

Characterizing the Impact of Network Substrate Topologies on Virtual Network Embedding

Marcelo Caggiani Luizelli, Leonardo Richter Bays, Luciana Saete Buriol,
Marinho Pilla Barcellos, Luciano Paschoal Gaspar
Institute of Informatics
Federal University of Rio Grande do Sul (UFRGS)
{mcluiuzelli,lrbays,buriol,marinho,paschoal}@inf.ufrgs.br

Abstract—Network virtualization is a mechanism that allows the coexistence of multiple virtual networks on top of a single physical substrate. One of the research challenges addressed recently in the literature is the efficient mapping of virtual resources on physical infrastructures. Although this challenge has received considerable attention, state-of-the-art approaches present, in general, a high rejection rate, *i.e.*, the ratio between the number of denied virtual network requests and the total amount of requests is considerably high. In this work, we investigate the relationship between the quality of virtual network mappings and the topological structures of the underlying substrates. Exact solutions of an online embedding model are evaluated under different classes of network topologies. The obtained results demonstrate that the employment of physical topologies that contain regions with high connectivity significantly contributes to the reduction of rejection rates and, therefore, to improved resource usage.

I. INTRODUCTION

Network virtualization is a mechanism that allows the coexistence of multiple heterogeneous virtual networks (VNs) sharing resources of the same physical substrate. The architectures, protocols, and topologies used in these VNs are unconstrained by the substrate network on which they are instantiated. Through network virtualization, Infrastructure Providers (InPs) are able to easily allocate and deallocate virtual networks with proper resource isolation. In other words, this mechanism enables InPs to support the creation of custom networks on demand, meeting different requirements imposed by requesters.

One of the major research challenges in network virtualization is the efficient mapping of physical resources to virtual networks (VNE – *Virtual Network Embedding*). The resource mapping process must consider the capacities of physical network devices, as well as the demands of virtual networks (for instance, virtual link bandwidth and processing capacity of virtual routers). Although previous work explores the problem of online virtual network embedding [1]–[4], considerably high rejection rates are commonly observed (as high as 53%). We assume that a subset of these rejections is caused by temporary resource exhaustion, *i.e.*, periods in which the available resources in the infrastructure as a whole are not sufficient to meet the demand. We theorize, however, that most rejections occur in situations in which a significant amount of resources is available, but a few saturated devices and links, depending on connectivity features of the physical substrate, hinder the acceptance of new requests.

Despite efforts to solve the virtual network embedding problem, we are not aware of previous attempts to investigate

the influence of network topologies in the process of virtual network embedding. Moreover, previous work in this area has considered topologies that often do not reflect those observed in commercial networks [5]. Understanding the relationship between the employment of different network topologies and the mapping process is important to determine how certain topological features influence this process. For example, topologies with higher connectivity in strategic regions may favor a better utilization of physical resources, which in turn may lead to lower rejection rates. Such outcomes have the potential to raise the profit obtained by InPs and, at the same time, reduce costs for virtual network requesters.

In this paper, we characterize the impact of different classes of topologies typically employed in commercial infrastructures on the quality of the virtual network embedding process. More specifically, we formalize an optimal virtual network embedding model and evaluate it on substrates with different types of topologies – namely, *star*, *ladder*, and *hub & spoke*. In this evaluation, we consider different metrics such as rejection rate and resource consumption of physical network devices. In summary, the main relevant contributions of this paper are: (i) the formalization of an *online* embedding model considering location constraints; (ii) the characterization of networks that are typically employed by infrastructure providers, and (iii) the evaluation and discussion of the impact of different types of topologies in the virtual network embedding process.

The remainder of this paper is organized as follows. Section 2 presents a discussion of related work in the area of virtual network embedding, highlighting the topologies considered in each work and the rejection rates obtained. In Section 3 we characterize the types of network topologies considered in this work. In Section 4 we formalize the *online* virtual network embedding model. In Section 5 we present and evaluate the obtained results. Finally, in Section 6 we conclude this paper with final remarks and perspectives for future work.

II. RELATED WORK

In this section, we present related work in the area of virtual network embedding. We briefly summarize the proposed solutions, highlighting the types of physical and virtual topologies employed, as well as rejection rates obtained by embedding methods, when available.

Yu *et al.* [1] present an online virtual network embedding model supporting path splitting and migration. Path splitting improves the utilization of physical resources by embedding a higher number of virtual networks on the substrate, while

the migration of virtual network elements aims to reoptimize physical resource usage. Router and link mapping are performed in distinct steps. The experiments employ randomly generated topologies for physical and virtual networks, with fixed connectivity of 50%. No rejection rates are presented.

