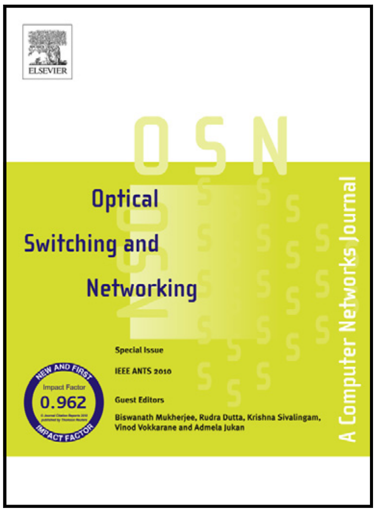

www.elsevier.com/locate/osn

Virtual Infrastructures as a Service enabling Converged Optical Networks and Data Centres

Jordi Ferrer Riera, Anna Tzanakaki, Alexandru-Florian Antonescu, Markos Anastasopoulos, Joan A. García-Espín, Eduard Escalona, Shuping Peng, Giada Landi, Giacomo Bernini, Bartosz Belter, Damian Parniewicz, Xavier Hesselbach, Sergi Figuerola, Dimitra Simeonidou

PII: S1573-4277(14)00061-7
DOI: http://dx.doi.org/10.1016/j.osn.2014.05.017
Reference: OSN313

To appear in: *Optical Switching and Networking*

Received date: 9 March 2014
Revised date: 27 April 2014
Accepted date: 13 May 2014

Cite this article as: Jordi Ferrer Riera, Anna Tzanakaki, Alexandru-Florian Antonescu, Markos Anastasopoulos, Joan A. García-Espín, Eduard Escalona, Shuping Peng, Giada Landi, Giacomo Bernini, Bartosz Belter, Damian Parniewicz, Xavier Hesselbach, Sergi Figuerola, Dimitra Simeonidou, Virtual Infrastructures as a Service enabling Converged Optical Networks and Data Centres, *Optical Switching and Networking*, http://dx.doi.org/10.1016/j.osn.2014.05.017

Virtual Infrastructures as a Service enabling Converged Optical Networks and Data Centres

Jordi Ferrer Riera[1*], Anna Tzanakaki[2], Alexandru-Florian Antonescu[3,4], Markos Anastasopoulos[2], Joan A. García-Espín[1], Eduard Escalona[1], Shuping Peng[2], Giada Landi[5], Giacomo Bernini[5], Bartosz Belter[6], Damian Parniewicz[6], Xavier Hesselbach[1,7], Sergi Figuerola[1], and Dimitra Simeonidou[2]

[1] Fundació i2CAT, C./ Gran Capità 2-4, Barcelona, Catalonia
[jordi.ferrer,joan.antoni.garcia,eduard.escalona, sergi.figuerola]@i2cat.net

[2] University of Bristol, High Performance Networks Group (HPNG), Department of Electrical and Electronic Engineering, United Kingdom.
[anna.tzanakaki,m.anastasopoulos,shuping.peng,dimitra.simeonidou]@bristol.ac.uk

[3] SAP Inc., Next Business and Technology, Althardstrasse 80, 8105 Regensdorf, Switzerland
[4] University of Bern, Communication and Distributed Systems (CDS), Neubrueckstrasse 10, 3012 Bern, Switzerland
[alexandru-florian.antonescu]@sap.com

[5] Nextworks, via Livornese 1027, Pisa, 56122, Italy
[g.landi,g.bernini]@nextworks.it

[6] Poznan Supercomptuing and Networking Centre, ul. Noskowskiego 10, Poznan, 61-704, Poland
[bartosz.belter,damianp]@man.poznan.pl

[7] Universitat Politècnica de Catalunya, Departament d'Enginyeria Telemàtica, C./ Gran Capità 2, Ed. C3, Barcelona
[xavierh@entel.upc.edu]

[*] Corresponding autor

Abstract

Cloud computing service emerged as an essential component of the Enterprise IT infrastructure. Migration towards a full range and large-scale convergence of Cloud and network services has become the current trend for addressing requirements of the Cloud environment. Our approach takes the infrastructure as a service paradigm to build converged virtual infrastructures, which allow offering tailored performance and enable multi-tenancy over a common physical infrastructure. Thanks to virtualization, new exploitation activities of the physical infrastructures may arise for both transport network and Data Centres services. This approach makes network and Data Centres' resources dedicated to Cloud Computing to converge on the same flexible and scalable level. The work presented here is based on the automation of the virtual infrastructure provisioning service. On top of the virtual infrastructures, a coordinated operation and control of the different resources is performed with the objective of automatically tailoring connectivity services to the Cloud service dynamics. Furthermore, in order to support elasticity of the Cloud services through the optical network, dynamic re-planning features have been provided to the virtual infrastructure service, which allows scaling up or down existing virtual infrastructures to optimize resource utilisation and dynamically adapt to users' demands. Thus, the dynamic re-planning of the service becomes key component for the coordination of Cloud and optical network resource in an optimal way in terms of resource utilisation. The presented work is complemented with a use case of the virtual infrastructure service being adopted in a distributed Enterprise Information System, that scales up and down as a function of the application requests.

Introduction

Cloud computing services are one of the fastest growing business opportunities for Internet service providers and telecom operators [1]. The emergence of even more resource demanding services, which hold high-performance, high-capacity, network-based applications with strict IT (e.g. computing and data repositories) resource requirements are driven by many technological advances. Distributed computing systems and large-scale computer networks supporting both communication and computation are able to run distributed high-performance applications. However, these applications require specific Cloud services that involve distributed IT resources interconnected through high-capacity, high-performance, and flexible networks, which cannot be intrinsically delivered by the current best-effort Internet [2]. In response to these needs, optical networking offers very high-capacity transport with increased dynamicity and flexibility through recent control planes, resource virtualisation and elastic mechanisms.

Optical networks enhanced with Generalized Multi-Protocol Label Switching (GMPLS) based protocols and Path Computation Element (PCE) offer the opportunity to automatically control, provision, and operate wavelength switched optical networks (WSON) connections [3]. GMPLS has proven itself to play an important role in realizing interconnections of a wide variety of resources, or even geographically distributed Data Centres (DCs). On the other hand, research and development on optical networks have matured considerably over the past decade, being deployed by telecom network operators all over the world.

Foster et al. stated in [4] that when plugging an electric appliance into an outlet, we care neither how electric power is generated nor how it gets to that outlet. This is possible because electricity is virtualized; that is, it is already available from a wall socket that hides power generation stations and a huge distribution grid. In fact, they claim, when extending this approach to information technologies, this concept means delivering useful functions while hiding how their internal works. Computing itself, to be considered fully virtualized, must allow computers to be built from distributed computing devices such as processing, data, and software resources. In this sense, Cloud computing, through fully virtualized environments, has emerged as a key paradigm providing services addressing user’s requirements over the Internet. However, one essential point that none of the countless definitions of Cloud computing addresses is the network availability or the Quality of Service (QoS) [5], which at the end impacts over the Quality of Experience (QoE), i.e. on how the end-user appreciates the Cloud services offered within the DCs. The Cloud computing paradigm typically considers the network to be always available and provisioned, which is not necessarily true, since applications or services running on a given instance of the distributed environment may be affected by network performance, throughput, or even delay. In order to dynamically provision Cloud resources located at the Data Centres and gain full benefit of these, it is crucial to have control over the quality of the network connections.

