Converged IT and Optical Network Virtualisation: the GEYSERS Virtual Infrastructure Service Provisioning Framework

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

Jens Buysse, Marc De Leenheer, Chris Develder
Ghent University - IBBT
Department of Information Technology (INTEC) - IBCN
{jens.buysse,marc.deleenheer,chris.develder}@intec.ugent.be

Shuping Peng
University of Essex
High Performance Networks Group
pengs@essex.ac.uk

Abstract—
The last evolutions of the Internet bring the fact that all the emergence of novel applications, requirements, services, roles, and the challenges associated to them is being built on top of the same Internet that was designed for handling completely different elements. Converged, both Information Technology (IT) and optical network, infrastructure resource virtualisation is currently one of the most promising approaches to face the Future Internet challenges. The major research problem associated to infrastructure resource virtualisation is the virtual resource embedding problem. This article presents the grouped Virtual Infrastructure (VI) mapping approach, contextualised within the Generalised Architecture for Dynamic Infrastructure Services (GEYSERS) virtual infrastructure service provisioning framework. Our findings show that batched VI mapping strategy enhances the amount of virtual entities to be allocated on top of the physical substrate. The technological solution presented and the simulation of potential benefits show a novel Information and Communication Technology (ICT) infrastructure control and management solution that is able to accomodate the optimisation requirements for the Future Internet, such as cost, energy, availability, or flexibility, in coordination with application deployments and cloud service models.

Index Terms—**Optical Network Virtualisation, Resource Abstraction, Virtual Infrastrucutre, Infrastructure Service**

I. INTRODUCTION

The current Internet has become a ubiquitous commodity to provide communication services to both enterprises and residential users [1]. Cloud computing has emerged as a key paradigm in order to provide computing services addressing users' requirements over the Internet. Cloud computing stands for transparent, on-demand access to IT hardware or software resources, which are geographically spread and interconnected by networks. In fact, analyses predict that in 2020 more than 80% of the infrastructure will be outsourced within the Cloud [2]. While there are countless definitions for the Cloud computing term, there seem to be common characteristics that a cloud infrastructure should have: (i) pay-per-use (no on-going commitment, utility prices); (ii) elastic capacity and the illusion of infinite resources; (iii) self-service interface; and (iv) resources that are abstracted or virtualised [3]. Others also argue that broad network access is the fifth essential characteristic, but this requires problem analysis and robust integration design of IT and networks resources and their associated service flows, and this is typically not addressed. None of the common points (i to iv above) mention the network, its availability or even the network resources composing it. This new IT provisioning paradigm considers that the network is always available and provisioned, which clearly is not necessarily true. Applications or services running in the cloud may be affected by network performance, throughput, delay, or any other Quality of Service (QoS) parameter: current applications, such as 3D-video streaming, hold high requirements in terms of network performance. Furthermore, bandwidth-provisioning systems typically do not take into account specific characteristics of the IT resources and services connected at the edge of the network. In other words, architectures for coordinated IT and network resource provisioning have been barely investigated [4].

This divergence has been present in the research community for several years, where provisioning optimisation has been scarcely addressed considering both realms at the same time. Coordinated IT and network infrastructure service provisioning is one of the main challenges to be faced. In order to dynamically provision IT resources and gain full benefit of these thanks to Cloud technologies, it is crucial to have control over the quality of the network connections used. In parallel to the emergence of this novel coordinated provisioning, ideas and concepts behind virtualisation have matured enough after being on the research arena for a while. The fact of sharing a common good in order to improve its efficiency, usage, and productivity has been a key goal in the research community, especially given the fact that the good, in our case, is an expensive and operationally costly ICT infrastructure. Therefore, since both IT and network realms have been totally independent, there have not yet been many approaches in the community considering combined resource virtualisation, using resources from both

domains.

The GEYSERS project aims at addressing coordinated IT and network resource virtualisation in order to address some of the Future Internet challenges [5]. The project aims to define a novel architecture capable of seamless and coordinated provisioning optical network and IT resources for end-to-end service delivery. However, there is an essential different between IT and optical network virtualisation: the first is totally deployed over commercial and production environments; while the lattest case, it is only deployed in testing environments or research projects' test-beds. Such a difference provides an ideal environment for the anlaysis and study of virtual optical networks within the GEYSERS converged virtual infrastructre service provisioning.

