Virtual Network Embedding in Optical Infrastructures

Albert Pagès, Jordi Perelló, Salvatore Spadaro, *Member, IEEE*
Advanced Broadband Communications Centre (CCABA), Universitat Politècnica de Catalunya (UPC)
Jordi Girona 1-3, 08034 Barcelona, Spain
Email: {albertpages, spadaro}@tsc.upc.edu, perello@ac.upc.edu

ABSTRACT
Optical network virtualisation has recently become a hot topic among the research community due to its potential to enable the shaping of the Infrastructure as a Service (IaaS) concept. A major challenge that arises of paramount importance when virtualising a network is the so called Virtual Network (VN) embedding problem. While this problem has been extensively treated for electrical networks, there is very little work regarding Virtual Optical Network (VONs) embedding. This paper addresses the VON embedding problem, giving an insight about the potential challenges and presenting solutions that allow to optimally solving this problem in dynamic network scenarios.
Keywords: Optical networks virtualisation, Network embedding, Optimization.

1. INTRODUCTION

1.1 Concept of Virtual Network

The concept of virtualization was introduced in the Information Technology (IT) realm in the 1960s [1], with the introduction of the Virtual Machine (VM), as result of placing a virtual layer between hardware and software layers. This allowed the abstraction of portions of an underlying physical resource, so that an end-user could directly interact with a portion of this resource, while perceiving it as a single real resource. As a result, it was possible to run multiple autonomous and isolated VMs on top of a single shared physical machine.

Network virtualization extended this concept from individual nodes or resources to data communication networks resulting in the advent of the so called Virtual Private Networks (VPNs) [2]. VPNs where initially thought as a way to provision logically separated private networks over a public infrastructure by tunnelling data traffic between geographically distant sites, offering connectivity services between these sites. The network operator owning the physical infrastructure establishes the connections composing the VPN and provides them together as a service to the client. Note that in VPNs all the connections are established and managed by the network owner, resulting in a client-server relationship.

An evolution of the VPN is the Virtual Network (VN) paradigm, which results in one step ahead on the evolutionary path towards the Future Internet. The VN paradigm pleads for a more disruptive approach, where network owners can offer parts of their infrastructures as services to external service providers, so that they are completely free to manage them to offer end-to-end services to final users. It aims at introducing the essentials of virtualization to network environments. VN are slices of the physical infrastructure, including link capacity and node resources, which are offered as a service to the client. In such a paradigm, the client is not restricted to only use the bandwidth of some connections, but can also interact with the underlying physical equipment to configure them. The main idea consists of creating several co-existing logical network instances, each one with its own topology and resources, over a shared physical substrate as shown in Fig. 1.

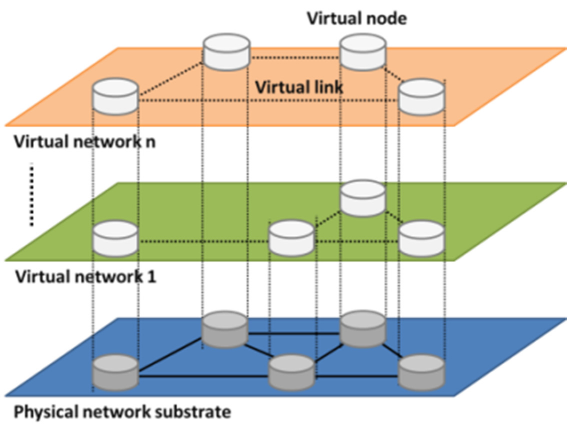

Figure 1. The virtual network paradigm.

1.2 Virtual Network Embedding

The VN paradigm opens many challenges that have to be addressed. Among them, the problem of how to map the requested resources by a VN into the current physical substrate arises of paramount importance in order to efficiently provision the requests and maximize the resource utilization of the underlying physical infrastructure. This problem has been addressed in the literature as the virtual network embedding problem [3].

Namely, given the topology and characteristics of a VN demand, it is necessary to map to actual physical resources the virtual nodes and virtual links of that demand. Essentially, the problem can be divided into two main sub-problems: node and link mapping. As its name suggests, the node mapping problem studies how to map the virtual nodes of a demand into a subset of physical nodes of the underlying substrate, while the link mapping problem studies the proper way of mapping the virtual links between virtual nodes into the physical network. Traditionally, both problems have been addressed separately, such as in [4], but this leads to suboptimal solutions. More refined studies propose that both stages should be attacked simultaneously or in a coordinated way to increase the optimality of the mapping solution [5].