The migration towards a full range and large-scale convergence of Cloud and network services has become the current trend, which implies the extension of the virtualisation concept from only computing to a joint computing and networks consideration. In fact, resource virtualisation is envisaged as the process that will homogenize both Cloud and network resources through the provisioning of combined IT and network Virtual Infrastructures (VIs). Converged virtualisation in multi-tenant environments allows the usage optimisation of the hardware devices, and therefore it actually avoids having an infrastructure with many similar devices performing much less than 100% just because they have to be under different administrative domains [6].

Network virtualisation is recognised as an enabling technology for the future Internet. Through dynamic mapping of virtual resources onto physical hardware, the benefit from the existing hardware can be maximized. Optimal dynamic resource allocation mechanisms, leading to the self-configuration and organisation of future networks, will be necessary in order to provide customised services to the end-users. However, several challenges emerge on the arena derived from the virtualisation environments, and even derived from the nature of the optical substrate itself, with new types of constraints when compared to electrical ones.

In [17, 23] the authors provided a general overview of the different challenges still to be solved in generic network virtualisation environments (NVEs), concluding that the materialisation of an NVE needs to satisfy the requirements sets by its characteristics and design goals. They claim that there is still research to go in order to achieve an open, flexible, and heterogeneous NVE. In [21] the authors considered that the application of the virtualisation technology relies on algorithms that can instantiate virtualised networks on a substrate infrastructure, optimising the layout for service-relevant metrics. They provided a complete survey of the current research in the network virtualisation area; based upon a novel classification scheme for the different algorithms, a taxonomy and classification of the current research approaches is provided in the manuscript.

Finally, in [22] the authors presented Integer Linear Programming (ILP) formulations to optimally allocate virtual optical networks over a given transparent optical physical substrate. The article focused on the different optical aspects of the substrate, although it provided formulations that serve the purpose of building either completely transparent virtual optical networks or opaque ones, where electrical termination capabilities were assumed at each virtual node.

Thus, as a primary conclusion, both IT and network virtualisation, through the abstraction of the physical devices as totally manageable, independent logical objects allow applications to easily deploy new services on top of such virtualised infrastructures. Major motivations to apply virtualisation for bringing together the clouds and the networks are cited in [6], and can be summarized as: (i) lower infrastructure operational costs; (ii) enable new business models; (iii) federate heterogeneous infrastructures; (iv) integrate different type of hardware in applications with service-oriented architectures; (v) scale infrastructure on-demand or elasticity; (vi) and reduce environmental impact.

In this context, the Generalised Architecture for Dynamic Infrastructure Services (GEYSERS) European project proposed the interconnection of IT resources through WSON networks in a converged infrastructure that can support delivery of end-to-end services through joint provisioning of Optical Network + IT resources at the edges. The project adopted the concept of the Infrastructure as a Service (IaaS) facilitated through virtualization of the combined DC and network infrastructure in order to offer performance advantages and enable sharing of physical resources, which brings new exploitation opportunities for the underlying physical infrastructures, both transport network and Data Centres. This approach makes network and IT resources dedicated to Cloud Computing converge on the same flexible and scalable layer.

The rest of the manuscript is structured as follows. Next section contains a brief description of the overall proposed architecture for coordinated Cloud and Network virtualisation and provisioning. Then, a first approach to virtual infrastructures provisioning, operation and dynamic re-planning is provided. Re-planning results presented in Section 3 comprehend one of the most representative results of the proposed solution. The next section provides an example of a large-scale Cloud Enterprise System deployment and how it is deployed over a managed virtual infrastructure. Finally, the article closes with conclusions and future research directions.

The GEYSERS approach: Bringing together Data Centres and Optical Networks

The Generalized Architecture for Dynamic Infrastructure Services (GEYSERS) project builds on the concept of Infrastructure-as-a-Service (IaaS). The project proposes a bottom-up architecture following an approach capable of providing dynamic, cost-efficient, and mission-specific virtual infrastructures that can be operated and managed by a virtual infrastructure composition layer and controlled by means of an enhanced IT- and energy-aware Network Control Plane. Within this concept of convergence and coordination, high-end IT resources such as Data Centres are fully integrated with the network service procedures, both at the infrastructure planning and connection provisioning stages.

This concept results in a new role for carriers that own their infrastructure by enabling them to offer their optical network integrated with IT infrastructures, either owned by them or by third-party providers, as a service. Resource sharing among virtual infrastructures allows for optimal utilization of the physical infrastructure; furthermore, the layered architecture allows virtual infrastructures to adapt to the effective load, and takes into account energy-efficiency considerations. Figure 1 depicts the layered architecture built on top of the heterogeneous physical infrastructure, composed of IT resources and optical network resources. The architecture provides converged planning and coordinated provisioning of both types of resources, enabling the new emerging application to run seamlessly over converged virtual infrastructures.

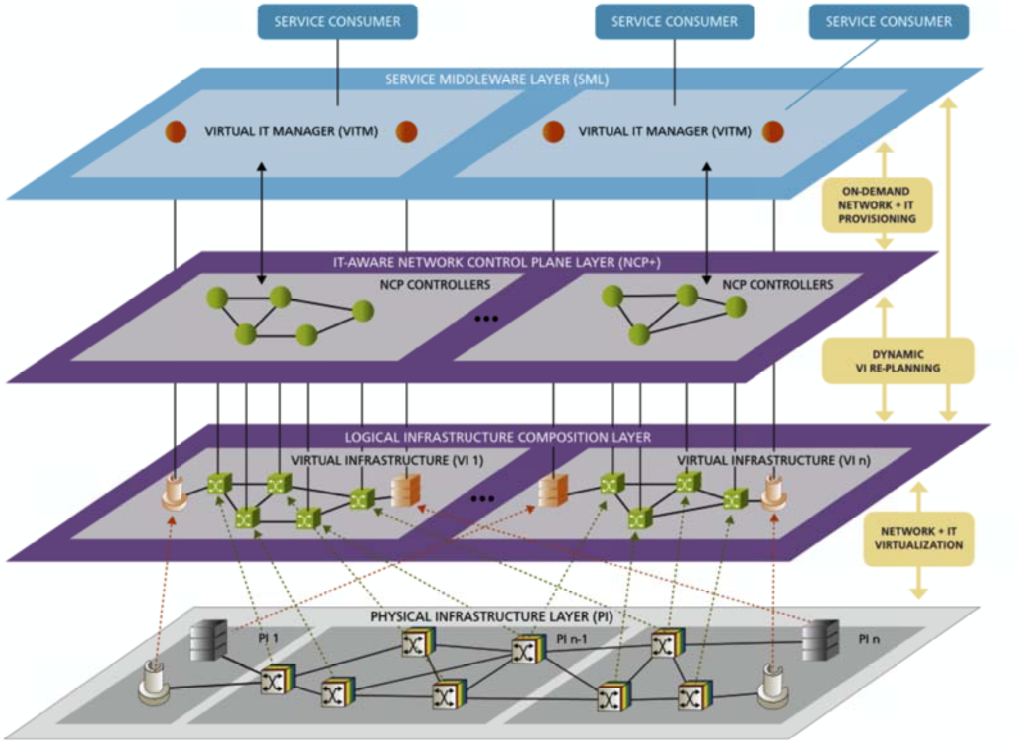

Figure 1: The proposed GEYSERS layered architecture [6]

In the layered architecture, physical devices populating the bottom layer – 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 Logical Infrastructure Composition Layer (LICL), the key element that provides converged infrastructure services. This component holds a uniform semantic resource description model for both network and IT resources, which unifies the abstraction process for the whole substrate. Therefore, at the LICL level, any resource is seen as a generic resource with a set of associated capabilities, such as switching capabilities or computing capabilities. This unified resource description model enables the virtual infrastructure provisioning processes, described in Section 3, to consider both optical networks and data centres in a converged manner. Further details on the unified resource modelling, the so-called LICL-IMF (LICL Information Modelling Framework), can be found in [8].

