CONTENT Project: Considerations towards a Cloud-based Internetworking Paradigm

Kostas Katsalis, Thanasis Korakis
University of Thessaly, Greece
kkatsalis,korakis@uth.gr

Giada Landi, Giacomo Bernini
Nextworks,Pisa, Italy
g.landi,g.bernini@nextworks.it

Bijan R. Rofoee, Shuping Peng
University of Bristol, UK
bijan.rofoee,shuping.peng@bristol.ac.uk

Markos Anastasopoulos, Anna Tzanakaki
Athens Information Technology, Greece
manast,atza@ait.gr

Dora Christofi, Michael Georgiades
PrimeTel PLC, Cyprus
dorac,michaelg@prime-tel.com

Renaud Larsen
Juniper Networks Inc.,Belgium
renaudlarsen@juniper.net

Jordi Ferrer Riera, Eduard Escalona, Joan A. Garcia-Espin
Fundació i2CAT, Spain
jordi.ferrer,eduard.escalona,joan.antoni.garcia@i2cat.net

***Abstract*—Although cloud computing and the Software Defined Network (SDN) framework are fundamentally changing the way we think about network services, multi-domain and multi-technology problems are not sufficiently investigated. These multi-domain, end-to-end problems concern communication paths that span from the wireless access and the wireless backhaul networks to the IT resources through optical networks. In this paper we present the CONTENT project approach to network and infrastructure virtualization over heterogeneous, wireless and metro optical networks, that can be used to provide end-to-end cloud services. The project goal is to drive innovation across multi-technology infrastructures and allow ICT to be delivered and consumed as a service by Virtual Network Operators. The communication mechanics between wireless and optical domains and the physical layer abstractions of a CONTENT Virtual Network are presented and the relation of the proposed approach with the SDN framework is investigated.**

***Keywords*—Network virtualization, inter-domain networking, wireless-optical networks, cloud computing**

I. Introduction

The mass adoption of the Internet is continually driving the technology in finding ways to face age-old obstacles, like user mobility, and physical infrastructure and network limitations. Network and infrastructure virtualization technologies give the opportunity to physical infrastructure providers to add value to their existing business propositions and generate an interesting new turnover, mainly because of the way they are able to handle these limitations [1]. To address the requirements of future virtual network operators, a combination of knowledge obtained by the recent advances in virtualization and Software Defined Networking (SDN) technology must be investigated, while in order to build and deploy efficient cloud based services, we must depart from models that consider intangible multi-domain end-to-end settings.

In this work we present the CONTENT project approach [2], an EU funded effort, for interconnecting geographically distributed computational resources through ubiquitous converged virtual network infrastructures. CONTENT project aims to deliver an SDN-enabled inter-domain cloud networking platform, engineered with a bottom up approach to meet the requirements of an end-to-end cloud computing environment. Thanks to the approach devised, the successful Mobile Virtual Network Operator (MVNO) model that is used to provide wireless services over the physical network provider, can be extended to a Mobile-Optical Virtual Network Operator (MOVNO) model. In this model a virtual operator, besides the wireless provider "sliced" resources, will be able to use virtualized resources on the optical metro network to seamlessly interface with virtualized resources in the Data Center.

The key motivation behind the CONTENT initiative is that multi-domain and multi-technology problems are not sufficiently investigated. Active research is striving to consolidate frameworks that, with respect to SDN philosophy, will be able to provide performance guarantees, not per segment but end-to-end. New protocols and frameworks, like OpenFlow [3] and OpenContrail [4], are emerging solutions, which seem promising to address these issues. However, global standards for Control Plane to Control Plane communications over diverse technologies are still missing, while the technology to support QoS and performance guarantees in SDN networks is still in an immature state. A main component of the CONTENT architecture is the communication framework that bridges the Control Planes residing in different technology domains. More specifically, it investigates how TSON (Time Shared Optical Network) virtualized optical networks [5], that utilize GMPLS Control Plane mechanics, will be able to communicate with virtualized converged WiFi/LTE networks, in such a way that traffic isolation between the virtual wireless networks is preserved and QoS guarantees are provided end-to-end.

In a nutshell, the CONTENT project is about building an end-to-end virtualization framework that allows for seamless orchestrated on-demand service provisioning across heterogeneous technology domains. The proposed architecture is based on the integration of metro optical-wireless access domains used to interconnect the end users with virtualized computational resources. We propose to leverage a solution that uses standards-based protocols and provides components for network virtualization, SDN-based network control, virtual routing functionalities and will be open to accept analytics engine and Northbound APIs, for enhanced network applications and cloud service orchestration. Two testbeds, the wireless NITOS testbed [6] and the optical TSON [5][7] will be the building blocks of a reference CONTENT testbed used for policies and framework evaluation and for proof of principle demonstrations.

