

A Service-Oriented Hybrid Access Network and Cloud Architecture

L. Velasco^{1*}, L. M. Contreras², G. Ferraris³, A. Stavdas⁴, F. Cugini⁵, M. Wiegand⁶,
and J. P. Fernández-Palacios²

¹ Optical Communications Group (GCO), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain.

² Telefónica Investigación y Desarrollo (TID), Madrid, Spain.

³ Telecom Italia, Torino, Italy.

⁴ University of Peloponnese (UoP), Tripolis, Greece.

⁵ Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT), Pisa, Italy.

⁶ Coriant, Munich, Germany.

*Corresponding author: lvelasco@ac.upc.edu

Abstract—Many telecom operators are deploying their own cloud infrastructure with the two-fold objective of providing cloud services to their customers and enabling network functions virtualization. In this article, we present an architecture named as SHINE, which focuses on orchestrating cloud with heterogeneous access and core networks; intra and inter data center (DC) connectivity is dynamically controlled, maximizing the overall performance in terms of throughput and latency whilst minimizing total costs. The main building blocks are: 1) a future-proof network architecture that can scale to offer potentially unlimited bandwidth based on an Active Remote Node (ARN) to interface end-users and the core network; 2) an innovative distributed DC architecture consisting of micro-DCs (μ -DC) placed in selected core locations to accelerate content delivery, reducing core network traffic, and ensuring very low latency; and 3) dynamic orchestration of the distributed DC and access and core network segments. SHINE will provide unprecedented Quality of Experience (QoE), greatly reducing costs by coordinating network and cloud and facilitating service chaining by virtualizing network functions.

Keywords: Telecom Cloud, Network Function Virtualization, Service Orchestration.

1 INTRODUCTION

A revolution in access networks is underway. The revolution is driven by the continued transformation of cellular networks offering to portable devices bit-rates and Quality of Service (QoS) comparable to those traditionally made available only through fixed networks. Driven by demand for video and the proliferation of data centers (DC), more than 75 percent of that traffic will stay in access/metro networks by 2017, as compared to 57 percent today, as forecasted in [1]. Accommodating the enormous traffic growth in a cost-effective and service-efficient way is essential for the viability of telecom operators and motivates a major network re-design. In fact, these shifting traffic patterns are the result of bringing content closer to the users to better manage quality of experience (QoE). For instance, the most popular video content can be cached and delivered to users locally over access/metro networks rather than being accessed from a central cache over the backbone network.

Coordinating these and new advanced services to be offered to a widely distributed number of customers requires building advanced service chains. Nevertheless, building service chains is very expensive and time consuming, since it requires, besides deploying dedicated hardware for each required network function, configuring each device using its proprietary command syntax, interfaces and protocols. Moreover, since loads often change over time, building a new service chain typically requires estimating future demands and over-provisioning IT and network resources to support growth. This prevents from reducing the final price that users pay undermining the average revenue per user (ARPU).

The continuous advances in computing hardware facilitate real-time processing to be performed on commodity hardware instead of specialized one. These advances enable network functions virtualization (NFV) [2]. By eliminating specialized network processors, multiple heterogeneous workloads can be consolidated onto a single architecture, thus reducing complexity and simplifying operation, leading to total cost of ownership (TCO) reductions.

Cloud technology offers numerous benefits including economies of scale, cost-effectiveness, efficient hardware utilization, and TCO reductions, both in capital and in operational expenditures [3]. These benefits are all key objectives for telecom operators, so the appeal of cloud technologies is clear. In fact, as revealed in a recent survey [4], many telecom operators are deploying cloud infrastructures. Notwithstanding, deploying the telecom cloud presents a differential set of challenges due to the industry's inherent requirements for availability (5-nines), very low latency, and complex networking (Ethernet, optical, wireless, etc.).

Scalability is also an issue since, in contrast to a small number of warehoused-sized DCs commonly used in public

clouds, telecom cloud must support a large number of small, distributed DCs to reduce traffic in the core network. A distributed DC architecture brings many benefits for network operators. By encapsulating workloads in virtual machines (VM) a cloud resource manager can migrate workloads from one DC to another looking for improving the perceived quality of experience (QoE), reducing energy consumption [5], or even in response to situations such as network failures or high-demand events. In addition, placing DCs closer to end-users enables the development of services and applications that can take advantage from very low latency.

