

Chapter 14

Converged Wireless Access/Optical Metro Networks in Support of Cloud and Mobile Cloud Services Deploying SDN Principles

Anna Tzanakaki, Markos Anastasopoulos, Bijan Rofoee, Shuping Peng, George Zervas, Reza Nejabati, Dimitra Simeonidou, Giada Landi, Giacomo Bernini, Roberto Monno, Nicola Ciulli, Gino Carrozzo, Kostas Katsalis, Thanasis Korakis, Leandros Tassiulas, Georgios Dimosthenous, Dora Christofi, Jordi Ferrer Riera, Eduard Escalona, Jacopo Pianigiani, Dirk Van Den Borne and Gert Grammel

Abstract This chapter proposes a next-generation ubiquitous converged infrastructure to support cloud and mobile cloud computing services. The proposed infrastructure facilitates efficient and seamless interconnection of fixed and mobile end users with computational resources through a heterogeneous network integrating optical metro and wireless access networks. To achieve this, a layered architecture which deploys cross-domain virtualization as a key technology is presented. The proposed architecture is well aligned and fully compliant with the Open Networking Foundation (ONF) Software-Defined Networking (SDN) architecture and takes advantage of the associated functionalities and capabilities. A modeling/simulation framework was developed to evaluate the proposed architecture and identify planning and operational methodologies to allow global

A. Tzanakaki (✉) · M. Anastasopoulos · B. Rofoee · S. Peng · G. Zervas · R. Nejabati · D. Simeonidou

High Performance Networks Group, University of Bristol, England, UK
e-mail: anna.tzanakaki@bristol.ac.uk

G. Landi · G. Bernini · R. Monno · N. Ciulli · G. Carrozzo
Nextworks, Busto Arsizio, Italy

K. Katsalis · T. Korakis · L. Tassiulas
University of Thessaly, Thessaly, Greece

G. Dimosthenous · D. Christofi
Primetel, Limassol, Cyprus

J.F. Riera · E. Escalona
I2CAT, Barcelona, Spain

J. Pianigiani · D. Van Den Borne · G. Grammel
JUNIPER, California, USA

optimization of the integrated converged infrastructure. A description of the tool and some relevant modeling results are also presented.

14.1 Introduction

Big data, Internet of Things, and Content Delivery Services are driving the dramatic increase of the global Internet traffic, which is expected to exceed 1.6 zettabytes by 2018. These new and emerging applications collect, process, and compute massive amounts of data, and require efficient management and processing of very large sets of globally distributed unstructured data. To support these services, the concept of cloud computing, where computing power and data storage are moving away from the end user to remote computing resources, has been introduced. To enable this, there is a clear need for a next-generation ubiquitous converged infrastructure. This infrastructure needs to facilitate the interconnection of data centers (DCs) with fixed and mobile end users.

In this environment, optical network solutions can be deployed to interconnect distributed DCs, as they provide abundant capacity, long reach transmission capabilities, carrier-grade attributes, and energy efficiency [1], while spectrum-efficient wireless access network technologies such as Long Term Evolution (LTE) can be effectively used to provide a variety of services to a large pool of end users. In this context, interconnection of DCs with fixed and mobile end users can be provided through a heterogeneous converged network integrating optical metro and wireless access network technologies.

However, to address limitations and inefficiencies of current optical and wireless network solutions optimized for traditional services, novel approaches that enable cost-efficiency and utilization efficiency through multi-tenancy and exploitation of economies of scale need to be developed. Toward this direction, infrastructure virtualization that allows physical resources to be shared and accessed remotely on demand has been recently proposed as a key enabling technology to overcome the strict requirements in terms of capacity sharing and resource efficiency in cloud and mobile cloud environments. Given the heterogeneous nature of the infrastructures suitable to support cloud and mobile cloud services described above, it is clear that cross-technology and cross-domain virtualization is required. In this context, the introduction of mobile optical virtual network operators (MOVNOs) facilitates the sharing of heterogeneous physical resources among multiple tenants and the introduction of new business models that suit well the nature and characteristics of the Future Internet and that enable new exploitation opportunities for the underlying physical resources.

Existing mobile cloud computing solutions commonly allow mobile devices to access the required resources by accessing a nearby resource-rich cloudlet, rather than relying on a distant “cloud” [2]. 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 [3], one-hop,




Fig. 14.1 **a** *Cloudlet* approach: micro-DCs in the wireless access to support mobile cloud traffic and large DCs in the core to support fixed cloud traffic, **b** *CONTENT* proposed approach: shared DC infrastructure fully converged with the broadband wireless access and the metro optical network

high-bandwidth wireless access to the cloudlet is required. In the case where a cloudlet is not available nearby, traffic is offloaded to a distant cloud such as Amazon's Private Cloud, GoGrid [4], or Flexigrid [5]. 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.

In response to these observations, this chapter is focusing on a next-generation ubiquitous converged network infrastructure [6]. The infrastructure model proposed is based on the Infrastructure as a Service (IaaS) paradigm and aims at providing a technology platform interconnecting geographically distributed computational resources that can support a variety of cloud and mobile cloud services (Fig. 14.1).

The proposed architecture addresses the diverse bandwidth requirements of future cloud services by integrating advanced optical network technologies offering fine (sub-wavelength) switching granularity with a state-of-the-art wireless access network based on hybrid LTE and Wi-Fi technology, supporting end user mobility. However, in these converged cloud computing environments, management and control information needs to be exchanged across multiple and geographically distributed network domains, causing increased service setup and state convergence latencies [7]. Given the complexity of the overall infrastructure, additional challenges that need to be addressed include the high probability of management/control data loss [8], often leading to violation of quality-of-service (QoS) requirements as well as impractical and inefficient collection of control data and management of network elements due to scalability issues [9]. Finally, these are highly dynamic environments where service requests can have big and rapid variations that are not known in advance.

To address the challenge of managing and operating this type of complex heterogeneous infrastructures in an efficient manner, Software-Defined Networking (SDN) and network function virtualization have been recently proposed [10, 11]. In SDN, the control plane is decoupled from the data plane and is moved to a logically centralized controller that has a holistic view of the network. At the same time, to enable sharing of the physical resources the concept of virtualization across heterogeneous technology domains has been recently proposed. Taking advantage

of the SDN concept and the benefits of cross-technology virtualization solutions, we propose and present a converged network infrastructure and architecture able to efficiently and effectively support cloud and mobile cloud services. The benefits of the proposed approach are evaluated through a modeling tool that allows the comparison of the performance of the proposed solution compared to alternative state-of-the-art solutions, extending the work presented in [6]. Our modeling results identify trends and trade-offs relating to resource requirements and energy consumption levels across the various technology domains involved that are directly associated with the services' characteristics.

The remaining of this chapter is structured as follows: Sect. 14.2 describes the relevant state-of-the-art. Section 14.3 provides a functional description together with a detailed structural presentation of the proposed architecture. This includes the details of the individual layers involved as well as a description of the interaction between the different layers. Section 14.4 includes a discussion on the modeling framework that is being developed with the aim to evaluate the architecture and propose optimal ways to plan and operate it, while Sect. 14.5 summarizes the conclusions.

14.2 Existing Technology Solutions Supporting Cloud and Mobile Cloud Services

14.2.1 Physical Infrastructure Solutions Supporting Cloud Services

14.2.1.1 State of the Art in Metro Optical Network Solutions

Recent technological advancements in optical networking are able to provide flexible, efficient, ultra-high data rate, and ultra-low-latency communications for DCs and cloud networks. Optical transmission solutions offering high data rate and low latency have been reported as field trial deployments of 400 Gb/s channels [12], while research on 1 Tb/s per channel is already in progress [13, 14]. However, beyond high capacity, optical networks need to address the requirement for fine granularity to enable efficient utilization of network resources both for service providers and end users. Optical orthogonal frequency division multiplexing (OFDM) [15], optical packet switched networks [16], and optical burst switching [17] technologies are examples of providing such fine granularity. These advanced novel optical network technologies offer the flexibility and elasticity required by the diverse, dynamic and uncertain cloud and mobile cloud services.

Time Shared Optical Networks (TSON) is a dynamic and bandwidth-flexible sub-lambda networking solution that has been introduced to enable efficient and flexible optical communications by time-multiplexing several sub-wavelength connections in a WDM-based optical network [17]. TSON allocates time slices over

wavelengths to satisfy the bandwidth requirements of each request, and statistically multiplexes them exploiting each wavelength channel capacity to carry data for as many requests as possible. In TSON, the electronic processing and burst generation is performed at the edge of the network, and the data of several users are transferred transparently in the core between TSON edge nodes in the form of optical bursts. The proposed infrastructure deploys TSON as it can be used to efficiently interconnect wireless networks with datacenters exploiting its fine granularity to accommodate the granularity mismatch inherent between the wireless and optical network domains.

14.2.1.2 State of the Art in Wireless Technologies Solutions

High-speed wireless access connectivity is provided by three prominent technologies: cellular LTE networks, WiMAX, and Wi-Fi. These technologies vary across a number of distinct dimensions [18], including spectrum, antenna characteristics, encoding at the physical layer, sharing of the available spectrum by multiple users, as well as maximum bit rate and reach. Femtocells appear to be a promising solution as they allow frequent spectrum reuse over smaller geographical regions with easy access to the network backbone. On the other hand, Wi-Fi networks are readily available and are easy to install and manage [19]. In this study, the proposed architecture relies on a converged 802.11 and 4G—Long Term Evolution (LTE) access network used to support cloud computing services.