This paper is organized as follows: in section II a description of the the envisioned architecture, the motivation behind CONTENT and related work is made; in section III the end-to-end virtualization strategy is presented; in section IV issues on multi-domain networking and the CONTENT testbed are described; and section VI concludes this work and presents the project's future plans.

II. Vision, Motivation and Related Work

An end-to-end infrastructure facilitates the interconnection of Data Centers (DCs) with fixed and mobile end users through a heterogeneous network, integrating optical metro and wireless access network technologies. The CONTENT project proposes an architecture design that integrates an advanced optical network solution, offering fine (sub-wavelength) switching granularity, with wireless WiFi and LTE access network technologies. To support the IaaS paradigm as well as the diverse and deterministic QoS needs of future Cloud and mobile Cloud services, the proposed architecture adopts the concept of virtualization across all technology domains (see Figure 1) and relies on a common DC infrastructure fully converged with the broadband wireless access and the metro optical network.

CONTENT defines a new stakeholder, the MOVNO, who will offer to the operators the ability to exploit new business opportunities. The MOVNO comprehends converged virtualization of WiFi, LTE, optical metro and IT resources for providing, voice, data and IT services to its customers. The MOVNOs can use the infrastructure of another provider, minimizing their technology systems to billing and customer care, content delivery management and business support systems. The end-user can experience a better quality of service in terms of improved reliability, availability and serviceability, whereas physical infrastructure providers can gain significant benefits through improved resource utilization and energy efficiency, faster service provisioning times, greater elasticity, scalability and reliability of the overall solution. In this context, the adoption of cross-domain and cross-technology virtualization facilitates migration towards a fully converged infrastructure.

The proposed architecture identifies three main actors, the Physical Infrastructure Provider (PIP), the Virtual Operator (VO) and the Service Provider (SP). The PIP owns the physical infrastructure and provides virtual infrastructure resources on top of its physical resources to the VO. The VO in turn has the ability to provide to the SP options for providing new services to its customers. The MOVNO will be able to reduce their CAPEX costs and allow a large percentage of population to be covered by the deployed infrastructure. This will also allow MOVNOs to spend more money on enhancing service provisioning. PIPs will have the opportunity to enhance their

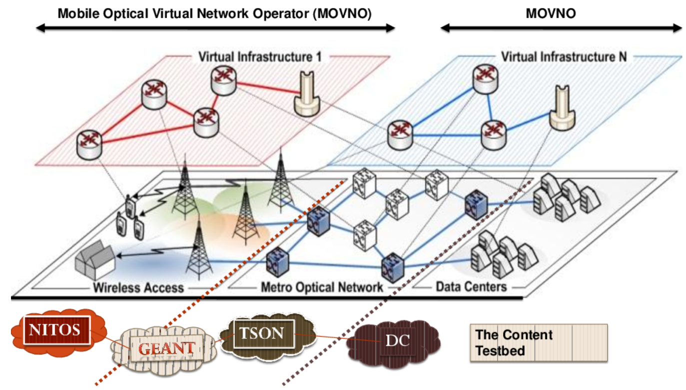

Figure 1. MOVNO and the Content testbed

infrastructure and generate revenues from MOVNOs including custom made solutions; the VO will be in position to provide end-to-end cloud services to its customers with guaranteed wireless broadband connectivity whilst being able to monitor and manage its own virtual network and cloud resources; and the SP will be able to offer to its customers cloud-based application services as part of the more general broadband network access services.

A. Related Work

Virtualization across multi-technology domains has gained significant attention over the last few years and several solutions already exist [8]. Typical examples include the pipe and the hose Virtual Private Network models [9]. Their main difference is that in the pipe model the bandwidth requirement between any two endpoints must be accurately known in advance, whereas the hose-model requires only the knowledge of the ingress and egress traffic at each endpoint having the advantages of ease of specification, flexibility and multiplexing gain [10].

Other similar research efforts have been focused on embedding Virtual Infrastructures (VI) over converged optical data center networks [11], [13] or across multiple substrate networks [8],[14]. A key assumption is that the details of each domain are not communicated to the other domains. Therefore, each infrastructure provider has to embed a particular segment of the VI without any knowledge of how the remaining VIs have already been mapped or will be mapped [8] (referred to as the PolyViNE) with the objective to unilaterally maximize its payoff. A typical example includes the case where an infrastructure provider wishes to optimally allocate its resources without considering the impact of its decision on the other domains.

