Design and Implementation of a Simulation Environment for Network Virtualization

Marc De Leenheer, Jens Buysse, Kevin Mets, Bart Dhoedt, and Chris Develder
Department of Information Technology
Ghent University - IBBT, Ghent, Belgium
Email: {marc.deleenheer, jens.buysse, chris.develder}@intec.ugent.be

***Abstract*—Network virtualization has been proposed as a key enabler of the future internet, as it allows multiple networks to coexist on a shared physical infrastructure, and as such overcomes the difficulties in deploying novel technologies in the current Internet. However, much confusion still exists on the impact virtualization will have on performance, since we lack the necessary tools to study the behaviour of a virtualized network infrastructure. We report on the architecture and implementation of a simulation environment for virtualizing both (optical) network and IT infrastructures. We provide a detailed overview of the layered architecture, give insight in the design and implementation of the simulator, discuss potential use cases and finally include some preliminary results.**

I. Introduction

Current projections indicate that at the end of this decade, the scale of information processing will scale from Petabytes to Exabytes of data [1]. Additionally, emerging paradigms such as cloud computing and IaaS, are driving profound transformations of networks' and users' capabilities [2]. Consequently, a new class of high-performance and high-capacity network-based applications are emerging, posing strict IT (e.g., computing and data storage) resource and service requirements. Due to its own succes and pervasiveness, the current best-effort Internet is unable to adapt to these novel service paradigms. Hence, there is an opportunity for operators/providers to create new services, especially integrated offerings of both optical network connectivity services and traditional IT services.

The GEYSERS project aims to design and showcase a novel architecture, able to provide network operators with an infrastructure composed of several optical network and IT resources in an on-demand fashion [3]. To this end, the physical resources can be partitioned and aggregated to create a virtual infrastructure (VI), which in turn can be controlled by a network operator without interference of other VIs [4]. To control this infrastructure on demand, GEYSERS' architecture deploys an enhanced Network Control Plane (NCP+) that can control both network and IT resources. This way, both network and IT resources can be seen as elements of one homogenous set, able to be provisioned on-demand.

Obviously, validating this architecture is not a straightforward task, given the software and protocol stack that the GEYSERS vision encompasses. As such, the project envisions experiments in a reasonably limited scale testbed comprising around 10–15 nodes. To perform full scale validation, and perform extensive testing of the architecture's scalability, experiments based on discrete event simulations have been identified as the most appropriate method to study the performance. The idea is to implement the full functionality of the layered architecture, and perform validation and testing on medium to large scale networks (consisting of hundreds of nodes). In this paper, we report on the design and implementation of this simulation environment, and demonstrate its features by way of a qualitative discussion of sample simulation scenarios.

The remainder of the paper is organized as follows: Section II discusses similar proposals, and Section III introduces the GEYSERS architecture on which our simulation environment is modeled, while Section IV describes the design and some implemention details of the simulator itself. The following Section V presents a number of use cases that will be validated, and finally Section VI summarizes the paper.

II. Related Work

The current interest in architectures for the future internet has led to substantial research on this topic [2]. For instance, sharing a physical infrastructure among multiple virtual networks (only considering networking elements, or, stated differently, disregarding IT end resources), is a topic well-studied and is referred to as Virtual Private Networks (VPN), overlay networks or even active networks.

The goal of a Virtual Private Network or VPN is to connect a number of known end-points over a dedicated communications infrastructure, usually by creating tunnels over a public medium (e.g. the Internet) [5]. These may exist on multiple layers of the network, as evidenced by the existence of either Layer 1, 2 or 3 VPNs. On the other hand, overlay networks are usually implemented on the application layer (L7), and are therefore aimed at providing specific services such as file sharing [6], multicasting or various other goals, including offering Quality of Service (QoS), protecting against Denial of Service (DoS) attacks and many others.