Another online model, formulated by Chowdhury *et al.* [2], also performs router and link mapping in distinct steps. However, location constraints are used to preselect physical routers on which virtual routers will be hosted. According to the authors, the preselection facilitates the subsequent stage of link mapping. This model also allows path splitting. Physical topologies used in the evaluation of this model have the same characteristics as those used by Yu *et al.* [1] (randomly generated with fixed connectivity of 50%). Experiments are performed with three types of VN topologies, namely *random*, *hub & spoke*, and *full mesh*. Rejection rates observed in experiments using *random* virtual topologies vary between 20% and 45%. *Hub & spoke* topologies led to a rejection rate of 35%, whereas the rejection rate of *full mesh* VNs varies between 40% and 45%.

Alkmim *et al.* [3] propose a model that focuses on minimizing the time needed to transfer binary images of virtual routers (stored in repositories connected to the network) to the physical routers that will host them. Their model considers requirements related to router and link capacity, as well as location constraints. The evaluation scenarios employ organic topologies created using the BA-2 model [6]. On average, rejection rates in experiments performed by the authors are approximately 53%.

Cheng *et al.* [4] perform the mapping of virtual network elements by ranking routers and links according to their own capacity as well as the capacity of their neighbors. Virtual routers and links are mapped to similarly ranked physical devices and paths. According to the authors, this strategy aims to avoid the formation of bottlenecks on the physical network. In the evaluated scenarios, physical and virtual topologies are randomly generated and rejection rates vary between 15% and 25%.

Davy *et al.* [7] present a model that builds virtual networks according to spanning trees and embeds them on a physical substrate. The model is able to prioritize either lower hosting costs or lower delay for embedded VNs. In contrast to the publications presented above, which use organic and randomly generated physical topologies, experiments are performed on a physical network created by the authors that is based on publicly available topologies used by Internet Service Providers (ISPs). However, it was out of the scope of this paper to compare different types of topologies used by ISPs. Additionally, no rejection rates are presented.

As previously stated, we are not aware of previous attempts to evaluate the results of virtual network embedding strategies considering different types of network topologies in a precise manner. Most publications in this area employ topologies that may not faithfully represent the topological properties of infrastructure provider networks, and to the best of our knowledge, there have been no attempts to evaluate the influence of different types of physical topologies. Moreover, as previously stated, considerably high rejection rates are often observed. For these reasons, this study aims to understand

how topologies that are typically employed in real physical substrates influence different aspects of virtual network embedding, such as the rejection of virtual network requests and physical resource usage.

III. INP NETWORK TOPOLOGIES

In the context of this paper, we assume, without loss of generality, that infrastructure providers (InPs) employ network topologies that are equivalent to those used by Internet Service Providers (ISPs). The most traditional ISP network topologies are known as: *ladder*, *star* and *hub & spoke*. Network topologies organized as *ladder* are characterized by the absence of hubs, *i.e.*, nodes with high connectivity and concentration of flows. Additionally, the infrastructure is formed by a set of loops. This type of topology tends to have lower cost regarding the deployment of links (due to its low connectivity) and the distance between nodes (in terms of number of hops) is typically high. *Star* networks have a low number of hubs connected to numerous nodes which, in turn, have low connectivity. In this type of network, the distance between nodes tends to be low, but traffic tends to become concentrated on the hubs. The *hub & spoke* class is characterized by a comparatively higher number of hubs, which tend to be interconnected. Additionally, a high number of nodes is connected to one or more hubs. Figure 1 illustrates examples of the three aforementioned topology classes.

Kamiyama *et al.* [8] conducted a study that formalizes the classification of ISP networks into the three previously described topology classes. In their study, the authors analyzed 23 commercial backbone networks (publicly available) with sizes ranging from 21 to 128 nodes. Through this analysis, the authors define a set of metrics that capture the main topological properties present in each infrastructure. Such metrics include, for example, the connectivity degree of the network and the presence of *hub* nodes. Thus, the authors map the relationship between these metrics and the type of network topology of the infrastructure, enabling the classification of ISP network topologies into one of the previously described classes.

Given the well accepted classification of ISP network topologies presented above and the systematic approach proposed by Kamiyama *et al.* [8] to characterize and, therefore, generate these topologies with high degree of fidelity (in relation to real networks), in this paper we consider these three classes as the basis of our investigation. Thus, we intend to analyze how the employment of these topological classes affects the mapping process in order to identify possible correlations between different topological factors, physical resource usage, and rejection rates.