Optical network virtualisation, coordinated with IT provisioning, is seen as solution for the operators in order to provide Future Internet services, which are characterized by global delivery of hihg-performance applications over a high-capacity dynamic optical network. With this new applications, service providers consider a key challenge for them the capability to deploy dynamic optical infrastructures at high data rates that are capable of supporting a wide range of application types. Each one of those applications holding their own access and network resource usage pattern [6]. Additionally, with this virtual optical networks need to be provided on the same physical substrate, selecting which virtual resource to be mapped to which physical resource becomes a major issue. Therefore, in this paper we present the potential benefits in terms of resource requirements that can arise from grouping several VI requests and provision them jointly.

The rest of the article is structured as follows. In Section II we present the related environment of the work we present; that is the GEYSERS project and the status of optical network virtualisation in the research community. Next section contains basic optical network definitions and features, and optical network composition methodologies. Section IV presents the proposed approach for solving the virtual infrastructure embedding problem, based on grouping the requests and provisioning them jointly. Finally, we present the conclusions of the work presented.

II. RELATED WORK

This section presents the environment within which the work for the paper has been performed. It is based on two major cornerstones; on the one hand, the GEYSERS project, which provides the environment to develop the work; and on the other hand, the optical network virtualisation, that represents the main research topic associated to the paper.

A. The GEYSERS project

The GEYSERS architecture presents an innovative approach by adopting the concepts behind the Infrastructure as a Service (IaaS) servicing model from cloud computing and service-oriented networking to enable infrastructure operators offering new IT and network converged services. In the GEYSERS layered architecture physical devices populating the bottom layer, i.e. the physical infrastructure layer, are abstracted and partitioned or grouped into virtual resources that can be selected to form the virtual infrastructures. This process takes place in the LICL, the key element of the GEYSERS stack in order to provide converged infrastructure services. On top of the virtual infrastructures, there is the Service Middleware Layer (SML) and the Network Control Plane+ (NCP+), responsible for configuring and managing virtual resources. Furthermore, the SML is responsible for translating the application requests and Service Level Agreements (SLAs) into technology specific requests in order to trigger the provisioning procedures at the NCP+ level [7].
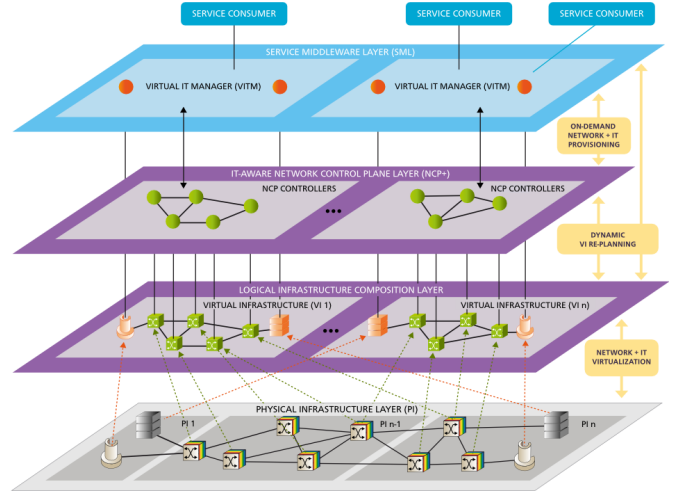