2. VIRTUAL OPTICAL NETWORK EMBEDDING: CHALLENGES AND OPEN ISSUES

Despite the vast research done regarding the topic in Layer-2/3 networks (e.g. IP networks), there exists very few works that focus on Virtual Optical Networks (VONs) embedding, due the fact that virtualization techniques in the optical domain are still at a very early stage of development [6].

The main challenges that the optical domain poses are derived from the analogue nature of the optical substrate in opposition to the electrical one, and, as a consequence, new constraints appear on the arena. When an optical network resource or element is virtualized, for example by partitioning, due to the Physical Layer Impairments (PLIs), different instances of virtualized optical resources sharing the same physical optical resource may interfere with each other. Therefore, any virtualization paradigm must take into account the physical characteristics of optical network resources [7]. Moreover, in an all-optical transparent scenario, where nodes are not equipped with wavelength conversion capabilities, the wavelength continuity constraint must be ensured for the physical paths that compose the virtual links of the VON.

Regarding the embedding problem, almost no work has been carried out. Focusing on the virtual links in VONs, each one of them can be understood as a lightpath demand between a source node and a destination node in the physical substrate. The problem of how to allocate resources to a lightpath in a Wavelength Division Multiplexing (WDM) network is known as the Route and Wavelength Assignment (RWA) problem, which is devoted to find a path and a wavelength to be used by the lightpath. In the context of VONs, the embedding problem includes an RWA problem but noteworthy differences exist, arising from the fact that VONs should be treated as entities on their own instead of a composition of lightpaths. Indeed a VON is successfully provisioned if, and only if, all composing lightpaths are provisioned too.

Depending on the services that are meant to be provided over a VON, the RWA problem can differ. If transparent services should be provisioned over a VON, i.e. no regeneration or Optical-Electrical-Optical (OEO) conversion is performed in the transit virtual nodes, it is necessary to exactly allocate the same subset of physical wavelengths in each of its virtual links, otherwise transparency will be unavailable for certain pair of source-destination nodes. This is what is known as transparent VONs.

Alternatively, in opaque VONs, it is assumed that electronic termination capabilities are physically present at each VON node and opaque transport services are provided from the VON viewpoint. In such a case, there is no need to allocate the same set of wavelengths for each virtual link, but can differ thanks to the OEO conversion stages. Note that in both type of VONs it is still necessary to preserve the wavelength continuity constraint for single virtual links. In [8] we have presented Integer Linear Programming (ILP) formulations that address the offline embedding of both types of VONs in a physical substrate with the aim of maximize the number of VONs that can be allocated in that substrate given a demand set. As for the current article, we will focus on more dynamic network scenarios. The next section details the considered scenario and the proposed solution to attack the problem.

3. DYNAMIC VON EMBEDDING

3.1 Problem formulation

Two main approaches exist in the context of VN embedding depending on the considered scenario. In the offline or static approach, all the VN demands are known in advance and the objective is to maximize the number of correctly mapped demands over the physical substrate given the scarcity of resources. In the online or dynamic approach, demands are not known in advance but they arrive dynamically at the network. Work regarding this scenario tries to embed a demand in a way that will favour the embedding of future demands.

In this section we focus on the dynamic scenario in the context of VONs. In this regard, we assume that VON demands arrive following a dynamic traffic profile at a physical substrate network with limited resources. The aim in such scenario, as it has been said before, is to embed the VONs in a manner that will favour the embedding of future incoming demands. To this end, we present an ILP formulation whose goal is to balance the load of the underlying physical resources. We also assume that demands may be served in batches to increase the optimality of the mapping solution provided by the ILP formulation. In this sense, the objective of the presented formulation is dual: maximize the number of successfully embedded VONs inside the batch while at the same time balance the load of the physical resources. Opaque VONs are assumed throughout this section.

Let the optical network substrate be characterized by a graph G = (N, E), where N denotes the set of nodes and $E = \{(i,j),(j,i): i, j \in N, i \neq j\}$ the set of physical links. Consider D as the batch of VON demands to be allocated over the optical network. Each demand $d \in D$, is characterized by a graph $G'd = (N'd, E'd)$, where $N'd$ denotes the set of virtual nodes and $E'd = \{(i,j),(j,i): i, j \in N'd, i \neq j\}$ the set of virtual links. Let W denote the set of available wavelengths per physical link and Wd denote the number of wavelengths per virtual link desired by demand $d \in D$.