IV. OPTIMAL MODEL FOR VN EMBEDDING

As our objective was to capture the impact of the underlying substrate topology in the network embedding problem, we formulated it as an Integer Linear Programming problem and solved it exactly by using an optimization software. The employment of such model means that the results are the best possible one would obtain when considering a substrate topology or another. Next, we describe the inputs, variables, and constraints of this model. Superscript letters are used to represent whether a set or variable is related to virtual (V) or




Fig. 1. Examples of InP topology classes.

physical (P) resources, or whether it is associated to routers (R) or links (L).

Both physical topologies and virtual network requests are represented by directed graphs $N = (R, L)$. Vertices R represent routers, while each edge L represents a unidirectional link. Bidirectional links are represented as a pair of edges in opposite directions (for instance, (a, b) and (b, a)). Thus, the model allows the representation of any type of physical and virtual topology.

Each physical router is associated with a location identifier stored in a set S^P . This enables virtual network requesters to indicate specific locations in which virtual routers must be instantiated (e.g., to ensure connectivity between two or more geographical locations). If a virtual router has a location requirement, it is stored in set S^V . Our approach of tagging routers with specific location identifiers differs from previous work [2], [9], which treat locations simply as regions defined by radius.

In real life, physical routers have limited CPU and memory capacities. In our model such capacities are represented, respectively, by C_i^P and M_i^P . Analogously, CPU and memory requirements of each virtual router on a network r are represented by $C_{r,i}^V$ and $M_{r,i}^V$. Likewise, physical links have limited bandwidth capacity, represented by $B_{i,j}^P$, while the bandwidth required by each virtual link is represented by $B_{r,i,j}^V$.

The model takes as input virtual network requests and embeds them in an online manner. Thus, it is necessary to consider virtual elements that were previously embedded on the substrate. Previously embedded virtual routers are stored in set $E_{i,r,j}^R$, while previously embedded links are stored in set $E_{i,j,r,k,l}^L$.

The variables are outputs of our model and represent the optimal solution of the virtual network embedding problem for the given set of inputs. These variables indicate in which location the requested virtual routers and links are allocated on the physical substrate. If a request is accepted, each of its virtual routers is mapped to a physical router, whereas each virtual link is mapped to one or more consecutive physical links (a path).

- $A_{i,r,j}^R \in \{0, 1\}$ – Router allocation, indicates whether physical router i is hosting virtual router j from virtual network r .
- $A_{i,j,r,k,l}^L \in \{0, 1\}$ – Link allocation, indicates whether physical link (i, j) is hosting virtual link (k, l) from virtual network r .

Based on the above inputs and outputs, we now present the

objective function and its constraints. The objective function of the model aims at minimizing the total bandwidth consumed by virtual networks embedded on the substrate. The purpose of each constraint is explained next.

Objective:

$$\min \sum_{(i,j) \in L^P} \sum_{r \in N^V, (k,l) \in L^V} A_{i,j,r,k,l}^L B_{r,k,l}^V$$

Subject to:

$$\sum_{r \in N^V, j \in R^V} C_{r,j}^V A_{i,r,j}^R \leq C_i^P \quad \forall i \in R^P \quad (C1)$$

$$\sum_{r \in N^V, j \in R^V} M_{r,j}^V A_{i,r,j}^R \leq M_i^P \quad \forall i \in R^P \quad (C2)$$

$$\sum_{r \in N^V, (k,l) \in L^V} B_{r,k,l}^V A_{i,j,r,k,l}^L \leq B_{i,j}^P \quad \forall (i, j) \in L^P \quad (C3)$$

$$\sum_{i \in R^P} A_{i,r,j}^R = 1 \quad \forall r \in N^V, j \in R^V \quad (C4)$$

$$\sum_{j \in R^V} A_{i,r,j}^R \leq 1 \quad \forall i \in R^P, r \in N^V \quad (C5)$$

$$\sum_{j \in R^P} A_{i,j,n,k,l}^L - \sum_{j \in R^P} A_{j,i,n,k,l}^L = A_{i,n,k}^R - A_{i,n,l}^R \quad \forall n \in N^V, (k, l) \in L^V, i \in R^P \quad (C6)$$

$$j A_{i,r,k}^R = l A_{i,r,k}^R \quad \forall (i, j) \in S^P, r \in N^V, (k, l) \in S^V \quad (C7)$$

$$A_{i,r,j}^R = E_{i,r,j}^R \quad \forall (i, r, j) \in E^R \quad (C8)$$

$$A_{i,j,r,k,l}^L = E_{i,j,r,k,l}^L \quad \forall (i, j, r, k, l) \in E^L \quad (C9)$$