Fig. 1.   GEYSERS architecture

The SML is a convergence layer for coordinating the management of IT resources that belong to an aggregate service. The SML contains the Virtual IT Manager, which is the element responsible of the end-to-end IT service management and the virtual IT resource configuration. The GEYSERS NCP+ performs all control and management functions necessary to operate the virtual network resources within the virtual infrastructure. The NCP+ also offers a set of functionalities towards the SML, in support of on-demand and coupled provisioning of the IT resources and associated network connectivity. The Logical Infrastructure Composition Layer (LICL) is a key component in the GEYSERS architecture and represents one of its architectural innovations. It is responsible for the creation and maintenance of virtual resources as well as virtual infrastructures composed of those virtual resources. Such a layer acts as a middleware on top of the physical infrastructure and offers a complete toolset to the involved entities in the infrastructure service provisioning workflow. More information of each layer presented can be found in [4], [8]. The LICL provides the infrastructure services within GEYSERS. This layer is the component in the GEYSERS architecture responsible for abstracting and virtualising the physical resources, and thus offering them as a service to the upper layers on the GEYSERS stack. Although having such other components and elements, the main pilars on
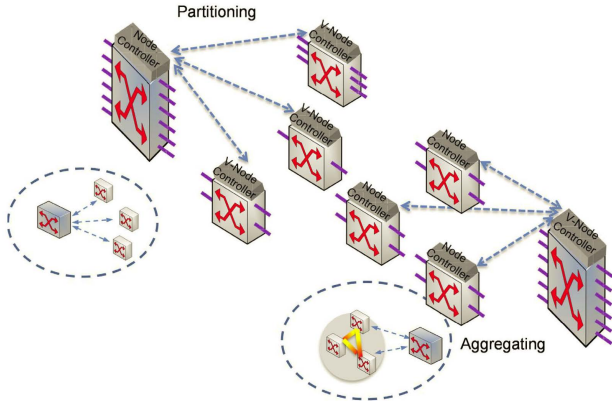
Fig. 2. Virtualisation of optical nodes

which the functionalities of the LICL rely are: (i) resource abstraction, (ii) the information modelling framework, and (iii) the synchronisation mechanisms.

B. Optical Network Virtualisation

The GEYSERS architecture enabled by optical network virtualization allows service providers to lease resources on-demand from infrastructure providers. Optical Network Virtualisation is defined as the composition of multiple isolated optical virtual networks (Optical VN) simultaneously coexisting over shared optical physical networks (Optical PN). The Optical PN can be provided by multiple Physical Infrastructure Providers (PIPs). Each of the optical VN is composed of a set of virtual optical nodes interconnected with virtual optical links, and managed by a single administrative entity. Virtual optical resources, including virtual optical nodes and virtual optical links, are often achieved by partitioning or aggregating physical optical resources, as shown in Fig. 2. The partitioning and/or aggregation of these resources empowers the creation of multiple simultaneous VIs, each with its own topology and QoS requirements, running over the same optical network infrastructure. Within an optical VN, a virtual link is defined as a connection between one port of a virtual network element to a port of another virtual network element. The granularity of virtual links in the virtualized optical network is inherited from the switching capabilities supported by physical devices (e.g., wavelength, or fiber). Generally, the virtualisation of optical switches should be considered together with that of optical links.

III. VIRTUAL OPTICAL NETWORK PROVISIONING

Requests for composing VIs are usually generated by service providers or operators. Each request has associated requirements that need to be fulfilled when composing the VI. Moreover, isolation and coexistence, the two most important characteristics of virtualized optical networks, need to be satisfied. Isolation implies that different VIs sharing the same underlaying substrate should be independent from each other. Virtual optical resources that are created by partitioning or aggregation from same physical optical network resources

should not interfere with each other. Coexistence means that different optical VIs sharing the same physical infrastructure can be supported and provisioned in parallel to different administrative entities. However, due to the analogue nature of optical networks, the optical layer constraints, such as wavelength continuity constraint and various physical layer impairments need to be considered when virtualizing optical networks. The wavelength continuity constraint will affect the network resource utilization, while the physical layer impairments impact the isolation of multiple coexisting VIs.

Given the information of physical infrastructures and the requirements of virtual network requests, an intelligent and dynamic composition mechanism is needed to create VIs on demand, utilizing the available physical resources. In [9], the authors have proposed a VI composition method considering the impact of the physical layer impairments to guarantee the isolation between multiple coexisting VIs. The studies in [10] show that the energy consumption of concurrent VIs is minimized over a shared IT and network infrastructure. A. Pages et al. studied in [11] the impact of the transport technology, that is, wavelength switching and spectrum switching, on the amount and characteristics of the virtual infrastructures that can be built on top of a physical infrastructure.

