the MCF formulation. We refer to these variables, which are defined outside of the sub-LP formulations, as *global variables* and do not index these. As we only consider the LP formulation, all variables are continuous.

The different sub-formulations are linked as follows. We employ Constraint 12 to enforce the setting of the (global) node mapping variables (cf. Constraint 2 of Formulation 1). By Constraints 13 and 14, the node mappings of the sub-LPs for mapping the subgraphs must agree with the global node mapping variables. With respect to cyclic subgraphs, we note that Constraint 14 allows for distributing the global node mappings to any of the $|V_{S,t}^{C_k}|$ formulations: only the sum of the node mapping variables must agree with the global node mapping variable. Constraint 15 is of crucial importance for the decomposability: considering the sub-LP for cycle C_k and target node $w \in V_{S,t}^{C_k}$, it enforces that the target node $t_r^{C_k}$ of the cycle C_k *must* be mapped on w . Thus, in the sub-LP $[C_k, w]$ both branches $\mathcal{B}_1^{C_k}$ and $\mathcal{B}_2^{C_k}$ of cycle C_k are *pre-determined* to lead to the node w . Lastly, for computing node allocations the global node mapping variables are used (cf. Constraint 16) and for computing edge allocations the sub-LP formulations' allocations are considered (cf. Constraint 17).

C. Decomposing Solutions to the Novel LP Formulation

We now show how to adapt the decomposition Algorithm 1 to decompose solutions to Formulation 2.

To decompose the LP solution for a request r , the acyclic reorientation G_r^A , which was also used for constructing the LP, must be handed over to the decomposition algorithm.

As the novel LP formulation does not contain (global) edge mapping variables, the edge mapping variables used in Lines 10 and 13 of Algorithm 1 must be substituted by edge mapping variables of the respective sub-LP formulations. Concretely, as each edge of the request graph G_r is covered exactly once, it is clear whether a virtual edge $(i, j) \in E_r$ is part of $G_r^{\mathcal{F}}$ or a cyclic subgraph $G_r^{C_k}$. If $(i, j) \in G_r^{\mathcal{F}}$ holds, then the edge mapping variables $z_{r,i,j}^{\mathcal{F}}$ are used. If on the other hand the edge $(i, j) \in E_r$ is covered in the cyclic subgraph $G_r^{C_k}$, then there exist $|V_{S,t}^{C_k}|$ many sub-LPs to choose the respective edge mapping variables from. To ensure the decomposability, we proceed as follows.

If the edge $(i, j) \in E_r^A$ is the first edge of $G_r^{C_k}$ to be mapped in the k -th iteration, the mapping variables $z_{r,i,j}^{\mathcal{F}}$ belonging to an arbitrary target node w , with $y_{r,i}^{m_r^V(i)}[C_k, w] > 0$, are used. Such a node w exists by Constraint 14.

Formulation 2: Novel LP for Cactus Request Graphs

$$\max \sum_{r \in \mathcal{R}} b_r x_r \quad (9)$$

$$\text{Cons. (2) - (7) for } G_r^{\mathcal{F}} \text{ on} \quad \forall r \in \mathcal{R} \quad (10)$$

variables $(x_r, \vec{y}_r, \vec{z}_r, \vec{a}_r)[\mathcal{F}_r]$

$$\text{Cons. (2) - (7) for } G_r^{C_k} \text{ on} \quad \forall r \in \mathcal{R}, C_k \in \mathcal{C}_r, w \in V_{S,t}^{C_k} \quad (11)$$

variables $(x_r, \vec{y}_r, \vec{z}_r, \vec{a}_r)[C_k, w]$

$$x_r = \sum_{u \in V_S^{r,i}} y_{r,i}^u \quad \forall r \in \mathcal{R}, i \in V_r \quad (12)$$

$$y_{r,i}^u = y_{r,i}^u[\mathcal{F}] \quad \forall r \in \mathcal{R}, i \in V_r^{\mathcal{F}}, u \in V_S^{r,i} \quad (13)$$

$$y_{r,i}^u = \sum_{w \in t_r^{C_k}} y_{r,i}^u[C_k, w] \quad \forall \left[\begin{array}{l} r \in \mathcal{R}, i \in V_r, u \in V_S^{r,i} \\ C_k \in \mathcal{C}_r : i \in V_r^{C_k} \end{array} \right] \quad (14)$$

$$0 = y_{r,t_r^{C_k}}^u[C_k, w] \quad \forall \left[\begin{array}{l} r \in \mathcal{R}, C_k \in \mathcal{C}_r, w \in V_{S,t}^{C_k} \\ u \in V_{S,t}^{C_k} \setminus \{w\} \end{array} \right] \quad (15)$$

$$a_r^{\tau,u} = \sum_{i \in V_r, \tau_r(i) = \tau} d_r(i) \cdot y_{r,i}^u \quad \forall r \in \mathcal{R}, (\tau, u) \in R_S^V \quad (16)$$

$$a_r^{u,v} = a_r^{u,v}[\mathcal{F}] + \sum_{C_k \in \mathcal{C}_r, w \in V_{S,t}^{C_k}} a_r^{u,v}[C_k, w] \quad \forall r \in \mathcal{R}, (u, v) \in E_S \quad (17)$$

$$\sum_{r \in \mathcal{R}} a_r^{x,y} \leq d_S(x, y) \quad \forall (x, y) \in R_S \quad (18)$$

If another edge (i', j') of the same cycle was already mapped in the k -th iteration, the *same* sub-LP before is considered. Accordingly, the mapping of cycle target nodes cannot conflict, and as these are the only nodes with potential mapping conflicts, the returned mappings are always valid.

To successfully iterate the extraction process, the steps taken in Lines 18 - 21 of Algorithm 1 must be adapted to consider the sub-LP variables. Again, as in each iteration at least a single variable of the LP is set to 0 and as the novel Formulation 2 contains at most $\mathcal{O}(|V_S|)$ times more variables than the MCF Formulation 1, the decomposition algorithm still runs in polynomial-time. Hence, we conclude that the result of Lemma 6 carries over to the novel LP Formulation 2 for *cactus* request graphs and state the following theorem.

Theorem 11. *Given a solution $(x_r, \vec{y}_r, \vec{z}_r, \vec{a}_r)$ to the novel LP Formulation 2 for a cactus request graph G_r , the solution can be decomposed into a convex combination of valid mappings $\mathcal{D}_r = \{(m_r^k, f_r^k) | m_r^k \in \mathcal{M}_r, f_r^k > 0\}_k$, such that:*

- *The decomposition is complete, i.e., $x_r = \sum_k f_r^k$ holds.*
- *The decomposition's allocations are bounded by \vec{a}_r , i.e., $a_r^{x,y} \geq \sum_k f_r^k \cdot A(m_r^k, x, y)$ holds for $(x, y) \in R_S$.*

IV. APPROXIMATION VIA RANDOMIZED ROUNDING

Above we have shown how convex combinations for the VNEP can be computed for cactus requests. Given these convex combinations, we now present our approximation algorithm (see Algorithm 2) and analyze its probabilistic guarantees in Sections IV-A to IV-C. Afterwards, we propose several rounding heuristics to apply our approximation in practice (see Section IV-D).

Algorithm 2: Randomized Rounding for the VNEP

```

1 foreach  $r \in \mathcal{R}$  do           // preprocess requests
2   compute LP Formulation 2 for request  $r$  maximizing
    $x_r$ 
3   if  $x_r < 1$  then remove request  $r$  from the set  $\mathcal{R}$ 
4   compute LP Formulation 2 for  $\mathcal{R}$  maximizing
    $\sum_{r \in \mathcal{R}} b_r \cdot x_r$ 
5   foreach  $r \in \mathcal{R}$  do           // perform decomposition
6   | compute  $\mathcal{D}_r = \{(f_r^k, m_r^k)\}_k$  from LP solution
7   do                               // perform randomized rounding
8   | foreach  $r \in \mathcal{R}$  select  $m_r^k$  with probability  $f_r^k$ 
9   while ( solution is not  $(\alpha, \beta, \gamma)$ -approximate and
   maximal rounding tries are not exceeded )

```

Synopsis of our Approximation Algorithm: The algorithm first performs a preprocessing in Lines 1 to 3 by removing all requests which cannot be fully (fractionally) embedded in the absence of other requests, as these can never be part of any feasible solution. In Lines 4 to 6 an optimal solution to the novel LP Formulation 2 is computed and afterwards decomposed into convex combinations. Then, in Lines 7 to 9, the rounding is performed: for each request $r \in \mathcal{R}$ a mapping m_r^k is selected with probability f_r^k . Importantly, the summed probabilities may not sum to 1, i.e., with probability $1 - \sum_k f_r^k$ the request r is not embedded.

