Optimal Allocation of Virtual Optical Networks for the Future Internet

Albert Pagès, Jordi Perelló, Salvatore Spadaro
Universitat Politècnica de Catalunya
Advanced Broadband Communications Center CCABA
{albertpages,spadaro}@tsc.upc.edu, perello@ac.upc.edu

Joan.A. García-Espín, Jordi Ferrer Riera, Sergi Figuerola
i2CAT Foundation, Catalunya
Distributed Applications and Networks Area
{jage,jordi.ferrer,sergi.figuerola}@i2cat.net

***Abstract*—Optical network infrastructures can be partitioned into multiple parallel, dedicated virtual networks for a physical infrastructure sharing purpose. However, different transport technologies may impact in both the amount and the characteristics of the different virtual instances that can be built on top of a single physical infrastructure. To analyse the impact of the transport technology in this regard, we present exact Integer Linear Programming (ILP) formulations that address the off-line problem of optimally allocate a set of virtual networks in two kind of substrates: wavelength switching and spectrum switching. Both formulations serve the purpose to provide opaque transport services from the virtual network point of view, where electronic terminations are assumed in the virtual network nodes. We carry out a series of experiments to validate the presented formulations and determine which is the impact of both substrates in the number of virtual networks that can be optimally allocated in the transport network.**

***Index Terms*—Optical Network Virtualisation, Integer Linear Programming.**

I. Introduction

Current service providers are focused on offering services on top of the infrastructures they own and manage. The end user has no control over these services and the provider-consumer relationship is far from being automated [1]. However, many emerging applications have increasing requirements in terms of bandwidth, Quality of Service (QoS), and manageability. Moreover, applications or even new paradigms such as Cloud computing o 3D-Video streaming require optimization and combined provisioning of different infrastructure resources and services that include both computational and networking realms [2]. As time goes by, demands are more and more variable and sporadic, driven by user behaviours and greedy applications. These new requirements are difficult to accommodate with the existing, rigid telecommunication operational models. However, through virtualisation, there exists the chance to allocate isolated instances from a given resource to different users or applications and manipulate it logically, before inferring changes to the real resource.

The concept of virtualisation in the IT realm was introduced by IBM in 1960s [3]. They introduced the virtual machine, as a result of introducing a virtual layer between hardware and software layers. Virtualisation technology can provide an isolated execution environment to applications, shield the dynamics and heterogeneous of hardware platform, and enable share and reuse of hardware resources. In order to create this virtual layer there exists the need for mapping the virtual elements on top of determined slices of the physical layer or substrate. Nowadays, several forms of virtualisation already exist, such as computational resources virtualisation (e.g. Cloud computing) or network link virtualisation basing on control services (e.g. Multi-Protocol Label Switching (MPLS) technology); or even router virtualisation using Virtual Router Forwardings (VRFs). Network virtualisation is a concept that precisely extends the idea of virtualisation from individual nodes or resources to entire networks [4]. The main idea consists of creating several co-existing logical network instances over a shared physical substrate.

The problem of how to map the requested resources along with a virtual network into a physical network substrate is referred in the literature as the virtual network embedding problem (cf. [5], [6]). Studies involving this problem seek finding the optimal mapping of the virtual network demands over a shared physical substrate with scarce network resources. Most of these studies are focused on layer-2/3 networks, such as Ethernet and IP networks, where virtualisation techniques are quite mature and well spread. In the optical realm, main challenges towards virtualising such kind of networks are derived from the hybrid nature of the optical substrate in comparison to the electrical one, and, as a consequence, the new constraints appearing on the arena. When an optical network resource or element is virtualised, for example by partitioning, due to the Physical Layer Impairments (PLIs), different instances of virtualised optical resources sharing the same physical optical resource may interfere with each other. Therefore, any virtualisation paradigm must take into account the physical characteristics of optical network resources [7].