In this paper, we focus on the grouped or clustered VI mapping in a Dense Wavelength Division Multiplexing (DWDM) network scenario. Several VI requests are grouped in a cluster wherein the bandwidth can be shared, whereas the isolation between different clusters is still guaranteed. The potential benefits in terms of the resource requirements of the joint VI provisioning is investigated.

IV. VIRTUAL INFRASTRUCTURE EMBEDDING

A virtualized Wavelength Division Multiplexing (WDM) network consists of several lightpaths (i.e. end-to-end wavelength connections) where each lightpath is completely owned and managed by one Virtual Infrastructure Provider (VIP). The aggregate of all lightpaths managed by the same VIP thus forms a (virtual) topology which is called a VI. Having a separate VI per VIP achieves complete isolation of the optical network resources (i.e. wavelengths) used by distinct VIs, it also implies a possible penalty in required network capacity to accommodate them all: given the coarse bandwidth granularity in current commercial DWDM products (each wavelength offers 10, 40 or 100 Gbps), the offered network capacity may be very high while only a small portion of this capacity is used. Therefore, in [12] we have proposed to cluster VI requests and introduce traffic grooming. As such, we do not offer de facto isolation within each cluster, although full isolation is enforced between different clusters.

Apart from the wavelength utilisation vs isolation trade-off, there is also a trade-off between resource utilization and the scalability of the Network Control Plane (NCP). Although a small number of VIs maximizes statistical multiplexing opportunities and hence increases resource utilization, it also leads to fairly large VI instances that, because of their size,
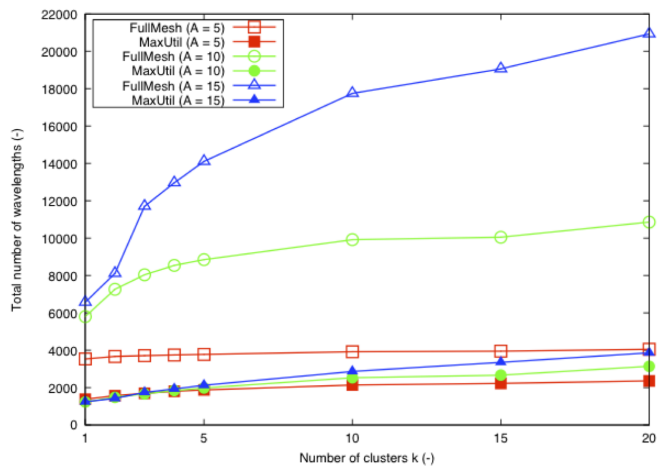
Fig. 3. The total number of wavelengths required to map a given set of VI requests depends of the number of clusters we partition them in, as well as the clustering strategy (FullMesh vs MaxUtil).

Fig. 4. The total number of wavelengths compared to that obtained by a random clustering.

may suffer from degraded control plane scalability: the number of control plane messages are proportional with the number of nodes in a network. A large number of isolated VIs however, would lead to poor resource utilization. This trade-off will be illustrated quantitatively in the following.

More formally, we solve the following VI mapping and clustering problem. Given: (i) the Physical Infrastructure (PI) topology, (ii) a set of traffic matrices representing the virtual infrastructure requests, and (iii) the number of isolated VIs that should be mapped on the PI. (Each isolated VI is composed of one or more virtual network requests) We should find: (i) the composition of the isolated VIs, i.e. which requests jointly form an isolated virtual etwork, and (ii) the mapping of the isolated virtual networks on the physical topology.

We aim to find solutions using a two-step algorithm: (i) we first perform clustering of individual VI requests into distinct virtual networks, after which (ii) we map these virtual network onto the PI, i.e. we select which the wavelength paths to use for each of the virtual network links.

For an optimal clustering solution, we turn to an ILP formulation as described in [12] and compare these results for benchmarking purposes with a random approach. We consider two alternatives for the actual mapping onto the PI: (i) a FullMesh strategy which minimizes the hop distance between virtual network nodes, and (ii) a MaxUtil strategy which maximally exploits statistical multiplexing to fill the available link capacity as efficiently as possible.

