# HIPER: Heuristic-based Infrastructure Expansion through Partition Reconnection for Efficient Virtual Network Embedding

Marcelo Caggiani Luizelli, Leonardo Richter Bays,
Marinho Pilla Barcellos, Luciano Paschoal Gaspary
Institute of Informatics – Federal University of Rio Grande do Sul (UFRGS)
{mcluizelli,lrbays,marinho,paschoal}@inf.ufrgs.br

*Abstract*—As research in the area of network virtualization continues to advance, there have been numerous efforts to solve the challenge of efficiently mapping virtual networks on top of physical structures. Despite these efforts, current state-of-the-art proposals still suffer from significant amount of rejection of virtual network requests in circumstances where overall resource availability would be sufficient to embed them. This is caused by the exhaustion of resources in certain key points of the infrastructure. In this paper, we propose HIPER – a strategy for expanding physical networks that suggests infrastructure upgrades with the objective of maximizing the acceptance of virtual network requests (and, as a consequence, physical resource utilization). This is achieved through the reconnection of strongly connected components (i.e., recurring partitions) of the infrastructure. Evaluated under realistic workloads, HIPER led to promising results. After the expansion of 10% to 20% of infrastructure resources, HIPER sustained an increase of up to 30% in virtual network acceptance, allowing an additional 52% in resource utilization.

## I. Introduction

Network virtualization enables the creation of multiple, isolated virtual networks (VNs) on top of physical infrastructures. By leveraging the ability to easily (de)instantiate VNs, Infrastructure Providers (InPs) are able to utilize their physical resources in a more effective manner, while offering VNs that are tailored to the needs of different types of customers.

Efficiently embedding virtual networks on top of physical infrastructures represents a major challenge in the area of network virtualization (known as VNE – Virtual Network Embedding). Recent research in VNE includes the proposal of models in the context of provider networks, aiming at satisfying distinct requirements such as load balancing [1] and privacy [2], as well as attempts to employ such models in the context of data center networks (as a means of providing traffic isolation). Despite the existence of a number of efforts to solve the aforementioned problem [1]–[6], we still observe significant amounts of rejection of virtual network requests. In a previous evaluation [7], we verified that an expressive number of rejections occur in situations in which there are sufficient resources available in the infrastructure as a whole, but the exhaustion of resources in a few key points of the infrastructure (due to connectivity-related attributes) hinders the fulfillment of additional requests. As we explain next, it is definitely important to revisit networking planning strategies so as to take full advantage of timely, dynamic and customized provision/management capabilities made possible by network virtualization.

We are not aware of previous attempts to investigate strategies for expanding the physical network of an infrastructure provider with the objective of enabling it to host a higher number of virtual networks. We emphasize that classic approaches to physical infrastructure expansion [8]–[10] are not suitable for network virtualization environments. These conventional (and widely used) approaches assume that pairs of devices that require a higher amount of resources are known beforehand, and expansion is planned according to a demand matrix (which indicates where network bottlenecks are). However, demand matrices observed as a result of the creation of slices in InP networks[1] reveal comparatively more homogeneous resource distribution among pairs of physical devices (this phenomenon is demonstrated in Section 3, right after the Related Work section). This is due to the fact that virtual routers and links, in general, may be hosted on (practically) any physical device with sufficient available resources. As the location of physical devices does not influence the virtual network embedding process in a significant manner, this hinders the identification of bottlenecks. In summary, due to the flexible allocation mechanisms allowed by virtual network infrastructures, it becomes a challenge to determine substrate regions that need to be replanned so as to increase the probability of successfully embedding higher amounts of virtual network requests.

In this paper, we propose a strategy based on the expansion of network elements that tend to partition the physical infrastructure when their resources are nearly or fully depleted. Through this, we are able to reconnect these partitions (identified as strongly connected components), ultimately leading to sustained increases in virtual network acceptance rates in the long term. To prove concept and technical feasibility of the proposed strategy, we carried out an extensive series of experiments in order to observe aspects such as increases in acceptance rates and physical resource usage. The main contributions of this paper are then twofold: *(i)* a novel definition of the problem of physical infrastructure expansion in the context of network virtualization; and *(ii)* a comprehensive evaluation of the proposed strategy, demonstrating how the strengthening of key parts of an infrastructure may lead to significant improvements in terms of resource usage.

The remainder of this paper is organized as follows. Section 2 presents a discussion of related work in the area of infrastructure expansion. In Section 3, we discuss the contrast between demand matrices observed in conventional infrastructures and network virtualization environments. In Section 4, we

---

[1]Throughout this paper, the expressions *physical substrate*, *physical network*, *physical infrastructure*, and *InP network* are used interchangeably.

formalize the InP network expansion problem and propose an algorithmic approach to solve it. In Section 5, we present and discuss the results of our evaluation. Last, Section 6 concludes the paper with final remarks and perspectives for future work.

## II. Related Work

Next, we review some of the main prominent approaches to the problem of network planning and expansion. We highlight the contexts in which these approaches were applied, as well as the methodologies adopted by them. We begin by summarizing planning and expansion approaches in the context of backbone infrastructures and, afterwards, move on to more recent approaches that focus on data center networks.

### A. Backbone Network Expansion

Mukherjee et al. [8] propose an expansion model for optical networks which aims at minimizing the average delay between infrastructure nodes. The proposed model is based on the Simulated Annealing metaheuristic and the flow deviation algorithm. Ramaswami et al. [9] and Krishnaswamy et al. [10] also approach the optical network expansion problem. However, their objective is to minimize congestion between infrastructure nodes. In both approaches, expansions are made considering maximum delay limits and existing demands between source-destination pairs. The authors proposed optimization models based on Integer Linear Programming (ILP) considering different constraints (e.g., link delay). These models were evaluated on small backbone networks. In all these approaches [8]–[10], expansions are made based on a static demand matrix, which represents the historic average traffic between source-destination pairs in the infrastructure.