On top of the virtual infrastructures, there are the Service Middleware Layer (SML) and the Network Control Plane (NCP+). The former is responsible for translating the application requests and Service Level Agreements (SLAs) into connectivity specific requests in order to trigger the provisioning procedures at the NCP+ level [7]. The latter is responsible thus for configuring, and managing virtual resources. These two components are responsible for the coordination of virtualised cloud computing resources and high-capacity network services. The inter-cooperation between the NCP+ and the SML allows the optical network connectivity to be automatically tailored to cloud service dynamics, guaranteeing high performances and reliability and, on the other hand, optimizing the efficiency of the resource provisioning utilization. The communication between the SML and the NCP+, performed through an enhanced User-to-Network Interface (UNI) involves two main procedures. On the one hand, advertisement of virtualised IT resources, that comprises the injection of attached virtual IT resources information (e.g. storage or computing capabilities) into the NCP+. This advertisement enables the utilization of the cloud-related information in the advanced computation processes that take place at the control plane. On the other hand, the interface supports the procedure of requesting coupled IT and network services. Thus, the SML is responsible for requesting the setup, modification and release of enhanced connectivity services. These connection services include the traditional unicast service and advance connectivity services such as the assisted unicast service, the restricted anycast service, and the full anycast service. Detailed information on the NCP+ and SML interactions, as well the advance connectivity services, can be found at [14].

Convergence and coordination of cloud computing and network resources in the proposed architecture takes place at the virtual infrastructure planning and virtual infrastructure operation stages respectively. Moreover, this layered architecture facilitates the emergence of new business entities, i.e. Virtual Infrastructure Operators (VIOs) and Virtual Infrastructure Providers (VIPs) that implement new behaviours depending on how they interact with the infrastructure. By means of decoupling the traditional infrastructure from the service provided on top of it, the architecture facilitates the appearance of new challenges on how to dynamically and optimally match the actual infrastructure with the service provided on top of them and avoid over-provisioning of resources. The proposed architecture introduces a new dynamism where ownership and operation of infrastructures are split and assumed by different players through the provisioning of virtual infrastructures as a service, allowing new actors to enter the market in a relatively short time and be capable of generating revenues with low initial investments.

Virtual Infrastructures as a Service

One of the key components of any on-demand cloud provisioning system is the Service Delivery Framework (SDF) [18], initially proposed by the TeleManagement Forum. We extend this conceptual framework connecting it to the proposed architecture and applying it to the dynamic and automated provisioning system. Therefore, a virtual infrastructure is considered and provisioned as a service, with its own lifecycle. In detail, virtual infrastructures are de-materialized resource aggregates, which provide a service over a time-limited period. Their lifecycle, as shown in Figure 2, involves their planning and creation phases, the service delivery phase, and finally their decommissioning.

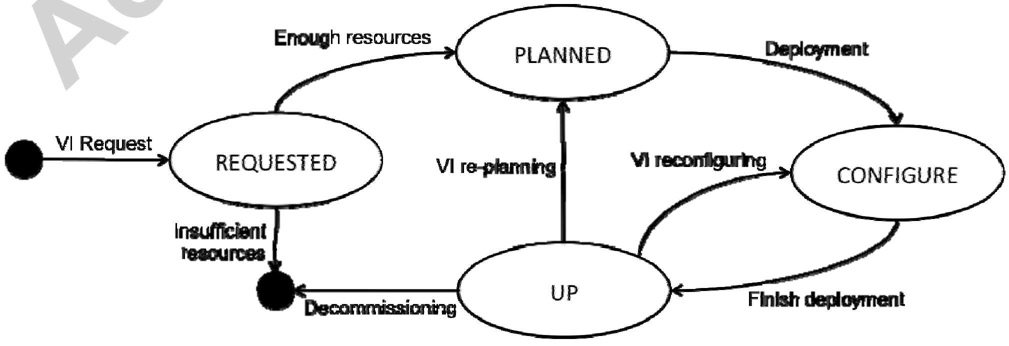

Figure 2: Lifecycle of a virtual infrastructure

In our architecture, Vis are requested from the virtual infrastructure operator, through the SML. Each request contains the description of the resources requested within the

VI, the topology describing how they should be connected, and the associated lifetime. When the operator requests a converged VI with both IT and network resources, the request is received at the internal LICL virtualization allocator through the LICL service interface. The algorithms utilized for planning the VI are explained later in this section. The state of a virtual infrastructure always depends on the different states of the virtual resources that compose it. Hence only once all resources of the virtual infrastructure could have been allocated, the VI moves to the planned state (PLANNED in Figure 2), otherwise the request is rejected. Once the VI is planned, its different virtual resources are created with the different parameters given in the VI request, and if the creation is successful, it moves to the configured state (CONFIGURE in Figure 2). The virtual infrastructure becomes up once all of its virtual resources have been instantiated on the physical resources to which they had been assigned (UP in Figure 2). It is then ready to be handed over to the users (i.e. virtual network operators), who can proceed with operating the network resources, installing applications or even create new virtual nodes. At the end of the reserved lifetime of a virtual infrastructure, all its virtual resources are decommissioned.

When providing virtual infrastructures composed of both IT and network resources as a service, the first problem in the management is the parameterization of the virtual infrastructure itself; i.e. the description of all the involved resources and their interconnections. This is performed through a semantic resource description-modelling framework, the aforementioned IMF, which converges cloud and optical network resources at the LICL level [8]. It becomes one of the basic components of the on-demand provisioning system, since it provides common description of Cloud and high-capacity network resources.

Furthermore, virtual infrastructures need to be appropriately designed and operated to address the very dynamic and unpredictable traffic profiles and service characteristics they are supposed to support. As an example, underestimating the required network and IT resources may lead to an inability to satisfy end-users requirements, whereas an overestimation may lead to over-provisioning of resources and hence increased operational and capital expenditures (OpEx and CapEx). In this context, optimal VI planning with respect to specific objectives of interest plays a key role in order to enable the IaaS paradigm into our proposal. Virtualisation, and, in detail, virtual network planning into the physical infrastructure is a well-studied topic in the literature [17, 21, 22, 23]. The VI planning algorithm through the GEYSERS architecture is realised into the LICL layer. Using as a basis the unified resource model that converges IT and network resources, several studies and analysis have been performed in order to benchmark the architecture [27]. In [9], Georgakilas et al. proposed, through detailed integer linear program modelling, one algorithm aiming at minimize the overall power consumption of the virtual infrastructure itself. In [10], Peng et al. considered that the composition method for virtual infrastructures should become aware of the different physical layer impairments that can later on affect onto the infrastructure operation. Finally, also through the GEYSERS architecture, in [11] the authors compared the performance of the planning algorithm in terms of the number of virtual infrastructures that can be provisioned depending on the switching capability of the substrate contained in the semantic description model (e.g. wavelength switching or spectrum switching).