In Fig.3 we show the total number of wavelengths necessary to instantiate a varying number of virtual network clusters (here denoted as k), using the ILP-based clustering algorithm. When comparing the increase in wavelength usage between MaxUtil and FullMesh, we note that MaxUtil has a very slow growth, in particular compared to a FullMesh virtual network design. Fig.4 indicates the effectiveness of the ILP-based cluster, where we show the ratio of the total number of wavelengths for the ILP-based algorithm over that from random clustering. First,
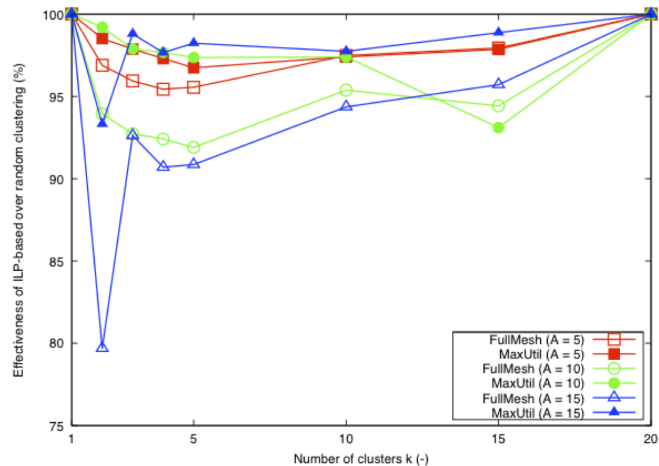
note that there is a gap of about 5 to 10%, showing that ILP performs better than Random. However, this relatively low improvement of ILP-based over random clustering indicates that there is a need for more advanced clustering (indeed we only incorporated node activity of VI requests, ignoring network grooming). Secondly, the effectiveness of the random clustering reaches a minimum around 2-3 clusters, indicating the region where intelligent clustering is most relevant.
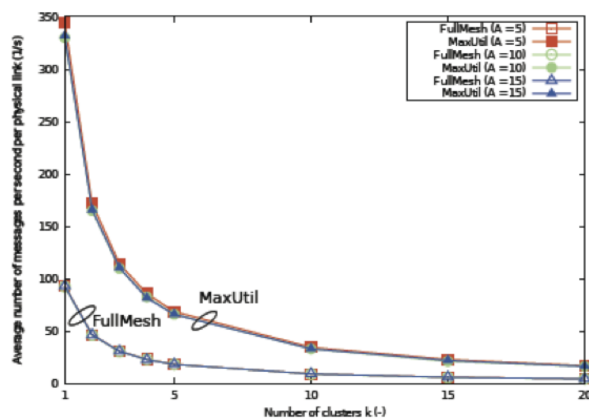

Fig. 5. Convergence of average message exchange rate for large number of virtual network clusters

Within each provisioned VI, connection requests are assumed to be issued at certain rates associated with the traffic matrix used for the VI mapping. In Fig.5 the control plane message exchange rate associated with those dynamic connection requests is averaged over all virtual networks. We only consider connection signalling traffic, as this forms the majority of control plane traffic, especially when introducing flooding reduction techniques for OSPF. Fig. 5 shows that both design techniques converge to approximately the same average message exchange rate, although the *MaxUtil* approaches a very high control plane load for a small number of virtual

networks. The size of the virtual network request does not influence the message exchange rates. Note than the average message exchange rate is a hyperbolic function and thus the total message exchange rate remains constant. However, the reduction in cotrol traffic within each cluster indicates that virtualisation offers operators the compelling advantage of control plane scalability, since the associated controllers can run independently from each other.

V. Conclusion

Cloud computing in essence has emerged thanks to the increased availability of network connectivity and bandwidth. However, despite the crucial role that networks play in making cloud services possible, network resource provisioning to date is not an integral part of the cloud service provisioning process. To alleviate this, and thus ensure that network performance is satisfactory in order to meet the specific characteristics of the cloud-bases applications, we present the virtual optical network provisioning within the GEYSERS project. We present the GEYSERS project, which proposes a holistic architecture, handling both IT and network resources in a converged manner, while exploiting virtualisation of both of them in order to maximize their efficient utilisation in an infrastructure as a service model.

