

Cost-Efficient Virtual Optical Network Embedding for Manageable Inter-Data-Center Connectivity

Jordi Perelló, Pablo Pavón-Mariño, and Salvatore Spadaro

Network virtualization opens the door to novel infrastructure services offering connectivity and node manageability. In this letter, we focus on the cost-efficient embedding of on-demand virtual optical network requests for interconnecting geographically distributed data centers. We present a mixed integer linear programming formulation that introduces flexibility in the virtual-physical node mapping to optimize the usage of the underlying physical resources. Illustrative results show that flexibility in the node mapping can reduce the number of add-drop ports required to serve the offered demands by 40%.

Keywords: VON, embedding, cost-efficiency, MILP.

I. Introduction

Network virtualization [1] allows network operators to capitalize on their network infrastructures by offering parts of the capacity in their physical nodes and links as infrastructure services to build virtual networks. Network virtualization fosters manageability and customizability, that is, service providers leasing a virtual network are free to configure and operate it to effectively deliver services to end users.

Network operators willing to strike into the virtual network service market seek to maximize profits from their network infrastructures by supporting as many virtual networks as possible, while minimizing equipment costs. With this in mind, the virtual network embedding problem [2] deals with the efficient mapping of virtual networks over a shared physical

network substrate. This problem has been generally addressed in the literature. However, very few works address this problem for transparent optical networks, the main candidates to implement the future metro-core network substrates.

This work concentrates on the cost-effective embedding of virtual optical networks (VONs) realizing the interconnection of data centers (DCs) inside a cloud. We present a novel mixed integer linear programming (MILP) model to optimally embed on-demand VON requests over a high-capacity transparent optical network. The model balances wavelength utilization at physical network links together with the add-drop port utilization at physical network nodes when allocating a VON request. This serves the purpose of maximizing link resource utilization, while minimizing the number of add-drop ports to be equipped per physical node to serve the incoming VON demands.

II. VONs for Data Center Interconnection

In this work, we envision the use of VONs as a promising solution for cloud service providers (CSPs) lacking physical network infrastructure to interconnect their DCs. Typically, CSPs emplace DCs in geographically distributed sites for high service reliability. In this way, they can replicate contents among them periodically, offering enhanced service availability upon failures. Additionally, for efficiency reasons, inter-DC workload migrations are very frequent, which can increase data flows among DCs up to daily totals of terabytes or petabytes. To realize such operations, a high-capacity optical network interconnecting DCs in a cloud is necessary. Moreover, such an optical network must provide flexibility and manageability, so as to let CSPs efficiently use it.

Thus, we believe that the VON solution perfectly suits CSPs' needs. A VON is composed of virtual nodes (that is, virtual ROADMs, in this letter), which are mapped over the physical

Manuscript received June 21, 2012; revised Aug. 2, 2012; accepted Aug. 8, 2012.

This work has been supported by the Spanish National projects ELASTIC (TEC2011-27310) and FIERRO (TEC2010-12250-E).

Jordi Perelló (phone: +34 93 405 40 65, perello@ac.upc.edu) is with the Computer Architecture Department, Universitat Politècnica de Catalunya, Barcelona, Spain.

Pablo Pavón-Mariño (pablo.pavon@upct.es) is with the Department of Information Technologies and Communications, Universidad Politécnica de Cartagena, Cartagena, Spain.

Salvatore Spadaro (spadaro@tsc.upc.edu) is with the Signal Theory and Communications Department, Universitat Politècnica de Catalunya, Barcelona, Spain.

<http://dx.doi.org/10.4218/etrij.13.0212.0271>




Fig. 1. VON for data center interconnection.

nodes of the underlying network substrate (that is, the physical ROADMs). Virtual ROADMs are interconnected by high-capacity virtual links, built on transparent lightpaths between the physical ROADMs over which the endpoints of the virtual links are mapped. In contrast to alternative services like Bandwidth on Demand or Layer-1 Virtual Private Networks, which only offer connectivity between remote sites, VONs also offer node reconfiguration capabilities. Indeed, CSPs leasing a VON can fully manage virtual nodes according to their needs, even from their own control plane. This allows them to dynamically set up and release connections between DCs at their convenience, optimizing the overall cloud performance.