While the aforementioned approaches focus on more traditional objectives (e.g, minimizing delay), there is also recent work on backbone network expansion considering other aspects such as resilience and energy consumption. For example, Curtis et al. [11] and Johnston et al. [12] present models that aim at optimizing infrastructure device usage considering aspects related to resilience (e.g., protection against link failures). The proposed models are based on ILP and have as their main objective minimizing financial investments while ensuring that the planned infrastructure is able to satisfy a number of requirements related to quality of service. These models determine which elements need to be added to the physical infrastructure (as well as the required capacity of each added element). With respect to expansion in the context of green networking, Gangxiang et al. [13] propose a model for infrastructure planning that attempts to ensure that future demands will be met while minimizing energy consumption in backbone networks. The model considers the energy consumption of network components to determine in which regions of the infrastructure new resources should be added. For scalability reasons, the authors also propose heuristics in order to enable the employment of their strategies in large scale networks. We emphasize that these approaches also rely on demand matrices in order to plan the capacity of infrastructure links.

As briefly introduced earlier (and more deeply analyzed in the next section), demand matrices observed in InP networks tend to exhibit comparatively more homogeneous resource usage among pairs of physical devices. As a consequence of this homogeneity, the employment of demand matrices is not sufficient for adequately determining where to expand or how much to invest. Therefore, topological characteristics of the infrastructure become considerably more influential in the expansion process than demand matrices themselves. In a previous work [14], we performed a first iteration to solve this problem. We now better characterize the difference in terms of resource consumption incurred by VNE, detail the proposed model and perform an extensive set of experiments to prove technical feasibility and better understand the trade-offs resulting from different instantiation parameters.

### B. Data Center Network Expansion

In addition to approaches focusing on the expansion of backbone networks, there have been also recent proposals focusing on data center network planning. Among these, we highlight the work of Curtis et. al. [15] and Gao et al. [16], which introduce optimization models and heuristics applied to the expansion of data center networks. Similarly to approaches focusing on backbone networks, the authors consider adding devices and/or increasing the capacity of existing ones in order to cope with rapidly increasing infrastructure demands.

Solutions focusing on planning and expansion of data center networks, in general, target at maximizing simultaneous data throughput without interference between pairs of devices in the infrastructure (bisection bandwidth). These solutions are typically designed for specific topologies (such as trees), which makes it difficult (or even impossible) to adapt them to backbone networks. Even if we consider the possibility of adapting such solutions, their employment in the context of backbone networks would be inefficient in terms of resource allocation, as the objectives of the expansion models are different. In light of this new reality, identifying substrate regions that need to be replanned in backbone networks so as to increase the likelihood of successfully embedding new virtual network requests is a challenge, which we address in the following sections.

## III. Why are Demand Matrices not Suitable for Virtual Network-Oriented Infrastructure Planning?

A demand matrix represents the volume of traffic flow between pairs of devices on a network infrastructure. In conventional networks (i.e., without network virtualization support), it is possible to know beforehand which pairs of devices require a higher amount of resources (e.g., bandwidth). More specifically, demands for resources in conventional infrastructures occur among specific regions, leading to significantly higher values in certain parts of the demand matrix. Moreover, abrupt, longstanding variations in demands among regions are uncommon (i.e., demand variations usually occur gradually over long periods of time). These aspects allow the identification of bottlenecks in a more straightforward manner, simplifying the planning of physical device capacities.

In contrast, as previously explained, demand matrices in network virtualization environments exhibit a comparatively more homogeneous resource distribution (as virtual routers and links may be potentially hosted on any physical device with sufficient available resources), making it difficult to identify bottlenecks. Although there may be some location constraints in virtual networking scenarios, they are substantially less restrictive and, in general, apply to a lower number of virtual network elements. For example, locations in network virtualization scenarios may be defined as sets of physical routers and

links, allowing some flexibility even when mapping elements with location constraints. Furthermore, virtual networks can be easily migrated among different regions, causing demand variations in a faster and more dynamic manner (as opposed to the behavior observed in conventional infrastructures).

As the aforementioned assumption represents a fundamental aspect of the proposed strategy, we deemed it important to confirm it. To this end, we performed a series of experiments. First, to analyze the characteristics of demand matrices in conventional infrastructures, we embedded a number of VN requests on top of a physical infrastructure (following the model proposed by Luizelli et al. [7]). We consider that these requests have strict location requirements – i.e., all requested routers must be mapped to specific physical locations. The purpose of these requirements is to emulate the behavior observed in conventional networks (e.g., ISP networks) with respect to demands in specific areas of the infrastructure. In order to analyze the demand matrix of infrastructures with network virtualization support, the same set of requests embedded in the conventional infrastructure was mapped to a virtualized infrastructure, disregarding strict location requirements. This enabled virtual routers, in this case, to be potentially mapped to a number of different areas in the physical infrastructure (as long as there was resource availability). We emphasize that the evaluated scenarios represent extreme cases (as, for example, virtual networks may not always be mapped disregarding all physical device locations) but nevertheless are able to express the expected behavior in a faithful manner.

Figure 1 illustrates the distribution of demand among pairs of physical devices in each evaluation scenario (conventional and virtualized infrastructures). Values pertaining the demand matrix of the conventional infrastructure are normalized in relation to the scenario with virtualization support and arranged in increasing order. In conventional networks, due to location-related demands, there are regions in the infrastructure with considerably higher resource consumption than others. In scenarios with network virtualization, as there are no strict location requirements, resource consumption tends to be more homogeneous. This figure reveals that 15 source-destination pairs in conventional networks exhibit higher demands than the ones observed in virtualized infrastructures (between 1.01 and 4.39 times higher), 27 pairs exhibit equal demands in both cases, and the remainder – 93 pairs – exhibit inferior demands (between 0 to 0.97 times lower). As such, the figure demonstrates that resource distribution among demand pairs in virtual scenarios tends to be homogeneous, instead of heterogeneous, thus hindering the possibility of adapting classic network expansion strategies.