However, those studies considered an offline approach, where the set of virtual infrastructures to be provisioned was known in advance and fixed. Given that the volume and type of service requests is not precisely known in advance, the required virtual infrastructure capacity may need to scale up and down on demand to ensure that all service requests can be supported in an efficient manner. This in practice can be performed through dynamic VI re-planning. Although the VI planning and re-planning algorithms are realised in the LICL layer, the triggering for the planning process comes from the upper layers into the presented architecture.

An issue of concern when dealing with dynamic re-planning of VIs is how to deal with existing service requests (i.e. already planned, and deployed virtual infrastructures) during the reconfiguration time of the given VI. From the pure business perspective, service disruption considerations are of capital importance, since there is no re-planning that can affect any of the currently provisioned service requests. This approach may not provide a globally optimal solution for the planning problem. However, it will ensure the infrastructure provider will meet the expected quality of service required in terms of disruptions. On the other hand, Dynamic Virtual Network Embedding approaches aim at reconfiguring the mapped virtual networks in order to recognize the resource allocation and optimise the global utilisation of the substrate resources.

This issue is analysed in [24]. The authors realize that most of the service requests rejections are caused by the bottlenecked substrate links. In order to improve the rejection ratio and the load balance in the substrate network, they propose a reactive and iterative algorithm (called virtual network reconfiguration). The algorithm just runs when a VI request is rejected. It works as follows. In first place it sorts the mapped virtual nodes by their suitability for migration, then it migrates the most suitable virtual node and its attached virtual links to another substrate node, and tries to map again the request. If the network cannot be mapped, the next iteration of the algorithm migrates the following virtual node and the process is repeated until the whole request is mapped or until a predefined number of iterations. Performance results presented a significant increase of mapped requests after the reconfiguration algorithm is applied. Dynamic re-planning is also considered in [25] by means of migrations when service access position changes, and in [26], where a heuristic uncoordinated is proposed to reduce the cost of periodic access position changes. Following the business considerations within the GEYSERS project, and the ecosystem of our proposal (service delivery framework), we consider that any virtual infrastructure (i.e. service) that is already provisioned cannot be disrupted by any dynamic re-planning procedure.

Dynamic virtual infrastructure re-planning

Information regarding the volume and type of service requests is not precisely available in advance of the requests to the VI providers. Cloud services can scale up and down on demand. Therefore, dynamic adaptation of the infrastructure to the elasticity of the Cloud services requires constant changes. The mechanisms to update a given virtual infrastructure can be either automatically or either manually triggered by various factors and events, having as main objectives: (i) to support the upcoming connectivity services requests that cannot be served by existing VIs; and (ii) to optimize the utilisation of network and IT resources. The automatically triggered are those described previously emerging from the cooperation of both the SML and the NCP+ with the LICL components.

However, for every planning period *t* the volume of the service requests can be described by a probability distribution function (pdf) that can be estimated based on history observations. In practise, this could be achieved by taking a weighted average of the traffic demand over the most recent time periods e.g. using the non-linear autoregressive analysis (NAR). For a detailed description on the subject the reader is referred to [28]. Once this information becomes available, an optimization criterion is selected and the optimal virtual infrastructures that can support the estimated services are identified in terms of both topology and resources. In this manuscript, we considered that the optimal virtual infrastructures are obtained minimising the energy consumption of the underlying substrate, through the following expected cost:

$$min \mathbb{E}_{\xi_{\mathrm{T}}} \left[\sum_t \mathcal{N}_t \left(\boldsymbol{y}, \xi \right) + \sum_t \mathcal{S}_t \left(\boldsymbol{u}, \xi \right) \right] \qquad (1)$$

whereby $\mathcal{N}_{\mathrm{t}}$ is the power consumption of the optical network resources y at time t, $\mathcal{S}_{\mathrm{t}}$ the power consumption cost of computing resources u at time t and ξ is a random vector that contains the uncertain parameters (i.e. traffic demands) that are involved in the planning process. Details regarding the power consumption models for the optical network and computing resources can be found in [6], [9]. However, for the sake of completeness these models are summarized as follows. The present paper is focusing on optical network technologies based on wavelength division multiplexing (WDM) utilizing Optical Cross-Connect (OXC) nodes to perform switching and facilitate routing at the optical layer. The overall network power consumption model is based on the power-dissipating (active) elements of the network that can be classified as switching nodes (OXC nodes), and transmission line related elements. More specifically the OXCs assumed are based on the Central Switch architecture using Micro-Electrical Mechanical Systems (MEMS), while for the fibre links a model comprising a sequence of alternating single mode fibre and dispersion compensating fibre spans together with optical amplifiers to compensate for the losses is employed. The details of these models are described in [29] with the only difference being that the current work assumes wavelength conversion capability available at the OXC nodes. For the computational resources, a linear power consumption model that mainly concentrates on the power consumption associated with the CPU load of IT resources is assumed and is described via the following linear equation [30]:

$$E_{st}(v_{st}) = P_s^i + P_s^b v_{st} \qquad (2)$$

where E_{st} that is the total power used for utilizing a portion v_{st} of the maximum processing capabilities of IT server s at time t and P_s^i, P_s^b are parameters describing the power consumption of the IT server s at idle state and per utilization unit, respectively [6]. In addition to the power consumption due to data processing, a 100% power overhead due to cooling has been incorporated in the power consumption model above described.

At the same time, a set of constraints should be taken into account including: (i) that the planned infrastructures have sufficient optical link capacity for all demands to be transferred to the IT servers, (ii) adequate IT server resources such as CPU, memory, disk storage to support all requested services.(iii) specific capabilities of the underlying physical infrastructure such as wavelength conversion, and (iv) protection from possible network or IT infrastructure failures, or specific security requirement through physical isolation.

In order to solve the above stochastic problem numerically, it is assumed that the random vector ξ has a finite number of possible realizations. Each one of these realizations is called scenario, and each scenario holds a known probability distribution function. Thus, in order to extract this information, the NAR method is adopted to predict the traffic demands for the upcoming time periods due to its inherently low computational complexity and high accuracy. However, due to the large number of scenarios involved in the optimisation, exact evaluation of eq. (1) is not possible. To address this issue, the *Sample Average Approximation* technique has been integrated with *Lagrangian Relaxation* and *Dual Decomposition* to achieve fast convergence to the optimal solution.

The performance of the proposed stochastic re-planning scheme is examined using the COST 239 reference topology [19] in which randomly selected nodes generate traffic demands that need to be served by a set of IT servers. The granularity of service duration for the generated services is one hour. Furthermore, we assume a single fibre per link, 40 wavelengths per fibre, wavelength channels of 10Gb/s each and that each IT server can process up to 2Tb/s and its power consumption ranges from 6.6 to 13.2KW, under idle and full load, respectively. The following scenario has been studied: i) 4 source nodes generate demands normally distributed, ii) the number of arrivals in any given time interval [0,t] follows the Poisson distribution with mean value 2 hours, iii) service times follows the exponential distribution with mean value 2 hours, iv) a single type of services has been considered that require instant access to the IT servers, v) each wavelength requires 10Tb/s of processing power.