The current manuscript addresses the off-line planning problem of optimally allocating a set of Virtual Optical Networks (VONs) over an all-optical network substrate, while accounting for the particularities of the optical transmission medium. We denote this problem as the Virtual Optical Network Allocation (VONA) problem. To the best of our knowledge, very few work in the literature has addressed this problem in the context of VONs.

We formulate the VONA problem as an ILP problem, studying the particularities that would entail the utilisation of diverse transport technologies for the optical network substrate. Furthermore, a series of experiments are carried out in order to demonstrate the validity of the ILP formulations proposed while giving an insight about the impact of the technology used in the transport plane in the number of VONs that can be successfully accommodated in a given optical network substrate.

The remainder of this paper is structured as follows: Section II gives an insight about the framework in which the work presented in this paper has been carried out, Section III explains the main assumptions taken about the working scenario, Section IV formally details the VONA problem and the solutions proposed to address it, Section V presents experimental results that validate the proposed solutions and discusses about them, and finally, Section VI states the main conclusions that can be extracted regarding the presented work.

II. Background

The VONA problem addressed in the manuscript is contextualised within the Generalised Architecture for Dynamic Infrastructure Services (GEYSERS)[1] EC-funded project. GEYSERS' vision is to qualify optical infrastructure providers and network operators with a new architecture, in order to enhance their traditional business operations by allowing a flexible role decoupling. Optical network infrastructure providers will compose virtual infrastructures and rent them out to network operators, who will run cost-efficient, dynamic and mission-specific networks by means of integrated control and management techniques. The GEYSERS solution is based on partitioning the physical optical network infrastructure in order to create VONs integrated with the IT resources connected to the edges of the network. This logical composition overcomes the limitations of networks and administrative domains segmentation [8].

Fig. 1 depicts the generalised architecture defined and designed within the GEYSERS project. The layered architecture is composed of the following elements: (i) the Service Middleware Layer (SML), which is a convergence layer in order to coordinate the management of IT resources that belong to an aggregate service; (ii) the Network Control Plane+ (NCP+) that is based on the Generalized Multi-Protocol Label Switching (GMPLS) and Path Computation Element (PCE) architecture and includes a set of extensions both in terms of architectural elements and protocol procedures in order to meet the emerging requirements; and (iii) the Logical Infrastructure Composition Layer (LICL) that is the layer responsible of abstracting and partitioning the underlying physical substrate and compose virtual infrastructures that are offered as a service.

In the GEYSERS architecture, the LICL acts as a middleware on top of the physical substrate and offers a set of tools that enable IT and optical network resource abstraction and virtualisation. Such an intermediate layer aims at decoupling infrastructure resource management from the actual service provisioning. LICL deals with the virtual infrastructure and virtual resource creation, management and handling. The business roles using the LICL (physical and virtual infrastructure providers) aim at maximising their revenues by means of efficiently improving the infrastructure performance and usage. As an

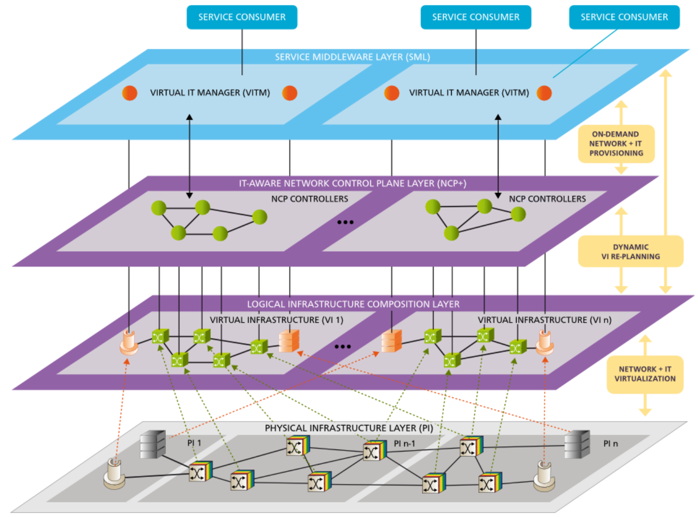

Fig. 1. EC FP7 GEYSERS project generic architecture

example of reference use case for multiple virtual infrastructure creation, one can consider the case that a company hosts an Enterprise Information System (EIS) externally on a Cloud rented from a provider. It relies on the resources provided by one or more IT and network infrastructure providers. It also connects data resources in an isolated virtual infrastructure.