Fig. 1. Demand matrices observed in conventional networks and in infrastructures with network virtualization support.

## IV. InP Network Expansion for Efficient Virtual Network Embedding

In this section, we describe the InP network expansion problem and introduce our proposed solution. We first formalize it as an Integer Linear Programming model, and afterwards, present an algorithmic approach, named HIPER.

### A. Problem Overview

One of the major causes of virtual network rejection in the context of VNE is the absence of a suitable partition in the physical infrastructure [7]. A partition is, in essence, a set of routers isolated from the rest of the infrastructure. The occurrence of partitions is directly correlated to the total (or near-total) exhaustion of available resources in specific devices (e.g., bridges or hubs) in the embedding process. As physical resource usage varies (as a result of resource allocation/reservation), the set of partitions changes (and, consequently, the devices that belong to each partition). Thus, it is necessary to identify which network elements cause the most impact on this partitioning – i.e., elements that frequently partition the substrate network when their resources are fully (or near-fully) depleted. Figure 2(a) illustrates the state of a physical infrastructure in a certain moment in time. Throughout a number of previously performed virtual network mappings, four recurring partitions (represented as light gray circles) were identified. These partitions are caused by the depletion of resources of links $(a, b)$, $(c, d)$, $(f, e)$, and $(f, g)$, represented as dashed lines in the figure.

Considering that partitioning is one of the main causes of virtual network rejection and that the physical infrastructure, as a whole, has a significant amount of available resources, we propose the reconnection of recurring partitions as a strategy for InP network expansion. This strategy goes a step further in comparison with traditional approaches (which rely only on demand matrices) by taking into account topological factors. The network expansion process is subject to a number of constraints related to the infrastructure provider, such as the number of physical elements (links and routers) that can be expanded (referred throughout the text as expansion coverage) and the available financial resources. In this context, we aim at investigating which physical elements in InP networks (prioritized according to relevant topological features such as cut-edges) need to be replanned with the objective of reconnecting recurring partitions. Figure 2(b) represents a possible solution based on partition reconnection. The structural reinforcement (i.e., the points where expansion investments will be performed) of the path that interconnects the main regions of the infrastructure is highlighted. This path is composed of links $(a, b)$, $(a, f)$, $(f, e)$, $(f, g)$, $(b, c)$, and $(c, d)$, as well as routers $(a, b, c, d, e, f, g)$. We expect the proposed strategy to be used with a certain frequency (e.g., annually or semiannually), depending on the investment policy of each InP. Additionally, we envision that different parts of the physical network would be expanded in each cycle, strengthening it in an organic manner.

### B. Definitions and Modeling

Next, we detail the inputs, variables, and constraints of our model. Superscript letters represent whether a set or variable refers to virtual ($V$) or physical ($P$) resources, or whether it relates to routers ($R$) or links ($L$).

(a) Physical infrastructure with four recurring partitions.

(b) Solution based on the context of partition reconnection. Elements selected for expansion are highlighted.

Fig. 2. Employment of the proposed strategy on a physical infrastructure.

**Substrate and virtual networks.** The topology of the substrate, as well as that of each requested VN, is represented as a directed graph $N = (R, L)$. Vertices $R$ represent routers, while each edge $L$ represents a unidirectional link. Bidirecional links are represented as a pair of edges in opposite directions (e.g., $(a, b)$ and $(b, a)$). Thus, the model allows the representation of any type of physical and virtual topologies.

In real environments, physical routers have a limited amount of resources, in addition to physical limits for resource expansion. In our model, CPU and memory capacities are represented, respectively, as $C_i^P$ and $M_i^P$, while the maximum expansion limits are represented as $EC_i^P$ and $EM_i^P$. Similarly, physical links have a given capacity, represented as $B_{i,j}^P$, and a physical expansion limit, represented as $EB_{i,j}^P$.

In order to take into account infrastructure resource consumption, the model needs to be aware of virtual networks already embedded on the substrate. Already embedded routers are represented by set $A_{i,r,j}^R$, which indicates whether virtual router $j$ from virtual network $r$ is hosted on physical router $i$. Embedded links, in turn, are represented by set $A_{i,j,r,k,l}^L$, which indicates whether virtual link $(k, l)$ from virtual network $r$ is mapped to physical link $(i, j)$. CPU and memory requirements of each router $i$ from virtual network $r$ are represented as $C_{r,i}^V$ and $M_{r,i}^V$, while $B_{r,k,l}^V$ represents the bandwidth requirement of each virtual link.

**Substrate partitioning.** Physical infrastructures may exhibit different topologies and degrees of connectivity. As a consequence of resource depletion in specific regions of the substrate (e.g., bridges or links connected to hubs), infrastructure partitioning may happen in different levels. In this context, a partition is defined as a Strongly Connected Component (SCC) comprising links with available bandwidth equal to or higher than the average requested by virtual networks. Cut-edges identified in the infrastructure are represented by set $X$, and for each $(i, j) \in X$ there is a corresponding value $I$, which represents the importance of each cut-edge in the context of virtual network embedding. The methodology used to calculate the importance of these links will be explained in Subsection IV-C.

**Expansion costs.** The expansion of device capacities incurs monetary costs for infrastructure providers. Moreover, required costs for increasing the capacity (e.g., memory or bandwidth) of physical devices – routers and links – are heterogeneous. These costs are modeled as $costCpu_i^P$ and $costMem_i^P$, related to the expansion of CPU and memory capabilities of router $i$, and $costBw_{i,j}^P$, related to bandwidth expansion of link $(i, j)$.