In this paper, we have focused on the grouped or clustered VI mapping in a DWDM network scenario. Several VI requests are grouped in a cluster wherein the bandwidth can be shared. We have presented illustrative results of such an approach. In particular, we showed that intelligently clustering the virtual infrastructure requests can attain non-negligible advantages in network capacity needed, order of 10%. Although promising, results obtained open a new door in order to enhance the embedding problem by means of clustering the requests and obtaining a trade-off complete isolation and network capacity required.

New clustering methodologies are left for future work. Initial work on the request clustering based on the IT server clustering in high-performance computing is being analysed. Also the impact of clustering VI requests on the control planes must be better characterised. Initial results show advantages to the operators in terms of control plane scalability, since the associated controllers to each cluster can run independently from each other. However, detailed characterisation of the control plane entities behaviour and its analysis is still missing in order to complete the analysis of the batched approach.

References

[1] P. Vicat-Blanc, S. Figuerola, X. Chen, G. Landi, E. Escalona, C. Develder, A. Tzanakaki, Y. Demchenko, J. Garcia Espin, J. Ferrer Riera *et al.*, "Bringing optical networks to the cloud: An architecture for a sustainable future internet," in *The Future Internet*, ser. Lecture Notes in Computer Science, J. Domingue, A. Galis *et al.*, Eds. Springer Berlin / Heidelberg, 2011, vol. 6656, pp. 307–320, 10.1007/978-3-642-20898-0-22.

[2] M. Nelson, "The next generation internet, e-business, and e-everything," http://www.aaas.org/spp/rd/ch20.pdf.

[3] R. Buyya, J. Broberg, and A. Goscinski, *Cloud computing: principles and paradigms*. Hoboken, New Jersey: Wiley, 2011.

[4] E. Escalona, R. Nejabati, J. Jimenez, A. Tovar, G. Landi, N. Ciulli, J. Garcia-Espin, J. Ferrer Riera, T. Ginkel, P. Robinson, M. Brogle, X. Chen, S. Greco, B. Belter, A. Tzanakaki, P. Donadio, Y. Demchemko, S. Soudan, and C. Develder, "Initial geysers architecture and interfaces specification," GEYSERS EC FP7-ICT-2009.1.1 No. 248657, Tech. Rep., September 2010.

[5] S. Figuerola, J. A. Garcia-Espin, J. F. 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.

[6] 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.

[7] A. Tzanakaki, M. Anastasopoulos, K. Georgakilas, J. Buysse *et al.*, "Energy efficiency in integrated it and optical network infrastructures: the geysers approach," in *Computer communications workshops (INFOCOM Workshops), IEEE Conference on*, 2011, pp. 343–348.

[8] E. Escalona, R. Nejabati *et al.*, "Geyses overall architecture, interfaces specification, and service provisioning workflow," GEYSERS EC FP7-ICT-2009.1.1 No. 248657, Tech. Rep., September 2010.

[9] S. Peng, R. Nejabati, S. Azodolmolky, E. Escalona, and D. Simeonidou, "An impairment-aware virtual optical network composition mechanism for future internet," in *Optical Communication (ECOC), 2011 37th European Conference and Exhibition on*, sept. 2011.

[10] M. Anastasopoulos, A. Tzanakaki, K. Georgakilas, and D. Simeonidou, "Energy aware planning of multiple virtual infrastructures over converged optical network and it physical resources," in *Optical Communication (ECOC), 2011 37th European Conference and Exhibition on*, Spet. 2011.

[11] A. Pages, J. A. Garcia-Espin, J. Perello, J. F. Riera, S. Spadaro, , and S. Figuerola, "Optimal allocation of virtual optical networks for the future internet," in *Proceedings of the 16th Conference on Optical Network Design and Modeling (ONDM 2012)*, April 2012.

[12] M. De Leenheer, J. Buysse, C. Develder, and B. Mukherjee, "Isolation and resource efficiency of virtual optical networks," in *Computing, Networking and Communications (ICNC), 2012 International Conference on*, 30 2012-feb. 2 2012, pp. 840 –844.