Considering the role of the virtual infrastructure provider, the business revenue behind the use case is to provide the maximum number of virtual infrastructures on top of the same physical infrastructure. In order to host the maximum number of distinct EIS, the virtual infrastructure provider aims at maximising the number of virtual networks used to interconnect the data resources of the corresponding EIS. The technological details of the underlying infrastructure will limit the number of virtual instances that can be built. Assuming the optical substrate, we provide an study and analysis of the influence of the physical substrate and the diverse transport technologies in the number of VONs that can be allocated on top of it.

III. Assumptions

Two alternatives enabling optical transport technologies have been considered as cases of study for determining the impact of the underlying optical network substrate technology on the number of VONs that can be allocated by means of the LICL planning system: (i) wavelength switching following the fixed-size spectrum grid defined by the International Telecommunication Union (ITU)[2], where the minimum granularity for allocating a connection is a full wavelength [9]; and (ii) spectrum switching, following a flexible spectrum grid as proposed in the Spectrm-slice Elastic Optical Path Network (SLICE) architecture [10]. In such a case, demands request a portion of spectrum, equivalent to a number of Frequency Slots (FSs) that can be efficiently allocated and switched thanks to the use of optical Orthogonal Frequency Division Multiplexing

[1] http://www.geysers.eu

[2] http://www.itu.int

(OFDM) [11] and Bandwidth Variable Wavelength Cross Connects (BV-WXC) [12], [13].

In both scenarios we assume that opaque transport services are being provisioned from the VONs point of view, that is, assuming that every node in the VON has electronic termination capabilities (e.g., is equipped with IP routers). Moreover, we assume an all-optical network substrate without wavelength/spectrum conversion capabilities, so that every virtual link of the VON must ensure the wavelength/spectrum continuity constraint. Note, however, that thanks to the Optical-Electrical-Optical (OEO) conversion stages, such a wavelength/spectrum continuity constraint can be relaxed among the virtual links composing the VON. Lastly, it shall be pointed out that the effects of the PLIs introduced in the optical network substrate have not been considered in this work to assess the feasibility of the provisioned virtual links. Provided that the proposed models would have to be applied to very large network scenarios, where the PLI could be a concern, the set of candidate paths connecting VON neighbouring nodes should have to be restricted, including only those paths whose physical distance would enable the desired bit-rate.

IV. Problem Formulation

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 set of VON demands to be allocated over the optical network. Each demand $d \in D$, is characterized by a graph $G'_d = (N'_d, E'_d)$, $N'_d \subseteq N$, $E'_d = \{(i,j),(j,i) : i,j \in N'_d, i \neq j\}$. The VONA problem consists in accommodating all or the maximum number of VONs from the demand set, given the limited capacity of the underlying optical network. VONs are treated as entities instead of a composition of lightpaths, which makes VONA differ from classical route and assignment problems with the objective to maximize the number of lightpaths established. Indeed, a specific demand $d \in D$ is accommodated if and only if all its virtual links in E'_d can be mapped over available resources. The rest of this section presents optimal ILP formulations for both fixed-size and flexible grid cases.

A. Fixed-VONA

This subsection presents an ILP model of the fixed-size grid VONA problem, hereafter referred as Fixed-VONA. For this, let W denote the set of available wavelengths per physical link and W_d denote the number of wavelengths per virtual link desired by demand $d \in D$. Also, we define P as the set of paths in the physical network, $P_{\{e',e\}}$ as the set of $p \in P$ associated with virtual link e' that traverse edge $e \in E$, and $P_{\{e',d\}}$ as the set of $p \in P$ associated with virtual link e' in demand d. The problem variables of Fixed-VONA are:

$x(d,e',p,w) = \{1$ if for demand d the virtual link e' is supported through path p and wavelength w, 0 otherwise$\}$