The rounding procedure is iterated as long as the constructed solution is not of sufficient quality or until the number of maximal rounding tries is exceeded. Concretely, we seek (α, β, γ) -approximate solutions which achieve at least a factor of $\alpha \leq 1$ times the optimal (LP) profit and exceed node and edge capacities by at most factors of $\beta \geq 1$ and $\gamma \geq 1$, respectively. In the following we derive parameters α , β , and γ for which solutions can be found with high probability.

Note that Algorithm 2 is indeed a polynomial-time algorithm, as the size of the novel LP Formulation 2 is polynomially bounded and can hence be solved in polynomial-time.

A. Probabilistic Guarantee for the Profit

For bounding the profit achieved by the randomized rounding scheme, we recast the profit achieved in terms of random variables. The discrete random variable $Y_r \in \{0, b_r\}$ models the profit achieved by the rounding of request $r \in \mathcal{R}$. According to our rounding scheme, we have $\mathbb{P}(Y_r = b_r) = \sum_k f_r^k$ and $\mathbb{P}(Y_r = 0) = 1 - \sum_k f_r^k$. We denote the overall profit by $B = \sum_{r \in \mathcal{R}} Y_r$ with $\mathbb{E}(B) = \sum_{r \in \mathcal{R}} b_r \cdot \sum_k f_r^k$. Denoting the profit of an optimal LP solution by B_{LP} , we have $B_{\text{LP}} = \mathbb{E}(B)$ due to the decomposition's completeness (cf. Theorem 11).

By preprocessing the requests and confirming that each request can be fully embedded, the LP will attain at least the maximal profit of any of the considered requests:

Lemma 12. $\mathbb{E}(B) = B_{\text{LP}} \geq \max_{r \in \mathcal{R}} b_r$ holds.

We employ the following Chernoff bound over continuous variables to bound the probability of achieving a small profit.

Theorem 13 (Chernoff Bound [20]). *Let $X = \sum_{i=1}^n X_i$, $X_i \in [0, 1]$, be a sum of n independent random variables. For any $0 < \varepsilon < 1$, the following holds:*

$$\mathbb{P}(X \leq (1 - \varepsilon) \cdot \mathbb{E}(X)) \leq \exp(-\varepsilon^2 \cdot \mathbb{E}(X)/2)$$

Theorem 14. *Let B_{IP} denote the profit of an optimal solution. Then $\mathbb{P}(B < 1/3 \cdot B_{\text{IP}}) \leq \exp(-2/9) \approx 0.8007$ holds.*

Proof. Let $\hat{b} = \max_{r \in \mathcal{R}} b_r$ be the maximum benefit among the preprocessed requests. We consider random variables $Y'_r = Y_r/\hat{b}$, such that $Y'_r \in [0, 1]$ holds. Let $B' = \sum_{r \in \mathcal{R}} Y'_r = B/\hat{b}$. As $\mathbb{E}(B) = B_{\text{LP}} \geq \hat{b}$ holds (cf. Lemma 12), we have $\mathbb{E}(B') \geq 1$. Choosing $\varepsilon = 2/3$ and applying Theorem 13 on B' we obtain $\mathbb{P}(B' \leq (1/3) \cdot \mathbb{E}(B')) \leq \exp(-2 \cdot \mathbb{E}(B')/9)$. Plugging in the minimal value of $\mathbb{E}(B')$, i.e., 1, into the equation we obtain $\mathbb{P}(B' \leq (1/3) \cdot \mathbb{E}(B')) \leq \exp(-2/9)$ and by linearity $\mathbb{P}(B \leq (1/3) \cdot \mathbb{E}(B)) \leq \exp(-2/9)$.

Denoting the profit of an optimal solution by B_{IP} and observing that $B_{\text{IP}} \leq B_{\text{LP}}$ holds as the linear relaxation yields an upper bound, we have $B_{\text{IP}}/3 \leq \mathbb{E}(B)/3$. Accordingly, we conclude that, $\mathbb{P}(B \leq (1/3) \cdot B_{\text{IP}}) \leq \exp(-2/9)$ holds. \square

B. Probabilistic Guarantee for Resource Augmentations

In the following, we analyze the probability that a rounded solution exceeds substrate capacities by a certain factor.

We first note that $d_{\max}(r, x, y) \leq d_S(x, y)$ holds for all resources $(x, y) \in R_S$ and all requests $r \in \mathcal{R}$. We model the allocations on resource $(x, y) \in R_S$ by request $r \in \mathcal{R}$ as random variable $A_{r,x,y} \in [0, A_{\max}(r, x, y)]$. By definition, we have $\mathbb{P}(A_{r,x,y} = A(m_r^k, x, y)) = f_r^k$ and $\mathbb{P}(A_{r,x,y} = 0) = 1 - \sum_k f_r^k$. Furthermore, we denote by $A_{x,y} = \sum_{r \in \mathcal{R}} A_{r,x,y}$ the random variable capturing the overall allocations on resource $(x, y) \in R_S$. As $\mathbb{E}(A_{x,y}) = \sum_{r \in \mathcal{R}} \sum_k f_r^k \cdot A(m_r^k, x, y)$ holds by Theorem 11, we obtain $\mathbb{E}(A_{x,y}) \leq d_S(x, y)$ for all resources $(x, y) \in R_S$.

We employ Hoeffding's inequality to upper bound $A_{x,y}$.

Theorem 15 (Hoeffding's inequality [20]). *Let $X = \sum_{i=1}^n X_i$, $X_i \in [a_i, b_i]$, be a sum of n independent random variables. The following holds for any $t \geq 0$:*

$$\mathbb{P}(X - \mathbb{E}(X) \geq t) \leq \exp(-2t^2 / (\sum_i (b_i - a_i)^2))$$

Lemma 16. *Consider a resource $(x, y) \in R_S$ and $0 < \varepsilon \leq 1$, such that $d_{\max}(r, x, y)/d_S(x, y) \leq \varepsilon$ holds for $r \in \mathcal{R}$. Let $\Delta(x, y) = \sum_{r \in \mathcal{R}: d_{\max}(r, x, y) > 0} (A_{\max}(r, x, y)/d_{\max}(r, x, y))^2$.*

$$\mathbb{P}(A_{x,y} \geq \delta(\lambda) \cdot d_S(x, y)) \leq \lambda^{-4} \quad (19)$$

holds for $\delta(\lambda) = 1 + \varepsilon \cdot \sqrt{2 \cdot \Delta(x, y) \cdot \log(\lambda)}$ and any $\lambda > 0$.