14.2.2 Infrastructure Management

Heterogeneous infrastructure management has been already addressed by several research projects and commercial systems. Traditionally, infrastructure management is vertically separated, i.e., each technological segment has its own management system. Thus, the management of the different essential operation components (e.g., policies, processes, or equipment) was performed in a per-domain basis. As such, there are clearly differentiated network management systems and cloud management systems.

Network management has been addressed following two approaches depending on the context and requirements of the network owners. On the one hand, centralized management assumes the existence of a single system that controls the entire network of elements, each of them running a local management agent. On the other hand, distributed management approaches introduce the concept of management hierarchies, where the central manager delegates part of the management load among different managers, each responsible for a segment of the network.

Most of the network infrastructure management systems are proprietary, since there are no common interfaces at the physical level. Each management solution depends on each equipment vendor. However, ISO Telecommunications

Management Network defines the Fault, Configuration, Accounting, Performance, Security (FCAPS) model [20], composed of fault, configuration, accounting, performance monitoring/management, and security components, which is the global model and framework for network management. Apart from the general network management model in [20], there are some standard network management protocols. Most relevant are the SNMP and NetCONF. SNMP is standardized by the IETF in RFCs 1157, 3416, and 3414. NetCONF is also an IETF network management protocol defined in RFC 4741 and revised later in RFC 6241. Its specifications mainly provide mechanisms to install, handle, and delete network device configurations. However, as the network grows, SNMP-based solutions that have limited instruction set space and rely on the connectionless UDP protocol, are inefficient. In the proposed cloud computing systems, information needs to be exchanged across multiple, heterogeneous, and geographically distributed network domains, causing increased service setup and state convergence latencies. At the same time, in such highly complex environments the probability that management/control data are lost is high, leading often to violation of quality-of-service (QoS) requirements. These challenges introduce a clear need for alternative management solutions.

To address these challenges in an efficient manner and also support multi-tenancy requirements inherent in cloud infrastructures, optical network virtualization becomes a key technology that enables network operators to generate multiple coexisting but isolated virtual optical networks (VONs) running over the same physical infrastructure (PI) [21]. Optical network virtualization in general adopts the concepts of abstraction, partitioning, and aggregation over node and link resources to realize a logical representation of network(s) over the physical resources [22]. Virtualization of optical networks is one of the main enablers for deploying software-defined infrastructures and networks in which, independently of the underlying technologies, operators can provide a vast array of innovative services and applications with a lower cost to the end users. Network management solutions focusing on optical network resources include the EU project GEYSERS (<https://www.geysers.eu>) which introduced the Logical Infrastructure Composition Layer (LICL) [23]. The LICL is a software middleware responsible for the planning and allocation of virtual infrastructures (VIs) composed of virtualized network and IT resources.

On the other hand, IT/server virtualization has reached its commercial stage, e.g., VMWare vSphere. In order to better serve the dynamic requirements that arise from the IT/data center side, coordinated virtualization of both optical network and IT resources in the data centers is desired. Joint allocation of the two types of resources to achieve optimal end-to-end infrastructure services has been investigated in the literature, but the reported work concentrates mostly on optical infrastructures supporting wavelength switching granularity [1, 24]. However, solutions addressing optical network technologies supporting sub-wavelength granularity, such as the CONTENT proposed approach [6], are still in their early stages.

In wireless networks, significant management challenges exist because of the system complexity and a number of inter-dependent factors that affect the wireless

network behavior. These factors include traffic flows, network topologies, network protocols, hardware, software, and most importantly, the interactions among them [25]. In addition, due to the high variability and dependency on environmental conditions, how to effectively obtain and incorporate wireless interference into network management remains an open problem. Similar to the optical domain, multi-tenancy in wireless networks can be provided through virtualization. Virtualization and slicing in the wireless domain can take place in the physical layer, data link layer (with virtual Mac addressing schemes and open-source driver manipulation), or network layer (VLAN, VPN, label switching). In [26], the SplitAP architecture is proposed in order for a single wireless access point to emulate multiple virtual access points. Clients from different networks associate corresponding virtual access points (APs) through the use of the same underlying hardware. The approach of creating multiple virtual wireless networks through one physical wireless LAN interface, so that each virtual machine has its own wireless network, is proposed in [27]. In [28], the CellSlice is proposed, focusing on deployments with shared-spectrum RAN sharing. The design of CellSlice is oblivious to a particular cellular technology and is equally applicable to LTE, LTE-Advanced, and WiMAX. CellSlice adopts the design of NVS for the downlink but indirectly constrains the uplink scheduler's decisions using a simple feedback-based adaptation algorithm.

The decoupling of the control from the data plane, supported through the emergence of the SDN paradigm, introduces the need for unified southbound interfaces and thus an abstraction layer that in turn helps to simplify network operation and management. In this context, suitable frameworks have been developed including OpenNaaS (<http://www.opennaas.org>) that provides a common lightweight abstracted model at the infrastructure level enabling vendor-independent resource management. Similarly, the OpenDaylight (<http://opendaylight.org>) platform aims at providing a software abstraction layer of the infrastructure resources that allows geographically distributed network domains to be unified and operated as a single network [29].

Management of IT resources in the context of a converged network/IT infrastructure is commonly performed using modern platforms based on web services (e.g., REST APIs). A typical example is the Open Grid Forum (OGF) Open Cloud Computing Interface (OCCI) [30], which comprises a set of open community-led specifications delivered through the OGF. Well-known cloud management system (CMS) implementations of OCCI-compliant application programming interfaces (APIs) include OpenStack, OpenNebula, and CloudStack.

14.2.3 Service Provisioning

The service provisioning and orchestration of IT resources (computing and storage) located in geographically distributed DCs, seamlessly integrated with inter-DC networking in support of a variety of cloud services, has been addressed in several

European projects. A number of technical solutions have been investigated and proposed for a variety of scenarios spanning from multi-layer architectures enabling the inter-cooperation between cloud and network domains, up to procedures, protocols, and interfaces allowing integrated workflows to support delivery and operation of joint cloud and network services.

The FP7 GEYSERS project has developed a framework for on-demand provisioning of inter-DC connectivity services, specialized for cloud requirements, over virtual optical infrastructures [1]. Following similar inter-layer approaches, some IETF drafts [31] have proposed cross-stratum solutions for the cooperation between application (service) and network layers in path computation for inter-DC network services, potentially combined with stateful path computation element (PCE) mechanisms [32]. Other relevant research efforts include the FP7 projects SAIL and BonFIRE.

In terms of infrastructure service provisioning, following the Network as a Service (NaaS) paradigm, the OpenNaaS framework has been recently proposed. It is an open-source framework that provides tools to manage the different resources present in any network infrastructure (including virtual infrastructures). The tools enable different stakeholders to contribute and benefit from a common software-oriented stack for both applications and services. Currently, OpenNaaS supports on demand both network and cloud service provisioning.

Emerging cloud applications such as real-time data backup, remote desktop, server clustering, etc., require more traffic being delivered between DCs. Ethernet remains the most used technology in the DC domain, but the end-to-end provisioning of Ethernet services between remote DC nodes poses a big challenge. Central to this capability is the SDN-based service and network orchestration layer [33]. This layer (a) is aware of the existing background and existing practices and (b) applies new SDN principles to enable cost reduction, innovation, and reduced time-to-market of new services, while covering multi-domain and multi-technology path—packet networks. This layer provides a network-wide, centralized orchestration. This high level, logically centralized entity exists on top of and across the different network domains and is able to drive the provisioning (and recovery) of connectivity across heterogeneous networks, dynamically and in real time.

14.3 Proposed Converged Network Architecture

14.3.1 *Vision and Architectural Approach*

As already discussed in the introduction, the infrastructure model proposed in this chapter is based on the IaaS paradigm. To support the IaaS paradigm, physical resource virtualization plays a key role and is enabled by using a cross-domain infrastructure management layer in the proposed architectural structure. Connectivity services are provided over the virtual infrastructure slices, created by

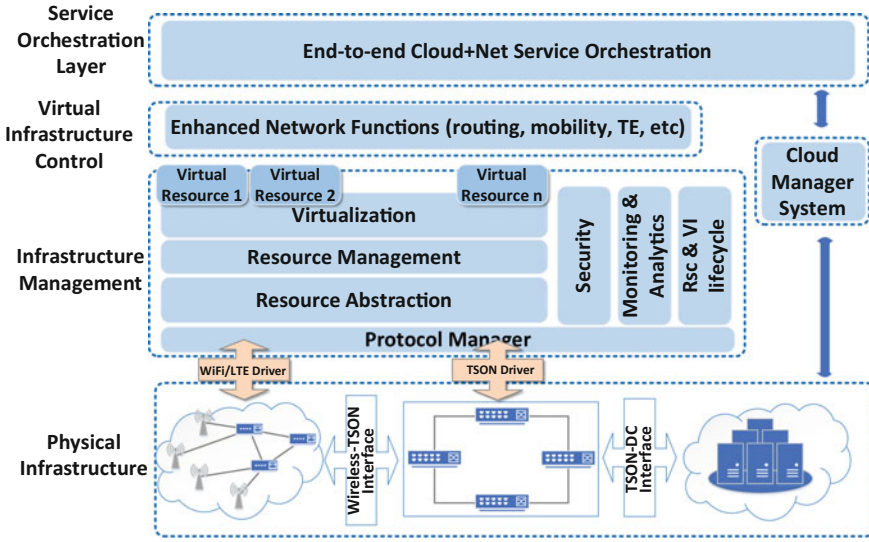


Fig. 14.2 Overall cross-domain, cross-layer architecture

the infrastructure management layer, through the virtual infrastructure control layer. Integrated end-to-end network, cloud, and mobile cloud services are orchestrated and provisioned through the service orchestration layer. The details of the proposed cross-domain, cross-layer architecture are discussed in Fig. 14.2.