We define P as the set of paths in the physical network, $P_{e'}^{d} \subseteq P$ as the set of candidate paths for virtual link e' in demand d and $P_{e'}^{e} \subseteq P$ as the set of candidate paths for virtual link e' that traverse physical link $e \in E$. Additionally, we denote R_p as the residual bandwidth of physical path p, being the residual bandwidth the number of wavelength with continuity between source and destination of the path, and W_u^p as the set of wavelength that are already in use in path p. We also assume that physical nodes have no limitations in the number of transponders that are using. The problem variables are:

$x(d, e', p, w)$ = {1 if virtual link e' in demand d is mapped over physical path p and wavelength w, 0 otherwise}.
$y(d, e', p)$ = {1 if virtual link e' in demand d is exclusively mapped over physical path p, 0 otherwise}.
$z(d)$ = {1 if demand d is successfully mapped, 0 otherwise}.

The ILP formulation is stated below:

$$\max\ \alpha \frac{1}{|D|}\sum_{d\in D} z(d) + \beta \frac{1}{|D||W|}\sum_{d\in D}\frac{1}{|E_d^{'}|}\sum_{e'\in E_d^{'}}\frac{1}{|P_{e'}^{d}|}\sum_{p\in P_{e'}^{d}} R_p \sum_{w\in W} x(d,e',p,w) \text{ s.t.} \tag{1}$$

$$\sum_{p\in P_{e'}^{d}}\sum_{w\in W} x(d,e',p,w) \le W_d,\ \forall\ d\in D,\ e'\in E_d^{'} \tag{2}$$

$$\sum_{d\in D}\sum_{e'\in E_d^{'}}\sum_{p\in P_{e'}^{e}} x(d,e',p,w) \le 1,\ \forall\ d\in D,\ e\in E,\ w\in W \tag{3}$$

$$x(d,e',p,w) = 0,\ \forall\ d\in D,\ e'\in E_d^{'},\ p\in P_{e'}^{d},\ w\in W_u^{p} \tag{4}$$

$$y(d,e',p) \le \frac{1}{|W_d|}\sum_{w\in W} x(d,e',p,w),\ \forall\ d\in D,\ e'\in E_d^{'},\ p\in P_{e'}^{d} \tag{5}$$

$$z(d) \le \frac{1}{|E_d^{'}|}\sum_{e'\in E_d^{'}}\sum_{p\in P_{e'}^{d}} y(d,e',p),\ \forall\ d\in D \tag{6}$$

Objective function (1) aims at maximizing the number of successfully served demands inside the batch while balancing the load in the physical network by choosing the less congested paths. The factors α and β are pondering factors that serve to put more or less weight to the terms. Constraints (2) ensure that at most W_d wavelength will be provisioned for every virtual link in demand d. Constraints (3) are the wavelength clashing constraints. Constraints (4) avoid using pairs (p, w) to map the demand that are already in use by previous allocated demands. Constraints (5) ensure that all wavelengths requested by a virtual link are mapped over the same path; this is to avoid delay problems at destination. Constraints (6) discriminate if demand d is successfully served.

3.2 Results and discussion

In order to test the presented formulation, we have performed a series of simulations. They are run over the 16-Node EON core network topology assuming 40 wavelengths per physical link. For the results, we assume that VON requests arrive at the network following a Poisson traffic profile with exponentially distributed mean Holding Times (HTs). Hence, different load scenarios are obtained modifying the demands' mean Inter-arrival Times (IATs) accordingly (load = HT/IAT).

We have generated 10^5 random VON demand following a 3 step process: firstly, 3 or 4 physical network nodes (with equiprobability) are randomly selected as virtual nodes for each demand. Next, the selected virtual nodes are then randomly connected with equiprobability and preventing the generation of non-connected graphs. Finally, the parameter W_d is set to {1, 2} with p = {80, 20}%. Moreover, we have limited the size of the candidate path set per virtual link to 10. Also, we have considered different size of the batch in order to analyse the impact of this factor on the Blocking Probability (BP) and execution time of the model. In this regard, a batch is constructed putting in a queue the demands until the desired size is reached. We do not consider any timeout for which the batch would be automatically processed if it has not reach the desired size after a set period of time.

Figure 2 (left) depicts the BP as a function of the network load for different batch sizes while Fig. 2 (right) displays the execution time of the model. Simulations are run in i5 Quad Core at 3.3 GHz PCs with 4 GB of RAM memory and using IBM Cplex v.12.2 as optimization software. For these results we have considered α = 1 and β = 0.5, that is, we prioritize the number of VONs that are successfully embedded inside the batch over balancing the load of the physical paths in the network.