Proof. We apply Hoeffding with $t = (1 - \delta(\lambda)) \cdot d_S(x, y)$:

$$\begin{aligned} & \mathbb{P}\left(A_{x,y} - \mathbb{E}(A_{x,y}) \geq (1 - \delta(\lambda)) \cdot d_S(x, y)\right) \\ & \leq \exp\left(\frac{-4 \cdot \varepsilon^2 \cdot \log(\lambda) \cdot \Delta(x, y) \cdot d_S^2(x, y)}{\sum_{r \in \mathcal{R}} (A_{\max}(r, x, y))^2}\right) \\ & \leq \exp\left(\frac{-4 \cdot \varepsilon^2 \cdot \log(\lambda) \cdot \Delta(x, y) \cdot d_S^2(x, y)}{\sum_{r \in \mathcal{R}: d_{\max}(r, x, y) > 0} (A_{\max}(r, x, y))^2}\right) \end{aligned}$$

$$\begin{aligned} &\leq \exp\left(\frac{-4 \cdot \varepsilon^2 \cdot \log(\lambda) \cdot \Delta(x, y) \cdot d_S^2(x, y)}{\sum_{r \in \mathcal{R}: d_{\max}(r, x, y) > 0} (\varepsilon \cdot d_S(x, y) \cdot A_{\max}(r, x, y) / d_{\max}(r, x, y))^2}\right) \\ &\leq \exp\left(\frac{-4 \cdot \log(\lambda) \cdot \Delta(x, y)}{\sum_{r \in \mathcal{R}: d_{\max}(r, x, y) > 0} (A_{\max}(r, x, y) / d_{\max}(r, x, y))^2}\right) = \lambda^{-4} \end{aligned}$$

The second inequality holds, as $A_{\max}(r, x, y) > 0$ implies $d_{\max}(r, x, y) > 0$. For the third inequality, $A_{\max}(r, x, y) \leq \varepsilon \cdot d_S(x, y) \cdot A_{\max}(r, x, y) / d_{\max}(r, x, y)$ is used, which follows from the assumption $d_{\max}(r, x, y) \leq \varepsilon \cdot d_S(x, y)$ and $d_{\max}(r, x, y) > 0$. In the next step, $\varepsilon^2 \cdot d_S^2(x, y)$ is reduced from the fraction. As the denominator equals $\Delta(x, y)$ by definition, the final equality follows. Lastly, we utilize that the expected allocation $\mathbb{E}(A_{x,y})$ is upper bounded by the resource's capacity $d_S(x, y)$ to obtain Equation 19. \square

Given Lemma 16, we obtain the following corollary.

Corollary 17. *Let $0 < \varepsilon \leq 1$ be chosen minimally, such that $d_{\max}(r, x, y) / d_S(x, y) \leq \varepsilon$ holds for all resources $(x, y) \in R_S$ and all requests $r \in \mathcal{R}$. Let $\Delta(X) = \max_{(x,y) \in X} \Delta(x, y)$,*

$$\begin{aligned} \beta &= (1 + \varepsilon \cdot \sqrt{2 \cdot \Delta(R_S^V) \cdot \log(|V_S| \cdot |\mathcal{T}|)}) , \text{ and} \\ \gamma &= (1 + \varepsilon \cdot \sqrt{2 \cdot \Delta(E_S) \cdot \log(|E_S|)}) . \end{aligned}$$

The following holds for all node resources $(\tau, u) \in R_S^V$ and edge resources $(u, v) \in E_S$, respectively:

$$\mathbb{P}(A_{\tau, u} \geq \beta \cdot d_S(\tau, u)) \leq (|V_S| \cdot |\mathcal{T}|)^{-4} \quad (20)$$

$$\mathbb{P}(A_{u, v} \geq \gamma \cdot d_S(u, v)) \leq |E_S|^{-4} \quad (21)$$

Proof. First, note that ε is chosen over all resources and requests and that $\Delta(R_S^V) \geq \Delta(\tau, u)$ and $\Delta(E_S) \geq \Delta(u, v)$ hold for $(\tau, u) \in R_S^V$ and $(u, v) \in E_S$, respectively. Equations 20 and 21 are then obtained from Lemma 16 by setting $\lambda = |V_S| \cdot |\mathcal{T}|$ for nodes and $\lambda = |E_S|$ for edges. \square

C. Approximation Result

Given the probabilistic bounds established above, the main approximation result is obtained via a union bound.

Theorem 18. *Assume $|V_S| \geq 3$. Let β and γ be defined as in Corollary 17. Algorithm 2 returns (α, β, γ) -approximate solutions for the VNEP (restricted on cactus request graphs) of at least an $\alpha = 1/3$ fraction of the optimal profit, and allocations on nodes and edges within factors of β and γ of the original capacities, respectively, with high probability.*

Proof. We employ the following union bound argument. Employing Corollary 17 and as there are at most $|V_S| \cdot |\mathcal{T}|$ node resources and at most $|V_S|^2$ edges, the joint probability that any resource exceeds their respective capacity by factors of β or γ is upper bounded by $1/(|V_S| \cdot |\mathcal{T}|)^3 + 1/|V_S|^2 \leq 1/27 + 1/9$ for $|V_S| \geq 3$. By Theorem 14 the probability of not finding a solution achieving an $\alpha = 1/3$ fraction of the optimal objective is upper bounded by $\exp(-2/9)$. Hence, the probability to not find a (α, β, γ) -approximate solution within a single round is upper bounded by $\exp(-2/9) + 1/9 + 1/27 \leq 19/20$. Hence, as the probability to return a suitable solution within a single round is at least $1/20$, the probability that the algorithm returns an approximate solution within $N \in \mathbb{N}$

rounding tries is lower bounded by $1 - (19/20)^N$. Thus, Algorithm 2 yields approximate solutions for the VNEP with high probability. \square

Note that while the theorem is restricted to $|V_S| \geq 3$, a similar result can be obtained for $|V_S| = 2$ by increasing the respective approximation factors slightly. Adapting Corollary 17 by using $\lambda = 2 \cdot |V_S| \cdot |\mathcal{T}|$ for nodes and $\lambda = 2 \cdot |E_S|$ for edges, the respective probabilities for violations become $(2 \cdot |V_S| \cdot |\mathcal{T}|)^{-4}$ and $(2 \cdot |E_S|)^{-4}$, such that the union bound yields $\exp(-2/9) + 1/2^7 + 1/2^6 \leq 0.85$ as a bound on the failure probability. Hence, the result holds with for $|V_S| = 2$ with $\beta = (1 + \varepsilon \cdot \sqrt{2 \cdot \Delta(R_S^V) \cdot \log(2 \cdot |V_S| \cdot |\mathcal{T}|)})$ and $\gamma = (1 + \varepsilon \cdot \sqrt{2 \cdot \Delta(E_S) \cdot \log(2 \cdot |E_S|)})$. For $|V_S| = 1$, the VNEP reduces to the knapsack problem and can be solved using respective algorithms.

D. Discussion & Proposed Heuristics

Theorem 18 yields the first approximation algorithm for the profit variant of the VNEP. However, the direct application of Algorithm 2 to compute (α, β, γ) -approximate solutions is made difficult by the cumbersome definition of the terms $\Delta(R_S^V)$ and $\Delta(E_S)$. Specifically, computing β and γ exactly requires enumerating all valid mappings, which is not feasible. Hence, to directly apply Algorithm 2, the respective values have to be estimated. The following upper bounds can be easily established:

$$\Delta(R_S^V) \leq |\mathcal{R}| \cdot \max_{r \in \mathcal{R}} |V_r|^2 \quad (22)$$

$$\Delta(E_S) \leq |\mathcal{R}| \cdot \max_{r \in \mathcal{R}} |E_r|^2. \quad (23)$$

Both inequalities follow from the observation that $A_{\max}(r, x, y) / d_{\max}(r, x, y)$ is upper bounded by the number of virtual nodes and edges, respectively, as $A_{\max}(r, \tau, u) \leq d_{\max}(r, \tau, u) \cdot |V_r|$ and $A_{\max}(r, u, v) \leq d_{\max}(r, u, v) \cdot |E_r|$ holds for node resources $(\tau, u) \in R_S^V$ and edge resource $(u, v) \in E_S$, respectively. Now, plugging these upper bounds into the definition of β and γ yields large resource augmentations of $\beta \in \mathcal{O}(\varepsilon \cdot \max_{r \in \mathcal{R}} |V_r| \cdot \sqrt{|\mathcal{R}| \cdot \log(|V_S| \cdot |\mathcal{T}|)})$ and $\gamma \in \mathcal{O}(\varepsilon \cdot \max_{r \in \mathcal{R}} |E_r| \cdot \sqrt{|\mathcal{R}| \cdot \log(|E_S|)})$, respectively.

To overcome estimating β and γ by the above (or similar) bounds, we propose the following to apply randomized rounding in practice. Concretely, we consider two types of adaptations to our Algorithm 2. The first adaptation is to discard the consideration of respective approximation factors α, β, γ and simply returning the best solution found within a fixed number of rounding iterations. We refer to this type of rounding heuristic as *vanilla rounding*, as the rounding procedure itself is not changed. The other type of adaptation is the avoidance of resource augmentations by considering *heuristic rounding*.