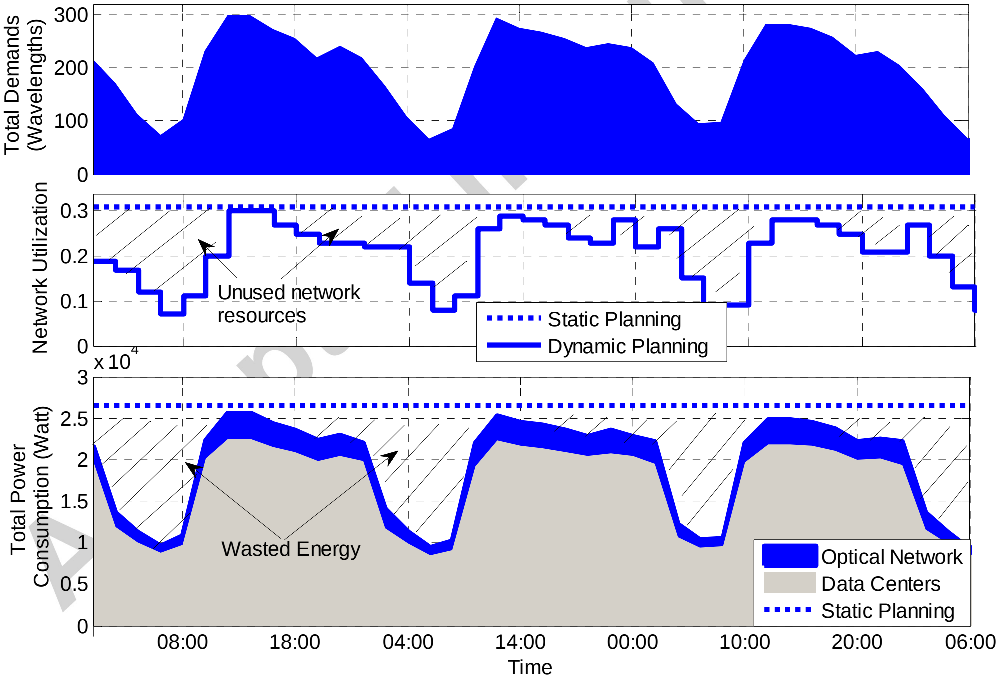

Figure 3: Lifecycle of a virtual infrastructure

Based on history observations for traffic demands recorded in the Pan-European network GEANT for a specific time period, the model is applied to estimate the traffic distribution for future time periods. Once this information is obtained, a set of traffic scenarios is generated using Monte Carlo simulations and the sample average approximated problem is solved to identify the optimal VIs for this predicted traffic distribution. The upper graph in Figure 3 illustrates the evolution of the traffic demands over time. The lower graphs of Figure 3 show the optical network resources allocated to the planned VIs and the total power consumption when applying traditional static planning and the proposed dynamic re-planning. Static planning refers to the case where virtual infrastructures are planned in advance for the highest volume of requests that has to be supported over time. As can be seen, when stochastic planning is adopted, the optical network resources allocated to the VIs are significantly lower (about 30%) than these required in case of static planning. Note that, utilization of optical network resources is defined as the ratio of the number of wavelength links that are used over the total number of available wavelength links. The power consumption share between the optical and the data centres is also depicted in the lower part of Figure 3 where it is seen that the optical network is responsible for 8% - 17% of the total power consumption.

The benefit of stochastic planning can be exploited in practice by adopting dynamic and periodic re-planning of the VIs over the PI. It should be noted that the benefits achieved through dynamic VI re-planning are very much dependent on the VI re-planning time granularity, the sensitivity of the triggering mechanism and the optimization objective chosen. Our modelling results show in Figure 4 that the higher the granularity of VI re-planning the lower the requirement for network resources and the lower the total power consumption.

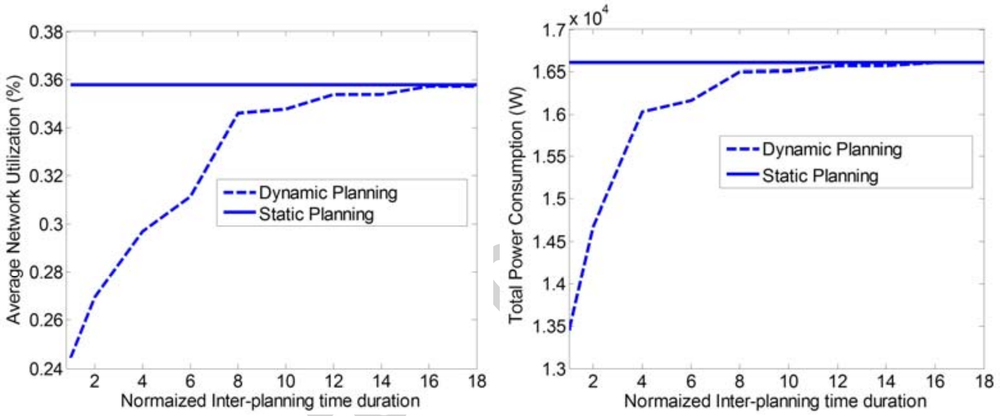

a) b)

Figure 4: Impact of inter-planning time duration on the: a) Utilization of Optical Network Resources and, b) Total Power Consumption

Distributed Enterprise Information System: a real use case

This section presents a practical example on how the proposed virtual infrastructure provisioning can be used to converge IT regions (i.e. DCs) and optical network resources, and how the NCP+ can be used by a real system deployed jointly with the SML to coordinate cloud and network resources composing a virtual infrastructure. The selected distributed application is representative for the enterprise-cloud class of appliances from the network and computing resource consumption. As the CPU utilization of the different VMs composing the distributed application varies from VM to VM, this creates the basis for optimizing the allocation of VMs to physical resources by taking into consideration the energy impact of oversubscribing the physical hosts. This is possible because some VMs consistently have a lower CPU utilization independent of the concurrent load. Thus, the presented application can be used for showing how energy-aware allocation policies can be used for optimizing both the virtual (software) and physical (servers and optical network connections) resources.

A Distributed Enterprise Information System (EIS) is composed of multiple load-balanced query-intensive application-servers and database systems, concurrently serving multiple users. In a dynamic EIS, the number of users and frequency of queries changes, such that the network and computational demand vary as well. In order to validate the response of the distributed system to the changing user loads, we specify what the users' performance expectations are. For this we use consumer-agreed SLAs [12] containing the performance invariants in terms of response times and load distribution. The SLAs are also used for scaling the computational and network resources of the virtual infrastructure through the re-planning functionality described in Section 3. The data used for generating the user load is obtained from empirical analysis of an existing EIS system [13,14] and from existing benchmarks [15] for Enterprise information systems including analytics and business information warehouses. The performed experiments show scaling under different types of load conditions.