We assume that the smallest possible units for expansion are measured in megabits per second for links, megabytes for memory, and processor cores for CPUs. This assumption allows us to generalize the expansion model. Equation 1 presents total estimated costs for expanding an infrastructure. Variables $xb_e$, $xm_i$, and $xc_i$ represent, respectively, the increased bandwidth, memory, and CPU capacities on physical devices.

$$ExpCost = \sum_{e \in L^P} xb_e \cdot costBw^P e$$
$$+ \sum_{i \in R^P} (xc_i \cdot costCPU^P i + xm_i \cdot costMem^P i) \quad (1)$$

**Variables.** The variables in this model represent the optimal solution to the expansion problem. Variables $z_e$ denote which cut-edges (stored in $X$) were selected for composing a new infrastructure core. Additionally, variables $x_e$ indicate whether link $e$ is part of the path built between cut-edges selected by $z_e$. Auxiliary variables $q_u$ indicate whether router $u$ from cut-edge $e$ belongs to the new core. All variables are restricted to the binary domain.

**Constraints.** In addition to expansion costs, the model considers a limited amount of physical resources to be expanded. For this reason, it is subject to a set of constraints, described next.

Equations 2, 3, and 4 ensure that the expanded capacity of physical links (bandwidth) and routers (CPU and memory) will not exceed their maximum expansion limits. We emphasize that Equations 3 and 4 are valid for both indexes – $i$ and $j$ – of variable $x$ (which represent the source and destination routers of links forming the path built between cut-edges). However, due to space restrictions, only equations relative to index $i$ are presented.

$$x_e \cdot \left( B_{i,j}^P + \frac{AvailableFunds \cdot Coverage}{costBw_{i,j}^P} \right) \leq EB_{i,j}^P \qquad \forall e \in L^P \tag{2}$$

$$x_{i,j} \cdot \left( C_i^P + \frac{AvailableFunds \cdot Coverage}{costCpu_i^P} \right) \leq EC_i^P \quad \forall (i, j) \in L^P \tag{3}$$

$$x_{i,j} \cdot \left( M_i^P + \frac{AvailableFunds \cdot Coverage}{costMem_i^P} \right) \leq EM_i^P \quad \forall (i, j) \in L^P \tag{4}$$

Equation 5 ensures that the the total resource investment (i.e., the amount of resources being added to the infrastructure) will not exceed the available funds of the InP.

$$ExpCost \leq AvailableFunds \qquad (5)$$

Equations 6 and 7 define the subset of physical devices that will form a new core, acting as a structural reinforcement for the network. This subset is represented by $R^C$, which is composed of all routers with one or more links in $X$. A path is built between the most relevant cut-edges selected by the objective function (Equation 9). $\delta^+(u)$ represents the outgoing links of router $u$, while $\delta^-(u)$ represents its incoming links.

$$\sum_{a \in \delta^+(u)} x_a^k - \sum_{a \in \delta^-(u)} x_a^k = \begin{cases} 1 \cdot q_u & \text{if } u = s \\ -1 \cdot q_u & \text{if } u = k \\ 0 & u \in R^P \setminus \{s,k\} \end{cases}$$
$$\forall u \in R^P, \forall k \in R^C \quad (6)$$

$$q_u \geq \frac{\sum_{a \in \delta^+(u)} z_a}{|\delta^+(u)|} \qquad \forall u \in R^P \quad (7)$$

In practice, it would be unlikely for an InP to consider expanding the capacity of its entire network at once. For this reason, Equation 8 guarantees that only a subset of physical devices will be affected by the expansion procedure. In other words, this equation sets the coverage limits for link expansion. The left side of the equation defines the percentage of links that will be expanded.

$$\frac{\sum_{\forall e \in L^P} x_e}{|L^P|} \leq Coverage \qquad (8)$$

**Objective.** The objective of the model, shown in Equation 9, is to maximize reconnection between the most important cut-edges. As a consequence, the VN acceptance rate should be maximized, and the usage of idle resources should be substantially improved.

$$Maximize \sum_{e \in X} I_e \cdot z_e \qquad (9)$$

### C. Proposed Expansion Strategy

In this subsection we present HIPER, our proposed strategy for solving the InP network expansion problem. We detail each specific procedure it uses to build a feasible solution, and present an overview of its algorithmic process. The proposed strategy is composed of two steps. The first step consists in identifying which devices need to be expanded, while the second step defines a strategy for resource distribution among the selected devices. We emphasize that we adopted a heuristic-based strategy due to the similarity of the approached problem with the minimum Steiner tree problem [17], which is known to be NP-Hard.

Algorithm 1 presents a simplified pseudocode version of HIPER, and its details are explained next. As previously described, a partition is defined in terms of physical resource usage. Thus, it is expected that a procedure for identifying infrastructure partitions and cut-edges (i.e., edges that cause partitioning due to the exhaustion of their resources) is executed periodically (e.g., once a day or week). In lines 5 and 6, we calculate partitions and cut edges observed in each time interval. There are polynomial algorithms for identifying

Strongly Connected Components (and, in consequence, cut-edges) in graphs, such as Tarjan's Algorithm [18], which can be computed with linear complexity $O(|R| + |L|)$.

A history of (previously computed) cut-edges is stored in set $X$ for the purpose of guiding the algorithm in the selection of the most important physical devices. As cut-edges change over time (due to variations in physical resource usage), the algorithm maintains set $I$, which contains information about the relevance of each link $(i, j)$ in $X$. The value stored for each cut-edge is a combination of frequency and coverage, which are cumulatively increased in each time interval when partitions and cut-edges are determined (lines 7–8). For frequency, the algorithm calculates how many times link $(i, j)$ caused infrastrucure partitioning, normalized in relation to the total amount of partitioning occurrences. For coverage, in turn, the algorithm calculates the percentage of routers that were disconnected from the rest of the infrastructure as a result of this partitioning.