The efficient integration of cloud-based services under a distributed DC architecture, including the interconnecting network, is a challenging task due to the required performance and high availability guarantees. The answer from network providers to the increasing traffic dynamism is to migrate their networks to a *cloud-ready transport network* [6], as a platform capable to handle dynamic traffic patterns and asymmetries. Although this approach enables a more elastic transport infrastructure, it has technical challenges on its own that must be addressed. In the recent work [7], the authors propose to use the dynamic connectivity provided by the flexgrid optical technology to improve resource utilization and save costs. The flexgrid technology enables a finer spectrum granularity adaptation and the ability to dynamically increase and decrease the amount of optical resources assigned to connections. The availability of flexgrid ready spectrum selective switches enables building bandwidth-variable Optical Cross-Connects (OXC), whereas the advent of sliceable bandwidth-variable optical transponders (SBVT), able to deal with several flows in parallel, adds even more flexibility and reduces costs [8].

In the access, higher speeds together with multiple data plane interfaces will drive the evolution of aggregation elements to multi-service nodes, abstracting capabilities from data plane specificities. The necessary support of legacy services and interfaces and the multi-service scope for those devices motivates the definition of programmable control, adapting the generic conception of the node to the specific need.

That control has to consider both, service and transport characteristics to orchestrate resources end-to-end. The advent of software-defined networking (SDN) is fueling the deployment of programmable control methods. In fact, several initiatives are currently under way for defining architectural frameworks for centralized control elements, such as the OpenDayLight project¹ or the Application Based Network Operations (ABNO) architecture [9]. The OpenFlow protocol² is well suited to handle transmission specifics and intra-DC connectivity [10]; extensions to OpenFlow can be defined to e.g., configure SBVTs. In contrast, some telecom operators might prefer using ABNO to control interconnection networks since it is based on working functional elements and facilitates network re-configuration [11].

In this article, we present a Service-oriented HybriD access Network and Cloud ArchitecturE (SHINE) that orchestrates cloud with heterogeneous access and core networks, dynamically controlling intra and inter DC connectivity. A number of challenges associated to end-to-end coordination and the migration from the existing networking framework need to be faced. Separation of service and transport oriented tasks is key to allow a scalable orchestration, facilitating its independent evolution; clear interfaces and taxonomy of functions is required.

A set of NFV use cases have been recently identified by the NFV group within ETSI [12] and several initiatives are being developed in that field, with a relevant number of Proofs-of-Concept in place³. In addition, the recently launched the mobile-edge computing (MEC) initiative aims at adding cloud-computing capabilities at the edge of the mobile network⁴. Notwithstanding, because of its versatile and adaptable architecture, SHINE offers a common infrastructure to deploy many different NFV scenarios.

2 SHINE ARCHITECTURE

SHINE proposes a new optical architecture capable of fulfilling the requirements in terms of capacity and dynamicity of future access networks bypassing metro aggregation layers. An Active Remote Node (ARN) node serves as a gateway for a number of heterogeneous networks and uses transmission and multiplexing to incorporate traffic from large geographic areas (rural and urban) directly to the core network. The ARN directly interfaces OXCs in the core network by means of point-to-point connections through dedicated links exploiting adaptive modulation formats to capitalize on their distance adaptive transmission properties (Fig. 1).

A number of μ -DCs are placed in some core locations to accelerate contents access times and to reduce core network traffic. μ -DCs are geographically distributed and connected through a flexgrid core network to behave as one single large DC. Large DCs can also co-exist to feed μ -DCs with contents.

SHINE also includes an orchestrated control and management plane to provide elastic and resilient cloud and network resource provisioning, combining resources in geographically separated μ -DCs. Dynamic network resource allocation will combine both, flexgrid core and access networks according to traffic needs.

¹ OpenDayLight: <http://www.opendaylight.org/>

² Open Networking Foundation: <https://www.opennetworking.org/>

³ http://nfvwiki.etsi.org/index.php?title=On-going_PoCs

⁴ <https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/MEC%20Executive%20Brief%20v1%2028-09-14.pdf>

The architecture of the SHINE's ARN and μ -DC is illustrated in Fig. 2.