$y(d,e',p) = \{1$ if in demand d all the wavelengths requested by virtual link e' use the same path p, 0 otherwise$\}$

$z(d) = \{1$ if demand d can be satisfied, 0 otherwise$\}$

The ILP formulation is stated below:

$$\max \sum_{d \in D} \alpha_d z(d), \text{s.t.} \tag{1}$$

$$\sum_{d \in D} \sum_{e' \in E'_d} \sum_{p \in P_{\{e',e\}}} x(d,e',p,w) \leq 1, \forall e \in E, w \in W \tag{2}$$

$$\sum_{p \in P_{\{e',d\}}} \sum_{w \in W} x(d,e',p,w) \leq W_d, \forall d \in D, e' \in E'_d \tag{3}$$

$$y(d,e',p) \leq \frac{1}{W_d} \sum_{w \in W} x(d,e',p,w), \forall d \in D, e' \in E'_d, p \in P_{\{e',d\}} \tag{4}$$

$$z(d) \leq \frac{1}{|E'_d|} \sum_{e' \in E'_d} \sum_{p \in P_{\{e',d\}}} y(d,e',p), \forall d \in D \tag{5}$$

Objective function (1) aims at maximizing the number of VONs to be allocated in the underlying optical network, where factors α_d are pondering factors used to put more or less weight in specific demands according to the network operator policies. Constraints (2) are the wavelength clashing constraints, avoiding that two virtual links are supported over the same wavelength in the same physical link. Constraints (3) ensure that at most W_d different wavelengths are assigned to every virtual link belonging to demand d. Constraints (4) ensure that every wavelength requested by virtual link e' is routed through the same path. Such constraints are adopted to avoid packet reordering problems that might appear at destination when sending data over different paths with different physical lengths and, thus, different end-to-end delays. Constraints (5) discriminate whether demand d is satisfied or not.

B. Flex-VONA

The flexible grid architecture allows to efficiently serve low data-rate sub-wavelength transmissions and ultra-high capacity super-wavelength transmissions onto the available network spectral resources, but poses new challenges compared to the classical Routing and Wavelength Assignment (RWA) problem applicable to Wavelength Switched Optical Networks (WSONs). Instead of wavelengths, a contiguous spectrum portion has to be allocated in flexible optical networks. Moreover, given a lack of spectrum conversion capabilities in the network, the assigned spectrum portion must show a continuity between the remote endpoints of the incoming connection requests (i.e., VON neighbouring nodes in this work). Both constraints, namely, spectrum contiguity and continuity constraints, must be ensured by the Routing and Spectrum Assignment (RSA) algorithm in the network.

In most works related to RSA, such as in [14], [15], it is assumed that the usable bandwidth of an optical fiber can be discretized into multiple FSs and so, the bandwidth requested by a demand can be converted into a number of FSs. Specifically, the authors in [15] proposed an ILP formulation for flexible optical networks targeting at the minimization of the number of FSs that must be provisioned per fiber link in order to serve the entire set of demands offered to the network. In contrast to this, we also depart, as in Fixed-VONA, from a capacitated optical network substrate, with a number of FSs per link instead of wavelengths, where we aim at maximizing the number of VONs that can be allocated. Hence, the ILP model for the flexible grid problem, called Flex-VONA, becomes a modification of the formulation presented in [15]. The definitions for the paths sets presented in Fixed-VONA are the same here. Besides, we define $F = \{f_1, f_2, ..., f_{|F|}\}$ as the ordered set of available FSs per physical link and F_d as the number of FSs per virtual link desired by demand $d \in D$. The problem variables of Flex-VONA are:

$x(d, e', p, f) = \{1$ if FS f in path p is selected to be the lowest indexed slot assigned to virtual link e' in demand d, 0 otherwise$\}$

$y(d, e', p, f) = \{1$ if FS f in path p is assigned to virtual link e' in demand d, 0 otherwise$\}$