Concretely, we consider two *vanilla rounding* heuristics, namely $\text{RR}_{\text{MinLoad}}$ and $\text{RR}_{\text{MaxProfit}}$, as well as one *heuristic rounding* variant $\text{RR}_{\text{Heuristic}}$, as described below:

$\text{RR}_{\text{MinLoad}}$ A fixed number of solutions is rounded at random as in Line 7 of Algorithm 2. The solution minimizing

resource augmentations (and among those the one of the highest profit) is selected.

RR_{MaxProfit} A fixed number of solutions is rounded at random as in Line 7 of Algorithm 2. The solution maximizing the profit (and among those the one of the least resource augmentations) is selected.

RR_{Heuristic} Line 7 of Algorithm is adapted in such a way that selected mappings are only accepted, when its incorporation into the solution does not exceed resource capacities. In other words, if a mapping m_r^k is selected for request $r \in \mathcal{R}$ whose addition would exceed any resource capacity, then this mapping is simply discarded and request r is not embedded. To increase the diversity of found solutions, the order in which requests are processed is permuted before each rounding iteration. A fixed number of solutions is rounded and the one maximizing the profit is returned.

Besides the above heuristics, we will also study the best solution obtainable via randomized rounding while respecting resource capacities. Specifically, similar to Jarray et al. [10], we propose to use a Multi-Dimensional Knapsack (MDK) Integer Programming formulation (see IP 3) as follows. Given the set of decomposed mappings $\mathcal{D}_r = \{(f_r^k, m_r^k)\}_k$ for each request, the MDK formulation introduces a single binary variable $x_r^k \in \{0, 1\}$ per mapping m_r^k of request r while ensuring that at most one of the mappings of a request is selected and that capacities are not violated. As each decomposed mapping m_r^k has a probability of $f_r^k > 0$ to be chosen, the MDK allows to compute the best possible solution that may be attained by rounding the solution while having an *exponential* runtime.

IP 3: Multi-Dimensional Knapsack – Optimal Rounding

$$\max \sum_{r \in \mathcal{R}, (\cdot, m_r^k) \in \mathcal{D}_r} b_r \cdot x_r^k \quad (24)$$

$$\sum_{(\cdot, m_r^k) \in \mathcal{D}_r} x_r^k \leq 1 \quad \forall r \in \mathcal{R} \quad (25)$$

$$\sum_{r \in \mathcal{R}, (\cdot, m_r^k) \in \mathcal{D}_r} x_r^k \cdot A(m_r^k, x, y) \leq d_S(x, y) \quad \forall (x, y) \in R_S \quad (26)$$

$$x_r^k \in \{0, 1\} \quad \forall (\cdot, m_r^k) \in \mathcal{D}_r \quad (27)$$

V. EXPLORATIVE COMPUTATIONAL STUDY

We now complement our formal approximation result in the standard multi-criteria model with resource augmentation with an extensive computational study. Specifically, we study the performance of vanilla rounding and heuristical rounding (i.e., without resource augmentations) as introduced above.

As we are not aware of any systematic evaluation of the profit maximization in the offline settings, we present a synthetic but extensive computational study. Following existing works [7], [8], we consider randomly generated (cactus) request graphs while employing real-world wide-area topologies as substrate networks. For our study, we have created 7,500 offline VNEP instances, consisting of up to 100 requests. To study the approximation guarantees and resource augmentations under varying demand-to-capacity ratios and under

TABLE I
SUMMARY OF USED SUBSTRATE NETWORKS

Name	Identifier	Type	Year	$ V_S $	$ E_S $
Deutsche Telekom	DT	Global	2010	30	110
NTT	NT	Global	2011	32	126
Geant	GE	Continent	2012	40	122
UUnet	UN	Country	2011	49	168
Surfnet	SN	Country	2010	50	136

different resource scarcities, the instances are generated according to a large parameter space. To analyze the performance of the proposed algorithm, baseline solutions were computed by solving the Formulation 1 for all instances.

We have implemented all presented algorithms in Python 2.7, employing Gurobi 8.1.1 to solve the Mixed-Integer Programs and Linear Programs. Our source code is freely available at [21]. All experiments were executed on a server equipped with Intel Xeon E5-4627v3 CPUs running at 2.6 GHz and reported runtimes are wall-clock times.

A. Instance Generation

a) *Substrate Graph*: We use five different substrate networks obtained from the Internet Topology Zoo [22]. Specifically, wide-area network topologies of between 30 and 50 nodes and between 110 and 168 edges are considered. The topologies span single countries, a continent, or the whole world (see Table I). We consider a single node type and set node and edge capacities uniformly to 100.

b) *Request Topology Generation*: Cactus graph requests are generated by (i) sampling a random binary tree of maximum depth 3, (ii) adding additional edges randomly unless they do refute the cactus property *as long as edges can be added*, and (iii) orienting the created edges arbitrarily.

Concretely, the sampling process of binary trees works as follows: starting with an initial root node, the number of children is drawn using the discrete distribution $\mathbb{P}(\#\text{children} = 0) = 0.15$, $\mathbb{P}(\#\text{children} = 1) = 0.5$, and $\mathbb{P}(\#\text{children} = 2) = 0.35$. For each (newly) generated node (of depth less than 3) further children are generated according to the same distribution. We discard graphs having less than 3 nodes. According to the above generation procedure, the

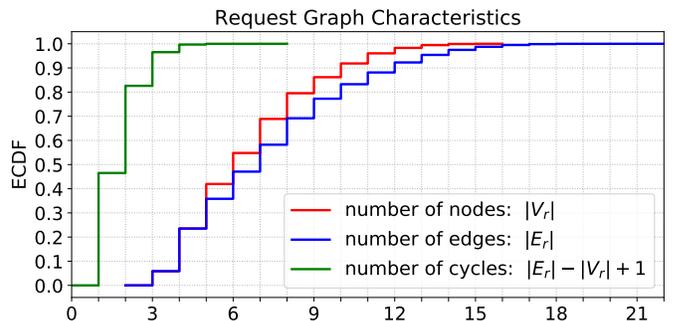


Fig. 4. Characteristics of the generated cactus requests graphs, namely the number of nodes, edges and number of cycles. The depicted empirical cumulative distribution function (ECDF) is based upon 100k sampled cactus requests graphs.

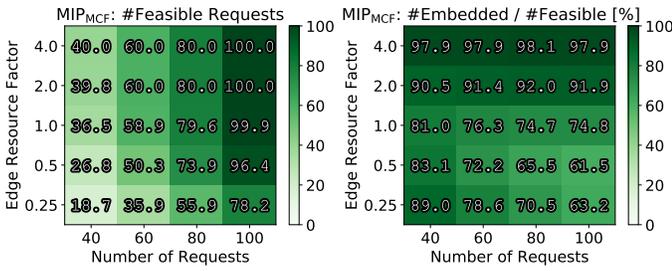


Fig. 5. Overview of the feasibility of generated requests and the baseline's acceptance ratio. Each cell averages the result of 375 instances.

Left: The feasibility of requests is obtained from optimally embedding the requests to compute the profit a priori. Note that absolute numbers are given. Right: The acceptance ratio of feasible requests by the baseline MIP_{MCF}.

expected number of nodes and edges is 6.54 and 7.28, respectively. On average, 61% of the edges lie on a cycle. Figure 4 offers an in-depth view on the characteristics of the requests.

c) Mapping Restrictions: To force the virtual networks to span across the whole substrate network, we restrict the mapping of virtual nodes to one quarter of the substrate nodes. Hence, the mapping of virtual nodes is restricted to ten substrate nodes. The mapping of virtual edges is not restricted.

d) Demand Generation: We control the demand-to-capacity ratio of node and edge resource using a node resource factor NRF and an edge resource factor ERF. The requests' demands are drawn from an exponential distribution and afterwards normalized, such that the following holds:

$$\sum_{r \in \mathcal{R}} \sum_{i \in V_r} d_r(i) = \text{NRF} \cdot \sum_{u \in V_S} d_S(u) \quad (28)$$

$$\text{ERF} \cdot \sum_{r \in \mathcal{R}} \sum_{(i,j) \in E_r} d_r(i,j) = \sum_{(u,v) \in E_S} d_S(u,v) \quad (29)$$