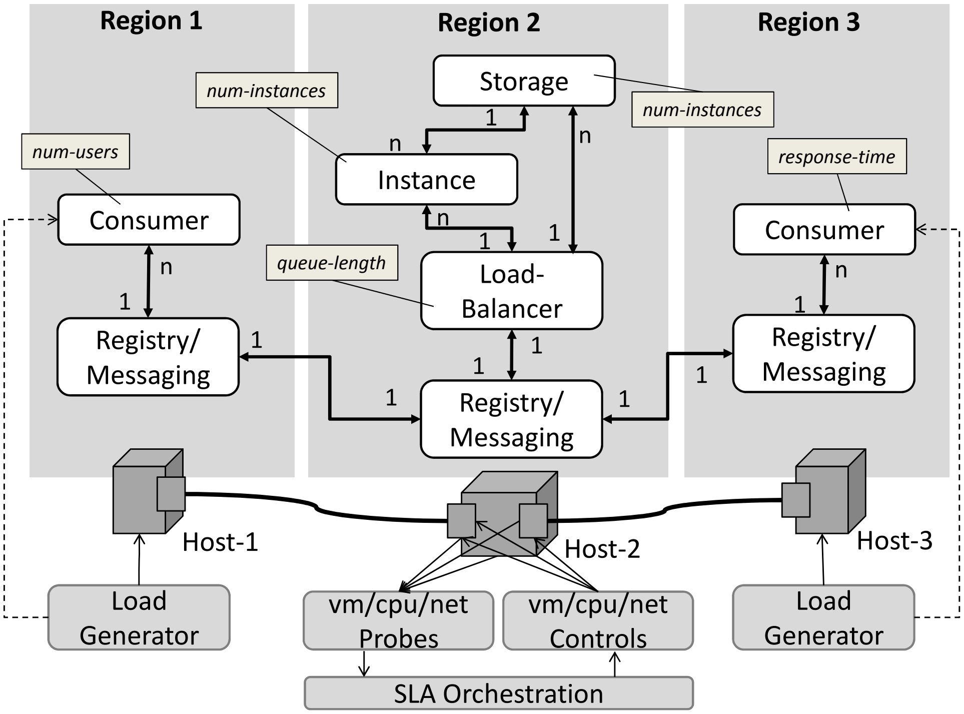

Figure 5: Distributed Enterprise Information System overview

The EIS is composed of four different logical entities: Load Balancer (LB), Worker Instance (WI), Storage (DB) and Consumer (EIS-C). The LB provides the logic by which Consumer requests are allocated to processing entities (WI). The algorithm used for load distribution is Power Saving, where requests are sent to a WI until the instance is full and a next one will be spawned and used. The WI executes requests either locally or on the DB, depending on the request type. The EIS-C simulates a variable number of parallel EIS users which are sending requests according to a given execution plan. All the service entities are implemented as Distributed OSGi (Open Service Gateway initiative) services and use a central registry for discovery.

The virtual infrastructure required for running the EIS is first created to allow the provisioning of the EIS services and then scaled based on the actual reported utilization levels of the resources. The scale process of the EIS is performed through the coordination of the SML and the NCP+ components.

In order to deploy the Distributed EIS, the complete GEYSERS service delivery framework is used, from actual infrastructure planning to control plane deployment, configuration and IT-advertisement. Therefore, first, a VI description containing the desired virtual resources is registered at the management component. Next, the VI planning algorithm takes place at the LICL and the VI is provisioned by configuring the various virtualization components in each administrative domain. This also creates resource reservations in each computational and network domain, bringing together in a single virtual infrastructure converged cloud and optical network resources.

After the VI has been provisioned, it is advertised at the Service Middleware Layer, which will then discover the available computational pools and the network resources. Then, the NCP+ is deployed and the IT advertisement process starts, coordinating the SML and the NCP+. The geographic locations for running the EIS services are selected at the GMPLS+ deployment phase. At this stage, the application description will be registered at the SML, containing the EIS service entities and the dependencies between them. Also, the monitoring metrics for each service are specified. These metrics will later be used for determining services' performance and the health of application SLAs.

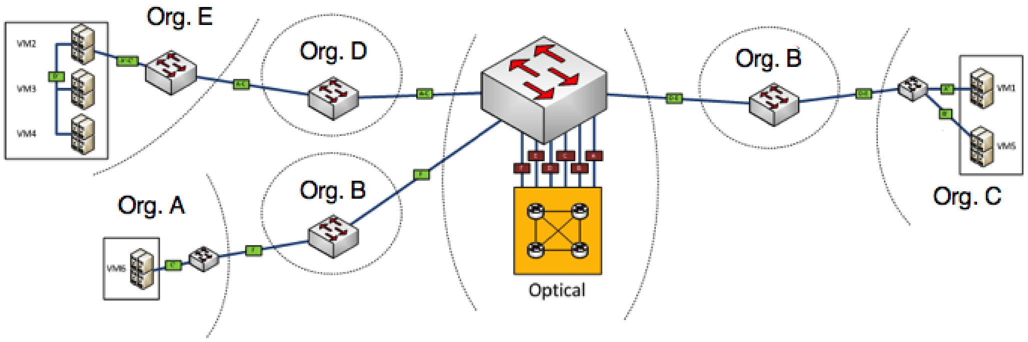

Figure 6: Distributed EIS test-bed

Figure 6 depicts how the distributed EIS application is deployed in a test-bed composed of network and computing resources belonging to different organizations. The computing sites are connected through two different networks: one deployed over optical equipment and one deployed over Ethernet devices. The former is used by the services composing the EIS for exchanging data payloads, while the latter is used for signalling between the different management components of the GEYSERS stack.

Once the GEYSERS components have been deployed and correctly configured to control and operate over the corresponding scenario, the SML will then request the creation of an optical connectivity service between each pair of IT regions. This network connectivity is realised through the NCP+ deployed also on top of the virtual infrastructure. The minimum bandwidth of each network connection is specified by the SML through the UNI service, and it will be later increased based on actual EIS consumers' activity. After the network connectivity has been established and signalled, the SML will begin the instantiation of EIS services in dedicated VMs according to the distributed application's blueprint specification. The list of service dependencies for the service is analysed to find the correct service instantiation order, then the SML will create and start the service VMs. Once that for each service its context has been resolved, a request for deploying a virtual machine with the service binaries will be sent to LICL. The SML will then wait for receiving the VM instantiation notification before marking the service as active and beginning its SLAs monitoring.

In Fig. 7 we display some experimental application performance and network utilization measurements gathered while running the EIS on a small-scale test-bed composed of dual core servers with 4GB of RAM memory and gigabit network interfaces. Each EIS service was running in its own virtual machine with one CPU and 1GB of RAM allocated. In the experiment shown, the concurrent load was maintained at 10 requests per second during a 10 minutes time window. The average network traffic generated by one EIS cloud tenant was 100 MB/minute per VM, distributed as EIS-C: 86.2 MB/min WI: 142.1 MB/min DB: 65.6 MB/min. According to [13] the maximum number of requests per VM before reaching the performance threshold is 50, which is equivalent to a VM network traffic of 500 MB/min or 5GB/min network traffic for a EIS cloud tenant with 10VMs. Considering a datacentre with 1000 quad core servers, the generated network traffic is approximately 2000 GB/min or 2.5Gb/sec.

Based on the reported number of active sessions at the LB, the average response time measured at the EIS Consumers and the aggregated size of the EIS responses' payload (as shown in Fig. 7), the SML will scale [13] the number of WI and DB service as well as the bandwidth of the virtual optical network circuits. Dynamic scale-up (or even scale-down if there are unused resources) of both IT and network resources of the virtual infrastructure over which the EIS is deployed is realised through the re-planning mechanisms aforementioned.