In Fig. 1, we depict a VON that interconnects four DCs in a cloud. We assume that a virtual ROADM with a number of virtual add-drop ports is requested to connect each DC to the VON. A main objective of this work is to evaluate the benefits that an operator can get by introducing flexibility in the virtual-physical node mapping when embedding a VON, which should yield optimal resource utilization in both physical nodes and links. Hence, we assume that a virtual ROADM can be mapped to any physical ROADM within a geographical area around each DC (denoted as $A(VR_n)$ in Fig. 1, where n stands for the virtual ROADM ID). The connectivity from a DC to any physical ROADM within its area can be achieved through a high-capacity access or metro network owned by the same operator provisioning the VON. Virtual links connecting virtual ROADMs are mapped on transparent lightpaths between the physical ROADMs on which the remote endpoints of the virtual links are mapped. Thus, for each virtual link, we need as many lightpaths as virtual wavelengths it requests. In this work, we assume that all virtual links of a VON request the same number of virtual wavelengths.

We evaluate two virtual-physical node mapping approaches: *unrestricted* 1:1 and *unrestricted* 1: N . In the former, a virtual ROADM can only be mapped to one physical ROADM. In the latter, the virtualization middleware allows splitting a virtual

ROADM over multiple physical ROADMs within its area.

III. Related Work

Even though the virtual network embedding problem has been largely addressed for electrical networks, very few works on this topic exist for transparent optical network substrates. For instance, [3] presented a VON composition mechanism to serve on-demand VON requests with physical layer impairment (PLI) awareness over an optical network substrate. In particular, fixed 1:1 virtual-physical node mappings were applied based on location requirements, and no solutions were given for performing node mappings if multiple physical nodes match the location requirements of a virtual node. Then, virtual links were mapped on the shortest paths between the physical nodes hosting the virtual link endpoints. This work was extended to mixed-line-rate networks in [4]. The work in [5] also focused on the VON embedding problem, differentiating between transparent and opaque VONs, meaning that VON-wide optical transparency was desired or not, respectively. However, fixed 1:1 virtual-physical node mappings were also assumed, prior to the receipt of the VON requests. Finally, although not strictly addressing the VON embedding problem, [6] focused on the energy-efficient design of integrated virtual infrastructures (VIs) involving both optical and IT resources, where a VI must support a set of computation demands to be served in the network IT resource locations.

In contrast to the references above, the goal of this letter is to introduce flexibility in the virtual-physical node mapping, quantifying its benefits in terms of VON request blocking probability and required add-drop ports per substrate node to serve the incoming demands. We do not consider the effects of PLIs, since we assume that the distances of physical paths onto which virtual links are mapped are within the transparent reach of the optical signals (for example, 3,000 km at 10 Gbps [7]). Nonetheless, PLI awareness could also be achieved by pruning the candidate physical paths whose end-to-end distance exceeds the transparent transmission reach.

IV. Model Formulation

We model the physical substrate network as a graph $G_i(N_i, E_i)$, where N_i is the set of physical nodes and E_i is the set of unidirectional fibers that interconnect them. The same spectral grid of W_i wavelengths is assumed for each physical fiber. We also denote P_i as the set of candidate paths over the physical network, namely, $p_i \in P_i$ is a candidate route to allocate a lightpath. A VON request is represented as a graph $G_v(N_v, E_v)$, where N_v represents the set of virtual nodes and E_v represents the set of virtual links. Each virtual link of the VON requests

W_v virtual wavelengths. We assume that a virtual node should be provided with as many add/drop ports as the total number of output/input virtual wavelengths from/to it. We use the terminology a and b to refer to the source and destination endpoints, respectively, of a physical link, path, or virtual link.

The decision variables of the problem are as follows:

- $y_{n_v n_i} = \{1 \text{ if node } n_v \text{ is mapped on node } n_i; 0 \text{ otherwise}\}$, $n_v \in N_v, n_i \in N_i$.
- $x_{e_v p_i w} = \{1 \text{ if } e_v \text{ is mapped on } p_i \text{ with wavelength } w; 0 \text{ otherwise}\}$, $e_v \in E_v, p_i \in P_i, 1 \leq w \leq W_i$.
- $t(n_i)$ = number of add-drop ports used in n_i once the VON request is allocated.
- $t_{\max}(n_v)$ = maximum add-drop port utilization among all $n_i \in A(n_v)$.
- $w(e_i)$ = number of wavelengths used at link e_i once the VON request is allocated.
- w_{\max} = maximum wavelength usage for all $e_i \in E_i$.