---

**Input**: Available funds, percentage of coverage, periodicity of expansion, InP physical infrastructure $N$, virtual network mappings $A_{i,r,j}^R$ and $A_{i,j,r,k,l}^L$
**Output**: Set of physical devices to be expanded
1   $X \leftarrow \emptyset$
2   $F \leftarrow \emptyset$
3   $C \leftarrow \emptyset$
4   **foreach** $TimeUnit$ **do**
5      $partitions \leftarrow obtainSetOfPartitions(N, E^R, E^L)$
6      $X \leftarrow updateCutEdges(partitions)$
7      $F \leftarrow updateFrequency(X)$
8      $C \leftarrow updateCoverage(X, partitions)$
9      **if** $ExpansionPeriodicity$ **then**
10         **foreach** $(i, j) \in X$ **do**
11            $I_{i,j} \leftarrow \frac{F_{i,j}}{\sum_{\forall(i,j)} F_{i,j}} \cdot C_{i,j}$
12         **end**
13         **while** $true$ **do**
14            $sortDesc(X, I)$
15            $list \leftarrow$ select % of edges in $X$ (% equal to coverage)
16            $R^c \leftarrow$ routers in $list$
17            $L^c \leftarrow$ links in least cost paths between each pair of routers $(a, b) \in R^c$
18            $N^c = (R^c, L^c)$
19            $sol^c \leftarrow mst(N^c)$
20            **if** $numberOfRouters(sol^c) > Coverage$ or $numberOfEdges(sol^c) > Coverage$ **then**
21               $X.removeItem()$
22            **else**
23               $stop$
24            **end**
25         **end**
26         $suggestLinkExpansion(AvailableFunds)$
27         $suggestRouterExpansion(AvailableFunds)$
28         $X \leftarrow \emptyset$
29   **end**

**Algorithm 1:** Overview of HIPER.

---

When the physical infrastructure is about to be expanded (which may be performed annually or semiannually, for example), the algorithm analyzes the history of cut-edges stored in $X$ and the importance $I$ of each one. A higher value in $I_{i,j}$ indicates a high probability that link $(i, j)$ is frequently causing the partitioning of a large set of physical routers. $I_{i,j}$ is calculated by multiplying the accumulated values for frequency and coverage. This procedure is performed in line

The proposed strategy takes into account the importance of topological elements of the physical infrastructure – unlike strategies applied in traditional approaches, which rely only on demand matrices for choosing devices to be expanded. To this end, first, physical links $(i, j)$ stored in $X$ are sorted according to their corresponding value in $I_{i,j}$ (line 14). Then, the algorithm builds a subgraph $N^c = (R^c, L^c)$, in which $R^c$ is a subset of routers (with size equal to the percentage of coverage) within set $X$ and $L^c$ contains all links that form the shortest (or least cost) paths between routers in $R^c$ (lines 15-18). A Minimum Spanning Tree (MST) algorithm is applied on subgraph $N^c$ in order to build a new core for the infrastructure, which will act as a structural reinforcement preventing the partitioning of critical regions (line 19). As there are polynomial algorithms for performing this procedure, such as Prim's algorithm, this operation can be calculated in $O(|L^c| + |R^c| log |R^c|)$ steps. It is important to note that the solution obtained from $mst(N^c)$ may contain more elements than the maximum coverage (constraint 7). This is due to the fact that $L^c$ is composed of a set of paths that interconnect all routers in $R^c$. As a result, there may be intermediate routers included in the solution that were not initially included in $R^c$. In this case, the process of rebuilding graph $N^c$ is iteratively repeated, eliminating the least relevant cut-edge (lines 20–24).

Last, after the creation of the reinforcement structure, the algorithm suggests the capacity expansion of each router and link (lines 26–27). Different strategies may be used for this purpose. For example, all available expansion funds may be uniformly invested among all routers and links. Alternatively, a probabilistic distribution based on the importance of each element may be adopted. Due to space constraints, in this paper we do not consider a prioritization strategy, assuming that expansion resources will be uniformly distributed among devices selected in the previous step.

## V. EVALUATING THE IMPACT OF INFRASTRUCTURE EXPANSION ON VIRTUAL NETWORK EMBEDDING

In order to assess the impact of expansions suggested by the proposed strategy on the virtual network embedding process, we implemented HIPER and subjected it to a systematic evaluation. Experiments were performed on a machine with four AMD Opteron 6276 processors and 64 GB of RAM, using the Ubuntu GNU/Linux Server 11.10 x86_64 operating system.

### A. Workloads and VNE Model

For the purpose of this evaluation, we developed a generator of virtual network requests. This generator is run for a total of 360 time units (we consider that each time unit corresponds to a day, comprising a total length of 360 days). In each time unit, three virtual network requests are generated. Each virtual network request mapped to the substrate has a limited duration – i.e., after a given number of time units, it is removed from the physical network. We emphasize that this form of instantiation is in line with related work in the area of virtual network embedding.

Virtual network requests are mapped using an online VNE model based on the one proposed by Luizelli et al. [7], which takes into account the main constraints observed in recent literature. In contrast to the original model, in this paper we do not consider location requirements, as our initial objective is to evaluate the impact of the expansion strategy considering a more general version of the VNE problem.

Networks used as physical substrate were generated with IGen[2]. The topology of these networks follows the hub & spoke model, as characterized by Luizelli et al. [7]. This type of topology was chosen as an approximation of those observed in real environments. Instantiated physical networks have a total of 50 routers, each with total CPU capacity of 100% and 256 MB of memory, while the bandwidth of physical links is 10 Gbps. Moreover, the maximum capacities of physical routers and links after expansion are defined as double their initial ones.