$z(d) = \{1$ if demand d can be satisfied, 0 otherwise$\}$

The ILP formulation is stated below:

$$\max \sum_{d \in D} \alpha_d z(d), \text{s.t.} \tag{6}$$

$$\sum_{p \in P_{\{e',d\}}} \sum_{f \in F} x(d, e', p, f) \leq 1, \forall d \in D, e' \in E'_d \tag{7}$$

$$x(d, e', p, f_i) \leq y(d, e', p, f_j), \quad \forall d \in D, e' \in E'_d, p \in P_{\{e',d\}}, f_i, f_j \in F, i = 1, ..., |F| - F_d + 1, j = i, ..., i + F_d - 1 \tag{8}$$

$$x(d, e', p, f_i) = 0, \quad \forall d \in D, e' \in E'_d, p \in P_{\{e',d\}}, f_i \in F, i = |F| - F_d + 2, ..., |F| \tag{9}$$

$$\sum_{d \in D} \sum_{e' \in E'_d} \sum_{p \in P_{\{e',e\}}} y(d, e', p, f) \leq 1, \forall e \in E, f \in F \tag{10}$$

$$\sum_{p \in P_{\{e',d\}}} \sum_{f \in F} y(d, e', p, f) \leq F_d, \forall d \in D, e' \in E'_d \tag{11}$$

$$z(d) \leq \frac{1}{|E'_d|} \sum_{e' \in E'_d} \sum_{p \in P_{\{e',d\}}} \sum_{f \in F} x(d, e', p, f), \forall d \in D \tag{12}$$

Objective function (6) seeks to maximize the number of VONs to be allocated in the underlying optical network. Factors α_d have the same role as before. Constraints (7) serve the purpose of selecting for every virtual link $e' \in E'_d$ a unique path from the candidate paths set and a FS to be the lowest indexed slot assigned to the virtual link. Constraints (8) are the contiguous FS assignment constraints. If slot f_i is selected as the lowest indexed slot for virtual link e', the consecutive $F_d - 1$ slots should be assigned to this virtual link. Constraints (9) ensure that any FS selection option will have enough space in the frequency spectrum, if chosen. Constraints (10) are the spectrum clashing constraints, avoiding that two virtual links are supported over the same FS in the same physical link. Constraints (11) ensure that at most F_d different FSs are assigned to every virtual link of demand d. Constraints (12) discriminate whether demand d is satisfied or not.

V. Results and Discussion

To analyse the impact of the underlying physical technology used by the substrate network on the number of VONs that can be allocated, we have executed series of experiments using both Fixed-VONA and Flex-VONA formulations. The experiments have been executed on the 16-Node EON core network topology [16], assuming that every physical link has a usable bandwidth of 400 GHz. In the fixed-size grid scenario, following a grid with a 50 GHz channel spacing, it results in 8 wavelengths per link; in the flexible grid scenario, considering a FS width of 6.25 GHz, it results in 64 FSs per link. In particular, $|D|$ sizes from 5 to 25, in steps of 5, has been considered, assuming for both models that all factors $\alpha_d = 1$, that is, regardless of its size or the spectral resources demanded, all VONs are treated equally. Moreover, we have fixed the number of candidate paths per virtual link to the first 6 shortest paths using the distance in hops as the metric, so as to avoid excessive execution times for the models. Although the presented results may not match the optimal ones in some occasions, the presented formulations are still valid and the absolute optimal could be obtained if the whole set of candidate paths per virtual link is considered.

The generation of the demand sets for all experiments throughout this section follows a 3-step process. Firstly, 3 or 4 physical network nodes (with equiprobability) are randomly selected as virtual nodes for each demand. In this way, we obtain reasonable medium-sized virtual networks compared to the underlying physical network size. Next, the selected virtual nodes are then randomly connected using the Erdős-Rényi algorithm [17], here slightly modified to prevent the generation of non connected graphs (any connected connectivity matrix is generated with equiprobability). Finally, the bandwidth requested by the demand in GHz is selected from the set {25, 50, 100} with probabilities of 0.4, 0.4 and 0.2, respectively. In the fixed-size grid scenario such bandwidth requests are translated to 1, 1 and 2 wavelengths. In the flexible grid scenario they are equivalent to 4, 8 and 16 FSs. We assume that any guard band needed by the physical equipment to perform correctly the switching between demands is included in the bandwidth requested by them.