Constraint C1 ensures that the CPU capacity of each physical router will not be exceeded, therefore assuring that the CPU requirement of each virtual router will be met. Constraint C2 applies the same restriction to the memory capacity of routers, and constraint C3, to link bandwidth. Constraint C4 ensures that all virtual routers will be mapped to a physical router. In turn, constraint C5 prevents multiple virtual routers that belong to a single virtual network from sharing the same physical router. As our objective function aims at minimizing bandwidth usage, the absence of this constraint would encourage a significant number of routers from a single virtual network to share the same physical router, which could lead to availability issues. Constraint C6 ensures that all virtual links will be mapped to a valid physical path. Thus, the physical path hosting a virtual link (a, b) is guaranteed to be a valid path between the physical router hosting virtual router a and the physical

router hosting virtual router b . Constraint $C7$ makes sure that all virtual routers with location requirements are mapped to physical routers at the required locations. Finally, constraints $C8$ and $C9$ ensure that all elements from previously embedded virtual networks remain hosted on the same physical elements. Router mappings are maintained by constraint $C8$, while link mappings are maintained by constraint $C9$.

V. EVALUATION OF THE IMPACT OF TOPOLOGIES ON VN EMBEDDING

In order to evaluate the impact of different network topologies in the process of virtual network embedding, the model formalized in the previous section was implemented and run in *CPLEX Optimization Studio*¹ version 12.3. All experiments were performed on a machine with four AMD Opteron 6276 processors and 64 GB of RAM, using the Operational System Ubuntu GNU/Linux Server 11.10 x86_64.

A. Workloads

To perform the experiments, we adopt a strategy in line with related work, such as the ones conducted by Yu *et al.* [1] and Houidi *et al.* [10]. Like them, we rely on time units and distribution models for the arrival and duration of requests.

We developed a virtual network request generator, which is run for a period of 500 time units for each experiment. Within each time unit, five requests are generated on average, according to a Poisson distribution. Each request has a limited duration, *i.e.*, after a number of time units, it is removed from the substrate. Requests have an average duration of five time units, following an exponential distribution.


Physical substrate topologies are created using the IGen² tool, as publicly available real substrate networks are not normalized in terms of network size. In order to generate networks with the topological features of the previously presented classes – *star*, *ladder* and *hub & spoke* – we use, respectively, the methods *Mentor*, *MultiTour* and *TwoTree*. In line with the topology characterization presented in Section III, *ladder* network nodes have an average degree of 3 and normalized maximum degree of 4. *Star* networks have a proportion of highly interconnected nodes (*hubs*) of less than 0.25, while in *hub & spoke* networks this proportion is greater or equal than 0.25. This ratio is defined as the number of nodes with connectivity degree greater than the average connectivity degree of the network, divided by the total number of network nodes. Besides these topological properties, physical networks have 50 routers, each with a total CPU capacity of 100% and 256 MB of memory. Routers are uniformly distributed among 16 locations, and the bandwidth of physical links is uniformly distributed between 1 and 10 Gbps.

The topology of each virtual network is generated using BRITE³ with the Barabasi-Albert (BA-2) [6] model. Each virtual network has between 2 and 5 routers. Virtual routers require between 10% and 50% of CPU and between 24 MB and 128 MB of memory. Both parameters follow a uniform distribution. Virtual link bandwidth is uniformly distributed


¹<http://www-01.ibm.com/software/integration/optimization/cplex-optimization-studio/>

²<http://igen.sourceforge.net/>

³<http://www.cs.bu.edu/brite/>



(a) Scenarios with location requirements.



(b) Scenarios without location requirements.


Fig. 2. Average percentage of rejected requests in all experiments.

between 1 and 5 Gbps. Two scenarios were evaluated on each physical topology. The distinctive feature of each scenario is the presence or absence of location requirements. In the first scenario, each virtual network has two routers (its end points) with location requirements, which are randomly selected among the 16 existing locations. In the second scenario, there are no such requirements. Each experiment was run 30 times, considering different instances for each type of network substrate.