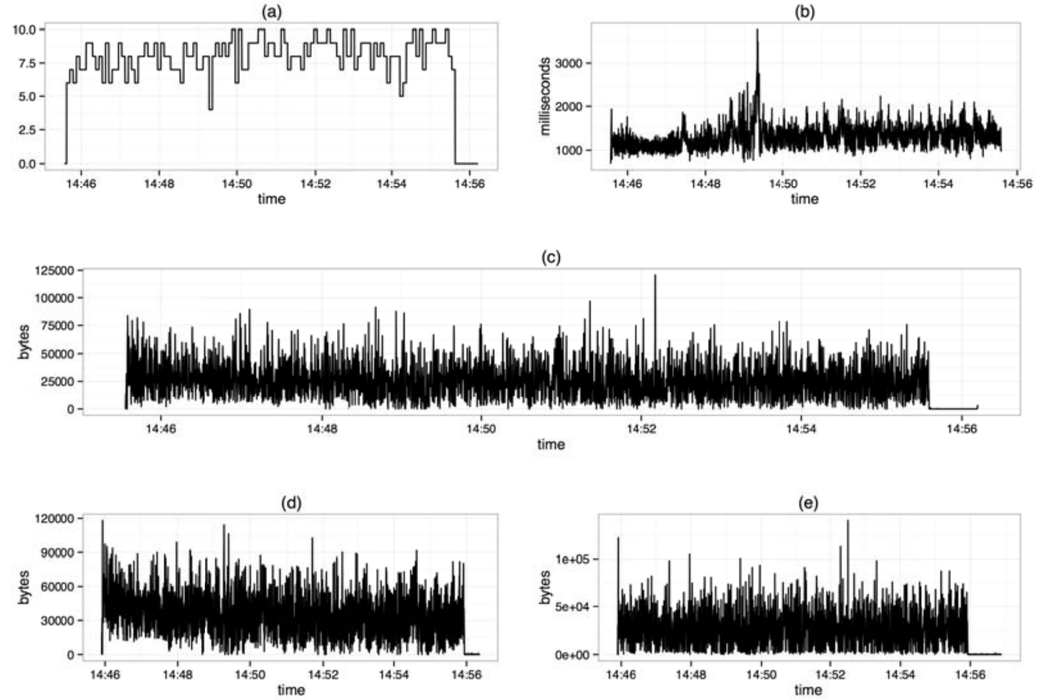

Figure 7 (a) Load Balancer number of concurrent requests, (b) Consumer request execution time, (c) Consumer network read bytes, (d) Worker network write activity, (e) Storage network write activity

In Figure 8 we display the simulated energy consumption behaviour for the two servers hosting the VMs described above. The simulation uses a linear power model based on Spec Power [20] benchmarking information using the servers' average CPU utilization. The actual server CPU utilization is derived using the VMs' CPU utilization trace measured in the previous step. According to the simulation, the average server power consumption is 4.2KW/minute or 250KWh per server.

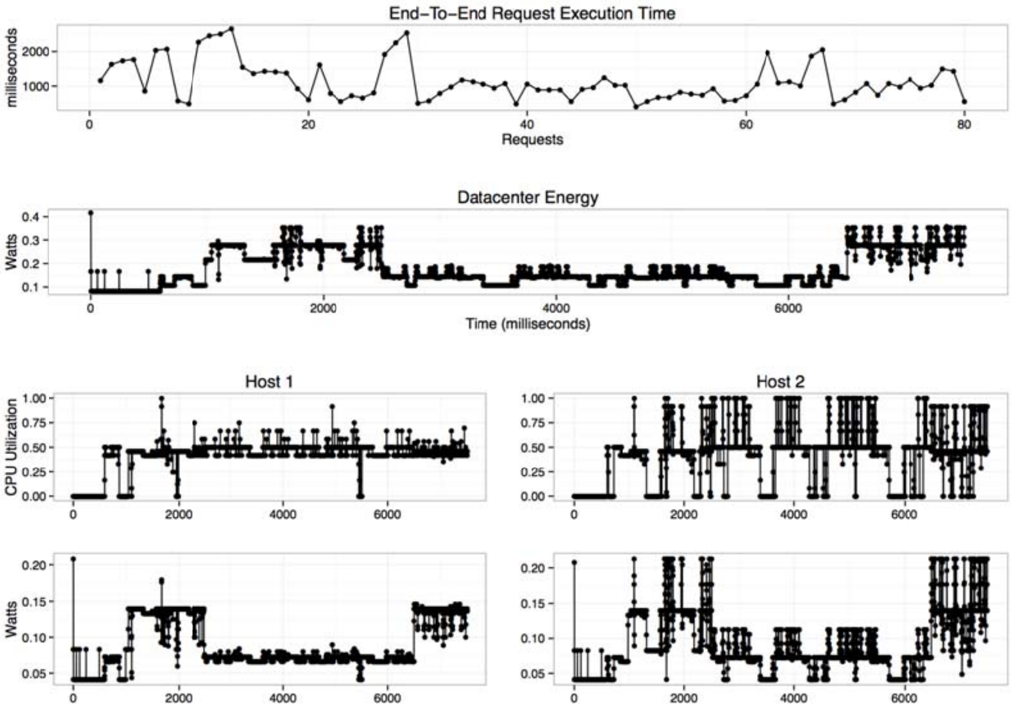

Figure 8. Simulated data centre power consumption versus CPU consumption

As an example, if the communication between region 1 and region 3 requires more network capacity, the SML communicates this to the NCP+, which is the element responsible of triggering the re-planning mechanism at the virtualisation layer, the LICL. Then, the LICL starts the planning process considering resources that are not allocated to any virtual infrastructure, e.g. free wavelengths in the transport network. Thus, the virtual link connecting different regions of the D-EIS is modified, and the NCP+ can now update the network service according to the application needs. It is worth mentioning that, additional to the dynamic modification of the virtual optical links, the re-planning mechanisms also enable adding or removing unused resources, as well as the dynamic modification of the virtual IT pools that the D-EIS can use to deploy services of the application. This deployment of the system becomes a clear and sound example on how virtualisation can be used to converge Cloud services with high-capacity optical network circuits.

Conclusion

This paper has presented an architecture emerging from the European GEYSERS project that combines both IT and optical network resources through virtual infrastructures. The virtualisation approach can potentially increase the usage of the infrastructure by sharing its resources, bringing higher profits from the same underlying substrate. It clearly reduces the need of infrastructure over-provisioning since it enables infrastructure providers to dynamically adapt the infrastructure to different operators' businesses at any moment of the service lifetime. We have presented how a converged virtualisation enables the creation of virtual infrastructures that address the high demanding requirements of dynamic, elastic cloud applications. The virtual infrastructure service proposed has been automated, in order to avoid manual presence in the provisioning workflow. Automation of the provisioning includes different possibilities for describing the resources that compose the virtual infrastructure and determining over which physical resources the virtual infrastructure is instantiated. Therefore, VI planning becomes a key element of the service workflow. Details of the planning algorithms are not included for the sake of simplicity, although typical optimal functions focus on the resource utilisation, or even the energy consumption of the resources.

Furthermore, re-planning features of the virtual infrastructure service have been introduced to address uncertainty of the behaviour of the Cloud applications. When stochastic planning is adopted, the optical network resources allocated to the VIs are significantly lower than these required in case of static planning. We have shown how the benefit of stochastic planning can be exploited in practice by adopting dynamic and periodic re-planning of the VIs over the physical substrate. Finally, in order to show a real use case for the virtual infrastructure service and its re-planning feature, we have presented the distributed Enterprise Information System, deployed over a distributed infrastructure composed of both IT resources and optical network resources. A virtual infrastructure is created on top of the different administrative domains, creating virtual resources to be used by the EIS. Up or downscaling of the infrastructure is performed as a function of the values monitored at the application level by the EIS.