From the assumptions above, one can conclude that the complexity of both models is closely related to the size of the set of candidate paths for the virtual links. In more detail, in the Fixed-VONA model, the number of decision variables is in the order of $\mathcal{O}(|D||\bar{E}'_d||\bar{P}_{\{e',d\}}||W|)$ and the number of constraints is in the order of $\mathcal{O}(|E||W| + |D||\bar{E}'_d||\bar{P}_{\{e',d\}}|)$, being $|\bar{E}'_d|$ the average number of virtual links per demand, and $|\bar{P}_{\{e',d\}}|$ the average number of candidate paths per virtual link. We can see that the main contributions to the complexity of Fixed-VONA are the wavelength clashing constraints (2) and the unique path constraints (4). In the Flex-VONA model, the number of decision variables is in the order of $\mathcal{O}(2|D||\bar{E}'_d||\bar{P}_{\{e',d\}}||F|)$ and the number of constraints is in the order of $\mathcal{O}(|D||\bar{E}'_d||\bar{P}_{\{e',d\}}||F||\bar{F}_d|)$, being $|\bar{F}_d|$ the average number of FSs requested by the demands. Therefore, the main

Tab. I
COMPLEXITY OF THE MODELS

	Fixed-VONA			Flex-VONA		
	Variables	Constraints	Time (s.)	Variables	Constraints	Time (s.)
$\|D\| = 5$	840	289	0.272	13440	53760	2.37
$\|D\| = 15$	2520	499	3.28	40320	161280	1.25×10^4
$\|D\| = 25$	4200	709	18.03	67200	268800	7.29×10^4

contributions to the complexity of Flex-VONA is the large number of FSs in some network scenarios, resulting in a huge number of contiguity constraints (8).

To exemplify the complexity of both formulations, table I displays the value of these expressions for the scenario considered in the executions plus the average execution time of both models. The results shown throughout this section have been averaged over 100 executions, randomly generating a new set of demands at the beginning of each execution. The experiments have been launched on Intel Core2 Quad at 2.66 GHz PCs with 4 GB RAM memory. The optimization software used for all executions is IBM ILOG CPLEX Optimizer v.12.2 [18].

Table I shows that the complexity of Flex-VONA, in terms of number of variables and constraints, is substantially greater than the one of Fixed-VONA. This comes from the fact that the number of FSs in the flexible grid scenario is substantially larger than the number of wavelengths in the fixed-size grid scenario. Besides this, the spectrum contiguity constraints add a considerable complexity to the problem. Focusing on the execution times of Flex-VONA, they notoriously increase with the size of D. Hence, an heuristic for the flexible grid scenario might be necessary when dealing with large scenarios as the the execution of the exact ILP formulation becomes impractical. The study and development of an heuristic for the flexible grid scenario is out of the scope of this paper and left for future work.

As for the results obtained through the executions of both models, Fig. 2 shows the Blocking Probability (BP) of the demands as a function of the size of the demand set. We can observe that the flexible grid scenario provides a lower BP figure (i.e. allocates a larger number of VONs) than the fixed-size grid scenario, being the difference more notorious as $|D|$ increases(e.g., we can observe a difference in BP of around 1.1% for $|D| = 10$ while a difference of around 8% is observed for $|D| = 20$).

This capacity of being able of allocating a larger number of VONs in the flexible grid scenario is due the fact that its granularity is more finner than the fixed-grid scenario's granularity, making it possible to adjust more to the bandwidth needs of the demands. Focusing on the traffic profile considered for the experiments, for the demands requesting 25 GHz, the flexible grid scenario allocates exactly 25 GHz of spectrum to these demands, while in the fixed-size grid scenario, due its coarser granularity, it allocates 50 GHz of spectrum to these demands, adding an overhead of 100%.