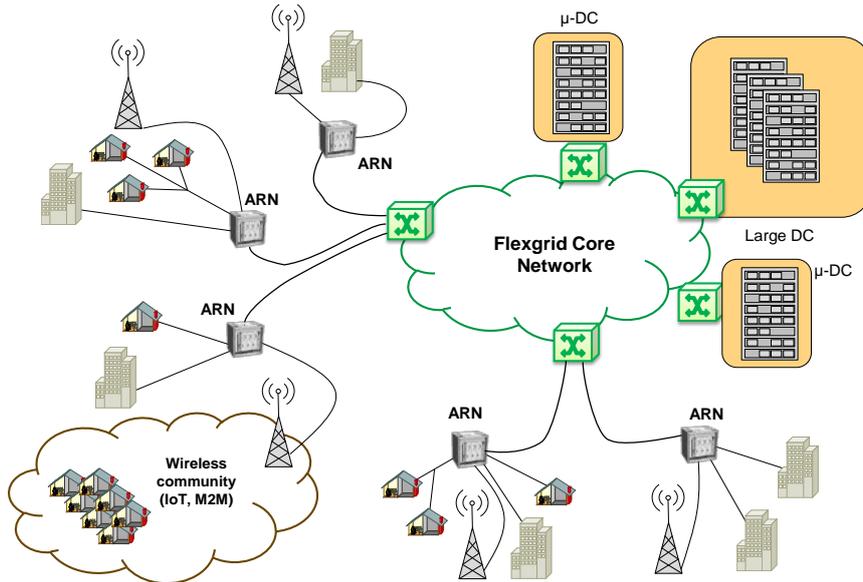


Fig. 1 Network scenario. SHINE combines hybrid access and a distributed DC connected through a flexgrid core network.

2.1 Access Network System

The SHINE's ARN works as protocol termination point where frame aggregation is implemented using either IP/MPLS, Ethernet or OTN platforms (Fig. 2a). The conceived functionality offers the potential for a service transparent solution whenever this is needed (e.g. for mobile front/back hauling). The main building blocks of the ARN include:

- Access interface implementation by means of transceiver modules. 10 GbE modules serving point-to-point connections from a multitude of services are envisaged, as well as 10 GbE PON for residential access.
- Upstream interface implementation by means of SBVTs allowing access to a number of client signals.
- OpenFlow switching and programmable network processing backplane.

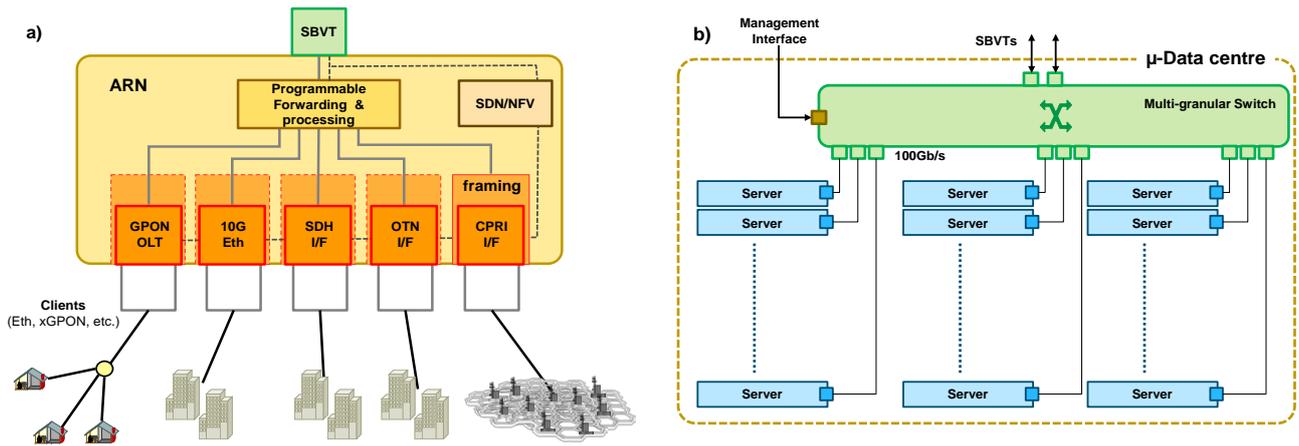


Fig. 2 SHINE's ARN (a) and μ -DC (b) architecture.

2.2 μ -DCs

The SHINE's μ -DC architecture aims at creating an energy-efficient cloud infrastructure while keeping latency ultra-low. To that end, one single multi-granular switch is used to connect every server in the μ -DC, using 100 Gb/s optical interfaces, to the interconnection network (Fig. 2b).