Although existing approaches (e.g. PolyViNE) target relatively straightforward solutions from an implementation perspective, they may lead to inability in meeting some cloud and mobile cloud service requirements (e.g. end-to-end latency) and inefficiencies in terms of resource requirements, energy consumption etc. [15]. Various aspects of SDN related to transport networking, and issues to be addressed, are presented in [16]. Existing mobile cloud computing solutions allow mobile devices to access the required resources by accessing a nearby resource-rich cloudlet, rather than relying on a distant cloud [17]. In order to satisfy the low-latency requirements of several content-rich mobile cloud computing services such as high definition video streaming, online gaming and real time language translation [18], one-hop, high-bandwidth wireless access to the cloudlet is required. In the case where a cloudlet is not nearby available, traffic is offloaded to a distant cloud such as Amazons Private Cloud, GoGrid [19] or Flexigrid [20]. However, the lack of service differentiation mechanisms for mobile and fixed cloud traffic across the various network segments involved, the varying degrees of latency at each technology domain and the lack of global optimization tools in the infrastructure management and service provisioning make the current solutions inefficient.

III. CONTENT END TO END VIRTUALIZATION APPROACH

In this section we present the challenges in virtualizing the optical and the wireless segments and a short description of the virtualization technologies available. Then we classify the virtualization approaches as a function of the layer where the virtualization happens followed by a discussion on the proposed scheme.

A. Challenges in wireless/optical segment virtualization

Optical Domain Virtualization: Optical network virtualization is defined as the composition of isolated virtual optical networks (VONs) simultaneously coexisting over shared optical physical network infrastructure. Within a VON, the granularity of virtual links is inherited from the switching capabilities supported by the physical devices (e.g. sub-wavelength, wavelength, waveband, or fiber). The two most important characteristics of VONs are isolation and coexistence.

Unlike Layer 2 and Layer 3, optical network resources and transport formats are characterized by their analogue nature. Optical layer constraints such as wavelength/spectrum continuity and physical layer impairments (PLIs) are examples of optical network differentiators. Hence, optical virtualization paradigms should take into account the physical characteristics of optical networks [21]. Linear and non-linear impairments in optical communication systems (noise, dispersion, or cross-talk) can affect the reach and quality of the optical signals, and hence limit the scale and flexibility of virtualization in optical networks. Apart of the wavelength continuity, time slice continuity should be also considered for the time-shared sub lambda networks. This means that if the burst of data is transmitted on a specific period of time, that period (considering the delays) should remain untouched by other signals otherwise contention occurs.

Virtualization techniques for different types of optical networks of different technologies and granularities may vary, according to the node and link characteristics in a particular network. For example, in a wavelength switched network with optical cross connects (OXCs) and Reconfigurable Optical Add-Drop Multiplexors (ROADMs) virtualization approaches are different from networks with sub-wavelength granularity of switching and control. Initial work on optical routers and switches was performed by Qiao [22] in the context of the Optical Burst Switching (OBS) proposal. This effort has lead to some very interesting research and prototype developments, such as an OBS ring testbed in [23], the OPST network solution by [24] and TSON sub wavelength switching [5]. Sub-wavelength switching also is enabled by using Optical Orthogonal Frequency Division Multiplex (OOFDMA) solutions.

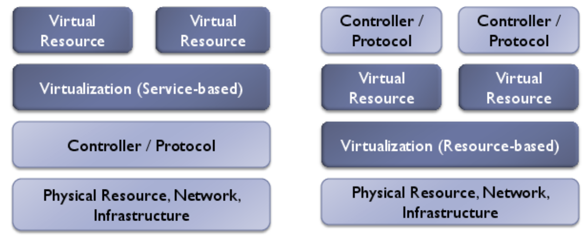

Figure 2. Service Virtualization (left) and Resource Virtualization (right)

The OOFDM sub-wavelength solution offers great flexibility in data rates and modulation levels, however the technology is still in its infancy to be properly deployed in a network environment.

Wireless Domain Virtualization: A number of challenges exist in the wireless domain that do not allow for a straightforward adoption of a specific virtualization strategy, if QoS and service differentiation must be provided, besides traffic isolation between the virtual networks. Like optical network that suffer from impairments and nonlinearity, wireless networks suffer from time varying channel characteristics. There is no well established theory for computing the channel capacity, or providing simple bounds to the maximum information rate based only on the channel impulse response. In addition, unlike wired networks, co-existence of slices with bandwidth-based and resource-based reservations is not straightforward, since the bandwidth achieved by a slice from a given amount of resources varies with the channel quality of the users, the number of users that are sharing the medium and the flow scheduling policies. Other concerning issues include non-location specific addressing, Association, Authentication (AuthN) and synchronization complexity and channel configuration [25].