The capacity of the flexible grid scenario to allocate more

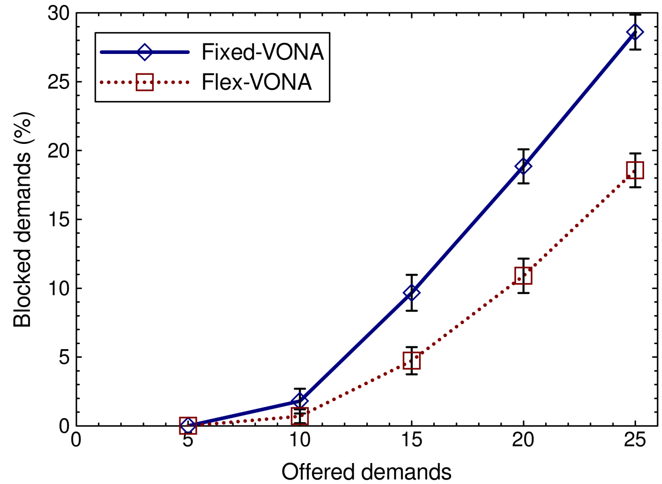

Fig. 2. Blocking probability as a function of the size of the demand set.

VONs than the fixed-size grid could be potentially greater for traffic profiles where the disparity between the requested bandwidth by the demands and the channel spacing in the fixed-size grid scenario is big, either for sub-wavelength or super-wavelength traffic demands.

VI. CONCLUSION

Optical network virtualisation is seen as a solution for the new emerging applications. GEYSERS project proposes a new architecture stack in order to qualify optical infrastructure providers and network operators with tools that enhance their traditional business operations. In detail, we have presented how the LICL layer is the corresponding element of the GEYSERS stack which is responsible to compose virtual infrastructures. The VONA problem has been introduced as a mean to describe formally the problem of how to map a set of VONs into a given optical network substrate with finite resources, while accounting for the particularities imposed by the optical medium. Depending on the technology used by the physical substrate, more or less VONs can be accommodated into the transport network.

To analyse the impact of this issue, we have considered two transport technologies as cases of study, wavelength switching and spectrum switching respectively. Exact ILP formulations for both scenarios have been proposed in aims to optimally allocate a set of VONs into a substrate using these technologies. We have validated the proposed formulations through a series of experiments, reaching the conclusion that for demands that require a fine granularity in terms of bandwidth, the spectrum switching alternative could be the more appropriate technology.

As for future work, the study and developing of an heuristic for the flexible grid scenario could be done, as the execution times for the exact ILP formulation become impractical. Moreover, we could extend the presented models to also incorporate IT nodes and the related information to the network models in aims to provide a joint virtualisation model for both network + IT resources, analysing how this new scenario impacts in the number of VONs that can be optimally allocated.

VII. Acknowledgements

The work in this paper has been supported by the European Commission through the FP7 funded GESYERS project (grant no. 248657), 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) and the CYCIT ref. TIN2010-20136-C03-03.

References

[1] S. Figuerola and M. Lemay, "Infrastructure services for optical networks [invited]," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 1, no. 2, pp. A247 –A257, july 2009.

[2] P. Vicat-Blanc, S. Figuerola, X. Chen, G. Landi, E. Escalona, C. Develder, A. Tzanakaki, Y. Demchenko, J. A. G. Espín, J. Ferrer, E. López, S. Soudan, J. Buysse, A. Jukan, N. Ciulli, M. Brogle, L. van Laarhoven, B. Belter, F. Anhalt, R. Nejabati, D. Simeonidou, C. Ngo, C. de Laat, M. Biancani, M. Roth, P. Donadio, J. Jiménez, M. Antoniak-Lewandowska, and A. Gumaste, "The future internet," in *The future internet*, J. Domingue, A. Galis, A. Gavras, T. Zahariadis, and D. Lambert, Eds. Berlin, Heidelberg: Springer-Verlag, 2011, ch. Bringing optical networks to the cloud: an architecture for a sustainable future internet, pp. 307–320.