The resource factors can be best understood under the assumption that all requests are embedded. Under this assumption, a resource factor NRF = 0.6 implies that the node load – averaged over all substrate nodes – equals exactly 60%. As virtual edges can be mapped on arbitrarily long paths (even of length 0), the edge resource factor should be understood as follows: the ERF equals 'the number of substrate edges that each virtual edge may use'. In particular, a factor ERF = 0.5 implies that if *each* virtual edge spans *exactly* 0.5 substrate edges (and all requests are embedded), then the (averaged) edge resource utilization equals exactly 100%. Hence, while increasing the NRF renders node resources more scarce, increasing the ERF increases the available bandwidth resources.

e) Profit Computation: To correlate the profit of a request with its size, its resource demands, and its mapping restrictions, we compute for each request its minimal embedding costs as follows. The cost $c(u, v)$ of using an edge $(u, v) \in V_S$ equals the geographical distance of its endpoints. The cost of nodes is set uniformly to $c(\cdot, u) = \sum_{(u,v) \in E_S} c(u, v) / |V_S|$ for all $u \in V_S$. Hence, the total node cost equals the total edge cost. Defining the cost of a mapping m_r to be $\sum_{(x,y) \in R_S} A(m_r, x, y) \cdot c(x, y)$, we compute the minimum cost embedding for each request $r \in \mathcal{R}$ using an adaption of Mixed-Integer Program 1 and set b_r accordingly.

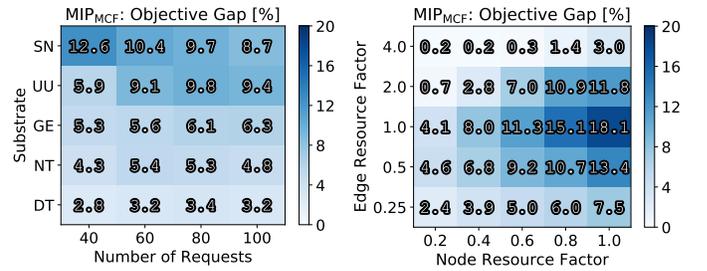


Fig. 6. Overview of the objective gap achieved by the baseline algorithm MIP_{MCF} after up to 2 hours of computation time. Note the different x-axes. The left plot averages the results of 375 and the right one the results of 300 instances per cell. The number of requests, i.e., the problem size, has a less distinct impact on the objective gap than the resource factors alone.

f) Parameter Space: Besides the five different substrate topologies (see Table I), the following parameters are used:

- $|\mathcal{R}| \in \{40, 60, 80, 100\}$,
- $\text{NRF} \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$,
- $\text{ERF} \in \{0.25, 0.5, 1.0, 2.0, 4.0\}$.

Generating 15 instances per parameter combination, 7,500 instances are obtained overall.

B. Computational Results

We first present our baseline results computed by solving the MCF Integer Programming Formulation 1 and then study the performance of vanilla rounding and heuristical rounding.

a) Baseline MIP_{MCF}: To obtain a near-optimal baseline solution for each of the 7,500 instances, we employ Gurobi 8.1.1 to solve the Mixed-Integer Programming Formulation 1 (using a single thread). We terminate the computation after 2 hours or when the *objective gap* falls below 1%, i.e., when the constructed solution is *provably* less than 1% off the optimum. On average the runtime per instance is 80.6 minutes (cf. Figure 8, right).

Figure 5 gives an initial overview of the number of requests for which feasible embeddings exist and the acceptance ratio of the best solution as a function of the number of requests and the edge resource factor. The number of feasible requests is determined during the a priori profit computation and may (on average) lie below 50% when edge resources are very scarce (ERF = 0.25) but lies otherwise consistently above 75%. Similarly, the acceptance ratio of the baseline solution highly depends on the edge resource factor, ranging from close to 63% to roughly 98% (on average).

Figure 6 depicts quality guarantees for the baseline solutions obtained during the solution process of IP Formulation 1. The formulation is solved using Gurobi's branch-and-bound implementation, consistently yielding upper bounds on the attainable profit (by LP relaxation). Accordingly, the objective gap depicted in Figure 6 gives guarantees on how far (at most) the found solutions are off optimality (on average). While increasing the number of requests does not increase the objective gap per se, both the node and edge resource factors have a distinct impact. Particularly, for the *maximal* node resource factor of 1.0 and *medium* edge resource factor 1.0, the averaged objective gap lies slightly above 18%. The

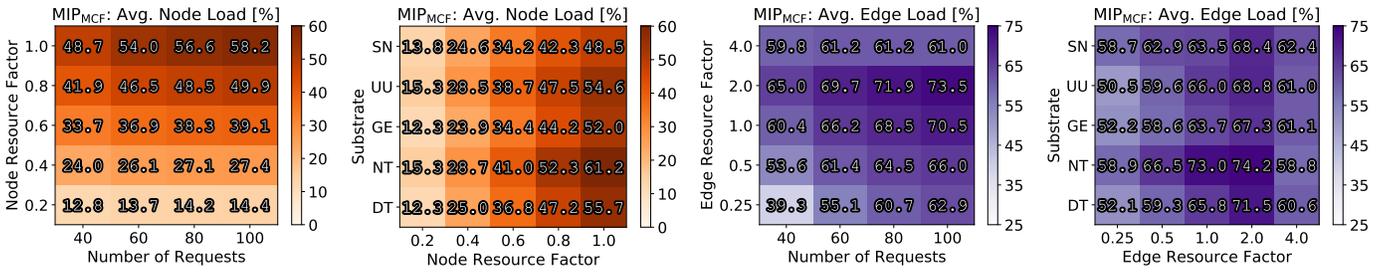


Fig. 7. Overview of the resource loads of the solutions computed by the baseline algorithm MIP_{MCF} . Depicted are the averaged node and edge loads as a function of the respective node and edge resource factors and the number of requests or the substrate topology. Each cell averages the results of 375 or 300 solutions, respectively. While the node factor has a distinct impact on the node loads, the averaged edge loads lie consistently between 55% and 75%, allowing the conclusion that edge resources often represent the bottleneck resource. Furthermore, the substrate topology has only a minor impact.

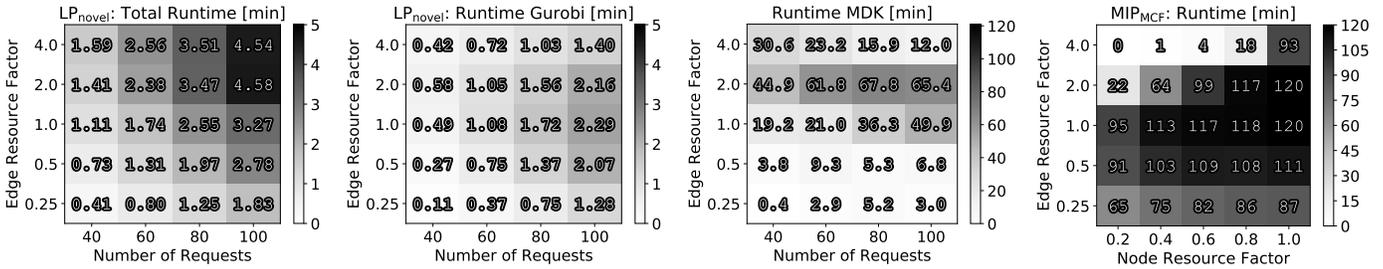


Fig. 8. Overview of the runtimes of the novel LP Formulation 2, the MDK IP 3, and of the baseline algorithm MIP_{MCF} . The three plots on the left average 375 results per cell while the right one averages 300 per cell. Note the different x-axes as well as the different scales.

Left: The runtime of the novel LP Formulation 2 including the creation of the LP and the decomposition of the solutions.

Center, left: The pure runtime of Gurobi 8.1.1 for solving the LP, i.e., pre- and post-processing times are excluded.

Center, right: Runtime for solving the MDK Integer Program 3 using Gurobi 8.1.1 Computations are terminated after 120 minutes.

Right: The runtime of the MCF Integer Program 1 being solved by Gurobi 8.1.1 Computations are terminated after 120 minutes.

mean objective value is 6.6% and the maximum observed gap across all 7,500 instances less than 62%.