The multi-granular switch is able to switch packets, flows and/or optical signals from/to the SBVTs, enabling configurable multiplexing towards the SBVT front-end cards. Based on the modulation formats supported by the SBVTs, the multi-granular switch can be configured to aggregate heterogeneous client lower granularity packet flows (by performing full electronic packet processing) or entire optical flows (by performing optical port-to-port forwarding) to a given tributary signal having a certain destination (e.g., a remote μ -DC). Such flexibility, orchestrated by the local SDN controller, allows the adoption of energy-efficient grooming strategies, aiming at reducing the impact of electronic processing only where and when needed. As an example, if enough amount of traffic is generated by a local server

toward a remote μ -DC, such traffic can be assigned to single or multiple flows. Such flows can be optically switched directly towards one or more tributary lines of the SBVT that reaches the remote μ -DC, thus bypassing electronic processing. Local SDN controller is in charge of such optimization by automatically and dynamically configuring the flow entries of the multi granular switch.

2.3 Orchestrated service management and control

SHINE considers the deployment of an orchestrated service management and control architecture spanning along the μ -DCs (Fig. 3); this architecture leverages on the ABNO framework for the interaction with the transport infrastructure. A parent module is in charge of the overall coordination of cloud and networking resources, being the common entry point for services. Specific management and control modules at μ -DC level are in charge of the resources internal to a given μ -DC, whereas ABNO coordinates both, optical nodes in the core and ARNs in the access network. See [13] for details on the interaction between components.

Components of *service management and control* are:

- The scheduler, which assigns VMs to servers looking at using resources effectively and achieving the target QoE. In addition, energy efficiency can result from energy-aware VM scheduling and server consolidation.
- The QoE estimation module estimates parameters related to the QoE experienced by end-users, mainly delay.
- The statistics and monitoring module gathers information regarding the use of resources and the performance of services to be used to predict likely scenarios.

Components of *network management and control* are:

- The network control module issues commands to μ -DC level modules and the ABNO looking to create virtual networks among VMs running in one or more μ -DCs.
- The SDN controller is in charge of intra μ -DC network resources controls both, the multi-granular switch and the SBVTs installed on it.
- The ABNO module is in charge of the connections among DCs and from them to ARNs in the access network.

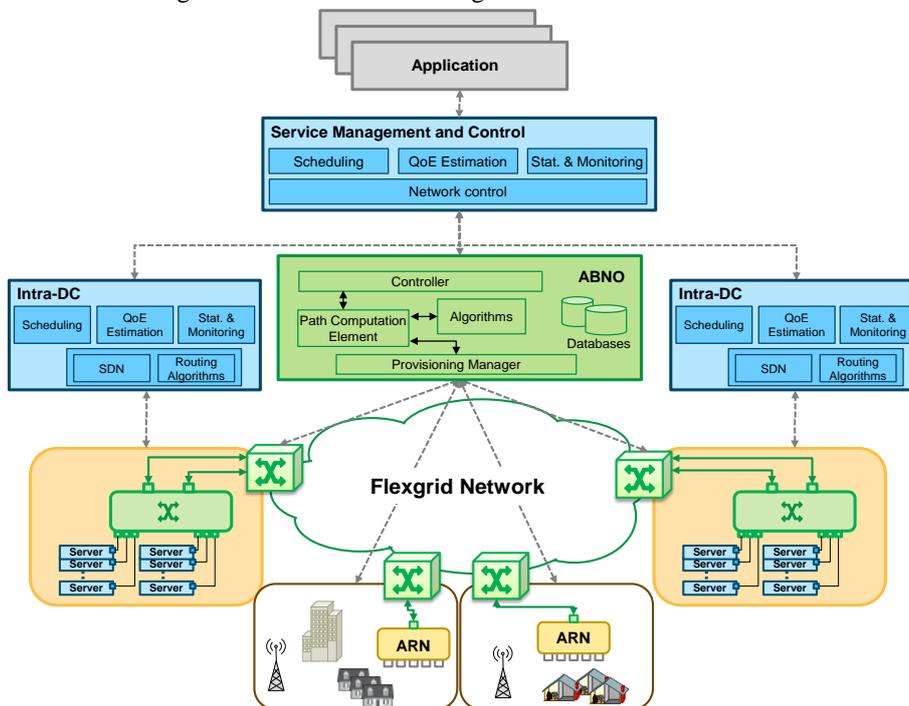


Fig. 3 SHINE's control and management architecture.