[3] Z. Jian, L. Xiaoyong, and G. Haibing, "The optimization of xen network virtualization," in *Computer Science and Software Engineering, 2008 International Conference on*, vol. 3, dec. 2008, pp. 431 –436.

[4] N. Chowdhury and R. Boutaba, "Network virtualization: state of the art and research challenges," *Communications Magazine, IEEE*, vol. 47, no. 7, pp. 20 –26, july 2009.

[5] N. Chowdhury, M. Rahman, and R. Boutaba, "Virtual network embedding with coordinated node and link mapping," in *INFOCOM 2009, IEEE*, april 2009, pp. 783 –791.

[6] Z. Cai, F. Liu, N. Xiao, Q. Liu, and Z. Wang, "Virtual network embedding for evolving networks," in *GLOBECOM 2010, 2010 IEEE Global Telecommunications Conference*, dec. 2010, pp. 1 –5.

[7] R. Nejabati, E. Escalona, S. Peng, and D. Simeonidou, "Optical network virtualization," in *Optical Network Design and Modeling (ONDM), 2011 15th International Conference on*, feb. 2011, pp. 1 –5.

[8] S. Figuerola, J. García-Espí and, J. Ferrer Riera, and N. Ciulli, "An optical network and it infrastructure virtualisation and provisioning framework," in *Optical Internet (COIN), 2010 9th International Conference on*, july 2010, pp. 1 –3.

[9] "Spectral grids for WDM applications: DWDM frequency grid," ITU-T Recomendation G.694.1, ITU-T, June 2002.

[10] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies," *Communications Magazine, IEEE*, vol. 47, no. 11, pp. 66 –73, november 2009.

[11] J. Armstrong, "OFDM for optical communications," *Lightwave Technology, Journal of*, vol. 27, no. 3, pp. 189 –204, feb.1, 2009.

[12] N. Amaya, I. Muhammad, G. Zervas, R. Nejabati, D. Simeodinou, Y. Zhou, and A. Lord, "Experimental demonstration of a gridless multi-granular optical network supporting flexible spectrum switching," in *Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011 and the National Fiber Optic Engineers Conference*, march 2011, pp. 1 –3.

[13] N. Amaya, M. Irfan, G. Zervas, K. Banias, M. Garrich, I. Henning, D. Simeonidou, Y. R. Zhou, A. Lord, K. Smith, V. J. F. Rancano, S. Liu, P. Petropoulos, and D. J. Richardson, "Gridless optical networking field trial: Flexible spectrum switching, defragmentation and transport of 10g/40g/100g/555g over 620-km field fiber," in *37th European Conference and Exhibition on Optical Communication (ECOC)*, September 2011, pp. 1 –3.

[14] Y. Wang, X. Cao, and Y. Pan, "A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks," in *INFOCOM, 2011 Proceedings IEEE*, april 2011, pp. 1503 –1511.

[15] M. Klinkowski and K. Walkowiak, "Routing and spectrum assignment in spectrum sliced elastic optical path network," *Communications Letters, IEEE*, vol. 15, no. 8, pp. 884 –886, august 2011.

[16] S. D. Maesschalck, D. Colle, I. Lievens, M. Pickavet, P. Demeester, C. Mauz, M. Jaeger, R. Inkret, B. Mikac, and J. Derkacz, "Pan-european optical transport networks: An availability-based comparison," *Photonic Network Communications*, vol. 5, pp. 203–225, 2003, 10.1023/A:1023088418684.

[17] D. Cordeiro, G. Mounié, S. Perarnau, D. Trystram, J.-M. Vincent, and F. Wagner, "Random graph generation for scheduling simulations," in *Proceedings of the 3rd International ICST Conference on Simulation Tools and Techniques*, ser. SIMUTools '10. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2010, pp. 60:1–60:10.

[18] "IBM ILOG CPLEX: high-performance mathematical engine," 2010. [Online]. Available: http://www.ibm.com/