In the wireless medium, radio resources can be shared and thus virtualized in different ways such as in time, space, and frequency. Splitting and slicing techniques of the wireless medium may use specific time slots (Time Division Multiplexing), space (Space Division Multiplexing), frequencies (Frequency Division Multiplexing) or combinations[26]. Ideas and techniques of virtualizing 802.11 access points are examined in [27] and [28], where we note that in [28] the proposed solution utilizes programmable data plane technologies to guarantee flow differentiation. Work on Base Station (LTE/WiMax) virtualization is presented in [29] and [30]. We also note that OpenFlow [3] and flow-based control offers the potential for end-to-end virtualization in the flow level. For example, OpenFlow technology can be used for slicing wireless backhaul networks and offer communication between controllers that reside in different domains.

1) Virtualization models: In [31], the authors provide a classification of the virtualization approaches as a function of the layer where the virtualization happens. Additionally, in [32], the authors focus on a theoretical approach. They provide a technical classification of the virtualization algorithms presents in the literature as a function of the different parameters optimized by the allocation algorithm, without considering specific technologies of the scenario where the algorithm is applied. From the CONTENT project perspective, and in order to align the virtualization to the software-defined networking requirements, we provide a generic definition of the different types of network virtualization, using as a basis the initial virtualization classification that has been done within the IT realm, with the virtual machines. Therefore, we distinguish between two different types of virtualization, resource-based virtualization (bare-metal), and service-based virtualization. Figure 2 depicts both approaches and their relationship with the SDN control planes.

Resource-based virtualization (or resource virtualization) becomes the most suitable option in order to have the finest feasible granularity in any case, as well as the maximum level of flexibility to control the different virtual resources. However, in terms of implementation, the complexity of this approach is higher, since it needs to completely control the physical resource itself. On the other hand, service-based virtualization loses part of the flexibility, since the virtualization system is located on top of the network service. As an example, the L3VPN, represents a case of service-based virtualization, since the virtualization happens of top of the IP protocol. Granularity is also limited for this approach, since it depends on the northbound interface of the service itself.

B. Virtualization in CONTENT

Applying different virtualization schemes on the different technological segments has different effects. In CONTENT a cross-domain virtualization approach is adopted that aims to overcome these difficulties in order to provision virtual infrastructures composed of resources coming from multiple heterogeneous domains. Based on the functional requirements, for each technology segment there is a specific virtualization system that follows the resource-based approach. On top of each domain-specific system, there exists the cross-domain component, which holds a holistic view of the infrastructure and which is responsible to perform the virtualization over the different domains.

In CONTENT, we adopt TSON as the underlying optical network technology. Its unique features bring more challenges on virtualization. The bursty nature of the signals is more vulnerable to the noise of optical components, such as Erbium Doped Fibre Amplifier (EDFA) and transceivers, and also other impairments, which limits the reach and OSNR. Besides the requirements on the wavelength/spectrum continuity, TSON also requires the time slice continuity. During the process of TSON virtualization, available time slots should also be abstracted and exposed as virtual resources. In TSON, global synchronization among the network nodes is needed due to lack of buffers. Therefore, in each virtual slice the synchronization should also be taken into account and guaranteed. Moreover, in order to flexibly accommodate with user's diverse demands, the virtualization of transceivers in the sub-wavelength network is essential and also a big challenge considering the limitations on the hardware.

In the wireless domain, programmable data planes, software routers and software switching technologies have already been proposed in the literature to perform network slicing or provide other network functionalities. Such as technology is the Click Modular router [33]. With respect to the SDN framework where data, control, management and service planes must be clearly separated, in order to provide a slicing mechanism over a converged WiFi/LTE network, a combination of technologies like OpenFlow in combination with programmable