B. Results

First, we analyze the rejection rate of virtual networks requests in the previously described scenarios. Virtual network requests are only rejected if it is not possible to map all of its routers and links on the physical substrate. Figure 2 depicts the average rejection rate in each scenario. Each point on the graph represents the average rejection rate since the beginning of the experiment until the current time unit. It is clear that when location requirements are considered, rejection rates are substantially higher (ranging from 65.38% to 83.71%) in all three physical topologies, in comparison to scenarios with no such requirements (in which rejection rates range from 0% to 41.32%). This behavior is influenced by the reduction in the exploration space of feasible solutions caused by the presence of location constraints.

The graph depicted in Figure 2 also reveals that there is considerable difference in rejection rates when using different physical topologies. *Hub & spoke* networks lead to a lower rejection rate in comparison to other topologies in both evaluated scenarios (68.44% in the scenario with location requirements and 0.53% in the scenario without such requirements). In



(a) Scenarios with location requirements.




(b) Scenarios without location requirements.


Fig. 3. Average bandwidth overhead needed to accommodate accepted requests.

contrast, *star* topology networks lead to the worst performance (rejection rate of 85.04% in the scenario with location requirements and 43.10% in the scenario without location requirements). *Ladder* topology networks present rejection rates of 75.63% and 24.66% for the scenarios with and without location requirements, respectively. *Hub & spoke* networks tend to cause the rejection of a lower number of requests because they have, on average, a greater number of highly interconnected nodes (hubs). The presence of multiple hubs lowers the probability that the depletion of the resources of one of these central nodes may cause a significant impact on the ability to embed future requests. In contrast, as *star* topology networks have a low number of central nodes, there is a high probability that these nodes may become a bottleneck in the process of virtual network embedding if their resources are depleted. In *ladder* networks, as there are no central nodes, the depletion of resources in some physical links may hinder the creation of virtual links that would use such physical links as “bridges” to interconnect certain points of the infrastructure.

Figure 3 illustrates the average overhead caused by virtual networks embedded in each experiment. This overhead is measured as the ratio between the effective bandwidth consumed by a virtual network hosted on the physical substrate and the bandwidth requested by such network. In general, the actual bandwidth consumption is greater than the total bandwidth required by virtual networks, due to the frequent need to map virtual links to paths composed of multiple physical links. The absence of overhead is only observed when each virtual link is mapped to a single physical link (ratio of 1.0). We emphasize that lower overhead rates directly favor



(a) Scenarios with location requirements.




(b) Scenarios without location requirements.


Fig. 4. Average CPU usage of physical routers.

the infrastructure provider, by sparing resources that may be used to embed future requests. Moreover, this lower resource consumption may lead to lower costs for virtual network requesters. *Ladder* topologies lead to higher average overheads in comparison to other topologies (94.59% in scenarios with locations requirements and 23.36% in scenarios without locations requirements), as the topological structure has, on average, longer distances (in terms of hops) between nodes, in addition to the absence of hubs. *Hub & spoke* networks achieve the lowest overhead rates (64.67% in the scenario with location requirements), as they have a higher number of hubs and interconnections between nodes. In the scenario without location requirements (Figure 3(b)), the average overhead rates observed in the *ladder* and *star* infrastructures are similar – 23.36% and 21.92%, respectively – whereas the overhead rate caused by the employment of the *hub & spoke* topology is only 1.89%. This demonstrates that, when location requirements are considered, virtual links tend to be mapped to longer paths in the substrate, due to a higher average distance between the locations where virtual routers are hosted. However, topologies that have a higher number of hubs tend to reduce the impact of these constraints.

Next, we evaluate the average consumption of physical resources – CPU and memory capacity of routers, as well as link bandwidth – in all experiments. Figure 4 shows the average CPU consumption of physical routers. Considering location requirements, we observe that the average CPU consumption when employing the *ladder* topology (11.49%) is approximately twice as high as the average CPU consumption in the experiment in which we use the *star* topology (5.36%). The *hub & spoke* topology led to an average CPU



(a) Scenarios with location requirements.




(b) Scenarios without location requirements.

Fig. 5. Average memory usage of physical routers.


consumption of 16.82%. CPU consumption in experiments without location constraints (Figure 4(b)) is noticeably higher. The average CPU consumption in the *star* network is 23.80%, whereas in the *ladder* and *hub & spoke* networks, 35.32% and 38.60% of CPU resources are consumed, respectively. Increased consumption of CPU resources is influenced by the number of virtual networks embedded on the substrate. This explains the higher consumption in scenarios that do not consider location requirements, as in these cases rejection rates are lower. However, the relatively low CPU consumption, which does not exceed 40% in any experiment, reveals that the rejection of virtual network requests is not caused by the global exhaustion of CPU resources in the infrastructure.