2.4 SHINE IaaS in support to NFV

IaaS-based cloud services can be offered on top of the resulting SHINE architecture, where VMs can run on the servers available in the μ -DCs and large DCs, whilst connectivity can be created to connect VMs belonging to the same client, disregarding their placement. The same infrastructure can be shared for NFV applications, where in addition to servers in DCs, VMs supporting NFV applications can run in ARNs. This creates three levels with different characteristics: *i)* ARNs can host those functions that require proximity to the end-users because of latency or to aggregate data from a reduced number of sources; *ii)* μ -DCs offer a good trade-off between latency and end-user proximity; *iii)* large DCs offer economies of scale and can be used for those delay tolerant services.

The service manager module is in charge of managing dynamic optimal VM placement and, once decided, optimal virtual networks are created or reconfigured using online optimization algorithms. This *self-management* allows applications to be deployed based on SLA agreements, including QoE parameters (e.g. max user delay, max average delay, etc.). For scalability, application administrators can request elastic operations being applied to *scale-out* by adding more VMs, *scale-up* adding more resources to a VM, and *scale-down* their cloud services.

All the above can be used to support NFV applications. Several use cases are presented next to illustrate the automated composition and allocation of computing and network resources and the interaction with the rest of elements in the SHINE architecture.

3 USE CASES

This section presents potential use cases addressing different markets of interest for any network operator, namely content delivery, and business, mobile and fixed broadband access. For that goal, we extract the corresponding use cases from the set defined in [12], aiming at briefly describing how these NFV scenarios fit into the proposed SHINE architecture.

3.1 Use case I: Content Delivery Network

Content Delivery Networks (CDN) incorporate a number of components, like cache nodes, which are orchestrated by a controller. The CDN controller is a centralized component that selects a cache node to serve an end-user request, and then redirect the end-user to the selected cache node; selecting nodes closer to the end user reduces traffic in the core network and enables delivering higher quality multimedia flows. CDN cache nodes are distributed within the network and currently deployed as dedicated physical appliances or software running on dedicated hardware.

CDN cache nodes can be virtualized to run on VMs placed in μ -DCs and/or ARNs, whereas the CDN controller can run in large DCs. Based on SLA agreements, VMs encapsulating CDN nodes are autonomously placed by SHINE's service manager to meet QoE parameters. In addition, the performance and load of the CDN nodes need to be monitored by the own CDN service administrator so as to elastically adapt the deployed nodes to the current service needs. In case the load of some cache nodes reaches an upper threshold, elastic operations to scale-up specific VMs or to scale-out to add new VMs encapsulating cache nodes can be requested to SHINE's service manager. On the contrary, when the load decreases, opportunities appear to reduce the resources (CPU or memory) available to some VMs or to consolidate workload in few VMs. As before, the CDN service should detect these opportunities to request the proper configuration to the SHINE's service manager. Finally, it is worth noting that each CDN service runs isolated from other services, so several CDNs can be deployed, where cache nodes share the underlying infrastructure.

Finally, one of the main characteristics of the SHINE service manager is its capability to reactively reconfigure deployed services in the event of QoE degradation and even proactively reconfigure them to improve overall performance. As an example, let us imagine that a failure in a link in the core network has triggered restoration and the length of the restoration path makes that suddenly the measured delay from the users to the serving application increases trespassing a given threshold. In that case, SHINE's service manager re-computes optimal VM placement to meet the committed QoE, which might result in decreasing the traffic through the restoration path, thus saving resources in the core network.

3.2 Use case II: Business access

The business connectivity market is typically characterized by the provision of isolated virtual private networks (VPN) to a variety of geographically dispersed access points. Even more, business services are demanding connectivity to some form of cloud networking, accessing either to private, public or hybrid clouds.

Specific network elements are deployed at customer premises, but also at the access point of presence (PoP) to collect such traffic; ideally aggregating a high number of enterprises demanding similar services. At the customer side, not only the customer edge (CE) equipment, but also some other devices like firewalls could be in place. In the network side, it is at least required the deployment of a provider edge (PE) router for customer access. The new trends in NFV can facilitate the virtualization of such network elements by instantiation of network functions. This can have a direct impact on service savings, since for managed services, both CE's cost and its operation are entirely allocated to the enterprise customer.