The proposed MILP formulation is

$$\min \alpha \sum_{n_v} t_{\max}(n_v) + (1 - \alpha) w_{\max}$$

subject to

$$\sum_{n_i} y_{n_v n_i} = 1, \forall n_v, \quad (1.a)$$

$$\sum_{n_i} y_{n_v n_i} \geq 1, \forall n_v, \quad (1.b)$$

$$y_{n_v n_i} = 0, \forall n_v, n_i \notin A(n_v), \quad (2)$$

$$x_{e_v p_i w} \leq y_{a(e_v) a(p_i)}, \forall e_v, p_i, w, \quad (3)$$

$$x_{e_v p_i w} \leq y_{b(e_v) b(p_i)}, \forall e_v, p_i, w, \quad (4)$$

$$\sum_{p_i} \sum_w x_{e_v p_i w} = W_v, \forall e_v, \quad (4)$$

$$\sum_{e_v} \sum_{p_i: e_i \in p_i} x_{e_v p_i w} \leq 1, \forall e_i, w, \quad (5)$$

$$t_p(n_i) + \sum_{e_v} \sum_{\substack{p_i: a(p_i)=n_i \\ p_i: b(p_i)=n_i}} \sum_w x_{e_v p_i w} = t(n_i), \forall n_i, \quad (6)$$

$$t(n_i) \leq t_{\max}(n_v), \forall n_v, n_i \in A(n_v), \quad (7)$$

$$w_p(e_i) + \sum_{e_v} \sum_{p_i: e_i \in p_i} \sum_w x_{e_v p_i w} = w(e_i), \forall e_i, \quad (8)$$

$$w(e_i) \leq w_{\max}, \forall e_i. \quad (9)$$

Constraints (1.a) and (1.b) ensure that each virtual node is mapped over only one (1.a) or multiple physical nodes (1.b), depending if unrestricted 1:1 or unrestricted 1: N node mapping is applied. Constraint (2) ensures that virtual nodes are only mapped on physical nodes within their area. Constraint (3) restricts virtual link mappings to paths connecting the physical nodes over which the remote endpoints of the virtual links are mapped. Constraint (4) ensures that W_v wavelengths are

assigned to all virtual links. Constraint (5) is the wavelength clashing constraint. Constraints (6) through (9) are used to store maximum add-drop port and wavelength usages once the VON is served. The number of add-drop ports and wavelengths used in physical node n_i and link e_i before serving the VON are $t_p(n_i)$ and $w_p(e_i)$, respectively.

V. Results and Discussion

Using the presented model, we quantify the benefits that a network operator can achieve by introducing flexibility in the virtual-physical node mapping with unrestricted 1:1 and unrestricted 1: N node mapping approaches. Simulations are run over the extended Abilene Internet2 network shown in Fig. 2, in which boxes denote areas where a DC can be located. In this network, 40 wavelengths per fiber link are assumed.

To benchmark the performance of the proposed approaches, a restricted node mapping approach has been also assumed, where the physical node on which a virtual node can be mapped is fixed. We do not show the MILP formulation behind this approach due to the lack of space. Note, however, that it is considerably simpler than the one presented in section IV. Indeed, variables $y_{n_v n_i}$ and constraints (1) through (3) are not needed here, as all $e_v \in E_v$ have already associated the physical nodes over which $a(e_v)$ and $b(e_v)$ must be mapped. Therefore, it is enough here to enforce $x_{e_v p_i w} = 0$ for all $p_i \in P_i$ not connecting such physical nodes, which can be specified as an additional set of constraints. In any case, execution times around 1 s are experienced in most occasions for all restricted and unrestricted approaches in a standard Intel Core i3 PC with 4 cores at 2.93 GHz and 2 GB RAM running CPLEX v.12.2, which highlights the practicality of the proposed formulations in a relatively large transport network, as the one represented in Fig. 2.




Fig. 2. Extended Abilene Internet2 network: 28 nodes, 37 bidirectional fiber links.




Fig. 3. (a) Average VON blocking probability and (b) peak add-drop port usage per node.