Figure 7 validates the impact the node and edge resource factors have on the respective resource loads. Depicted are the averaged node and edge loads as a function of the respective resource factors or the topologies and the number of requests. As can be seen, increasing the node resource factor increases the averaged node resource loads. For the edge resource loads, the picture is a different one. In particular, the averaged edge load lies with 40% to 75% significantly above the averaged node load. Interestingly, the edge loads increase when making more edge resources available, indicating that for very low edge resource factors only some edges can at all be fully utilized. Furthermore, as is shown in Figure 10 (left), both the maximal node and edge resource loads lie above 99% for more than 90% of the instances, i.e., for each instance there exist both some bottleneck node resource as well as some bottleneck edge resource.

b) Solving LP Formulation 2: To apply the rounding algorithms presented in Section IV-D, we solve our novel LP Formulation 2 by employing again Gurobi 8.1.1, specifically its Barrier algorithm. Figure 8 (left) depicts the averaged runtime to solve the LP, including the time to construct the (potentially large) LP and the time to decompose the solution into convex combinations. This pre- and post-processing is not negligible as the formulation contains up to 1,000k variables. In fact, solving the LP itself is on average faster than the respective processing steps (see Figure 8, center left). The runtime increases from less than half a minute for $|\mathcal{R}| = 40$ to

around 4.5 minutes for $|\mathcal{R}| = 100$. The maximally observed runtime in our experiments amounted to roughly 15 minutes.

c) Vanilla Rounding: With respect to the results of our rounding heuristics, we first discuss the results of our vanilla rounding heuristics $RR_{MinLoad}$ and $RR_{MaxProfit}$. Concretely, we report on the best solution found within 1,000 rounding iterations. Figure 11 depicts the respective results as a scatter plot while Figure 10 depicts empirical cumulative distribution functions (ECDF) for both the achieved profit and the maximal resource loads. As can be seen, for $RR_{MinLoad}$ the algorithm achieves a profit between 50% and 140% compared to the best solution constructed by the MIP, while exceeding resource capacities mostly by 150% of the resource's capacity. For $RR_{MaxProfit}$, the achieved profit always exceeds the baseline's profit. This is to be expected based on our analysis in Section IV-A, as the expected benefit is at least as large as the optimal Mixed-Integer Programming formulation's value. The resource loads mostly lie below 500% with the maximum being 871% (single data point).

For both selection criteria, the edge resource factor has a distinct impact on the (overall) maximum load. This can be explained as follows. As each request edge may use any of the substrate edges (compared to the restricted node mappings), the chances of 'collisions', i.e., multiple request edges being mapped on the same substrate edge, is higher. Additionally, the fact that the number of virtual edges is (always) at least as large as the number of nodes (cf. Figure 4) and that a single request edge may use *multiple* substrate edges, complicates the rounding of edge mappings. This observation is further

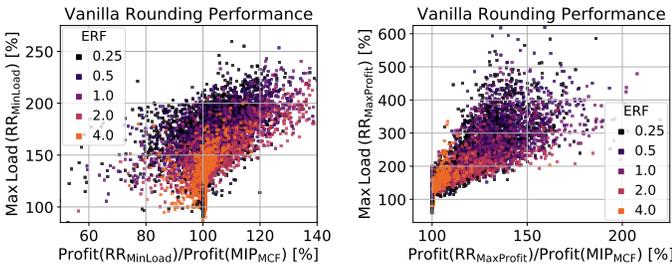


Fig. 9. Depicted are the solutions obtained via the two *vanilla* rounding schemes RR_{MinLoad} (left) and $RR_{\text{MaxProfit}}$ (right). Each point corresponds to a single instance and is colored according to the instance’s edge resource factor. The left and right plots shows results for 7,458 and 7,499 of the 7,500 instances, respectively; the other results lie outside the depicted area.

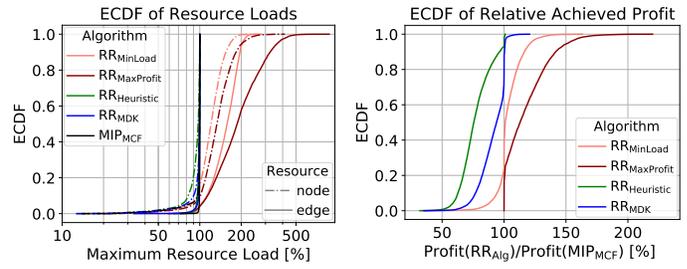


Fig. 10. Comparison of the different algorithms in terms of (maximal) resource usage and the relative achieved profit for all 7,500 results. Left: ECDF of maximal node and edge resource loads. Right: ECDF of the profit achieved by the randomized rounding algorithms compared to the profit achieved by the baseline solution.

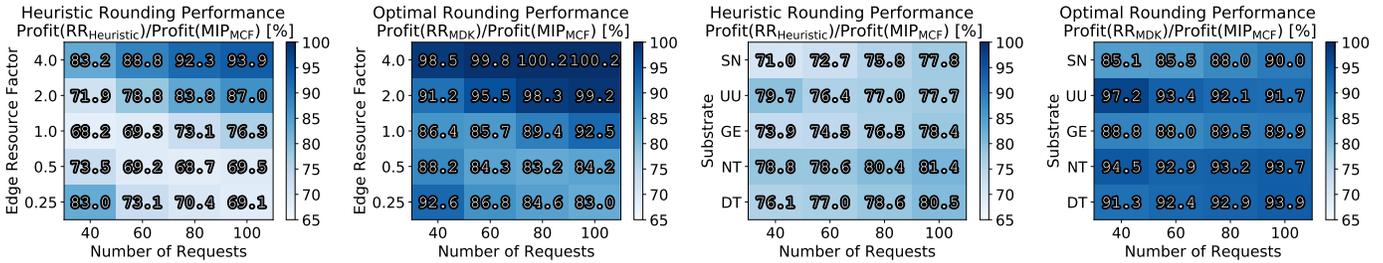


Fig. 11. Overview of the averaged relative performance achieved by the heuristic rounding algorithm and the MDK Integer Program 3 as a function of the edge resource factor and the number of requests (left) and the substrate topologies (right). Each cell averages the result of 375 results. Left: Both the performance of the heuristical rounding and the MDK generally depend on the edge resource factor as well as the number of requests. Right: The performance of the heuristical rounding and the MDK are consistent among the different substrates.

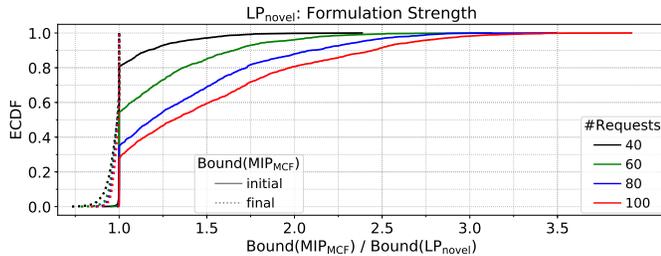


Fig. 12. Comparison of the bounds on the profit computed by the novel LP Formulation 2 and the classic MCF Formulation 1. As objective bounds are continuously improved during the solution process of the Mixed-Integer Program MIP_{MCF} , we report on the initial (weakest) bound and the final (best) bound. The initial bound is essentially the objective value of the LP relaxation of Formulation 1, but might be improved by the solver Gurobi based on the introduction of cutting planes valid only for the integer variant.

substantiated by the ECDF presented in Figure 10: the maximal edge loads are consistently larger than the maximal node resource load for all randomized rounding algorithms.