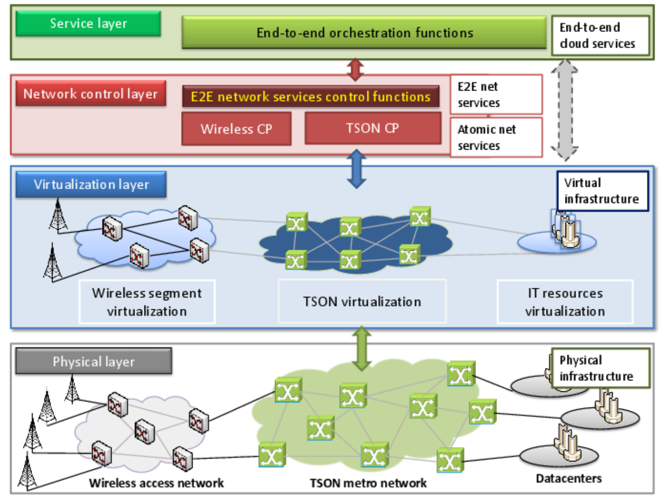

Figure 3. CONTENT Layering Structure

data plane technologies, is adopted. These technologies can be used to perform traffic shaping and slice isolation from the edge network to the wireless backhaul/access network. We note that in works like [28], programmable data plane technologies, besides providing potential technologies for traffic shaping and policy enforcement in the edge network, they are also used to virtualize a single Access Point itself. The combination of OpenFlow (or Openflow like) technologies with programmable data planes gives us the ability to perform flow isolation, traffic shaping and QoS policies, for both the ingress traffic that comes from the optical network and the egress traffic flows that span from the wireless access networks through the wireless backhaul to the wireless edge.

C. CONTENT Control Plane

The CONTENT Control Plane (CP) operates on top of the hybrid virtual infrastructures composed of wireless and TSON domains and virtual data centers. Following the SDN paradigm and the four planes approach (data, virtualization, control & management, service layers) it is fully decoupled from the physical layer and interacts with the underlying virtual infrastructure through open and vendor-independent interfaces.

The CONTENT control plane relies on an SDN-like approach with functional entities operating as controllers of the virtualized networks. These entities provide management, control and analytics procedures through three main internal functions: *Configuration functions* responsible for translating the high-level data model into a lower level form suitable for interacting with network elements; *Control functions* responsible for propagating this low level state to and from network elements; *Measurement functions* responsible for capturing real-time data from network elements, abstracting it and presenting it in a form suitable for administration purposes or to consume by applications.

The CONTENT CP is structured in two hierarchical layers. The lower layer is related to the control of technology-specific domains through dedicated CPs responsible for providing connectivity within the wireless and the TSON networks respectively (see also Figure 3). The upper layer uses the functionalities exposed by wireless and TSON CPs in order to provide end-to-end and multi-layer network services to interconnect the data centres and the mobile users, as required by the cloud orchestration functions deployed on top of the architecture.

In this context the wireless CP handles the connectivity in the wireless domain and it manages the intra- and inter-technology handovers. Also it is responsible for the configuration of the Ethernet backhaul network, which is composed of programmable switches, through an SDN controller. On the metro side, the TSON CP implements all the basic functionalities to provide intra-domain network services with flexible capacity (through the support of the sub-wavelength bandwidth granularity), dynamicity and resiliency guarantees. Both the wireless and the TSON CPs expose at the north-bound side an interface to access the basic service primitives implemented in each domain, so that enhanced network control applications with end-to-end scope can be easily developed on top of them.

These applications will operate on the entire virtual infrastructure through the abstracted services exposed by the wireless and TSON CP north-bound interface, without dealing with the specific network-oriented semantics, internal technologies and protocols implemented in each domain. These end-to-end network applications manage the cross-domain connectivity taking into account the requirements of the cloud services in mobile environments. As an example, they can dynamically modify the resource allocation in the two network segments depending on different parameters or network conditions, such as the real-time traffic load in the metro between the data centers, the network utilization, the scheduled cloud services, the characteristics of the traffic generated by the mobile users.

IV. The CONTENT Testbed

The CONTENT testbed aims to evaluate the integrated solution and the proposed CONTENT framework, though performance metrics that quantify the efficiency of ubiquitous fast broadband access, the flow isolation, the QoS and QoE that the end users enjoy and the signaling overhead. As shown in Figure 4 the wireless segment of the CONTENT testbed is the NITOS testbed [6] and the Optical segment is the TSON testbed [5]. Custom virtualized infrastructure or cloud based resources will be used to support the necessary back-end DC infrastructure.