To the best of our knowledge, this paper is the first to report on simulation activities on virtualized networking architectures that comprise combined network+IT virtualization, and comprise both control plane and a virtualization layer. Nevertheless, some research has already appeared on the topic of simulation of service-oriented networks. For instance, [7] presents an extensible toolkit for the modelling and simulation of cloud computing environments, while [8] does the same for

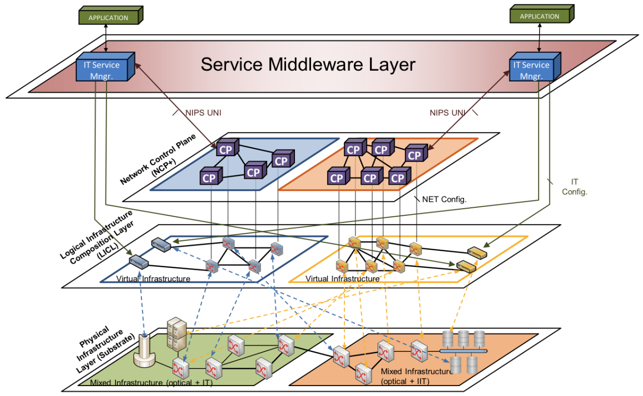

Fig. 1. GEYSERS layered architecture

Grid computing infrastructures. Similarly, an example of pure network control plane simulations, esp. GMPLS-based, is [9]. Finally, we mention some work on scalability testing of large networks, as can be found for example in [10], where various approaches are taken to study the performance of different aspects of communication networks.

Complementary to pure simulation studies (as we envisage), also emulation approaches have been proposed to study scalability of large scale networks. For instance, in [11] the authors describe an emulation environment to study fault behaviour and network behaviour in an environment modeled after the Internet.

III. The GEYSERS Layered Architecture

The Generalised Architecture for Dynamic Infrastructure Services (GEYSERS) is a European FP7 project, that will design a novel architecture for seamless and coordinated provisioning of both optical and IT resources, and to develop the necessary tools and software to realize this objective. In particular, virtualization is one of the key goals in this project: adequate mechanisms for abstraction, partitioning and aggregation will be provided. The resources which we consider include optical networking nodes and links, and IT resources such as computational- and data storage equipment. Another point of focus is the inclusion of energy efficient mechanisms on all levels of the architecture.

The architecture is detailed in Fig. 1, and is basically composed of four layers. First, devices in the Physical Infrastructure (PI) layer are abstracted and partitioned or grouped into virtual resources that can be selected to form the Virtual Infrastructures (VI) in the Logical Infrastructure Composition Layer (LICL). Within each VI, controllers in the IT-aware network control plane (NCP+) layer configure and manage virtual network resources. The Service Middleware Layer (SML) is responsible for translating the application requests and service level agreements (SLAs) into technology specific requests to trigger the provisioning procedures at the NCP+. Refer to [3] for a more detailed discussion on the different components in the layered architecture.

Our aim is to validate the overall GEYSERS architecture, and in particular the end-to-end service provisioning workflow across the various layers and associated interfaces. To this end, we are developing a simulation environment to evaluate the performance of a large-scale network, to complement the relatively small-scale tests of the actual implementation in a testbed. The simulation framework will be used to evaluate the performance and scalability of the architecture and its workflows, as well as associated algorithms for routing, allocation, dynamic partitioning, etc. Ultimately, the outcome of our research will be used to refine and validate the overall architecture.

IV. Simulator Architecture

We model the major novel components of the aforementioned GEYSERS architecture, specifically the LICL and the NCP+; we do not elaborate on the SML as this component is already in existence in current service-oriented networks.

The main objectives in developing this simulator are to:

- demonstrate the *feasibility* of the GEYSERS architecture (e.g. in terms of achieving energy efficiency)
- identify which potential *bottlenecks* may exist within the architecture
- verify whether the novel components can *scale* towards large networks:
  - comprising a large number of physical resources,
  - supporting a large number of virtual infrastructures,
  - performing as expected under highly dynamic network conditions and user demand.

This requires a thorough investigation of the scalability, overhead, response times and blocking behaviour of the mechanisms, protocols and interfaces.

A. Overview

The simulator is built on the OMNeT++ simulation framework [12], which is an extensible and highly scalable [13] C++ discrete-event simulation environment aimed at building network simulators. Of particular interest is the INET Framework, which offers implementations for a variety of both wired and wireless networking protocols (covering most of the TCP/IP stack). It also includes an incomplete (but useable) implementation of the (G)MPLS protocol, and forms the basis of our implementation of the NCP+ functionality. The relevant standards include Resource Reservation Protocol (RSVP), Label Distribution Protocol (LDP), and Contrained Shortest Path First (CSPF) routing.