The **physical infrastructure layer (PIL)** comprises an optical metro network and a hybrid Wi-Fi/LTE wireless access system. The optical metro network supports frame-based sub-wavelength switching granularity and is used to interconnect geographically distributed DCs.

The **infrastructure management layer (IML)** is responsible for the management of the network infrastructure and the creation of virtual network infrastructures over the underlying physical resources. This involves several functions, including resource representation, abstraction, management, and virtualization across the heterogeneous TSON and the wireless access network domains. An important feature of the functionalities supported is orchestrated abstraction of resources across domains, involving information exchange and coordination across domains. The IML functional architecture can be mapped over the SDN architecture as defined by the ONF. The IML comprises mainly the infrastructure segment of the SDN architecture and the associated management.

The **virtual infrastructure control layer (VICL)** is responsible to provision IT and (mobile) connectivity services in the cloud and network domains, respectively. The VICL is structured in two main levels of functionality. The lower-layer functions, implemented within one or more SDN controllers, deal with the details of each technology deployed at the data plane in the wireless access and TSON-based metro network segments, as exposed in the associated virtual resources through the

IML. These functions provide elementary services with a per-domain scope and are specialized to operate on top of specific (virtual) technologies. On top of the basic services offered by the SDN controller(s), further enhanced functionalities are developed in order to operate the entire heterogeneous infrastructure in a unified manner. These functions manage more abstracted entities (e.g., “per-domain TSON paths”) that summarize the results of the configuration actions performed by the lower-level functions directly on the virtual resources. The VIACL architecture is fully compliant with the SDN architecture defined by the ONF.

The *service orchestration layer (SOL)* is responsible for the converged orchestration of cloud and network services, with the automated and cloud-aware setup and adjustment of the user-to-DC and inter-DC network connectivity. Furthermore, it is used for the composition and delivery of multi-tenant chains of virtualized network functions from a service provider’s perspective.

The proposed architecture is fully compliant with the SDN architecture defined by the ONF shown in Fig. 14.3. As it can be seen, the ONF architecture is structured in a data plane, a controller plane, and an application plane together with a transversal management plane. More detailed discussion on the commonalities of the two architectures is provided in Sects. 14.3 and 14.4 that follow.

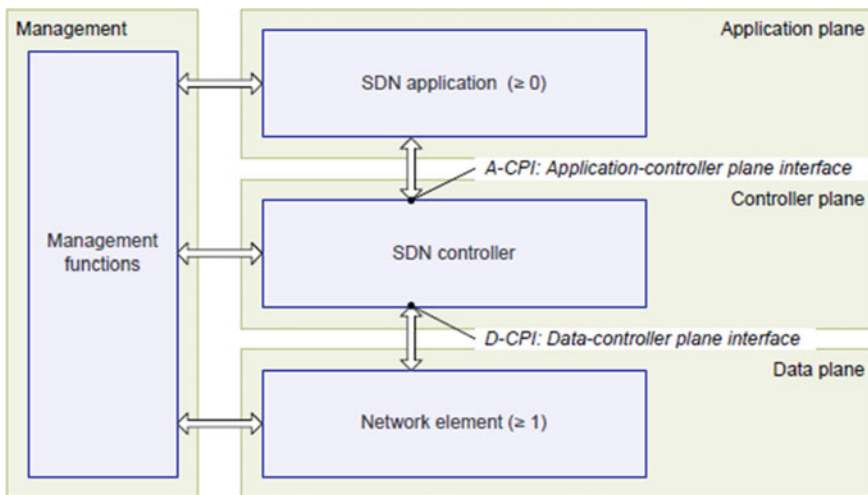


Fig. 14.3 SDN architecture defined by ONF

14.3.2 Physical Infrastructure Layer

14.3.2.1 Optical Physical Infrastructure

The optical metro network solution adopted supports frame-based sub-wavelength switching granularity and implements time-multiplexing of traffic flows, offering dynamic connectivity with fine granularity of bandwidth in a connectionless manner [17]. This technology can efficiently support connectivity between the wireless access and DC domains by providing flexible rates and a virtualization-friendly transport technology. It can also facilitate the seamless integration of wireless and IT resources to support a converged infrastructure able to provide end-to-end virtual infrastructure delivery.

A typical scenario of an optical network interconnecting the wireless access with distributed DCs is shown in Fig. 14.4. The optical network comprises FPGA platforms (to receive the ingress traffic and convert it into optical bursts), and 10-ns optical switches (to route them within the TSON core to the final destination all-optically). The routing at each node is performed using the allocation information displayed as matrices on top of each link.

The operational architecture of the optical nodes follows the generic approach presented in [34] and can be structured in three layers. The optical network data plane consists of FPGA nodes for high-speed data processing in rates of 10 Gb/s/λ. Ethernet traffic from the client sides in the wireless network and in the data centers is the ingress traffic to the optical network. At the optical network edge, this tributary traffic is processed and converted into optical bursts, which are then sent to fast switches and, from there, they are routed all optically to the destination nodes. The optical network uses 10GE SFP+ modules, and, by using extended FPGA features, the edge nodes can process any combination of Ethernet header information including MAC address and VLAN IDs.

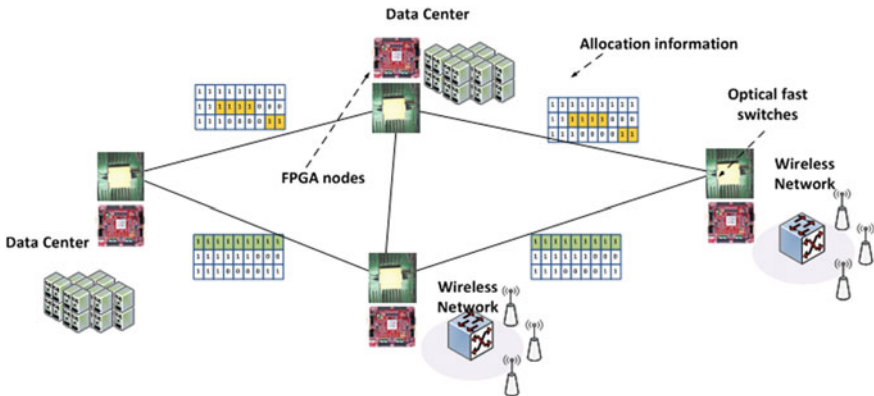


Fig. 14.4 Optical network interconnecting DC and wireless access networks

14.3.2.2 Wireless Physical Infrastructure

The wireless network consists of a hybrid IEEE 802.11 and LTE network and a wired backhaul OpenFlow-enabled packet core network that will interact with the optical domain through a Wireless Domain Gateway. A high-level overview of the wireless network architecture is presented in Fig. 14.5. In this work, the LTE network utilizes a commercially available EPC solution and an Orbit Management Framework (OMF) Aggregate Manager service that enables control of the LTE femtocells and the EPC network.

14.3.2.3 Integrated Infrastructure

The integration of the two physical infrastructures creates a datapath across multiple technology domains including DC, optical, and wireless access networks (Fig. 14.6). The wireless network is equipped with OpenFlow-based backhaul switches which can apply policies to the ingress traffic and redirect them to different access points. Wi-Fi and LTE access points are available in the test bed and provide connectivity to the end users.

The horizontal connectivity between the different technologies is based on the Ethernet protocol and uses MAC addressing, VLAN IDs, and also specific priority bits [35]. For instance, in a simple case VLAN IDs can be identified end-to-end when transporting different flows across the optical and wireless backhaul. When more complicated grooming is the objective, the optical network can classify and transport traffic based on a set of Ethernet bits, and the traffic on the wireless side

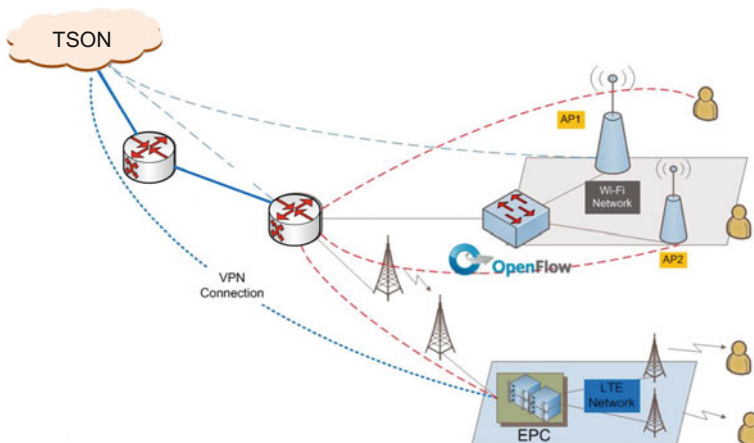


Fig. 14.5 Wireless access network. The OMF aggregate manager service is running at a dedicated server for each AP




Fig. 14.6 Integrated physical infrastructures with end-to-end datapath

can apply its own forwarding policies using, e.g., OpenFlow layer 2 on either the downlink or uplink.

14.3.3 Infrastructure Management

The IML is devoted to the converged management of resources of the different technology domains and is responsible for the creation of isolated virtual infrastructures composed of resources from those network domains.

The overall IML architecture is shown in Fig. 14.7. The bottom part of the IML contains the components (i.e., drivers) that are responsible for retrieving information and communicating with each one of the domains. Once the information has been acquired, the resources are abstracted and virtualized. Virtualization is the key functionality of the infrastructure management layer. This layer of the architecture provides a cross-domain and cross-technology virtualization solution, which allows the creation and operation of infrastructure slices including subsets of the network connected to the different computational resources providing cloud services. From the architectural and functional perspective, the design of the IML is capable of covering and addressing the set of specific requirements identified for physical resource management and virtualization.