A. Description of the CONTENT Testbed

TSON, is designed and implemented as a novel time multiplexing metro network solution, which uses multiple wavelength for transportation of time multiplexed data, while offering dynamic connectivity with fine granularity of bandwidth in a contention-less manner. TSON data plane is augmented with an extended GMPLS control plane for signaling and reservation of network resources for establishing lightpaths. At the heart of the control plane, two centralized routing and resource allocation drivers of Path Computation Element (PCE) and Sub Lambda Assignment Engine (SLAE) manage the network resources for the incoming lightpath requests and the existing ones. TSON uses high performance FPGA platforms of XILINX V6HT for high speed 10GE frame processing and EOC conversion of the frames into bursts of traffic at the edge nodes, whilst TSON core nodes use the high speed FPGA platforms to transparently bypass the traffic,

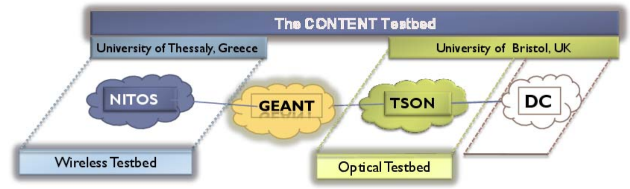

Figure 4. The CONTENT Testbed

via controlled advanced PLZT based fast switches (10ns). Therefore the TSON data plane is capable of offering up to 10 Gb/s data rate over each wavelength with the granularity of 100 Mb/s. The interconnection and communication between the TSON data plane nodes and control plane elements take place through web service interfaces; this approach decouples the functionalities of the two layers from each other, therefore enhances the manageability and extensibility of the network.

The wireless segment of the CONTENT testbed architecture is provided by an heterogeneous wireless testbed called NITOS, offered by University of Thessaly (UTH) in Greece. NITOS testbed is a FIRE facility and the current deployment includes more than 50 WiFi enabled nodes, a wide-range WiMAX/LTE Base Station, a WiMAX/LTE experimental network and an extended wireless sensor network. Furthermore, wireless nodes feature USRP boards that are software defined radio systems. Through these, a user can build from scratch PHY layer interfaces as well as the needed MAC layer. NITOS is based on open-source software and drivers allowing experimenters to modify any layer of the network protocol stack, while the control and management of the testbed is based on the OMF/OML framework. In addition, NITOS offers experimentation capabilities using fully programmable OpenFlow switches and allows for the design, build and evaluation of SDN scenarios that span both the wireless access and the wireless backhaul network. This is very important for the aims of the CONTENT project where a convergent WiFi/LTE network is examined.

B. Physical interconnection between the different testbeds

The physical interconnection between the wireless testbeds (NITOS in Volos, Greece) and the optical testbed (TSON in Bristol, UK) is made through the GEANT network [34]. VLAN-based connections are currently used, while the GEANT network robustness guarantees the delivery of speed and services to enable greater collaboration. CONTENT project will exploit all the advances and updates on the GEANT network and the services offered, like for example the Bandwidth on Demand service (BoD). The involved NRENs (GRNET in Greece and JANET in UK), are already carrying out pilots for the introduction of the BoD service. With the BoD service, multi-domain provisioning is offered to users that can make reservations in advance, or provision in real-time point-to-point circuits with the capacity needed.

The interconnection of the two technology domains includes interfacing of the testbeds both at the physical layer and also at the logical and software layer. For the physical interconnections, the two technology domains need to support common protocols to enable seamless traffic interchange between them. Ethernet 802.1Q VLAN tagging and techniques like QnQ tagging and EVPN have been considered to interface the two networks, which enable inter domain traffic transport, are compatible with other Ethernet based technologies (thus enhancing the scalability of the network), and facilitate the establishment of end-to-end virtual networks as services.

For connecting the two testbeds (TSON and NITOS), data rate characteristics of both testbeds are considered. TSON as the metro backhaul solution provides from 100 Mbps up to 20 Gbps connectivity for a variable number of LTE/WiFi cells, depending on the rates required at each access point and the level of aggregation to be performed. The rate matching between the two domains can cause inefficiency in using the available bandwidth. However TSON as the sub-wavelength switching technology with fine data rate granularities and dynamic setup and tear-down of connections can adopt to the requirements in the access quite swiftly.

The control of the physical setup and interconnection of the two networks needs to be addressed in the control layer. This way the allocation/optimization mechanisms will have global understanding of the resources, can have access to these resources (for e.g. nodes, ports, spectral, time units) and so create optimum end-to-end slices of the network. In this regard, the TSON system exposes the information in its edge nodes out to the controller through a web service interface, while NITOS does it via both a web based interface and SFA/mySlice [35]. This enables the controller to take advantage of the programmability and flexibility of the TSON and NITOS technologies to establish optimum paths stretching from users in wireless access areas to the distributed data centers through the metro/core network.