Acknowledgements

This work has been founded by the European Commission ICT-2009.1.1 Project #248657 Generalised Architecture for Dynamic Infrastructure Services (GEYSERS). We thank our project partners for their valuable contributions.

References

[1] Verchere, D., "Cloud computing over telecom network," Optical Fiber Communication Conference and Exposition/National Fiber Optic Engineers Conference (OFC/NFOEC), 6-10 March 2011.

[2] Vicat-Blanc, P., Figuerola, S., et al. “Bringing optical networks to the cloud: an architecture for a sustainable future internet”. The Future Internet Assembly Book. pp 307-320, Springer-Verlag 978-3-642-20897-3

[3] Lee, Y., et al. “Framework for GMPLS and Path Computation Element (PCE) control of Wavelength Switched Optical Networks (WSONs)”. RFC 6163.

[4] Ian Foster. “The grid: computing without bounds”. Scientific American, 288(4):78– 85, April 2003.

[5] Figuerola, S.; Lemay, M., "Infrastructure services for optical networks [Invited]," IEEE/OSA Journal of Optical Communications and Networking, vol.1, no.2, pp.A247-A257, July 2009. 10.1364/JOCN.1.00A247

[6] Tzanakaki, A., Anastasopoulos, M., Georgakilas, K., et al. “Energy efficiency in integrated IT and optical network infrastructures: the GEYSERS approach”. Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on; 10-15 April 2011, pp. 343 - 348, 978-1-4577-0249- 5, 10.1109/INFCOMW.2011.5928835.

[7] Bohnert, T.M.; Figuerola, S.; Ciulli, N.; Simeonidou, D.; Vicat-Blanc, P., "Optical networking for cloud computing," Optical Fiber Communication Conference and Exposition/National Fiber Optic Engineers Conference (OFC/NFOEC), 6-10 March 2011.

[8] Ghijsen, M.; Van der Ham, J.; Grosso, P.; De Laat, C., "Towards an infrastructure description language for modeling computing infrastructures," Parallel and Distributed Processing with Applications (ISPA), 2012 IEEE 10th International Symposium on , vol., no., pp.207,214, 10-13 July 2012. 10.1109/ISPA.2012.35.

[9] K. Georgakilas, A. Tzanakaki, M. P. Anastasopoulos and Jens M. Pedersen, "Converged optical network and data center virtual infrastructure planning", IEEE/OSA Journal of Optical Communications and Networking, Vol. 4, No. 9, pp. 681–691, Sep. 2012.

[10] S. Peng, R. Nejabati, D. Simeonidou, “Impairment-aware optical network virtualization in single-line-rate and mixed-line-rate WDM networks”, IEEE/OSA Journal of Optical Communications and Networking, 5(4), pp. 283-293, April 2013.

[11] A. Pages, J. Perelló, S. Spadaro, J.A. Garcia-Espin, J.F. Riera, and S. Figuerola, "Optimal allocation of virtual optical networks for the future internet", in Proc. ONDM, 2012, pp.1-6.

[12] A.-F. Antonescu, P. Robinson, and T. Braun, “Dynamic topology orchestration for distributed cloud-based applications,” in Network Cloud Computing and Applications, IEEE 2nd Symposium on, 2012.

[13] A.-F. Antonescu, A.-M. Oprescu, Y. Demchenko, C. de Laat, T. Braun "Dynamic optimization of SLA-based services scaling rules", 5th IEEE International Conference on Cloud Computing Technology and Science (CloudCom), Bristol, UK, December 2 - 5, 2013

[14] A.-F. Antonescu, P. Robinson, T. Braun "Dynamic SLA management with forecasting using multi-objective optimizations" in IFIP/IEEE Integrated Network Management Symposium (IM 2013) , Ghent, Belgium, May 27 - 31, 2013, ISBN 978-1-4673-5229-1

[15] SAP Standard Application Benchmarks for 3-tier systems. Available Online: http://www.sap.com/solutions/benchmark/sd3tier.epx

[16] Landi, G.; Ciulli, N.; Buysse, J.; Georgakilas, K.; Anastasopoulos, M.; Tzanakaki, A.; Develder, C.; Escalona, E.; Parniewicz, D.; Binczewski, A.; Belter, B., "A network control plane architecture for on-demand co-provisioning of optical network and IT services," *Future Network & Mobile Summit (FutureNetw), 2012* , vol., no., pp.1,8, 4-6 July 2012.

[17] N. M. M. K. Chowdhury. "Network virtualization: state of the art and research challenges". IEEE communications magazine, 47(7): 20–26, July 2009.

[18] TR139, Service Delivery Framework Overview, Release 2.0". http://www.tmforum.org/ *(Available online)*

[19] O'Mahony, M.J.: "Results from the COST 239 project. Ultra-high capacity optical transmission networks". *Optical Communication, 1996, ECOC '96. 22^nd^ European Conference on*, vol. 2, no., pp11,18, vol.2, 19-19. Sept. 1996.

[20] SpecPower (online, Apr 2014) HP ProLiant ML110 G5 http://www.spec.org/power_ssj2008/results/res2011q1/power_ssj2008-20110124-00339.html

[21] A. Fischer, et al. "Virtual network embedding: a survey". Communications Surveys Tutorials, IEEE, pp (99):1 –19, 2013

[22] A. Pagès, et al. "Strategies for virtual optical network allocation", IEEE Communications Letters, vol. 16, no.2, pp. 268-271, Feb. 2012

[23] N.M. Mosharaf Kabir Chowdhury and Raouf Boutaba. "A survey of network virtualization". Computer Networks, 54(5):862 – 876, 2010. ISSN 1389-1286.

[24] Fajjari, N. Aitsaadi, G. Pujolle, and H. Zimmermann. "Vnr algorithm: a greedy approach for virtual networks reconfigurations". In Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE, pages 1 – 5, dec. 2011

[25] Marcin Bienkowski, Anja Feldmann, Dan Jurca, Wolfang Kellerer, Gregor Schaffrath, Stefan Schmid, and Joerg Widmer. "Competitive analysis for service migration in vnets". In Proceedings of the second ACM SIGCOMM workshop on Virtualized infrastructure systems and architectures, VISA '10, pages 17–24, New York, NY, USA, 2010. ACM.

[26] Y. Zhu and M. Ammar. "Algorithms for assigning substrate network resources to virtual network components". In INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings, pages 1–12, 2006.

[27] Weigang Hou; Lei Guo; Yejun Liu; Qingyang Song; Xuetao Wei, "Virtual network planning for converged optical and data centers: ideas and challenges," IEEE Network, vol.27, no.6, pp.52,58, 2013

[28] A. Tzanakaki, M. P. Anastasopoulos and K. Georgakilas, “Dynamic virtual optical networks supporting uncertain traffic demands,” *IEEE/OSA Journal of Optical Communications and Networking*, OFC/NFOEC 2013 Special Issue (invited), Vol. 5, No. 10, Oct. 2013

[29] A. Tzanakaki et al, “Power considerations towards a sustainable pan-european network”, JWA061, in Proc. of OFC 2011

[30] V. Valancius et. al, “Greening the internet with nano data centers”, In Proc. of CoNEXT ’09, 37-48, 2009.