The IML functional architecture can be mapped over the SDN architecture defined by the ONF as depicted in Fig. 14.8. The IML comprises mainly the infrastructure segment of the SDN architecture and the associated management. The virtual slices created by the IML are seen as actual network elements by the corresponding controller deployed on top of the IML. The IML is also responsible for composing the virtual resources created and offering them as isolated virtual infrastructures. The virtual resources will hold operational interfaces, which will be used by the VICL in order to configure and provision dynamic connectivity services.

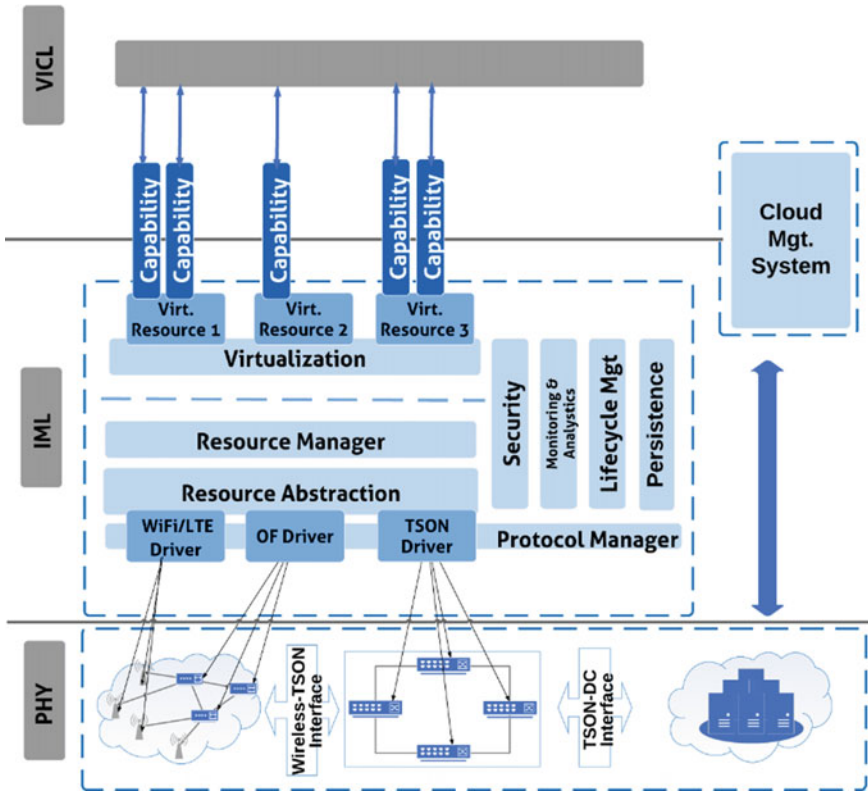


Fig. 14.7 IML overall architecture

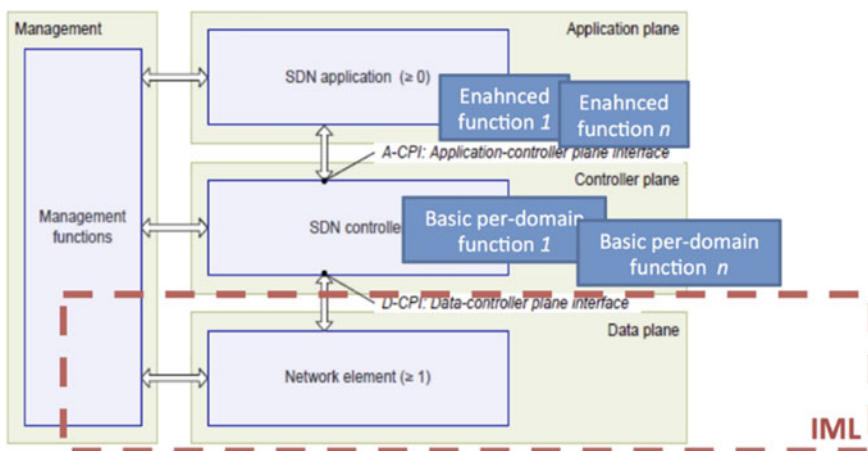


Fig. 14.8 IML mapping over the SDN ONF architecture

14.3.4 Virtual Infrastructure Control Layer

The architecture design of the VICL is structured in two main levels of functionalities, as shown in Fig. 14.9. Lower-layer functions, implemented within one or more SDN controllers, deal with the details of each technology deployed at the data plane in the wireless access and optical metro network segments, as exposed in the associated virtual resources through the IML. These functions provide elementary services with a per-domain scope and are specialized to operate on top of specific (virtual) technologies: radio, L2 wireless backhaul, and optical virtual resources. On top of the basic services offered by the SDN controller(s), further enhanced functionalities are developed in order to operate the entire heterogeneous infrastructure in a unified manner, without dealing with the specific constraints of the virtual resources exposed by the IML. These functions manage more abstracted entities (e.g., “per-domain TSON paths”) that summarize the results of the configuration actions performed by the lower-level functions directly on the virtual resources. Some examples of enhanced functions are the provisioning of end-to-end, multi-domain connectivity crossing wireless and metro domains or the re-optimization of this connectivity based on current traffic load and characteristics of the cloud services running on top. The cooperation of these enhanced functions allows MOVNOs to operate efficiently their entire infrastructures and dynamically establish the connectivity between users and DCs, or between DCs, that is required

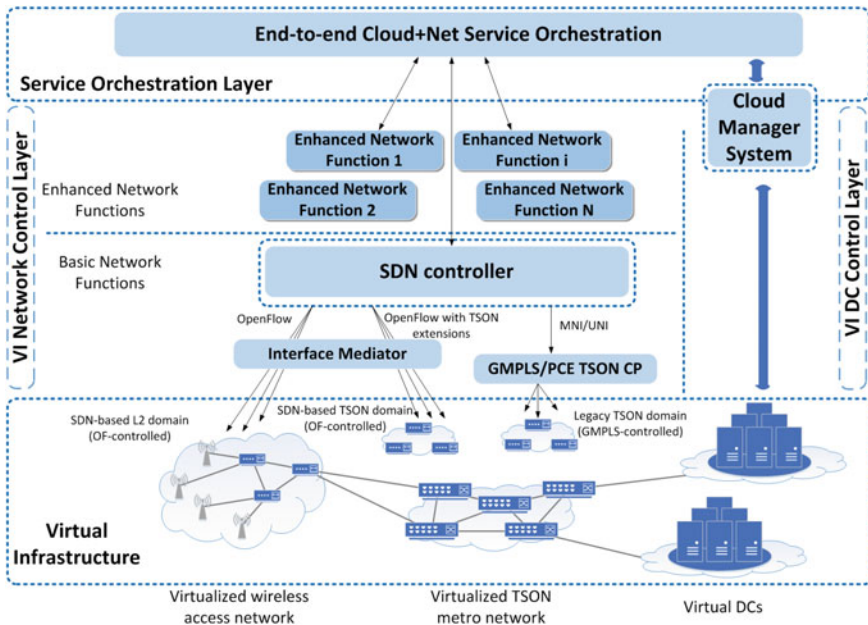


Fig. 14.9 High-level architecture of virtual infrastructure control layer

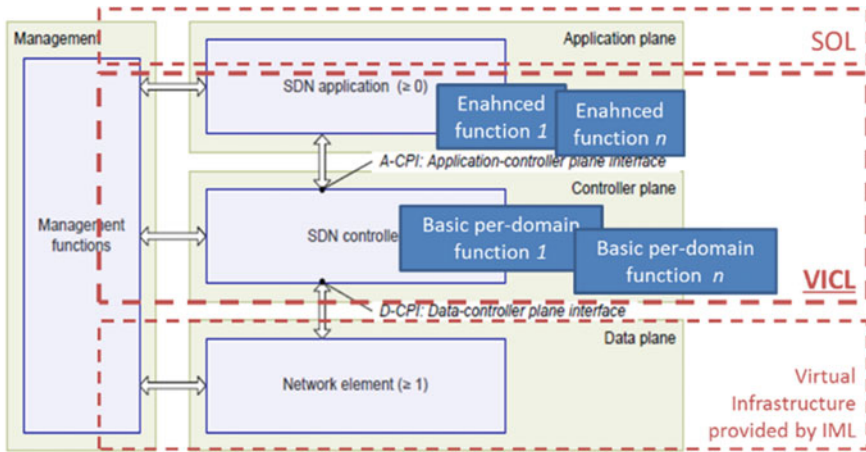


Fig. 14.10 Mapping between the proposed VICL architecture and SDN architecture defined by ONF

to properly execute and access the cloud applications running on the distributed IT resources at the virtual DCs.

The VICL architecture is fully compliant with the SDN architecture defined by the ONF shown in Fig. 14.3. Following the VICL perspective, the data plane corresponds to the virtual infrastructure exposed by the IML, the SDN controller offers the basic per-domain functions, while the enhanced application for provisioning and optimization of cloud-aware multi-domain connectivity operates at the application plane (Fig. 14.10).

Here, a single SDN controller is deployed to control the whole infrastructure, with specific plug-ins for the different virtual technologies. A unified resource control layer is responsible to create an abstracted internal information model, properly extended to cover the technological heterogeneity of the virtual infrastructure. However, the MOVNO may decide to deploy the VICL following a distributed approach, with different SDN controllers dedicated to each single technology domain and an upper layer SDN controller implementing the cross-domain enhanced functions. Other applications may run on top of this controller in order to manage the interaction with the cloud environments.