V. Concluding remark: CONTENT project and the SDN framework

The CONTENT layered architecture adopts the SDN paradigm since the control functions are separated from the physical devices and operate on top of programmable entities through open interfaces. The SDN approach is reflected in several architectural choices, from the introduction of virtualization mechanisms between the physical infrastructure and the control plane, up to the internal design of the control plane itself.

The virtualization layer in CONTENT generates isolated virtual infrastructures and provides an abstracted view of the different domains, hiding the vendor-dependent details while exposing the full technology capabilities through a unified interface. This design is aligned with the SDN approach, where control plane functions operate on top of programmable devices, in this specific case represented by the virtual network resources (TSON nodes in the metro networks, base stations and backhaul L2 switches in the wireless segment). The interfaces enabling the control of the virtual nodes follow the same design principles of SDN protocols, like OpenFlow. In particular, they are conceived to support synchronization, configuration and monitoring actions through commands and asynchronous notifications expressed in an open language, independent from the specific protocols and interfaces implemented on the corresponding physical devices.

At the control plane level, on top of the virtual infrastructures, the wireless and TSON virtual domains are managed through dedicated lower-layer control frameworks: the wireless and the TSON CPs. These are customized to deal with the specific technology constraints (e.g. management of sub-wavelength switching capability in the TSON metro network) and implement a set of basic intra-domain mechanisms, like connectivity setup and monitoring. With reference to SDN architectures, the local CPs act as a sort of macro SDN controllers responsible of the basic network functions in each segment. On the other hand, all the end-to-end enhanced functionalities for on-demand provisioning, maintenance, resiliency, monitoring and dynamic re-configuration of network connectivity in support of mobile cloud services are delegated to upper layer network applications.

It is clear that this internal organization of the CONTENT CP is strongly based on the programmability concept, that represents one of the pillars of SDN architectures. In the CONTENT architecture, specialized network applications implement cloud-oriented functions over generic controllers offering simpler network service primitives. In other terms, the controllers provide a further level of service abstraction on top of the resource virtualization layer, while the dynamic programmability of the network is handled through the upper layer applications, tightly integrated with the orchestration functions at the cloud service layer.

VI. Conclusions & Future work

This paper presented the CONTENT project approach for building virtual networks over converged multi-domain network infrastructures. The proposed architecture is based on a sample environment with distributed data center sites interconnected through a TSON metro network, while a hybrid wireless network, based on LTE and Wi-Fi technologies, provides the access for mobile cloud users. The joint orchestration and cooperation between cloud and network services through an enhanced cloud-to-network interface, will allow for the strict requirements imposed by cloud services to be met, in terms of dynamicity, resiliency, end-to-end QoS and performance.

The CONTENT testbed has been designed to validate the proposed architectural solution using relevant cloud-based applications. Trials and experimental tests are planned during the next year of the project, to evaluate the performance of different types of cloud services on top of a CONTENT-enabled virtual infrastructure. A "Mobile broadband services by MOVNO" scenario will be the pilot use case for which the CONTENT testbed will provide the virtual infrastructure used for the provisioning of cloud services.

The project consortium closely follows the advancements in the cloud computing technology, in the SDN framework and the Network Functions Virtualizaton (NFV) that are related to the project's goals and can potentially be exploited by the proposed framework.

Acknowledgements

This work has been funded by the EU Project 318514 "Convergence of wireless Optical Network and IT rEsourcesiN support of cloud services" (CONTENT). We thank all our project partners for their valuable contributions.

References

[1] M. Armbrust, et al., “A view of cloud computing”. Commun. ACM 53, 4 (April 2010)

[2] CONTENT Project: http://content-fp7.eu/

[3] OpenFlow Switch Specification, version 1.3.1, Open Networking Foundation, September 2012

[4] The OpenContrail Project: www.opencontrail.org/

[5] G. Zervas, J. Triay, N. Amaya, Y. Qin, C. Cervell-Pastor, and D.Simeonidou, Time Shared Optical Network (TSON): a novel metro architecture for flexible multi-granular services, Optics Express, 19 (26).

[6] NITOS test-bed, http://nitlab.inf.uth.gr/NITlab/

[7] TSON testbed http://www.bris.ac.uk/engineering/research/pho/hpn/

[8] M. Chowdhury, F. Samuel, R. Boutaba, PolyViNE: policy-based virtual network embedding across multiple domains, in Proc. of ACM SIGCOMM, workshop on Virtualized infrastructure systems and architectures (VISA ’10), 2010.