We considered ring and random as the solicited topologies for the virtual network requests. Ring topologies were chosen because they are inherently resilient against individual link failures and this property is specially important in the context of virtual networks. The choice for random topologies, in turn, is due to being widely employed in related work and exhibiting a certain variability in terms of topological structure. All virtual networks have 5 routers each, and random topologies have fixed connectivity degree of 50%. Virtual routers require 20% CPU and 48 MB of memory, while virtual links require 2.5 Gbps bandwidth. This choice of parameters aims at reproducing the same phenomena observed in previous evaluations, namely high virtual network rejection rates in infrastructures with significant amounts of available (residual) resources. In [7], for example, we observed that the described workload led to an acceptance rate of 66.60% and a global resource usage of 60.03% in terms of bandwidth, 52.43% in terms of CPU, and 55.93% in terms of memory. In other words, while overall resource usage remained under 60%, we still observed rejection rates in the order of 30%. It is important to note that similar behavior is observed in other related VNE proposals, such as those mentioned in the Introduction. Through precise physical infrastructure expansions, we expect to reach significantly higher acceptance levels and, most importantly, greatly improve the usage of resources that would otherwise remain idle.

We initially evaluate the impact of a single expansion, performed on the 180th time unit (i.e., after six months). Then, we evaluate the proposed solution considering consecutive expansions performed on the same infrastructure. Expansion costs are defined in a homogeneous manner – i.e., we consider that devices are identical and, therefore, expansion costs are the same. Further, our experiments employ two variable parameters: expansion (the total amount of resources added to existing infrastructure devices) and coverage (the percentage of devices to be expanded). Each experiment was performed 30 times, considering different physical network instances.

### B. Results

First, we analyze the average increase in terms of virtual network acceptance after the expansion of physical infrastructure resources. We emphasize that virtual network requests are only embedded if it is possible to map all of its virtual routers and links on the physical network.

Figure 3 depicts average gains in terms of additional accepted virtual network requests considering variations in the percentage of expanded resources (expansion), the number of devices affected by the expansion (coverage) and virtual

---

[2]http://igen.sourceforge.net/

network topologies (ring and random). When considering virtual networks with ring topology, the graph in Figure 3(a) reveals that ideally there should be a balance between expansion and coverage. Observed averages demonstrate that, in order to maximize the acceptance of additional virtual networks, coverage should remain between 20% and 30%. With these values, we observe an increase of 31.06% and 31.76%, respectively, in terms of additional VN acceptance for an expansion of 20% of infrastructure resources. For an expansion of 10%, the observed benefits were, respectively, 14.48% and 12.87%. With respect to virtual networks with random topologies, Figure 3(b) reveals that greater benefits (in terms of virtual network acceptance) are obtained when the expansion is concentrated on a lower number of devices (10% to 20% coverage). These percentages of coverage lead to increases of, respectively, 23.26% and 27.06% in terms of additional virtual network acceptance (for an expansion of 20% of infrastructure resources). For an investment of 10%, observed gains are, respectively, 12.91% and 15.48%.
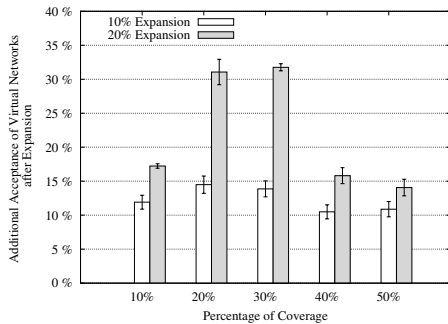


(a) Experiments considering virtual networks with ring topology.



(b) Experiments considering virtual networks with random topologies.

Fig. 3. Average increase in virtual network acceptance after the employment of the expansion strategy.

As one can observe, the topology (ring or random) of the mapped virtual networks plays an important role in the outcome of the expansion strategy. First, there are differences in terms of additional increase in virtual network acceptance. This is due to varying amounts of resources requested by VNs. In scenarios that consider ring topologies, the workload is homogeneous (i.e., VNs have a fixed size), since they always contain 5 links. However, in random topologies, although the number of routers is always the same, the number of links is uniformly varied. In a best case scenario, a VN with random topology contains only 4 links (star topology) and, in the worst case, 20 links (full mesh). Therefore, as random VNs consume

more resources on average, comparatively lower acceptance rates are expected. Second, there is a difference in terms of the ideal coverage for the expansion process for each type of virtual topology. As the number of links requested by VNs varies, the connectivity degree of virtual routers also changes. VNs with random topologies require, on average, routers with higher connectivity degree (up to a degree of 4) than those with ring topologies (always a degree of 2). Due to the tendency of virtual elements that are more interconnected to be mapped to substrate regions with greater connectivity (hubs and nearby links), the exhaustion of resources in these physical devices tends to lead to higher infrastructure partitioning levels. Since hub & spoke topologies, in general, tend to contain a high number of hubs – which are interconnected and close to each other – the solution to the expansion problem is generally composed of a subset of these elements (with higher topological importance values, $I_{i,j}$). As these elements are usually close to each other, fewer elements need to be expanded (i.e., lower coverage) in order to efficiently reconnect infrastructure partitions. When considering VNs with ring topology, due to the higher flexibility of embedding such topologies, the set of candidate devices for expansion is larger (containing, for example, intermediate links between main cut-edges). Therefore, expansions that lead to higher benefits tend to contain a larger set (in the order of 10% additional coverage) of links and routers.

Figure 4 and Figure 5 outline the average usage of physical links after the expansion process for the different types of virtual topologies considered in our experiments. Each point in the graph represents the average usage since the beginning of the experiment until the time unit shown in the horizontal axis. The black horizontal line in each graph represents the performed resource expansion. Values above this line indicate that the InP is taking advantage of a portion of previously unused (or underused) resources, beyond the added resources (10% or 20%). In cases where the average is below the horizontal line, expanded resources are underused.