The simulator is composed of two major blocks (Fig. 2): one portion is implemented in OMNeT++, while the other portion makes use of a relational database. This design choice reflects the rather static behaviour of the physical infrastructure, while more dynamic components such as the network control objects are modelled in OMNeT++. This has the advantage of freeing up more memory space for the upper layers of the architecture. Note that as scalability is of major concern, offloading parts of the modeling to a database is preferred, even though a minor penalty can be expected due to database retrieval operations.

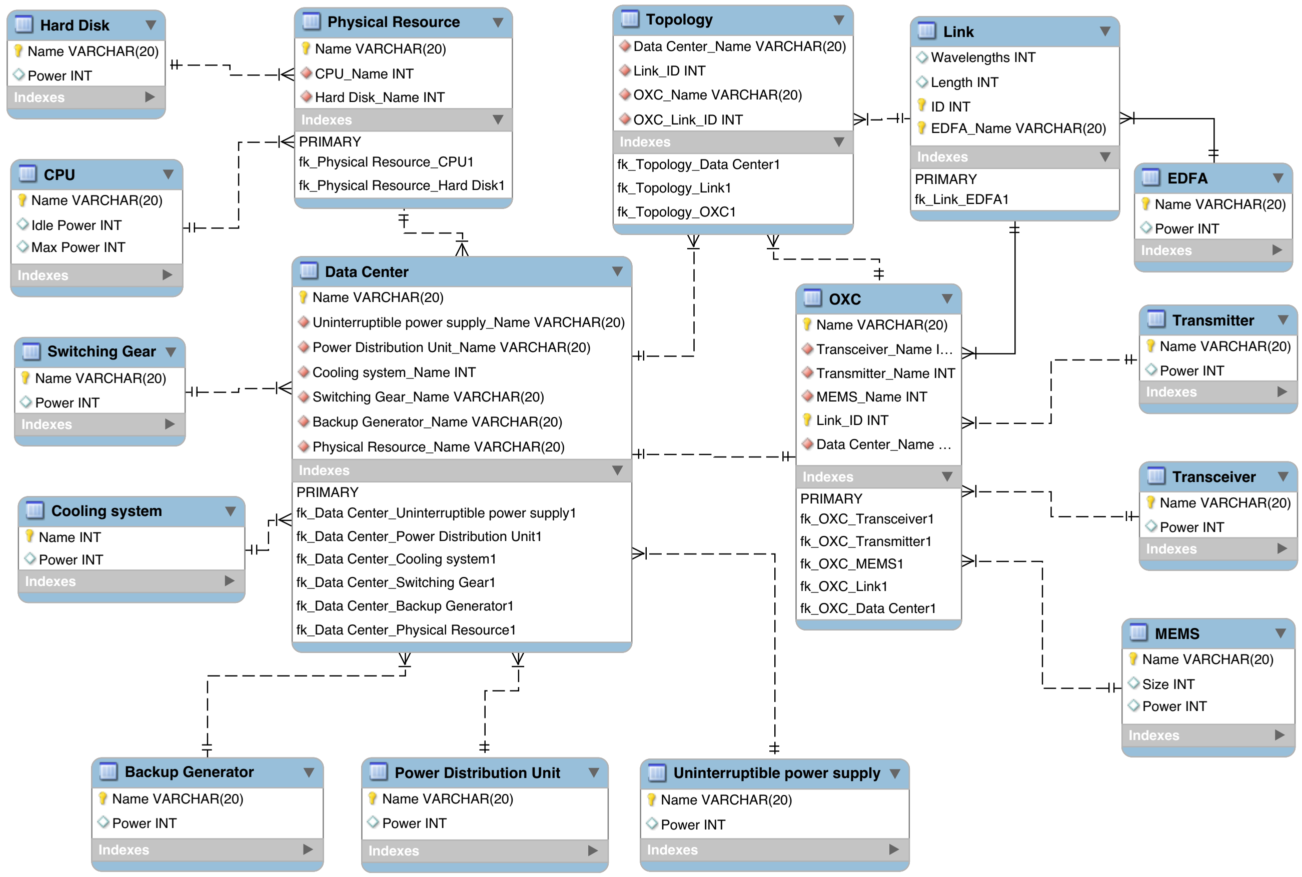

Fig. 3. Entity relationship model for the physical infrastructure

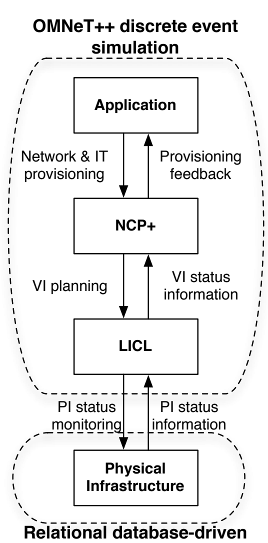

Fig. 2. Overview of the simulation environment

B. Entity Relationship Diagram

As shown in Fig. 2, a simulation is constructed by running both a discrete event simulator, and a relational database-driven component. The data model for the physical infrastructure is shown as an entity relationship diagram in Fig. 3. Of note is the detailed description of the optical networking components, as evidenced by the inclusion of transmitter, transceiver, optical switching fabric (MEMS) and the optical amplifiers (EDFA). Also observe the equipment that contributes to the energy consumption of the architecture, in particular the cooling, uninterruptible power supple (UPS) and backup generators, all of which are common in today's communication networks.

C. UML Models

The detailed designs of the LICL and NCP+ layers are depicted in Fig. 4 and 5. The LICL component is composed of virtual devices that correspond to either networking or IT equipment. The key networking entity is the optical cross-connect (`OXC`), which serves as a switching device and is composed of ports (`Port`). Each port can be either