This model is compliant with the recursive hierarchical distribution of SDN controller defined by ONF (Fig. 14.11), with two levels of controllers. On the lower side, the “child” controllers are able to manage the specific granularity of each technology and have a geographical scope limited to the topology of the wireless access or metro domains. On the upper side, the “parent” controller operates on more abstracted resources, as exposed by the child controllers, but with a wider geographical scope, managing the entire virtual infrastructure (Fig. 14.12).

The ONF architecture defines two main interfaces, the Application-Controller Plane Interface (A-CPI) and the Data-Controller Plane Interface (D-CPI). In the

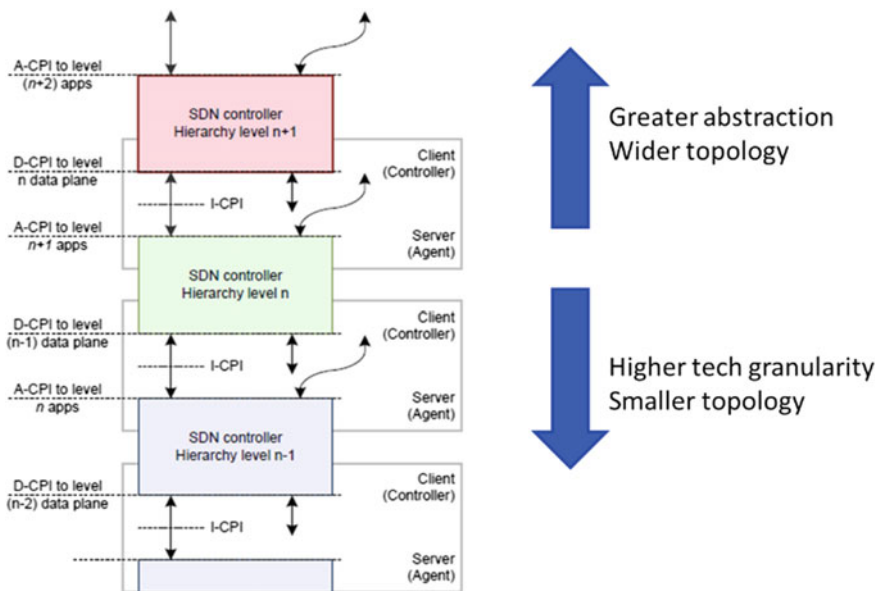


Fig. 14.11 Recursive hierarchical distribution of SDN controller, as defined by ONF

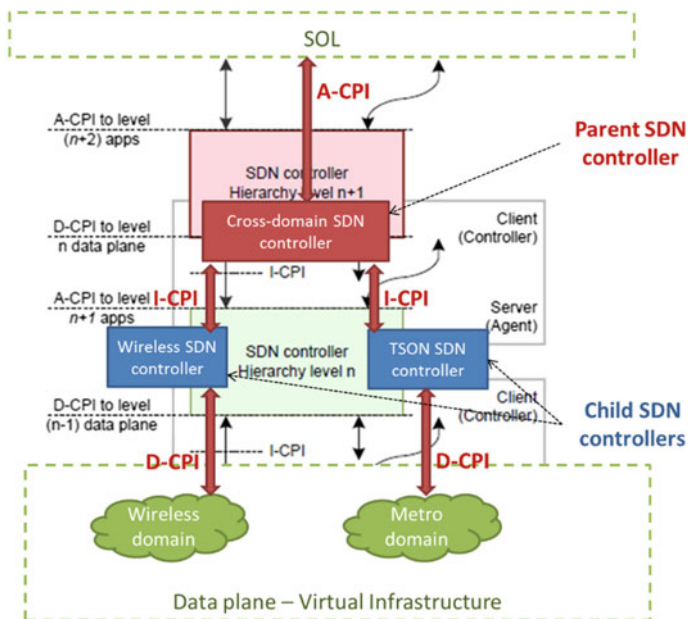


Fig. 14.12 Mapping of VICAL distributed deployment model with the recursive hierarchical distribution of SDN controllers defined by ONF

proposed architecture, the D-CPI corresponds to the interface between the IML and the VICL and provides all the functionalities related to discovery, configuration, and monitoring of the virtual resources belonging to the virtual infrastructure under control of a given SDN controller (of the lower layer) deployed in the VICL. The protocol adopted at this interface is OpenFlow for the operation of layer 2 resources and an extended version of OpenFlow for the operation of optical resources. The extensions introduced in OpenFlow allow handling the specific granularity and characteristics of optical resources, through modeling of wavelengths and the related time slots available for each port. Similar extensions could not be adopted in the OF protocol in the case of wireless resources, since real-time response to channel configuration request is not feasible, in terms of acceptable time, in reaction to packet-in events. Thus, in the wireless domain, the OF protocol is used for the operation of OF systems (e.g., Open vSwitch (OVS) instances in the access points and the backhaul network) and Representational State Transfer (REST)-based management interfaces for the management and control of other types of resources.

The A-CPI is considered at two different levels. The former is the interface between the SDN controller providing the elementary per-domain functionalities and the cross-domain applications running on top of that for the operation of the entire infrastructure. A second level of A-CPI is the interface between the entire VICL and the SOL, responsible for the management of the end-to-end cloud services. It should be noted that, from the SOL perspective, the actual deployment of the VICL is completely transparent and the VICL itself can be considered as a single SDN controller that exposes some services on its A-CPI to control the end-to-end multi-domain connectivity. In this sense, the components of the SOL act as SDN applications which operate on the most abstracted level of the resources managed at the VICL (i.e., the end-to-end connections). The A-CPIs in the proposed architecture are based on the REST paradigm and allow to manage the associated resources through create, read, update, and delete (CRUD) operations (Fig. 14.13).

14.3.4.1 Provisioning of Cross-Domain Connectivity

The dynamic provisioning of cross-domain connectivity is handled through a set of cooperating enhanced functions, implemented as SDN applications that run on top of the SDN controller and makes use of the basic per-domain services offered through the A-CPI REST interface of the controller (see Fig. 14.10). Part of these applications offers a north-bound interface to be invoked by the SOL at the upper level, while other applications are related to the internal management of the virtual infrastructure and are not consumed by external entities.

Figure 14.14 shows an example of interaction between some cross-domain enhanced services at the SDN application plane and per-domain basic services implemented at the SDN controller plane. When a new multi-domain connection,

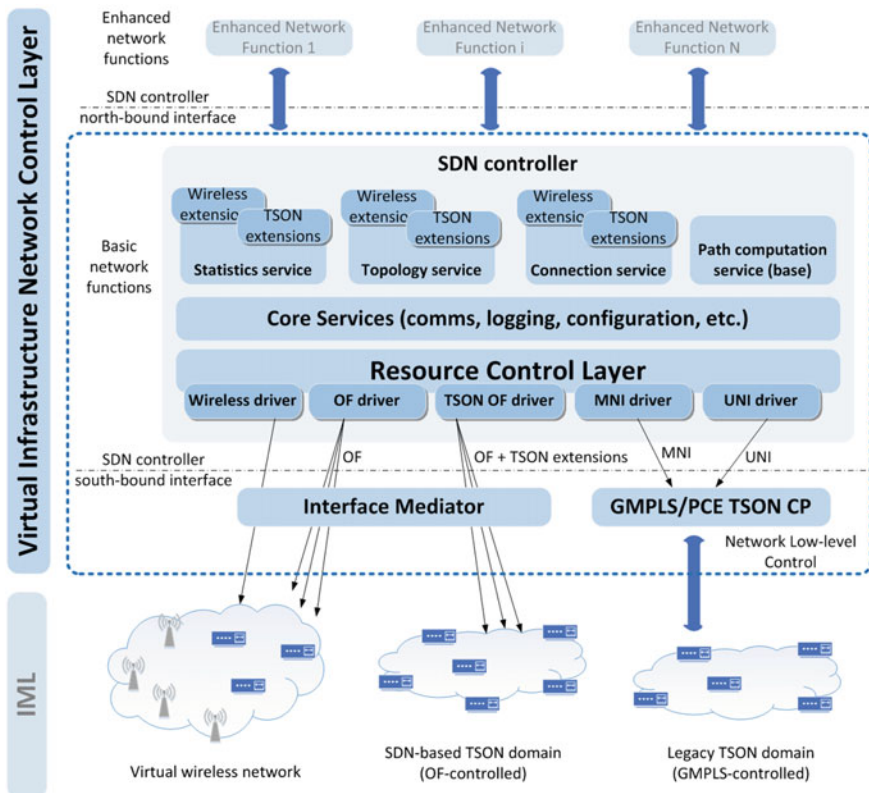


Fig. 14.13 Low-level control functions in the SDN controller

spanning from the wireless access to the metro optical domain, must be established to connect a cloud mobile user to a DC, the SOL sends a request to the end-to-end (E2E) Multi-domain Connection Service. This service acts as a consumer of the E2E Multi-layer Path Computation Service, which is a sort of parent PCE specialized in the elaboration of multi-layer, multi-technology paths. The E2E Multi-layer Path Computation Service elaborates a path which specifies just the edge nodes of the domains traversed by the flow, together with the ports and resources to be used at the inter-domain links. This information is used by the E2E Multi-domain Connection Service to trigger the creation of per-domain connections between these edge nodes, consuming the wireless access and TSON connection services provided by the SDN controller.