In Figure 5 we present the average memory consumption of physical routers. The behavior of memory usage in physical routers is similar to that of CPU usage. In experiments that consider location requirements, the average memory usage is of 5.65% in the *star* topology, 11.21% in the *ladder* topology, and 15.93% in the *hub & spoke* topology. In contrast, experiments that do not consider such requirements lead to an average memory utilization of 23.49% in the *star* topology, 35.52% in the *ladder* topology, and 38.03% in the *hub & spoke* topology. Note that memory resources are not fully depleted in any of the experiments. Thus, we can state that this factor, similarly to CPU usage, is not the cause of the rejection rates observed in these scenarios.

The average bandwidth consumption in physical links is shown in Figure 6. In experiments in which virtual networks have location requirements, average link utilization is of 18.01% in the *star* topology, 40.63% in the *ladder* topology, and 35.89% in the *hub & spoke* topology. In the experiments



(a) Scenarios with location requirements.




(b) Scenarios without location requirements.

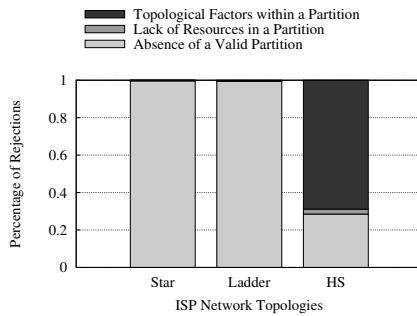
Fig. 6. Average bandwidth usage of physical links.

without location requirements, the average consumption is 44.93% in the *star* topology, 66.50% in the *ladder* topology, and 54.08% in the *hub & spoke* topology. These results show that the overall bandwidth consumption in the physical network, in an analogous manner to CPU and memory usage, is not responsible for the rejection rates observed in the experiments. The saturation of specific points in the physical infrastructure makes it impossible to embed a higher number of virtual networks, even though a global view of the physical network reveals a considerable amount of available resources.

ISP network topologies may have different connectivity degrees and, due to the depletion of resources in specific parts of the substrate (e.g. links connected to hubs or bridges), infrastructure partitioning occurs at different levels. In this context, a partition is defined as a strongly connected component (SCC) with residual bandwidth equal to or greater than the average link bandwidth requested by VNs. We observed that, in scenarios with location requirements, *star* topologies have, upon the embedding of a new VN, 13.60 partitions with an average of 3.63 routers per component, while *ladder* topologies have, on average, 12.30 partitions with 4.38 routers per component. In turn, *hub & spoke* topologies present lower levels of partitioning and a higher number of routers per component, i.e. 7.22 partitions and 7.52 routers, respectively. In scenarios without location requirements, infrastructure partitioning levels are higher due to higher link bandwidth consumption. *Star* and *ladder* topologies have similar average partitioning levels; 19.20 and 19.67 partitions with 2.61 and 2.48 routers, respectively. *Hub & spoke* topologies have, on average, 12.65 partitions and 3.94 routers per component. Next, we analyze how this and other factors influence the



(a) Scenarios with location requirements.



(b) Scenarios without location requirements.

Fig. 7. Average percentage of rejection causes observed in all experiments.

rejection rates observed.

Figure 7 shows results regarding causes of rejection in the performed experiments. Three main causes of rejection have been observed: the absence of a valid partition, the lack of physical router resources within a partition, and other topological factors. The first cause is related to the unavailability of a proper partition on the infrastructure with size equal to or greater than the number of virtual routers requested. This means that it is not possible to map the VN request on the topology due to the lack of connectivity between physical routers. In other words, there is no SCC that has, for each pair (a,b) of routers, a path between “a” and “b” with bandwidth greater than or equal to the maximum bandwidth requested. In scenarios with location requirements, 59.45% of rejections on *star* topologies as well as 54.44% on *ladder* topologies are caused by this kind of connectivity problem. In scenarios without location requirements, 99.62% of rejections on *star* topologies and 99.47% of rejections on *ladder* topologies are associated with this cause. In *hub & spoke* topologies, rejections associated with this cause amount to 20.85% in scenarios with location requirements and 28.50% in scenarios without such requirements. These results point to a direct correlation between higher partitioning levels and higher rejection rates.