an input (`Inport`) or output (`Outport`) port. As we mainly focus on optical networking, each wavelength (`Lambda`) is part of a (`Port, Phy Link`) pair, in which the latter represents a physical link. The central object for IT equipment is the physical resource (`Phy Server`), which contains one or more processing units (`CPU`) and storage disks (`Disk`). In its turn, physical resources can be grouped into a cluster or datacenter environment (`Datacenter`). Finally, note the

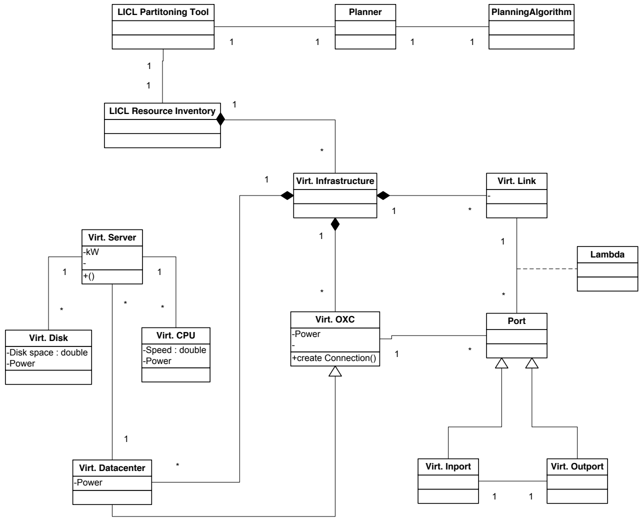

Fig. 4. LICL UML diagram

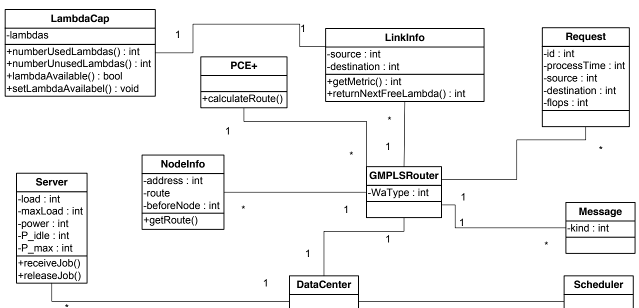

Fig. 5. NCP+ UML diagram

EnergyController which provides the energy consumption of various devices (refer to Section V-C for more detail). A number of additional classes are necessary to drive both the planning of the VI process and maintain the mapping between VI and PI. Specifically, the information of the VI-to-PI mappings is stored in the LICL Resource Inventory, while the LICL Partitioning Tool is responsible for planning. In its turn, the Planner can choose between different objectives by selection of an appropriate planning algorithm (PlanningAlgorithm, see Section V-A).

Finally, the NCP+ simulation component draws from the basic concepts in the GMPLS protocol. A demand for a connection is modeled by the Request object, containing the relevant connection parameters. The main element of the network is the GmplsRouter, which forms the start and endpoint of an optical connection, represented by a Route. This route is calculated by the PCE+ class, passed on to the GMPLSRouter, which then stores this information in the NodeInfo object. Each router has a database, consisting of LinkInfo objects, to track connections and the wavelengths they use (stored in the LambdaCap) for each link. This information is then exchanged through the OSPF-TE and RSVP-TE protocols.