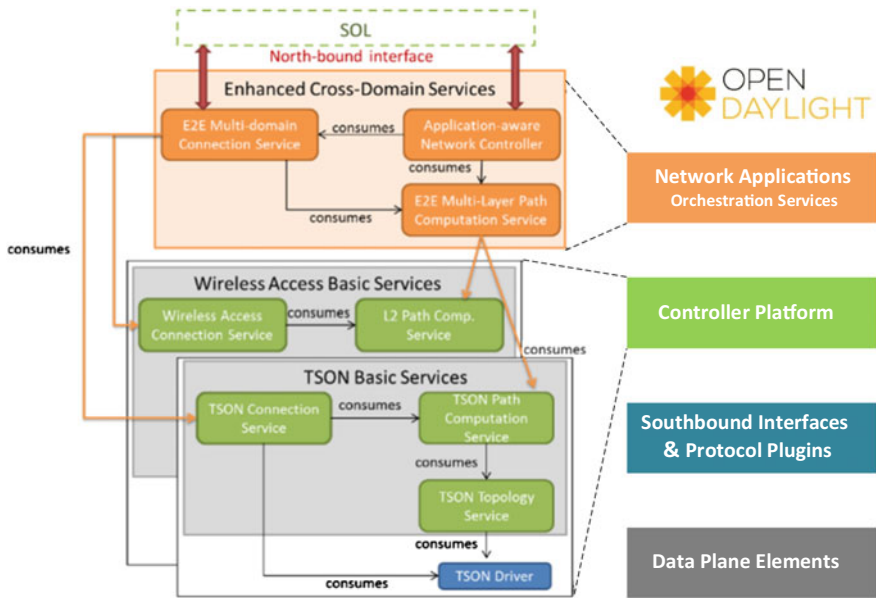


Fig. 14.14 Example of interaction between cross-domain enhanced services and per-domain basic services

14.3.5 Converged Service Orchestration

The mechanisms for the converged orchestration of cloud and network services, with the automated and cloud-aware setup and adjustment of the user-to-DC and inter-DC network connectivity, are implemented at the SOL.

Figure 14.15 shows the SOL architecture together with the interaction between the SOL and the entities responsible to manage, control, and operate cloud and network resources. As analyzed in previous sections, the network control layer is provided through the combination of one or more SDN controllers (based on the OpenDaylight platform) and additional SDN applications developed on top of them. These applications provide enhanced functions for end-to-end provisioning of multi-domain connectivity, as required to support a suitable access and execution of distributed cloud-based applications. REST APIs are available to allow the interaction between SOL and VICL, from the basic requests for creation, tear-down, and modification of user-to-DC and inter-DC connections to more advanced bidirectional communications to collect monitoring metrics or cross-layer information to take joint decisions about the evolution of the cloud service. The cloud management system can adopt several open-source solutions to control the different types of cloud resources (computing and storage), together with the internal connectivity of the data center. An example is OpenStack, a flexible cloud management platform widely used in cloud environments to deliver multi-tenant services.

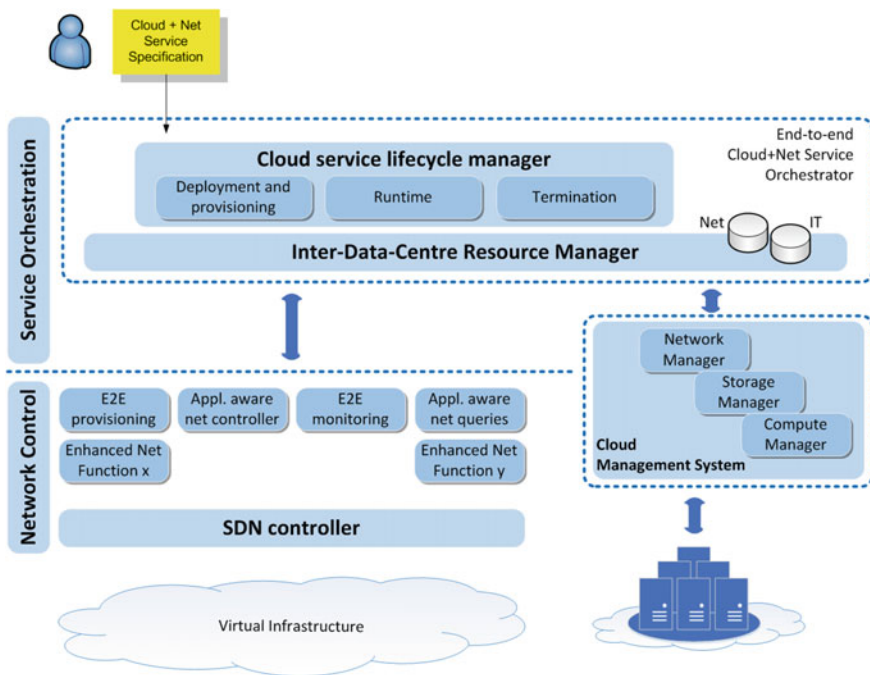


Fig. 14.15 Service orchestration layer high-level architecture

14.4 Architecture Evaluation

14.4.1 Network Scenario and Related Work

As already discussed in detail, one of the main innovations of the reference architecture proposed (Fig. 14.2) is that of cross-domain virtualization that allows the creation of infrastructure slices spanning the converged optical and wireless network domains as well as the IT resources in the DCs. Cross-domain virtualization is performed by the IML that is responsible for abstracting the physical infrastructure, composing the virtual resources created, and offering them as isolated MOVNOs exploiting the functionalities available through the adoption of the SDN principles. An important aspect of this process is that the formed MOVNOs are able to satisfy the virtual infrastructure operators' requirements and the end users' needs, while maintaining cost-effectiveness and other specific requirements such as energy efficiency. This involves the identification of the optimal virtual infrastructures that can support the required services in terms of both topology and resources and includes mapping of the virtual resources to the physical resources. To achieve this goal, the IML plays a key role as it allows taking a centralized approach with a complete view of the underlying infrastructure, thus allowing the identification of globally optimal MOVNOs.

In order to assess the performance of this type of infrastructures, we have developed a modeling framework based on integer linear programming (ILP) for the integrated wireless, optical, and DC infrastructure. In terms of wired technologies, computing resources are interconnected through the TSON WDM metro network with frame-based sub-wavelength switching granularity incorporating active nodes that will also interface the wireless microwave links supporting backhauling from the wireless access segment. This infrastructure is suitable to support traffic that is generated by both traditional cloud and mobile cloud applications. In this context, traffic demands corresponding to traditional cloud applications are generated at randomly selected nodes in the wired domain and need to be served by a set of IT servers. Mobile traffic on the other hand is generated at the wireless access domain and in some cases needs to traverse a hybrid multi-hop wireless access/backhaul solution before it reaches the IT resources through the optical metro network.

Using this modeling framework, an extended set of results is produced that are used to evaluate the performance of the proposed architecture in terms of metrics such resource and energy efficiency. Our study focuses on optimal planning of MOVNOs in terms of both topology and resources over the physical infrastructure. To identify the least energy-consuming dependable MOVNOs, detailed power consumption models taken from [36] for the optical metro, the wireless access networks and IT resources are considered. In the general case, the MOVNO planning problem is solved taking into account a set of constraints that guarantee the efficient and stable operation of the resulting infrastructures. The following constraints are considered:

- Every demand has to be processed at a single IT server. This allocation policy reduces the complexity of implementation.
- The planned MOVNO must have sufficient capacity for all demands to be transferred to their destinations.
- The capacity of each link in the MOVNO should be realized by specific PI resources.
- A critical function is the mapping and aggregation/de-aggregation of the traffic from one domain to the other. Given that the PI consists of a heterogeneous network integrating optical metro and wireless access domains, each physical link g is assumed to have a different modular capacity M_κ , ($\kappa = \text{wireless, optical}$). For example, the wireless backhaul links are usually treated as a collection of wireless microwave links of 100 Mbps while the TSON solution may offer a minimum bandwidth and granularity of 100 Mbps. Therefore, during the MOVNO planning process the *conservation of flow* between the different technology domains is also considered.
- The planned MOVNO must have adequate IT server resources such as CPU, memory, and disk storage to support all requested services.
- As a final consideration, depending on the type of service that has to be provided, the end-to-end delay across all technology domains should be limited below a predefined threshold.

The objective of our formulation is to minimize the total cost during the planned time frame of the resulting network configuration that consists of the following components: (a) k_g the cost for operating capacity u_g of the PI link g in the optical domain, (b) w_l the cost for operating capacity u_l of the PI link l in the wireless domain, and (c) σ_{sr} the total cost of the capacity resource r of IT server s for processing the volume of demand h_d .

$$\text{Minimize } F = \sum_g k_g(u_g) + \sum_l w_l(u_l) + \sum_s \sum_r \sigma_{sr}(h_d)$$

The costs considered in our modeling are related to the energy consumption of the converged infrastructure comprising optical and wireless network domains interconnecting IT resources. This objective has been chosen as it addresses the energy efficiency target of the proposed solution, but also because it can be directly associated with the operational expenditure (OpEx) of the planned MOVNO. Apart from minimizing energy consumption, the planned MOVNOs can be also designed (based on the same formulation) taking into account other optimization objectives such as distance (or number of hops) between source and destination nodes, end-to-end service delay, load balancing, resource utilization, or number of active IT servers. More details on the modeling framework can be found in [6].

14.4.1.1 Numerical Results

To investigate the performance of the proposed MOVNO design scheme across the multiple domains involved, the architecture illustrated in Fig. 14.16 is considered: The lower layer depicts the PI and the layer above depicts the VIs. For the PI, a macro-cellular network with regular hexagonal cell layout has been considered, similar to that presented in [36], consisting of 12 sites, each with 3 sectors and 10 MHz bandwidth, operating at 2.1 GHz. The inter-site distance (ISD) has been set to 500 m to capture the scenario of a dense urban network deployment. Furthermore, a 2×2 MIMO transmission with adaptive rank adaption has been considered while the users are uniformly distributed over the serviced area.