The second cause of rejection – the lack of physical router resources within a partition – occurs when there is an appropriate partition (*i.e.*, there is a SCC that contains enough physical routers to map all virtual routers in the VN request), but it is not possible to map the requested VN due to the depletion of physical router resources (CPU and memory) in this partition. In scenarios with location requirements, rejection rates caused

by insufficient router resources are 3.2%, 1.8% and 14.7% on *star*, *ladder* and *hub & spoke* topologies, respectively. In scenarios without such requirements, these rates are, respectively, 0.33%, 0.37% and 2.58%. *Star* and *ladder* topologies present similar percentages in both scenarios. In scenarios without location requirements, an insignificant number of VNs are rejected due to this cause, as the main cause of rejection is infrastructure partitioning. In *hub & spoke* topologies, when considering location requirements, the percentage of VN rejections associated with this cause is higher in contrast to other topologies. This behavior is related to lower rejection rates associated with the absence of valid partitions in this topology, which leads to greater resource usage in physical routers.


The third cause of rejection refers to VN requests that are denied despite the existence of an adequate partition and available resources on its routers. In this case, rejections are a result of topological factors within a partition that make it impossible to map an incoming VN request. This occurs either when location constraints can not be met or when VNs have a significantly different topology in relation to the topology available in the partition. In the latter situation, a greater amount of resources will be required to embed these VNs, which may not be available. When considering location requirements, *star* and *ladder* topologies have, respectively, 37.30% and 43.71% of rejection rates associated with topological properties. In scenarios without location requirements, these percentages are 0.04% and 0.15%, respectively. Rejections related to topological factors in *hub & spoke* networks amount to 64.39% in scenarios with location requirements, and 68.90% in scenarios with no such requirements. When analyzing rejections associated with this cause, *star* and *ladder* topologies present a similar behavior, as partitioning levels and overall rejection rates are similar in both topologies. Nevertheless, minimal differences exist, which can be explained by topological factors such as the absence of hub nodes in *ladder* networks. In *hub & spoke* networks, however, topological factors are the most significant cause of rejection rates. Although this type of topology presents higher connectivity, the rejection of a lower number of VNs leads to higher bandwidth usage in physical links, which leads to reduced connectivity within network partitions.

Last, in Figure 8 we present the average time needed to find the optimal mapping of each accepted request. We emphasize that, in this graph, the vertical axis is shown in logarithmic scale, as results differ significantly among results. In all scenarios with location requirements, the time needed to optimally embed virtual network requests remains under 1 second. The use of *star* topology leads to an average solution time of 0.055 seconds, the *ladder* topology to an average time of 0.20 seconds, and the *hub & spoke* topology, to an average time of 0.44 seconds. In the remaining scenarios, which do not consider location requirements, all averages remain under 3 seconds. The *star* topology leads to an average solution time of 0.13 seconds, the *ladder* topology to an average time of 0.92 seconds, and the *hub & spoke* topology, to an average time of 2.41 seconds.

The relatively high solution times observed in scenarios that employ the *hub & spoke* topology are explained by the



(a) Scenarios with location requirements.



(b) Scenarios without location requirements.

Fig. 8. Average time needed to find the optimal solution.

presence of a higher number of links in this topology. This characteristic leads to a larger set of possible mappings for each virtual network request, which tends to increase the time needed to find the optimal mapping. Similarly, the removal of location requirements also leads to a larger space of feasible solutions, which explains the higher average time observed in such scenarios. Furthermore, a number of peaks are observed near the beginning of the experiments, reaching a maximum of 24.81 seconds in the scenario that employs the *ladder* topology and 9.99 seconds in the scenario in which the *hub & spoke* topology is employed. This behavior can be explained by the larger amount of unused resources in the beginning of these experiments (when the availability of resources on physical networks is substantially high), which also increases the number of possible mappings for virtual networks on the substrates.

The results presented in this section show that the employment of different types of physical network topologies in network virtualization environments causes a significant impact on rejection rates and physical resource usage. This impact is even more pronounced when location requirements of virtual networks are considered. Such experiments reveal that the rejection of virtual network requests is not caused by the overall depletion of resources in the infrastructure. Instead, it is caused by factors related to certain topological features.

The main factor that influences the rejection of virtual networks is resource depletion in specific regions of the substrate, which leads to higher partitioning levels in the infrastructure. For example, in *star* and *hub & spoke* topologies, the exhaustion of resources in physical hubs, as well as links connected to them, tends to be the main cause that

makes it impossible to map new virtual network requests even when there are partitions with sufficient size and connectivity. In *ladder* topologies, in which there are no hubs, the main cause for the increase in rejected requests is the depletion of resources in specific connections between nodes. Links used as “bridges” to interconnect different points of the infrastructure can become bottlenecks and, if the bandwidth of one of these links is sufficiently depleted, the infrastructure is partitioned into two groups of routers with no connectivity between them. Moreover, in all studied topologies, partitioning may lead to the existence of segments that, despite having adequate size and sufficient resources, are unable to accommodate certain VNs due to topological incompatibilities.