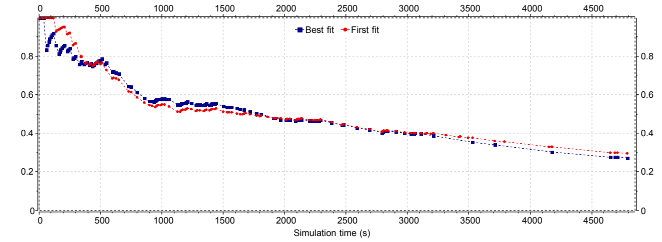

Fig. 6. VI request acceptance ratio, when mapping of resources is performed in a first fit or best fit fashion

Regarding the IT functionality of the NCP+, the DataCenter groups a number of Server objects, that are controlled by a Scheduler. Finally, the Message class is used to exchange information between the different objects in the NCP+.

V. Use Cases

In the following, we describe three use cases which will be studied using our simulation environment. These scenarios have served as guidelines during development and are thus the minimal functionalities the simulator supports. The cases are the result of a consultation of the GEYSERS' partners, comprising representatives from both academia and industry.

A. LICL Scalability

One of the fundamental issues in virtualized network environments is how to perform the mapping of virtual infrastructure (composed of both network and IT resources) requests on a shared physical infrastructure [15]. A virtual infrastructure (VI) is in essence a subset of the underlying physical infrastructure, and relevant objectives include maximization of the number of accepted VI requests and energy efficient mapping (see Section V-C). This is one of the key roles of the LICL, and thus a number of VI mapping algorithms will be developed and evaluated in the simulation environment.

Referring back to Fig. 4, the Planner and PlanningAlgorithm classes are responsible for this functionality, while the LICL Partitioning Tool will make the necessary changes to the LICL Resource Inventory based on the VI mapping algorithms' outcome (see Fig. 6 for a sample result).

Additionally, we will investigate the overhead introduced by the LICL layer, for instance when the NCP+ must provision a new network path, how much delay is added by going through the LICL layer? Finally, an evaluation of the different architectural options for implementing the LICL layer (centralized, distributed, or hybrid) will be performed.

B. NCP+ Scalability

This activity explores the scalability of the proposed NCP+, by evaluating different architectural options for PCE+ inter-domain path computation, and the overhead in terms of, among others, message exchange and signaling delay caused by introduction of IT resource state information in the GMPLS+

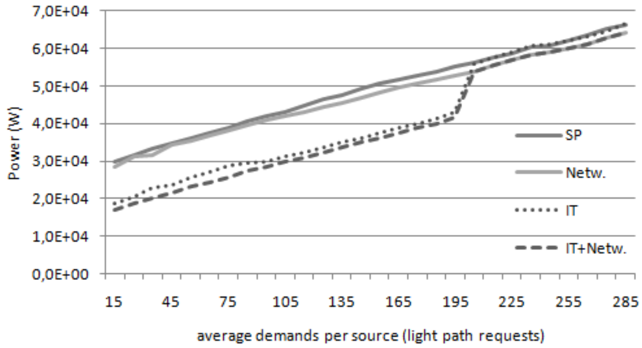

Fig. 7. Total energy consumption of the COST32 network. The different lines represent the used strategy: SP is shortest path routing, Netw. only considers network resources' power consumption, IT only considers IT resources' power consumption, while IT + Netw. considers both.

protocol. Several strategies are available and will be investigated:

- in the centralized approach, one `PCE+` will do the path computation for all GMPLS+ controllers.
- a hierarchical design, in which a number of GMPLS+ controllers share a `PCE+` object. In turn, these `PCE+` objects share a parent `PCE+`, which performs the path computation on an abstracted topology. A number of alternative tree designs will be studied to investigate the scalability of this approach.
- one `PCE+` object per `GmplsRouter`, such that the path computation process is performed in a fully distributed way.