d) Heuristical Rounding $RR_{\text{Heuristic}}$: The results of the heuristical rounding, which does not exceed resource capacities, are presented in Figure 10 (ECDFs) and Figure 11 (heatmaps). Again, 1,000 rounding iterations were considered. While for low edge resource factors, i.e., scarce edge resources, the solutions achieve around 68% of the profit of the MIP baseline, for larger edge resource factors, the relative performance exceeds 80%. For edge resource factors above 0.5, the performance of the heuristic notably improves with the number of requests. This can be explained as follows. For any combination of resource factors, the virtual resource demands are computed according to Equations 20 and 21

independently of how many requests are considered. Hence, when increasing the number of requests, the resource demands of each single request become smaller. Hence, the maximal demand-to-capacity ratio $d_{\text{max}}(r, x, y) / d_S(x, y)$ decreases for *all* substrate resources $(x, y) \in R_S$ when increasing the number of requests. Accordingly, the value of ε (cf. Lemma 16 and Corollary 17) can be chosen smaller and the expected resource augmentations decrease. Thus, the heuristical rounding algorithm is able to construct better solutions more easily. Overall, the average relative performance with respect to the baseline solutions is 77.2%, with the minimal one being 32.1% and less than 2% of constructed solutions achieving less than 50% of the baseline’s profit (cf. Figure 10).

e) Optimal Rounding Solution RR_{MDK} : We lastly discuss the results of executing the Multi-Dimensional Knapsack (MDK) Integer Program 3. As for the baseline IP, solutions were computed using Gurobi 8.1.1 and computations were terminated when an objective gap of less than 1% was reached or after 2 hours. As shown in Figure 11, the optimal rounding solution improves upon the heuristical rounding solutions significantly: the average relative profit is 91.2%, yielding an average improvement over the heuristical rounding by 14%. Interestingly, for the maximal edge resource factor and 80 and 100 requests, the solutions found by the MDK slightly improve upon the baseline solutions, with the maximal improvement over the baseline being 20.6% (cf. Figure 10). Since the runtime of the MDK (cf. Figure 8) consistently and at times significantly lies beneath the runtime for constructing the baseline solutions, the MDK may pose an interesting alternative to computing solutions using the MCF IP formulation.

C. Comparison of Formulation Strengths

Lastly, we empirically study the strength of our novel LP Formulation 2 and compare it with the classic MCF Formulation 1. Concretely, as noted in Theorems 8 and 9, the classic MCF formulation has an unbounded or very large integrality gap in general, i.e., the bound on the profit returned by the solution can be arbitrarily far off the optimal attainable profit. Our novel LP formulation is provably stronger than the old formulation, as it only allows for (fractional) solutions, which can be decomposed into valid mappings. Thus, it will always yield (equal or) better bounds on the attainable profit.

Figure 12 presents the experimental comparison of both formulations. In particular, for each of the 7,5000 instances we compare the objective of our novel *Linear Programming* Formulation 2 to *two* bounds computed during the solution process of the baseline MIP_{MCF} : the initial bound, i.e., the objective of the *Linear Programming* Formulation 1, and the final (best) bound computed during the solution process of the (Mixed-)Integer Program.

As can be seen, the initial LP bounds of the classic formulation at times exceeds our formulation's objective by more than 300%, i.e., the classic formulation 'overestimates' the maximal attainable profit by more than a factor of 3. For roughly 60% of the instances, the novel LP improves the bounds by a factor of 1.5. Clearly, the more requests are considered, the less accurate the classic MCF formulation is. Considering the final bound computed during the execution of the (Mixed-)Integer Program, we see that these bounds always improve upon the novel LP's bound. However, for 90% of the instances, the novel LP's bound is only improved by less than 10% with the maximal improvement being less than 28%.

Concluding, we note that our novel LP formulation is much *stronger* than the classic formulation: it consistently yields significantly better bounds in practice compared to the classic LP formulation and comes close to the bounds obtained by solving the (Mixed-)Integer Program for up to 2 hours.

VI. CONCLUSION

This paper has initiated the study of approximation algorithms for the Virtual Network Embedding Problem supporting arbitrary substrate graphs and supporting *cyclic* request graphs, specifically cactus request graphs. To obtain the approximation result, we have derived a novel strong LP formulation and have mathematically proven performance guarantees in the resource augmentation model. Our computational evaluation shows the practical significance of our work: solutions obtained by our derived heuristic achieve (on average) 77.2% of the baseline's profit while *not augmenting capacities*. Furthermore, our computational study has shown that optimally rounded solutions achieve on average 91.2% of the baseline's profit, while even exceeding it at times. Hence, future work might investigate improvements of the presented heuristical rounding procedure, to come closer to the optimal rounding.

Lastly, we note that the developed approximation framework – including the proposed rounding heuristics – is independent of how LP solutions are computed and decomposed. In particular, while the LP formulation presented in this paper is

only applicable for cactus request graphs, our approach can be generalized to arbitrary request graphs as shown in [17], [18], albeit coming at the cost of potentially exponential runtimes.

Acknowledgements. This work and its dissemination efforts were supported in part by the BMBF Software Campus grant 01IS1205 and the European Research Council (ERC) grant ResolutioNet (ERC-StG-679158). We thank Elias Döhne, Alexander Elvers, and Tom Koch for their contribution to our implementation [21].

REFERENCES

- [1] M. Rost and S. Schmid, "Virtual Network Embedding Approximations: Leveraging Randomized Rounding," in *Proceedings of the 2018 IFIP Networking Conference (IFIP Networking)*, May 2018.
- [2] J. C. Mogul and L. Popa, "What we talk about when we talk about cloud network performance," *SIGCOMM Computer Communication Review*, vol. 42, no. 5, pp. 44–48, Sep. 2012.
- [3] J. Napper, W. Haeffner, M. Stiernerling, D. R. Lopez, and J. Uttaro, "Service Function Chaining Use Cases in Mobile Networks," Internet-Draft, Apr. 2016. [Online]. Available: <https://tools.ietf.org/html/draft-ietf-sfc-use-case-mobility-06>
- [4] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron, "Towards predictable datacenter networks," in *Proceedings of the ACM SIGCOMM 2011 Conference*. New York, NY, USA: ACM, 2011, pp. 242–253.
- [5] S. Mehraghdam, M. Keller, and H. Karl, "Specifying and placing chains of virtual network functions," in *Proc. IEEE CloudNet*, October 2014.
- [6] M. Rost and S. Schmid, "Charting the Complexity Landscape of Virtual Network Embeddings," in *Proceedings of the 2018 IFIP Networking Conference (IFIP Networking)*, May 2018.
- [7] A. Fischer, J. F. Botero, M. T. Beck, H. De Meer, and X. Hesselbach, "Virtual network embedding: A survey," *IEEE Communications Surveys Tutorials*, vol. 15, no. 4, pp. 1888–1906, Fourth 2013.
- [8] M. Chowdhury, M. R. Rahman, and R. Boutaba, "Vineyard: Virtual network embedding algorithms with coordinated node and link mapping," *IEEE/ACM Trans. Netw.*, vol. 20, no. 1, pp. 206–219, Feb. 2012.
- [9] R. Hartert, S. Vissicchio, P. Schaus, O. Bonaventure, C. Filsfil, T. Telkamp, and P. Francois, "A declarative and expressive approach to control forwarding paths in carrier-grade networks," in *Proceedings of the ACM SIGCOMM 2015 Conference*. New York, NY, USA: ACM, 2015, pp. 15–28.
- [10] A. Jarray and A. Karmouch, "Decomposition approaches for virtual network embedding with one-shot node and link mapping," *IEEE/ACM Trans. Netw.*, vol. 23, no. 3, pp. 1012–1025, Jun. 2015.
- [11] M. Rost, S. Schmid, and A. Feldmann, "It's about time: On optimal virtual network embeddings under temporal flexibilities," in *Proceedings of the International Parallel and Distributed Processing Symposium*, May 2014, pp. 17–26.
- [12] M. Rost, C. Fuerst, and S. Schmid, "Beyond the stars: Revisiting virtual cluster embeddings," *SIGCOMM Computer Communication Review*, vol. 45, no. 3, pp. 12–18, Jul. 2015.
- [13] G. Even, M. Medina, and B. Patt-Shamir, "Online path computation and function placement in sdn," in *Proc. International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, 2016.
- [14] G. Even, M. Rost, and S. Schmid, "An approximation algorithm for path computation and function placement in sdn," in *Structural Information and Communication Complexity*, J. Suomela, Ed. Cham: Springer International Publishing, 2016, pp. 374–390.
- [15] T. Lukovszki and S. Schmid, "Online admission control and embedding of service chains," in *Structural Information and Communication Complexity*. Cham: Springer International Publishing, 2015, pp. 104–118.
- [16] N. Bansal, K. Lee, V. Nagarajan, and M. Zafer, "Minimum congestion mapping in a cloud," *SIAM Journal on Computing*, vol. 44, no. 3, pp. 819–843, 2015.
- [17] M. Rost and S. Schmid, "(FPT-)Approximation Algorithms for the Virtual Network Embedding Problem," Tech. Rep. arXiv:1803.04452 [cs.NI], March 2018.
- [18] M. Rost, E. Döhne, and S. Schmid, "Parametrized complexity of virtual network embeddings: Dynamic & linear programming approximations," *SIGCOMM Computer Communication Review*, vol. 49, no. 1, pp. 3–10, Feb. 2019.
- [19] P. Raghavan and C. D. Thompson, "Provably good routing in graphs: Regular arrays," in *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, ser. STOC '85. New York, NY, USA: ACM, 1985, pp. 79–87.