In summary, mapping virtual network in *star* topologies leads to low solution times but also high rejection rates, and consequently, to low physical resource usage. In *ladder* topologies, rejection rates and solution times have intermediate values in relation to other topology classes, but embedded virtual networks tend to consume a greater amount of bandwidth. Finally, *hub & spoke* topologies lead to low rejection rates and average bandwidth overhead, but the time needed to find the optimal solution is comparatively higher.

VI. CONCLUSION

Network virtualization is a topic that has received considerable attention both from the scientific community and the Industry, resulting in a series of studies involving, predominantly, issues related to virtual network embedding. However, as far as we are aware, there have been no previous attempts to evaluate the result of mapping strategies considering different network topologies in a precise manner. Previous work, in general, use organic or generic topologies, which do not faithfully represent the topological properties present in real InP networks.

After formalizing an optimal *online* virtual network embedding model and applying it on substrates with different topological features that are typically present in InP networks, we characterized the impact of different types of topologies regarding rejection rates and physical resource usage. The obtained results evidence the significant impact caused by embedding virtual networks on physical substrates with different topological features. The ability to embed virtual networks is hindered by resource depletion in some specific points of the physical infrastructure, although a global view of the network reveals that there are still resources available in the remainder of the substrate. This impact is even more expressive when the embedding model considers location requirements of virtual networks.

As future work, we intend to investigate strategies for reducing rejection rates observed in typical InP networks. More specifically, we intend to investigate ways to expand or scale specific points of physical substrates in order to reduce the rejection of virtual networks, considering the necessary costs to make such changes in the infrastructure. The main idea is to identify problematic regions of the substrate based on the size and total number of SCCs and, through specific adjustments in the infrastructure, to increase the acceptance of future virtual network requests in a consistent manner.

REFERENCES

- [1] 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. [Online]. Available: <http://doi.acm.org/10.1145/1355734.1355737>
- [2] M. Chowdhury, M. Rahman, and R. Boutaba, "Vineyard: Virtual network embedding algorithms with coordinated node and link mapping," *Networking, IEEE/ACM Transactions on*, vol. 20, no. 1, pp. 206–219, 2012.
- [3] G. Alkimi, D. Batista, and N. da Fonseca, "Mapping virtual networks onto substrate networks," *Journal of Internet Services and Applications*, vol. 4, no. 1, p. 3, 2013.
- [4] X. Cheng, S. Su, Z. Zhang, H. Wang, F. Yang, Y. Luo, and J. Wang, "Virtual network embedding through topology-aware node ranking," *SIGCOMM Computer Communication Review*, vol. 41, no. 2, pp. 38 – 47, Apr. 2011. [Online]. Available: <http://doi.acm.org/10.1145/1971162.1971168>
- [5] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network topologies: inference, modeling, and generation," *Communications Surveys Tutorials, IEEE*, vol. 10, no. 2, pp. 48 – 69, quarter 2008.
- [6] R. Albert and A.-L. Barabási, "Topology of evolving networks: Local events and universality," *Physical Review Letters*, vol. 85, pp. 5234 – 5237, Dec 2000.
- [7] S. Davy, J. Serrat, A. Astorga, B. Jennings, and J. Rubio-Loyola, "Policy-assisted planning and deployment of virtual networks," in *Proceedings of the 7th International Conference on Network and Service Management (CNSM)*, Paris, France, oct. 2011, pp. 1–8.
- [8] N. Kamiyama, R. Kawahara, T. Mori, S. Harada, and H. Hasegawa, "Impact of topology on parallel video streaming," in *Network Operations and Management Symposium (NOMS), 2010 IEEE*, april 2010, pp. 607 – 614.
- [9] I. Fajjari, N. Aitsaadi, G. Pujolle, and H. Zimmermann, "Vne-ac: Virtual network embedding algorithm based on ant colony metaheuristic," in *Communications (ICC), 2011 IEEE International Conference on*, 2011, pp. 1–6.
- [10] I. Houidi, W. Louati, W. B. Ameer, and D. Zeglache, "Virtual network provisioning across multiple substrate networks," *Computer Networks*, vol. 55, no. 4, pp. 1011 – 1023, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128610003786>