As indicated in [36], each site can process up to 115 Mbps and its power consumption ranges from 885 to 1087 W, under idle and full load, respectively. For the computing resources, three Basic Sun Oracle Database Machine Systems have been considered where each server can process up to 28.8 Gbps of uncompressed flash data and its power consumption ranges from 600 to 1200 W, under idle and full load, respectively [37]. For the metro network, the TSON solution has been adopted assuming a single fiber per link, 4 wavelengths per fiber, wavelength channels of 10 Gb/s each, minimum bandwidth and granularity 100 Mbps, and maximum capacity of the link 40 Gbps. Finally, cloud and mobile cloud computing traffic is generated at randomly selected nodes in the optical metro network and LTE base stations and needs to be served to the three IT servers. Note that in order

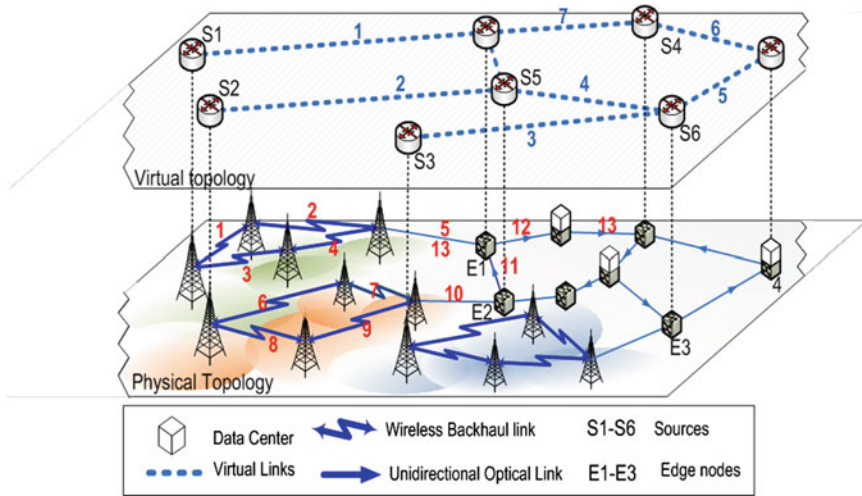


Fig. 14.16 Virtualization over heterogeneous network infrastructures

to study how the end user mobility and the traffic parameters affect the utilization and power consumption of the planned VI the service-to-mobility factor [38] is introduced. The service-to-mobility factor is defined as the fraction of the service holding time over the cell residence time. It is assumed that the cell residence time is exponentially distributed with parameter η and the service holding time is Erlang distributed with parameters (m, n) .

Figure 14.17 illustrates the total power consumption of the converged infrastructure (wireless access, wireless backhaul, optical network, and IT resources) when applying the proposed approach optimizing for energy efficiency or network resources. Comparing these two schemes, it is observed that the energy-aware MOVNO (EA MOVNO) design consumes significantly lower energy (lower operational cost) to serve the same amount of demands compared to the closest IT scheme providing an overall saving of the order of 37 % for low traffic demands. This is due to the fact that in the former approach fewer IT servers are activated to serve the same amount of demands. Given that the power consumption required for the operation of the IT servers is dominant in this type of infrastructures, switching-off the unused IT resources achieves significant reduction of energy consumption. Furthermore, it is observed that for both schemes the average power consumption increases almost linearly with the number of demands. However, the relative benefit of the energy-aware design decreases slightly with the number of demands, when approaching full system load. It is also observed that the wireless access technology is responsible for 43 % of the overall power consumption, while the optical network consumes less than 7 % of the energy. Note that the results have been produced assuming that the traffic generated in the wireless access (mobile cloud) is 2 % of the traffic generated by fixed cloud services in the optical domain (cloud traffic).

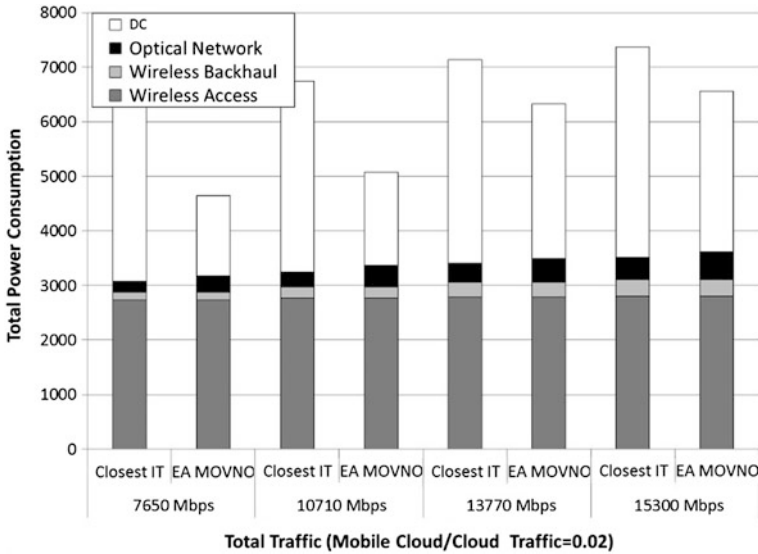


Fig. 14.17 Impact of traffic load on power consumption when the VI is planned for energy or distance

To further investigate how the service characteristics affect the optimum planned MOVNO, the impact of both traffic load and end user mobility on the total power consumption (operational cost) of the MOVNO is studied. Figure 14.18 shows that the total power consumption increases for higher end user mobility, as expected. More specifically, when mobility is higher (lower service-to-mobility factor) additional resources are required to support the MOVNO in the wireless access domain. However, it is interesting to observe that this additional resource requirement also propagates in the optical metro network and the IT domain. The additional resource requirements, across the various infrastructure domains, are imposed in order to ensure availability of resources in all domains involved (wireless access and backhauling, optical metro network, and DCs) to support the requested service and enable effectively seamless and transparent end-to-end connectivity between mobile users and the computing resources.

14.4.1.2 Impact of the Virtualization Solution on the Proposed Architecture

Network and cloud convergence can bring significant benefits to both end users and physical infrastructure providers. In particular, end users can experience a better quality of service in terms of improved reliability, availability, and service ability, whereas physical infrastructure providers can gain significant benefits through improved resource utilization and energy efficiency, faster service provisioning

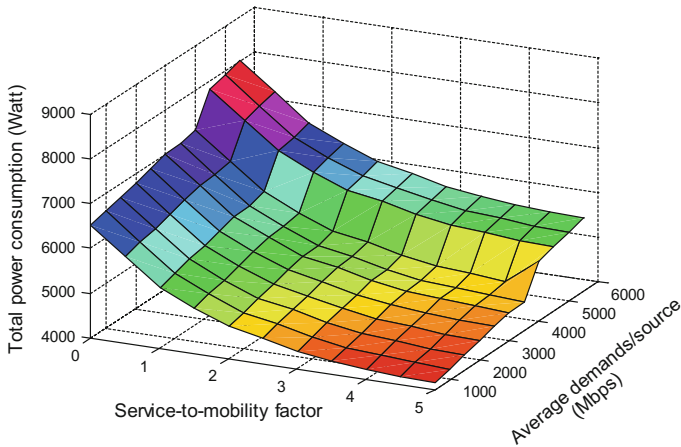


Fig. 14.18 Impact of mobility and traffic load on the total power consumption ($m = 5$)

times, greater elasticity, scalability, and reliability of the overall solution. In view of these, virtualization across multi-technology domains has recently gained significant attention and several solutions already exist (see [39]). Typical examples include the *pipe* and the *hose* virtual private network models [40]. 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 traffics at each endpoint having the advantages of ease of specification, flexibility, and multiplexing gain [41] (Fig. 14.19).

Other similar research efforts have been focused on embedding virtual infrastructures (VI) over multi-domain networks [39]. 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

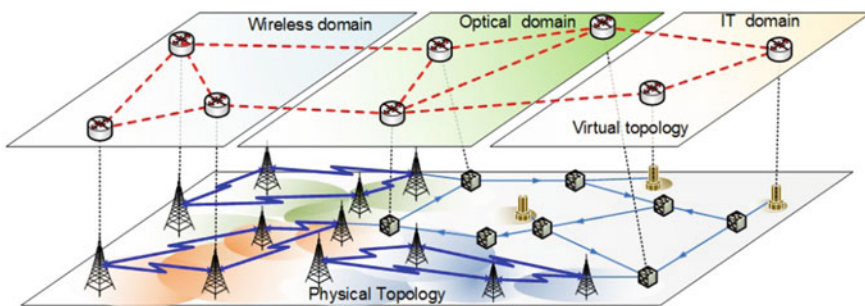


Fig. 14.19 Segmented multi-domain virtualization

(see [39] where this approach is 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. In the present study, it is argued that in order to optimally exploit the benefits of a converged cloud, some level of information exchange across the multi-technology domains must be provided. To facilitate this while offering a feasible solution, it is proposed to develop suitable abstraction models of the physical infrastructure that allow MOVNOs to maintain the information needed to be exchanged between technology domains, in order to overcome the issues described above. To give an example, the information required to be maintained may include the status of Rx/Tx queues of the TSON nodes, the physical layer impairments in the wireless and optical domains, the power consumption levels of the equipment, and the mobility profiles/characteristics of the end users. Preliminary modeling results (see Figs. 14.20 and 14.21) indicate that by means of maintaining and exchanging the required information across technology domains, significant performance improvements can be achieved. In more detail, our modeling results show that by maintaining this type of information through the relevant parameters the end-to-end delay for the mobile traffic has been reduced by up to 40 % (see Fig. 14.20) compared to [39]. In addition, approximately 20 % overall power savings