On the other hand, the cost of a PE can be shared among the customers connected to it. However, in practice, overprovisioning is required to provide the needed service coverage, thus deploying a huge number of PE equipment for enterprise access, which in reality exceeds the connectivity demand per area. This complex trade-off between service footprint and adequate platform dimensioning could be highly optimized by rolling-out virtualized PE functionality.

Considering the SHINE architecture, CEs can be deployed at the ARNs, together with additional functions if needed (e.g. firewalls as mentioned before). As for the PE function, it can be located deeper in the network, at core level, or it can be even distributed to the ARNs in case of scalability concerns. In the centralized case, the underlying flexgrid

transport network can guarantee the required SLAs for the service. Regarding the cloud resources complementing the business service, they can be placed and moved among DCs (large and μ) according to actual service needs. Fig. 4 shows a potential deployment scenario.

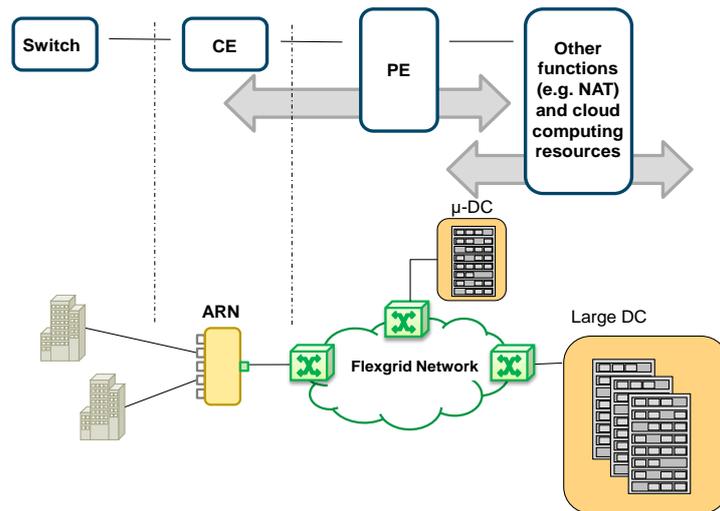


Fig. 4 Business access scenario.

From the operation viewpoint, a VPN application should be placed on top of the SHINE’s Service Management and Control module for *programming* forwarding rules among virtual CE functions residing in ARNs across the network, with the necessary isolation among customers. The intelligence needed for routing among customer branches or some other rich functions, e.g. network address translation (NAT), will be part of the VPN application, which interacts properly with the central ABNO controller for accomplishing end-to-end services. Additionally, in case hybrid cloud services, the needed orchestration with the DC infrastructure can be managed from the Service Management and Control module in a transparent way to the customer. This interaction would allow for elastic cloud services and to isolate between customer and operator cloud management, e.g. when moving VMs among DCs [14], [15].

3.3 Use case III: Mobile broadband access

Mobile access networks are of particular interest for network operators because the high capacity and capillarity they require to satisfy end-user expectations; this will become even more evident with the advent of 5G wireless networks. This scenario forces to explore new ways of deploying the necessary infrastructure to fulfil end-users requirements in a cost effective manner.

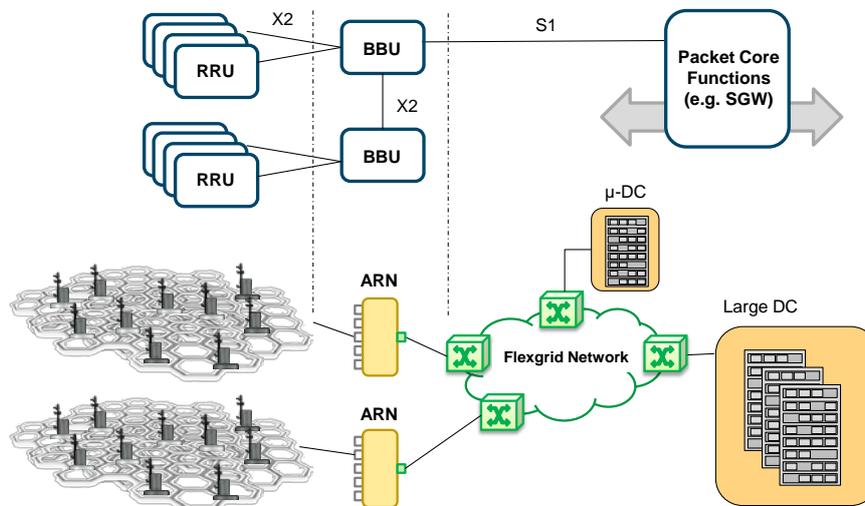


Fig. 5 Mobile broadband access scenario.