[9] J. Chu, Chin-Tau Lea, New architecture and algorithms for fast construction of hose-model VPNs, IEEE/ACM Trans. Netw. 16, 3, 670-679, June 2008.

[10] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. V. der Merwe, A flexible model for resource management in virtual private networks, in proc. ACM Sigcomm , San Diego, California, USA, August 1999.

[11] D.A. Schupke, Multi-Layer and Multi-Domain Resilience in Optical Networks, Proceedings of the IEEE, vol. 100, no. 5, May 2012.

[12] A. Tzanakaki et. al., Planning of Dynamic Virtual Optical Cloud Infrastructures: The GEYSERS approach, IEEE Communications Magazine, Special Issue on Advances in Network Planning, Jan. 2014.

[13] 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.681691, Sep.2012.

[14] D. Dietrich, A. Rizk, P. Papadimitriou, AutoEmbed: Automated Multi-Provider Virtual Network Embedding, Demo, ACM SIGCOMM 2013

[15] A. Tzanakaki, M.P. Anastasopoulos, G. Zervas, B. Rofoee, R. Nejabati and D. Simeonidou, Virtualization of Heterogeneous Wireless-Optical Network and IT infrastructures in support of Cloud and Mobile Cloud Services, IEEE Communications Magazine, Aug. 2013.

[16] Mcdysan, D., “Software defined networking opportunities for transport”, Communications Magazine, IEEE , vol.51, no.3, pp.28,31, March 2013.

[17] M. Satyanarayanan, P. Bahl, R.Caceres and N. Davies, The Case for VM-Based Cloudlets in Mobile Computing. IEEE Pervasive Computing 8 (4), pp.14-23, Oct. 2009.

[18] K. Mun, Mobile Cloud Computing Challenges, TechZine Magazine, http://www2.alcatel-lucent.com/techzine/mobile-cloud-computing-challenges/

[19] http://www.gogrid.com/

[20] http://www.flexiscale.com/

[21] S. Peng, R. Nejabati, D. Simeonidou, Impairment-Aware Optical NetworkVirtualization in Single-Line-Rate andMixed-Line-Rate WDM Networks, IEEE/OSA Journal ofOptical Communications and Networking, 5(4), pp. 283-293, April 2013.

[22] C. Qiao, M.Yoo, Optical burst switching (OBS) - a new paradigm for an optical Internet, Journal of High Speed Networks - Special issue on optical networking

[23] Ning Deng; et al ”A novel optical burst ring network with optical-layer aggregation and flexible bandwidth provisioning,” Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011 and the National Fiber Optic Engineers Conference , vol., no., pp.1-3, 6-10 March 2011

[24] Verismaivx 8000, Optical Packet Switch and Transport Intune Networks, 2011

[25] Hong-jiang Lei, et al., “Survey of multi-channel MAC protocols for IEEE 802.11-based wireless Mesh networks”, The Journal of China Universities of Posts and Telecommunications, v.18, Issue 2, April 2011,

[26] Aljabari G., Eren E., Virtualization of Wireless LAN Infrastructures, thThe 6 IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications, 2011

[27] Aljabari G., Eren E., “Virtual WLAN: Extension of Wireless Networking into Virtualized Environments”, International Journal of Computing Research Institute of Intelligent Computer Systems, 2010

[28] Bhanage, G.; Vete, D.; Seskar, I.; Raychaudhuri, D.; , SplitAP: Leveraging Wireless Network Virtualization for Flexible Sharing of WLANs, IEEE GLOBECOM, 2010

[29] R Kokku, R Mahindra, H Zhan et al,Remote Virtualization of a Cellular Basestation, US Patent , US 2012 0002620 A1

[30] Bhanage G., Seskar I., and Raychaudhuri D, “A virtualization architecture for mobile WiMAX networks.” SIGMOBILE Mob. Comput. Commun. Rev. 15, 4 (March 2012)

[31] Chowdhury, N.M.M.K.; Boutaba, R., ”Network virtualization: state of the art and research challenges,” Communications Magazine, IEEE , vol.47, no.7, pp.20,26, July 2009

[32] Fischer, A.; Botero, J.; Beck, M.; De Meer, H.; Hesselbach, X., “Virtual Network Embedding A Survey,” Communications Surveys & Tutorials, IEEE , vol.PP, no.99, pp.1,19 ,2013

[33] Kohler, Eddie, et al. ”The Click modular router.” ACM Transactions on Computer Systems (TOCS) 18.3 (2000): 263-297.

[34] GEANT network, http://www.geant.net

[35] http://sfawrap.info/