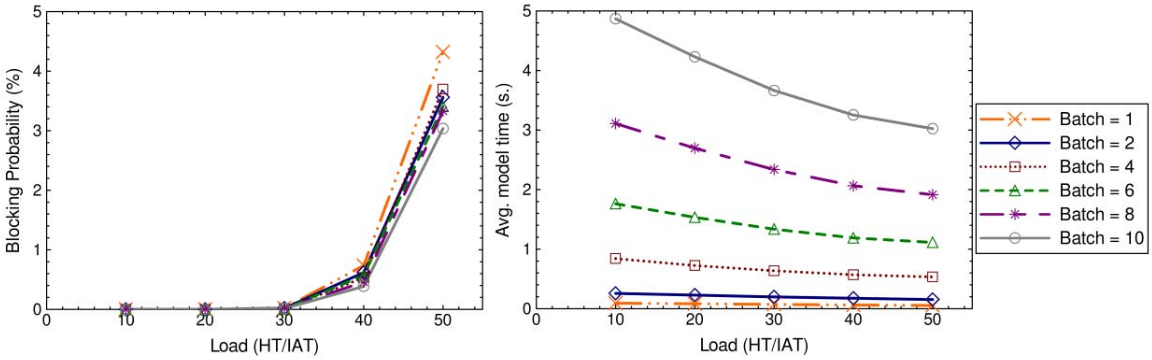

Figure 2. Blocking probability (left) and average model execution time (right) as a function of the load.

From the results above, we can observe that using bigger batches does not present a clear advantage in front of smaller batches. Focusing in the BP, for a load equal to 40 the difference between a batch of size 1 and 10 is about 0.35 % while this difference rises up to 1.3 % for a load equal to 50. On the other hand, for a batch equal to 1, the model execution times remains in the order of few milliseconds while for a batch equal to 10 this time ranges from about 5 to 3 seconds. In view of this, due the little improving in BP that offers the use of bigger batches compared to smaller batches, it may not justify the extra amount of time needed to construct the batch and execute the model, so smaller batches or even serve the demands one by one would be the preferred option.

4. CONCLUSIONS AND FUTURE WORK

We have discussed the problem of VN embedding in the context of optical infrastructures, presenting potential challenges and offering solutions for different network scenarios. Specifically, the current paper has focused on providing an ILP formulation for the dynamic scenario. The presented formulation aims at optimally allocate the VON demands in order to reduce the BP probability.

As for future work, and following the trend of the cloud computing paradigm, the problem of how to combine optical resources with IT resources in order to compose Virtual Infrastructures (VIs) with both networking and computational capabilities still has to be studied by the research community. In such a context, due the heterogeneity of the underlying physical substrate (optical network + IT resources), solutions jointly accounting for the particularities of each world should be considered when trying to embed such a VI.

ACKNOWLEDGEMENTS

The work in this paper has been supported by the Government of Catalonia and the European Social Fund through a FI-AGAUR research scholarship grant and the Spanish Science Ministry through the ELASTIC project (TEC2011-27310).

REFERENCES

[1] Zhang Jian, Li Xiaoyong, Guan Haibing: The optimization of Xen network virtualization, in *Proc. of International Conference on Computer Science and Software Engineering*, vol. 3, pp. 431-436, Dec. 2008.

[2] S. Khanvilkar, A. Khokhar: Virtual private networks: an overview with performance evaluation, *IEEE Communications Magazine*, vol. 42, no. 10, pp. 146-154, Oct. 2004.

[3] Zhiping Cai, *et al.*: Virtual network embedding for evolving networks, in *Proc. of Global Telecommunications Conference*, Dec. 2010.

[4] Y. Zhu, M. Ammar: Algorithms for assigning substrate network resources to virtual network components, in *Proc. of IEEE International Conference on Computer Communications*, Apr. 2006.

[5] N.M.M.K. Chowdhury, M.R. Rahman, R. Boutaba: Virtual network embedding with coordinated node and link mapping, in *Proc. of IEEE International Conference on Computer Communications*, Apr. 2009.

[6] R. Nejabati, *et al.*: Optical network virtualization, in *Proc. of 15th International Conference on Optical Networking Design and Modeling*, Feb. 2011.

[7] Shuping Peng, *et al.*: An impairment-aware virtual optical network composition mechanism for future Internet, in *Proc. of 37th European Conference and Exhibition on Optical Communication*, Sept. 2011.

[8] A. Pagès, *et al.*: Strategies for virtual optical network allocation, *IEEE Communication Letters*, vol. 16, no. 2, pp. 268-271, Feb. 2012.