One of the recent trends in the mobile industry is the centralization of some functions of the Radio Access Network (RAN), named as the Centralized-RAN (C-RAN) approach. C-RAN proposes allocating common radio-access processing resources, Base-Band Units (BBUs), currently deployed in mobile stations, in a central node, while just keeping remote only the infrastructure strictly needed to provide the wireless connectivity, i.e. the Radio Remote Units (RRUs).

The flexibility of C-RAN can be further extended by virtualizing the BBU functionality. Fig. 5 suggests a mapping of the C-RAN approach to the SHINE architecture. In this case, the BBU is deployed inside the ARNs and connected to the RRUs in the coverage area defined for this service. Such connection is implemented by means of high-speed Common Public Radio Interface line cards, supported by the ARN architecture.

Communication between two mobile stations allocated to the same BBU is performed through an X2 interface implemented directly in the ARN. When mobile stations are attached to distinct ARNs, connectivity is performed through the flexgrid core network available in the SHINE architecture, thus minimizing latency and guaranteeing quality indicators.

The SHINE's Service Management and Control module will be responsible for handling all the necessary connections in the access to ensure service provision. This involves not only the connections for X2 interface, but also the connections needed for the S1 interface that allows communication with the mobile packet core elements placed deep in the network. In fact, such core elements, e.g. Mobility Management Entity (MME), Serving Gateway (SGW), or Packet Data Network Gateway (PGW) in the LTE architecture, can also be deployed in the form of virtualized functions to run either in the large DC or the μ -DCs depending on the required scalability.

An application running on top of the SHINE's Service Management and Control module requests function deployment and the overall connectivity to the system.

3.4 Use case IV: Fixed broadband access

Similarly as for the mobile access, network operators are considering the viability of centralizing certain fixed broadband processing capabilities in access nodes looking for simplifying those network elements that provide connectivity to the end users. This simplification, ideally implementing just programmable forwarding capabilities, will have clear impacts on cost savings and service flexibility.

By doing so, access service provisioning can be highly simplified; end user connectivity will be just a matter of how much bandwidth the available infrastructure (e.g., xDSL, FTTx) provides, whilst centralizing all service logic. Aspects like QoS configuration, traffic filtering and prioritization, etc., can be governed independently of the underlying technology. New services creation would only imply updating service logic in a central point, whereas connectivity upgrades would become just a question of migrating to an access transport technology supporting more capacity.

Besides, this approach would also simplify the way access infrastructure is shared among operators; the flexibility introduced by separating forwarding and control planes in the access allows slicing of network assets, facilitating the control of dedicated portions of the network to different operators. Thus, infrastructure deployed by different operators can be shared, reducing costs and time to market.

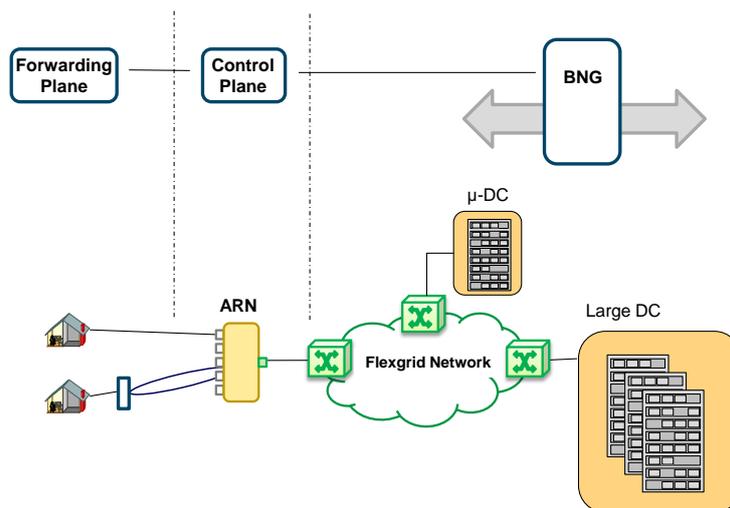


Fig. 6 Fixed broadband access scenario.