We assume that VON requests arrive at the network in a Poisson traffic process with exponentially distributed holding times (HTs). VON topologies are built upon arrival as follows. For each VON request, the number of virtual ROADMs is randomly chosen from 3 to 5. Each virtual ROADM pair is connected by a bidirectional virtual link with a probability of 0.5. Once the VON topology is built, we ensure its connectivity. Otherwise, we discard it, and a new topology is built. VONs request W_v wavelengths for all virtual links, where W_v is randomly chosen from $\{1, 2, 4\}$. Virtual nodes are assigned to the geographical areas with equiprobability, meaning that the DCs that should be connected to the VON are located in these areas. We consider that an area can allocate at most one virtual node of a VON.

Figure 3 depicts the VON blocking probability and average peak add-drop port usage per physical node for the evaluated virtual-physical node mapping approaches. Graphs are plotted as a function of the normalized VON HT, fixing the interarrival time of VON requests to 1 time unit. In the restricted approach, the physical node over which a virtual node must be mapped is chosen randomly among all nodes in its geographical area. In the model, $\alpha=0.5$ has been fixed. Moreover, P_i size has been limited to 6 paths for every physical node pair (that is, the shortest ones in terms of hops).

As seen in Fig. 3(a), both unrestricted node mapping approaches yield substantially lower VON blocking for medium and high loads, as flexibility in virtual node mappings allow better resource utilization in the network. However, almost no differences are observed between them in terms of blocking probability. This is due to the fact that good wavelength usage can already be achieved by mapping virtual to physical nodes 1:1 in the geographical area. In fact, the most significant benefits from applying unrestricted 1:N are observed in the physical node add-drop port usage. Figure 3(b) plots the average peak add-drop port usage to ensure zero blocking due to port unavailability among all network nodes. As shown, unrestricted node mapping approaches clearly

lower the number of add-drop ports needed per physical node in the restricted case, which is directly translated into CAPEX savings for network operators. Among them, unrestricted 1:N interestingly decreases the number of add-drop ports needed in the restricted and unrestricted 1:1 approaches by around 40% and 20%, respectively.

VI. Conclusion

In this letter, we addressed the efficient embedding of VONs for interconnecting geographically distributed data centers. Specifically, we investigated the benefits that a network operator can get by allowing flexibility in the virtual-physical node mapping process. We proposed two alternative flexible node mapping approaches: unrestricted 1:1 and unrestricted 1:N. The obtained results show that both approaches clearly reduce the VON request blocking and number of add-drop ports per physical node to serve the offered demands. Particularly, unrestricted 1:N reduces the number of add-drop ports per physical node by up to 40% against a restricted (that is, fixed) node mapping, arising as an attractive solution to be adopted by future network virtualization technologies, as it can significantly reduce equipment costs to network infrastructure providers.

References

- [1] N.M. Chowdhury and R. Boutaba, "Network Virtualization: State of the Art and Research Challenges," *IEEE Commun. Mag.*, vol. 47, no. 7, July 2009, pp. 20-26.
- [2] N.M. Chowdhury, M.R. Rahman, and R. Boutaba, "ViNEYard: Virtual Network Embedding Algorithms with Coordinated Node and Link Mapping," *IEEE/ACM Trans. Netw.*, vol. 20, no. 1, Feb. 2012, pp. 206-219.
- [3] S. Peng et al., "An Impairment-Aware Virtual Optical Network Composition Mechanism for Future Internet," *OSA Optics Express*, vol. 19, no. 26, Dec. 2011, pp. B251-B259.
- [4] S. Peng et al. "Virtual Optical Network Composition over Single-Line-Rate and Mixed-Line-Rate WDM Optical Networks," *Proc. OFC*, Mar. 2012.
- [5] A. Pagès et al., "Strategies for Virtual Optical Network Allocation," *IEEE Commun. Lett.*, vol. 16, no. 2, Feb. 2012, pp. 268-271.
- [6] M.P. Anastasopoulos et al., "Energy Aware Planning of Multiple Virtual Infrastructures over Converged Optical Network and IT Physical Resources," *OSA Optics Express*, vol. 19, no. 26, Dec. 2011, pp. B503-B508.
- [7] O. Rival and A. Morea, "Cost-Efficiency of Mixed 10-40-100Gb/s Networks and Elastic Optical Networks," *Proc. OFC*, Mar. 2011.