Figures 4(a) and 5(a) depict the average resource usage after an expansion of 10% of physical resources considering VNs with ring and random topologies, respectively. In these cases, improved resource usage is observed when the expansion coverage is lower. With a coverage of up to 20%, it is possible to obtain a sustained increase of up to 21% (considering VNs with ring topology) and 29% (considering VNs with random topologies) in terms of resource usage in comparison with the scenario in which no expansion is made. In other cases, with coverage equal to or greater than 30%, the benefits obtained in terms of resource usage remain below the invested 10%. This is directly related to the fact that a low amount of resources is distributed among a high number of devices, leading to minor increases in the capacity of each device. Figures 4(b) and 5(b) show the average bandwidth usage after an expansion of 20% of physical resources considering VNs with ring and random topologies, respectively. We observe that expansions with 20% and 30% coverage lead to increased resource usage of up to 45% (considering VNs with ring topology) and 52% (considering VNs with random topology) compared to the scenario with no expansion.

Observed variations in terms of average resource usage in the recently discussed results are related to the ability of VN requests to use residual substrate bandwidth in different scenarios. Since all VNs require the same amount of resources in scenarios that consider VNs with ring topology, it becomes
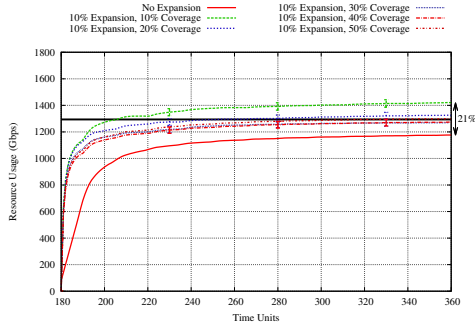
(a) Experiments considering an expansion of 10%.
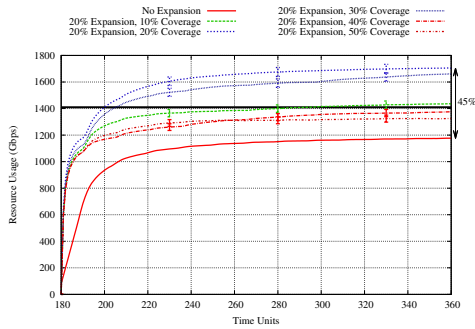


(b) Experiments considering an expansion of 20%.

Fig. 4. Overall physical bandwidth usage after the expansion, considering virtual networks with ring topology.



(a) Experiments considering an expansion of 10%.



(b) Experiments considering an expansion of 20%.

Fig. 5. Overall physical bandwidth usage after the expansion, considering virtual networks with random topologies.

more difficult to take advantage of residual resources after a certain level of exhaustion. In scenarios where VNs have random topologies, the variation of topological characteristics

among distinct requests allows VNs with lower numbers of links (e.g., star or line topologies) to take advantage of these residual resources. As a result, with proper resource expansion (i.e., adequate levels of coverage), average resource usage tends to be improved further when considering varying virtual topologies.

Benefits obtained in terms of resource usage are mainly due to the structural reinforcement performed in certain regions of the infrastructure. This reinforcement leads to lower partitioning levels and directly favors a lower fragmentation of idle resources. It is important to note that these benefits may lead to lower costs for VN requesters, as the InP is able to take advantage of previously unused resources. Similarly to previously discussed results (depicted in Figure 3), we observe a clear correlation among resource expansion, coverage, and improved resource usage.

Next, Figure 6 presents a qualitative view of the proposed strategy through a graphical representation of average resource usage in a physical infrastructure before and after an expansion. In this figure, darker shades represent higher resource usage. First, it is clear that resource usage in physical links and routers before the expansion is concentrated in certain regions of the infrastructure. The overuse of resources in certain physical devices (mainly hubs and cut-edges) leads to poor overall infrastructure resource usage. The expansion strategy creates a reinforcement structure based on the reconnection of recurring partitions. Because resource distribution becomes significantly more homogeneous as a whole, our expansion strategy leads to improved usage of physical resources (larger number of dark nodes).



(a) Resource usage before expansion. (b) Resource usage after expansion.

Fig. 6. Graphical representation of average infrastructure resource usage before (a) and after (b) an expansion of 20% of resources with coverage of 20%. Darker shades represent higher resource usage.

An additional set of experiments was performed in order to provide a better understanding of the relationship between consecutive employments of the expansion procedure. Two consecutive expansions were performed in intervals of 180 days (i.e., semianually). Due to space constraints, the results presented in this paper are limited to a subset of all performed experiments in which the first expansion (applied in the 180th time unit) was performed with 20% total resource expansion and 20% coverage (shown in Figure 3(b)). Moreover, this additional set of experiments was performed considering VNs with random topologies (as ring topologies represent a subset of random topologies). Figure 7 illustrates additional gains in terms of VN acceptance after the second expansion (i.e., performed in the 360th time unit). In contrast to previous results, we observe that resource expansions are more effective when coverage is higher. More specifically, a coverage of 50% was needed in order to ensure that the obtained benefits would be superior to the expansion investments. Considering this level of coverage, expansions of 10% and 20% led to 12.79% and

22.06% additional VN acceptance, respectively. This higher coverage requirement is mainly due to the substantially higher and more homogeneous resource usage achieved after the first expansion (depicted in Figure 6(b)). In this situation, a larger set of routers and links need to be expanded in order to significantly improve virtual network acceptance. As a result of this behavior, we envision that expansions with higher and lower levels of coverage should be performed in alternation in order to maximize long term benefits for the infrastructure provider.
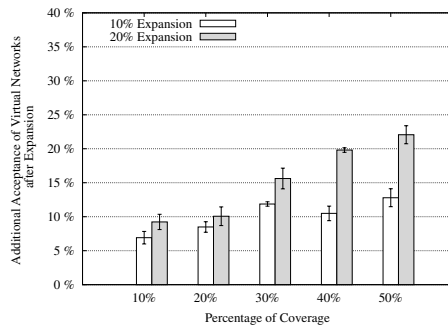


Fig. 7.    Average increase in virtual network acceptance after consecutive employments of the expansion strategy.