Apart from the access, network functions like the Broadband Network Gateway (BNG) can be virtualized and instantiated on cloud infrastructures, scaling according to real needs. Nowadays, monolithic BNGs are rolled-out per PoP considering a maximum user forecast. The reality is that the load of those BNGs is far lower than its maximum capacity, with the constraint that the vacant resources cannot attend (at least in an easy and optimal way) the demand in other PoPs, while consuming scarce resources, like IPv4 addressing. By deploying virtualized instances of BNGs, the right scale for accommodating the actual demand, which changes along the day, can be ensured.

Fig. 6 presents how the SHINE architecture can address this scenario. ARNs host control plane capabilities of fixed access network nodes, with the local instantiation for handling a variety of access technologies (G.PON, Ethernet, etc).

Each ARN supports client interfaces for all “last drop” technologies, abstracting control from data-plane characteristics.

On the other hand, the BNG function is deployed in the large DC or the μ -DCs, where the optimal placement depends on the actual demand. The SHINE architecture facilitates two ways of scaling: *i*) when a BNG is attending a huge number of dispersed customers across the network (i.e., accessing from different PoPs) the BNG can be deployed in the large DC and its network functionality scaled-out/down producing a bigger/smaller BNG; *ii*) in the case a punctual high demand arise in an specific area, a new instance of the BNG can be deployed in a nearby μ -DC.

As in previous use cases, an application running on top of the SHINE’s Service Management and Control module requests function deployment.

4 SUMMARY

In this article, the SHINE architecture has been presented; it orchestrates cloud with heterogeneous access and core networks to provide cloud services, being the base to support NFV. The SHINE approach incorporates: ARN nodes to interface end-users directly to the core, a set of μ -DC placed close to the access to reduce network traffic whilst ensuring low latency, and a Service Management and Control module to dynamically orchestrate cloud and network.

Four use cases addressing different markets of interest for any network operator have been used to illustrate how the SHINE architecture can be used to facilitate virtualizing network functions and orchestrating services.

The major challenges in the deployment of SHINE are associated to end-to-end coordination and the migration from the existing networking framework. In that regard, further specific studies are needed and migration steps need to be considered.

REFERENCES

- [1] Alcatel Lucent White Paper, “Bell Labs Metro Network Traffic Growth: An Architecture Impact Study,” 2013.
- [2] NFV White paper, “Network Functions Virtualisation,” (http://portal.etsi.org/NFV/NFV_White_Paper.pdf), 2012.
- [3] Ericsson White Paper “The Telecom Cloud Opportunity: How telecom operators can leverage their unique advantages in the emerging cloud market.” 2012.
- [4] IBM Institute for Business Value, “The natural fit of Cloud with Telecommunications,” 2012.
- [5] L. Velasco et al., “Elastic Operations in Federated Datacenters for Performance and Cost Optimization,” Elsevier Computer Communications, vol. 50, pp. 142-151 2014.
- [6] L. Contreras et al.: “Towards Cloud-Ready Transport Networks,” IEEE Communications Magazine, vol. 50, pp. 48-55, 2012.
- [7] L. Velasco et al., “Cross-Stratum Orchestration and Flexgrid Optical Networks for Datacenter Federations,” IEEE Network Magazine, vol. 27, pp. 23-30, 2013.
- [8] M. Jinno et al., “Multiflow Optical Transponder for Efficient Multilayer Optical Networking,” IEEE Communications Magazine, vol. 50, pp. 56-65, 2012.
- [9] D. King, A. Farrel, “A PCE-based Architecture for Application-based Network Operations,” IETF draft, work in progress, October 2014.
- [10] ONF Solution Brief, “OpenFlow-enable Transport SDN,” May 2014.
- [11] L. Velasco et al., “In-Operation Network Planning,” IEEE Communications Magazine, vol. 52, pp. 52-60, 2014
- [12] ETSI GS NFV 001, “Network Functions Virtualisation (NFV): Use Cases”, V1.1.1, October 2013.
- [13] Ll. Gifre et al., “Experimental Assessment of ABNO-driven Multicast Connectivity in Flexgrid Networks,” IEEE J. of Lightwave Technology (JLT), 2015.
- [14] M. Mishra, et al., “Dynamic resource management using virtual machine migrations,” IEEE Comm. Mag., vol. 50, pp. 34-40, 2012.
- [15] J. Barrera, M. Ruiz, L. Velasco, “Orchestrating Virtual Machine Migrations in Telecom Clouds,” in Proc. OFC, 2015.