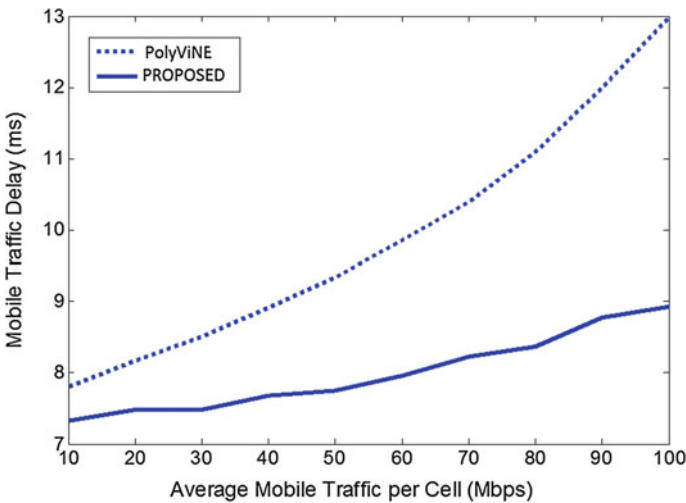


Fig. 14.20 Comparison of the proposed converged approach with PolyViNE in terms of mobile cloud computing traffic delay (fixed traffic 800 Mbps/source, source nodes = 3)




Fig. 14.21 Impact of the proposed solution on the total power consumption

(see Fig. 14.21) for various loading conditions can be achieved, when compared to the existing approaches [39]. Details of the relevant models used are available in [6].

14.5 Conclusions

This chapter focused on a next-generation ubiquitous converged network infrastructure involving integration of wireless access and optical metro network domains to support a variety of cloud and mobile cloud services. The infrastructure model proposed is based on the IaaS paradigm and aims at providing a technology platform interconnecting geographically distributed computational resources that can support these services. The proposed architecture addresses the diverse bandwidth requirements of future cloud and mobile cloud services by integrating advanced optical metro network technologies offering fine granularity with state-of-the-art wireless access network technology (LTE/Wi-Fi), supporting end user mobility. The concept of virtualization across the technology domains is adopted as a key enabling technology to support our vision. This study also provided a description of the proposed architecture including the details of the individual architectural layers, i.e., physical infrastructure layer, infrastructure management layer, virtual infrastructure control layer, and orchestrated end-to-end service layer. In addition, a detailed description of the interaction between the different layers is provided. The proposed architecture is well aligned and fully compatible with the OGF SDN architecture and takes advantage of the associated functionalities and capabilities. A modeling/simulation framework was developed to evaluate the proposed architecture and identify planning and operational methodologies to allow global optimization of the integrated converged infrastructure. A description of the tool and some relevant modeling results were also presented.

Acknowledgments This work was carried out with the support of the CONTENT (FP7-ICT-318514) project funded by the EC through the 7th ICT Framework Program.

References

1. Tzanakaki A et al (2014) Planning of dynamic virtual optical cloud infrastructures: the GEYSERS approach. *IEEE Commun Mag* 52(1):26–34
2. Satyanarayanan M, Bahl P, Caceres R, Davies N (2009) The case for VM-based cloudlets in mobile computing. *IEEE Pervasive Comput* 8(4):14–23
3. Mun K Mobile cloud computing challenges. *TechZine Mag*. <http://goo.gl/AuI5ZM>
4. <http://www.gogrid.com/>
5. <http://www.flexiscale.com/>
6. Tzanakaki A et al (2013) Virtualization of heterogeneous wireless-optical network and IT support of cloud and mobile services. *IEEE Commun Mag* 51(8):155–161
7. Yeganeh SH et al (2013) On scalability of software-defined networking. *IEEE Commun Mag* 51(2):136–141
8. Mateos G, Rajawat K (2013) Dynamic network cartography. *IEEE Signal Process Mag* 129–143
9. Rajawat K (2012) Dynamic optimization and monitoring in communication networks. University of Minnesota, Ph.D. dissertation
10. Channegowda M et al (2013) Software defined optical networks technology and infrastructure: enabling software-defined optical network operations (invited). *IEEE/OSA JOCN* 5(10): A274–A282
11. Zhang S, Tornatore M, Shen G, Zhang J, Mukherjee B (2014) Evolving traffic grooming in multi-layer flexible-grid optical networks with software-defined elasticity. *J Lightwave Technol* 32:2905–2914
12. France Telecom-Orange and Alcatel-Lucent deploy world's first live 400 Gbps per wavelength optical link, Alcatel Lucent press release Feb 2013. <http://goo.gl/xtvLBa>
13. Gerstel O, Jinno M, Lord A, Yoo SJB (2012) Elastic optical networking: a new dawn for the optical layer? *IEEE Commun Mag* 50(2):s12–s20
14. IEEE 802.3 Ethernet Working Group Communication (2012) IEEE 802.3™ industry connections ethernet bandwidth assessment. <http://goo.gl/Zp2Xym>
15. Jinno M et al (2009) Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies. *IEEE Commun Mag* 47(11):66–73
16. Verisma ivx 8000, Optical Packet Switch and Transport. Intune Networks, 2011
17. Zervas GS, Triay J, Amaya N, Qin Y, Cervelló-Pastor C, Simeonidou D (2011) Time shared optical network (TSON): a novel metro architecture for flexible multi-granular services. *Opt Express* 19:B509–B514
18. Lehr W, Sirbu M, Gillett S (2004) Municipal wireless broadband policy and business implications of emerging technologies. In: *Competition in networking: wireless and wireline*. London Business School
19. [RFIC2013] Cellular versus Wi-Fi: future convergence or an utter divergence? In: *Radio frequency integrated circuits symposium (RFIC), 2013 IEEE*, pp I, II
20. Castelli MJ (2002) *Network sales and services handbook*. Cisco Press. ISBN 978-1-58705-090-9
21. Peng S, Nejabati R, Azodolmolky S, Escalona E, Simeonidou D (2011) An impairment-aware virtual optical network composition mechanism for future internet. *Opt Express* 19(26):B251–B259
22. Jinno M et al (2013) Virtualization in optical networks from network level to hardware level. *IEEE/OSA J Opt Commun Netw* 5(10):A46–A56

23. García-Espín JA, Ferrer Riera J, Figuerola S, Ghijsen M, Demchenko Y, Buysse J, De Leenheer M, Develder C, Anhalt F, Soudan S (2012) Logical infrastructure composition layer, the GEYSERS holistic approach for infrastructure virtualisation. In: Proceedings of TNC
24. Tzanakaki A, Anastasopoulos MP, Georgakilas K (2013) Dynamic virtual optical networks supporting uncertain traffic demands. *IEEE/OSA J Opt Commun Netw* 5(10) (Invited)
25. Li Y, Qiu L, Zhang Y, Mahajan R, Rozner E (2008) Predictable performance optimization for wireless networks. In: Proceedings of ACM SIGCOMM, Seattle, WA, USA
26. Bhanage G, Vete D, Seskar I, Raychaudhuri D (2010) SplitAP: leveraging wireless network virtualization for flexible sharing of WLANs. In: Proceedings of IEEE GLOBECOM 2010, pp 1–6
27. Aljabari G, Eren E (2010) Virtual WLAN: extension of wireless networking into virtualized environments. *Int J Comp Res Inst Intell Comput Syst* 10(4):1–9
28. Kokku R et al (2013) CellSlice: cellular wireless resource slicing for active RAN sharing. In: Proceedings of COMSNETS
29. <http://goo.gl/oT5E5t> <http://www.contextream.com/news/press-releases/2014/contextream-announces-first-commercially-available-opensdaylight-based-carrier-grade-sdn-fabric-for-nfv/>
30. Nyren R, Edmonds A, Papaspyrou A, Metsch T (2011) Open cloud computing interface—core. GFD-P-R.183, OCCI WG
31. Dhody D, Lee Y, Ciulli N, Contreras L, Gonzales de Dios O (2013) Cross stratum optimization enabled path computation. IETF Draft
32. Crabbe E, Medved J, Minei I, Varga R (2013) PCEP extensions for stateful PCE. IETF Draft
33. Tzanakaki A et al (2014) A converged network architecture for energy efficient mobile cloud computing. In: Proceedings of ONDM 2014, pp 120–125
34. MAINS Deliverable D2.3. OBST mesh metro network design and control plane requirements
35. Rofoee BR et al (2015) First demonstration of service-differentiated converged optical sub-wavelength and LTE/Wi-Fi networks over GEANT. In: Proceedings of OFC 2015, paper Th2A.35
36. Auer G, Giannini V (2011) Cellular energy efficiency evaluation framework. In: Proceedings of the vehicular technology conference
37. Valancius V, Laoutaris N, Massouli L, Diot C, Rodriguez P (2009) Greening the internet with nano data centers. In: Proceedings of the 5th international conference on emerging networking experiments and technologies (CoNEXT'09), ACM, New York, NY, USA, pp 37–48
38. Fang Y, Chlamtac I (2002) Analytical generalized results for handoff probability in wireless networks. *IEEE Trans Commun* 50(3):369–399
39. Chowdhury M, Samuel F, Boutaba R (2010) PolyViNE: policy-based virtual network embedding across multiple domains. In: Proceedings of ACM SIGCOMM workshop on VISA'10
40. Chu Jian, Lea Chin-Tau (2008) New architecture and algorithms for fast construction of hose-model VPNs. *IEEE/ACM Trans Netw* 16(3):670–679
41. Duffield NG, Goyal P, Greenberg A, Mishra P, Ramakrishnan KK, der Merwe JEV (1999) A flexible model for resource management in virtual private networks. In: Proceedings of ACM SIGCOMM, San Diego, California, USA