- [20] D. P. Dubhashi and A. Panconesi, *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [21] E. Döhne, A. Elvers, T. Koch, and M. Rost, "Github repository," <https://github.com/vnep-approx/evaluation-ieee-acm-ton-2019>, 2019.
- [22] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The internet topology zoo," *Selected Areas in Communications, IEEE*, vol. 29, no. 9, 2011.

APPENDIX

Theorem 8. *The integrality gap of the MCF formulation is unbounded. This even holds under infinite substrate capacities.*

Proof. Consider the example of Figure 3. We introduce the following restrictions for mapping the virtual links: $E_S^{r,i,j} = \{(u_1, u_2), (u_4, u_5)\}$, $E_S^{r,j,k} = \{(u_2, u_3), (u_5, u_6)\}$, $E_S^{r,k,i} = \{(u_3, u_4), (u_6, u_1)\}$. Note that the LP solution depicted in Figure 3 is still feasible and hence the LP will attain an objective of b_r .

On the other hand, there does not exist a valid mapping of request r : assume i to be mapped on u_1 , then j must be mapped on u_2 and k must be mapped on u_3 due to the node and edge mapping. However, the edge mapping restrictions for (k, i) do not allow the establishment of a path from u_3 to u_1 . By the same argument, the mapping of i on u_4 is not feasible. Accordingly, as the optimal solution achieves a profit of 0 while the LP solution yields a profit of b_r , the integrality gap of the MCF LP Formulation 1 is unbounded. \square

Theorem 9. *The integrality gap of the MCF formulation lies in $\Omega(|V_S|)$, when only considering node mapping restrictions.*

Proof. Consider the following instance. The substrate is a cycle of an even number of n nodes u_i , with $1 \leq i \leq n$ and edges $\{(u_1, u_2), (u_2, u_3), \dots, (u_{n-1}, u_n), (u_n, u_1)\}$. We consider unit edge capacities, i.e., we set $d_S(e) = 1$ for $e \in E_S$. Consider now a request r with $V_r = \{i, j\}$ and edges $E_r\{(i, j), (j, i)\}$ and unit bandwidth demands $d_r(i, j) = d_r(j, i) = 1$. Assume that i can only be mapped on substrate nodes having an uneven index and that j can only be mapped on substrate nodes of even index. Clearly, any valid mapping of this request will use *all* edge resources. Consider now the following MCF solution with $x_r = 1$: nodes i and j are, together with the respective edges, mapped in an alternating fashion: $y_{r,i}^u = y_{r,j}^v = 2/n$ holds for all $u \in V_S^{r,i} = \{u_1, u_3, \dots\}$ and all $v \in V_S^{r,j} = \{u_2, u_4, \dots\}$, respectively. Similarly, $z_{r,i,j}^e = z_{r,j,i}^{e'} = 2/n$ is set for all edges e and e' originating at uneven and even nodes, respectively. Hence, the allocation equals $2/n$ on each edge. Hence, $n/2 \in \Theta(|V_S|)$ many copies of this request may be embedded in the MCF solution, while the optimal solution may only embed a single request and the integrality gap therefore lies in $\Omega(|V_S|)$. \square

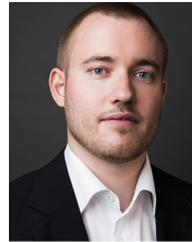
Lemma 10. *Consider a cactus request graph G_r and an acyclic reorientation G_r^A of G_r . The graph G_r^A can be uniquely partitioned into subgraphs $\{G_r^{A,C_1}, \dots, G_r^{A,C_n}\} \sqcup G_r^{A,\mathcal{F}}$, such that the following holds:*

- 1) *The subgraphs $\{G_r^{A,C_1}, \dots, G_r^{A,C_n}\}$ correspond to the (undirected) cycles of G_r and $G_r^{A,\mathcal{F}}$ is the forest remaining after removing the cyclic subgraphs. We denote the index set of the cycles by $\mathcal{C}_r = \{C_1, \dots, C_n\}$.*

- 2) *The subgraphs partition the edges of E_r^A : an edge $(i, j) \in E_r^A$ is contained in exactly one of the subgraphs.*
- 3) *The edge set E_r^{A,C_k} of each cycle $C_k \in \mathcal{C}_r$ can itself be partitioned into two branches $\mathcal{B}_1^{C_k}$ and $\mathcal{B}_2^{C_k}$, such that both lead from $s_r^{C_k} \in V_r^{A,C_k}$ to $t_r^{C_k} \in V_r^{A,C_k}$.*

Proof. We prove that the lemma holds independently of the chosen acyclic reorientation. Accordingly, let G_r^A be an arbitrary acyclic reorientation of G_r . We first show that $|\delta^-(i)| \leq 2$ holds for each virtual node $i \in V_r$ with respect to the edge set E_r^A , i.e., any virtual node has at most 2 incoming edges in G_r^A . If $|\delta^-(i)| > 2$ held, then by the definition of the acyclic reorientation there must exist at least 3 paths P_1, P_2, P_3 from the root r_r to i . Let $p_{1,2}, p_{2,3} \in V_r$ be the last common nodes lying on both P_1 and P_2 and P_2 and P_3 , respectively. Clearly, from $p_{1,2}$, there exist two (otherwise) node-disjoint paths to i and from $p_{2,3}$ there exist two (otherwise) node-disjoint paths towards i . Accordingly, there exist at least two cycles intersecting either in i and $p_{1,2}$ or i and $p_{2,3}$, which contradicts the assumption on G_r being a cactus graph. Accordingly, $|\delta^-(i)| \leq 2$ holds for all virtual nodes $i \in V_r$ according to E_r^A .

Now, consider any node $i \in V_r$ with $|\delta^-(i)| \leq 2$. By performing a graph-search in the opposite direction of edges in E_r^A , a unique common ancestor node i' can be determined, such that there exist exactly two paths \mathcal{B}_1 and \mathcal{B}_2 (branches) from i' to i . The union of these branches represents a single 'cycle' (cf. Statement 3). Removing the identified cycle from the graph, the cactus graph property still holds, as removing edges can never refute it. Accordingly, 'cycles' in the acyclic reorientation $G_r^{A,C_1}, \dots, G_r^{A,C_n}$ can be uniquely identified (decomposed) and after repeated removal only the forest $G_r^{A,\mathcal{F}}$ remains. Hence, the first statement of the lemma follows. Lastly, the second statement of the lemma holds trivially as each edge is either contained in any of the cycles or is part of the remaining forest. \square



Matthias Rost is a postdoctoral researcher at Technische Universität Berlin, Germany, from which he also received his MSc (2014) and his PhD (2019). He was awarded the KuVS prize for his master thesis on computing virtual aggregation and multicast trees by the German Informatics Society. His main research focus is the theoretical design of provably good algorithms for service orchestration in networks and their application in practice.



Stefan Schmid is a Professor at the University of Vienna, Austria. He received his MSc (2004) and PhD (2008) from ETH Zurich, Switzerland. Subsequently, Stefan Schmid worked as postdoc at TU Munich and the University of Paderborn (2009). From 2009 to 2015, he was a senior research scientist at the Telekom Innovations Laboratories (T-Labs) in Berlin, Germany, and from 2015 to 2018 an Associate Professor at Aalborg University, Denmark. His research interests revolve around the algorithmic problems of networked and distributed systems.