Last, we discuss the time needed to find a feasible solution to the expansion problem. In all experiments, the average time taken by the proposed algorithm remains in the order of few seconds, as the proposed expansion strategy has polynomial complexity. These results indicate that the proposed solution may be applied to larger infrastructures while still being able to generate timely solutions.

## VI. Conclusion

Despite the increasing prominence of network virtualization in both academia and the Industry, efficient embedding of virtual networks on top of physical infrastructures remains an enduring challenge. While there is a sizable body of work in this area, state-of-the-art proposals still suffer from significant amounts of rejection of virtual network requests due to the exhaustion of resources in certain key points of the infrastructure. Moreover, although network expansion procedures could be used to mitigate this issue, we demonstrated that classical, demand matrix-based expansion strategies are not suitable for virtual network-oriented infrastructure planning.

In this paper, we formalized a heuristic-based expansion strategy for physical InP infrastructures geared towards efficient virtual network embedding. Our strategy, named HIPER, achieves this by reconnecting recurring partitions in the infrastructure in order to fulfill demands imposed on InPs by the process of virtual network embedding. We applied this strategy on typical provider network topologies, evaluating its ability to improve virtual network acceptance rates and physical resource usage. The obtained results demonstrate that the expansion of infrastructure resources using the proposed strategy significantly contributes to a sustained increase of up to 30% in virtual network acceptance, as well as up to 52% in resource utilization compared to the original network.

As perspectives for future work, we envision extending the evaluation of HIPER by applying it to other types of backbone topologies, as well as conducting an in-depth analysis of the inter-relationships between its parameters. Moreover, we intend to propose and evaluate novel strategies for resource distribution among infrastructure devices selected for expansion.

## References

[1]   M. Chowdhury, M. R. Rahman, and R. Boutaba, "Vineyard: Virtual network embedding algorithms with coordinated node and link mapping," *IEEE/ACM Transactions on Networking*, vol. 20, no. 99, pp. 206–219, 2012.

[2]   L. R. Bays, R. R. Oliveira, L. S. Buriol, M. P. Barcellos, and L. P. Gaspary, "A heuristic-based algorithm for privacy-oriented virtual network embedding," in *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Krakow, Poland, May 2014.

[3]   M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: substrate support for path splitting and migration," *SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 17–29, Mar. 2008.

[4]   X. Cheng, S. Su, Z. Zhang, K. Shuang, F. Yang, Y. Luo, and J. Wang, "Virtual network embedding through topology awareness and optimization," *Computer Networks*, vol. 56, no. 6, pp. 1797 – 1813, 2012.

[5]   G. P. Alkmim, D. M. Batista, and N. L. S. Fonseca, "Mapping virtual networks onto substrate networks," *Journal of Internet Services and Applications*, vol. 3, no. 4, pp. 1–15, 2013.

[6]   Q. Zhang, M. Zhani, M. Jabri, and R. Boutaba, "Venice: Reliable virtual data center embedding in clouds," in *IEEE International Conference on Computer Communications (INFOCOM)*, Toronto, Canada, April 2014.

[7]   M. Luizelli, L. Bays, L. Buriol, M. Barcellos, and L. Gaspary, "Characterizing the impact of network substrate topologies on virtual network embedding," in *International Conference on Network and Service Management (CNSM)*, Zurich, Switzerland, 2013, pp. 42–50.

[8]   B. Mukherjee, D. Banerjee, S. Ramamurthy, and B. Mukherjee, "Some principles for designing a wide-area wdm optical network," *IEEE/ACM Transactions on Networking*, vol. 4, no. 5, pp. 684–696, 1996.

[9]   R. Ramaswami and K. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, pp. 840–851, 1996.

[10]   R. M. Krishnaswamy and K. N. Sivarajan, "Design of logical topologies: a linear formulation for wavelength-routed optical networks with no wavelength changers," *IEEE/ACM Transactions on Networking*, vol. 9, no. 2, pp. 186–198, Apr. 2001.

[11]   A. Curtis and A. Lopez-Ortiz, "Capacity provisioning a valiant load-balanced network," in *IEEE International Conference on Computer Communications (INFOCOM)*, Rio de Janeiro, Brazil, 2009, pp. 3006–3010.

[12]   M. Johnston, H.-W. Lee, and E. Modiano, "A robust optimization approach to backup network design with random failures," in *IEEE International Conference on Computer Communications (INFOCOM)*, Shanghai, China, 2011, pp. 1512–1520.

[13]   G. Shen and R. Tucker, "Energy-minimized design for ip over wdm networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 1, no. 1, pp. 176–186, 2009.

[14]   M. C. Luizelli, L. R. Bays, L. S. Buriol, M. P. Barcellos, and L. P. Gaspary, "Reconnecting partitions on physical infrastructures: Towards an expansion strategy for efficient virtual network embedding," in *Brazilian Symposium on Computer Networks and Distributed Systems (SBRC)*, Florianopolis, Brazil, May 2014, pp. 335–343.

[15]   A. Curtis, T. Carpenter, M. Elsheikh, A. Lopez-Ortiz, and S. Keshav, "Rewire: An optimization-based framework for unstructured data center network design," in *IEEE International Conference on Computer Communications (INFOCOM)*, Orlando, USA, 2012, pp. 1116–1124.

[16]   P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav, "It's not easy being green," *SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 211–222, Aug. 2012.

[17]   M. X. Goemans and Y.-S. Myung, "A catalog of steiner tree formulations," *Networks*, vol. 23, no. 1, pp. 19–28, 1993.

[18]   D. D. Sleator and R. E. Tarjan, "A data structure for dynamic trees," *Journal of Computer and System Sciences*, vol. 26, no. 3, pp. 362 – 391, 1983.