C. Energy Efficient Design and Operation

The GEYSERS architecture incorporates both energy efficient design and operation of infrastructures, whereby the joint optimization of both IT and optical networking resources is considered [14]. This implies both the LICL, which is reponsible for the VI planning phase (design), and the NCP+, responsible for the VI service provisioning (operation), should incorporate energy efficiency parameters. As illustrated in Fig. 4, these energy-related parameters will be handled by the `EnergyController` in the physical layer. For instance, when two different VI's have virtual resources derived from the same physical resource, the power consumption of each virtual resource is dependent on the total load of the physical resource. Generating the power consumption will be based on both experimental values and appropriate models of the relevant devices in the physical infrastructure.

We performed an initial case study (outlined in [14]) to evaluate potential energy savings, by considering both network and IT resource power consumption. Fig. 7 shows a number of strategies to reduce energy consumption, and demonstrates that the joint consideration of both network and IT resources can achieve a considerable decrease in energy consumption. Ongoing work involves simulation studies on various online mechanisms to assess the achievable energy savings.

VI. Conclusion

We presented an overview of the layered GEYSERS architecture, which aims to introduce virtualization of both optical network and IT resources. Furthermore, the design and implementation of a simulation environment, to accurately evaluate the feasibility of the architecture, was presented. The simulator allows extensive scalability testing of all relevant layers of the proposed architecture. Finally, a number of use cases were discussed that demonstrated the functionalities of our simulation environment. A preliminary case study identified substantial potential energy saving opportunities that could be achieved by the Geysers framework.

Acknowledgment

The work described in this paper was partly funded by the European Commission through the 7th ICT-Framework Programme (FP7) project GEYSERS (www.geysers.eu). M. De Leenheer is a post-doctoral fellow of the Research Foundation Flanders (FWO), and is currently working as a visiting postdoc at the University of California, Davis, supported by a Fellowship of the Belgian American Educational Foundation. J. Buysse is supported by the agency for Innovation by Science and Technology. C. Develder is funded in part by the FWO as a post-doctoral fellow.

References

[1] Cisco white paper, *Cisco Visual Networking Index: Forecast and Methodology, 2009-2014*, June 2010.
[2] J. Schonwalder et al., *Future Internet = content + services + management*, IEEE Communications Magazine, 47(7), July 2009.
[3] E. Escalona et al., *GEYSERS: A novel network architecture for "Optical + IT" service provisioning*, Submitted to IEEE Communications Magazine
[4] T. Anderson, L. Peterson, S. Shenker, J. Turner, *Overcoming the Internet Impasse through Virtualization*, Computer, 38(4):34–41, April 2005.
[5] E. Rosen, and Y. Rekhter, *BGP/MPLS IP Virtual Private Networks (VPNs)*, RFC 4364, February 2006.
[6] E.K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, *A survey and comparison of peer-to-peer overlay network schemes*, IEEE Communications Surveys & Tutorials, 7(2):72–93, 2005.
[7] R. Buyya, R. Ranjan, and R.N. Calheiros, *Modeling and simulation of scalable Cloud computing environments and the CloudSim toolkit: Challenges and opportunities*, Proc. Int. Conf. on High Performance Computing & Simulation, June 2009.
[8] S. Bagchi, *Simulation of grid computing infrastructure: challenges and solutions*, Proc. of the Winter Simulation Conference, December 2005.
[9] J. Korniak, P. Rozycki, *GMPLS simulation tools*, Proc. 1st Conf. "Tools of Information Technology", September 2006.
[10] G.R. Yaun et al., *Large-scale network simulation techniques: examples of TCP and OSPF models*, ACM SIGCOMM Computer Communication Review, 33(3), July 2003
[11] A. Vahdat et al., *Scalability and accuracy in a large-scale network emulator*, Proc. of the 5th Symp. on Operating Systems Design and Implementation, 2002.
[12] A. Varga, and R. Hornig, *An overview of the OMNeT++ simulation environment*, Proc. 1st Int. Conf. on Simulation Tools and Techniques for Communications, Networks and Systems, pp. 1–10, 2008.
[13] E. Weingärtner, H. Vom Lehn, and K. Wehrle, *A performance comparison of recent network simulators*, Proc. IEEE Int. Conf. on Communications (ICC), pp. 1287–1291, June 2009.
[14] A. Tzanakaki, et al., *Energy Efficiency in integrated IT and Optical Network Infrastructures: The GEYSERS approach*, Accepted Proc. IEEE INFOCOM 2011 Workshop on Green Communications and Networking.
[15] N. Chowdhury, and R. Boutaba, *A survey of network virtualization*, Comput. Netw., 54(5):862–876, 2010.