



## Towards Flexible SDN-based Management for Cloud-based Mobile Networks

**Artuso, Matteo; Caba, Cosmin Marius; Christiansen, Henrik Lehrmann; Soler, José**

*Published in:*

Proceedings of IEEE/IFIP Network Operations and Management Symposium

*Link to article, DOI:*

[10.1109/NOMS.2016.7502846](https://doi.org/10.1109/NOMS.2016.7502846)

*Publication date:*

2016

*Document Version*

Peer reviewed version

[Link back to DTU Orbit](#)

*Citation (APA):*

Artuso, M., Caba, C. M., Christiansen, H. L., & Soler, J. (2016). Towards Flexible SDN-based Management for Cloud-based Mobile Networks. In *Proceedings of IEEE/IFIP Network Operations and Management Symposium* [7502846] IEEE. <https://doi.org/10.1109/NOMS.2016.7502846>

---

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

# Towards Flexible SDN-based Management for Cloud-Based Mobile Networks

Matteo Artuso, Cosmin Caba, Henrik Lehrmann Christiansen & José Soler  
DTU Fotonik, Department of Photonics Engineering  
Technical University of Denmark  
2800 Kongens Lyngby, Denmark  
E-mail: {matart, cosm, hlch, joss}@fotonik.dtu.dk

**Abstract** — New technologies and architectures arise in the telecommunications industry in order to cater to the ever growing demands in terms of resource utilization, manageability and user experience. C-RAN and SDN represent two such novel paradigms, both advocating for centralization of a set of resources or control capabilities respectively. The C-RAN architecture requires a significant amount of link capacity which may be a prohibitive factor in its adoption hence an obvious solution is to intelligently share the physical infrastructure among several virtual operators. In this context, a new challenge is to flexibly manage the sharing of the infrastructure. This paper argues that a centralized, SDN-based approach can bring the needed flexibility in the management of the C-RAN. More specifically, this paper proposes a policy-centric management framework, which uses the SDN architecture to enforce various rules for sharing the physical infrastructure. A testbed based on Floodlight and Mininet has been implemented to show the benefits of using this automatic management tool for sharing the mobile site capacity.

**Keywords**—C-RAN, SDN, Network Slicing

## I. INTRODUCTION

The increasing capacity demands in mobile networks are seriously challenging for network operators, which are facing shrinks in revenues due to the high costs of the network and the limited profitability of being “just bit pipes”. To overcome these challenges, a new architecture is being considered in the research and industrial communities that inherits some of the concepts that have been successful in cloud computing systems. This architecture is referred to as Cloud-Radio Access Network (C-RAN) [1]. In C-RAN, the typical functionalities of a base station are split into two entities: a Remote Radio Head (RRH) and a Baseband Unit (BBU). The former remains on the site of the cell next to the antennas to perform digital/analogue conversion and power amplification, while the latter is moved away from the site. Several BBUs are grouped to form a BBU pool (Figure 1). The BBU pool takes care of most of the functionalities of the base station, in particular the digital processing of the signal. These pools therefore resemble a data center that is typical in cloud IT infrastructures. This decoupling process allows for generic nodes on the sites that are “technology-agnostic” and allows for savings both in Capital Expenditure (CAPEX) and Operational Expenditure (OPEX), when compared to traditional mobile network architectures [1].

This is thanks to the use of simpler and cheaper sites and to the co-location of intelligence and processing power in the pool. It also opens up new perspectives and challenges for network operators, infrastructure vendors and even new players. While this new architecture appears interesting at a first glance, it introduces a number of challenges, as the capacity required in the fronthaul network, which is the network interconnecting the BBU pool with the RRHs. The fronthaul transports In-phase/Quadrature (I/Q) data which implies that around 10x higher capacity is needed on the link compared to the capacity offered from the cell at the air interface [2]. Therefore, even though C-RAN can introduce several opportunities for CAPEX and OPEX reductions, the transition cost from a traditional, distributed network, can be prohibitive, especially because of the fronthaul. This transition can be eased if the newly deployed infrastructure is shared by its owner with other network operators, possibly not owning their own physical infrastructure. Another type of operator, often referred to as Mobile Virtual Network Operator (MVNO), is introduced with this business model. An MVNO has a Service Level Agreement (SLA) with the infrastructure owner capturing the relation between the two in terms of allocated resources, prices and services, among others. In this scenario, it is beneficial for the infrastructure owner to maximize the network utilization to profit from the investment made. Therefore, it must be able to flexibly manage the network and the allocation of resources to various MVNOs.

Software Defined Networking (SDN) proposes the decoupling of the control and data planes in the network devices, such that the control plane logic is centralized in a logical entity named SDN Controller (SDNC). The network devices are thus simple forwarding elements instructed by the SDNC on how to process the data packets. Due to its centralization of the control plane logic, SDN has been proposed as a tool to ease the management of networks and potentially lead to better resources utilization, by enabling the use of algorithms that globally optimize the network resources consumption [3]. Furthermore, SDN can also be applied as a tool in mobile networks to facilitate the management of the physical network and sharing of resources between MVNOs. Nevertheless, some caveats must be considered when evaluating the feasibility of applying SDN concepts in cloud-based mobile networks as there are some inherent limitations to the potential flexibility that SDN aims to provide as will be presented in the following.

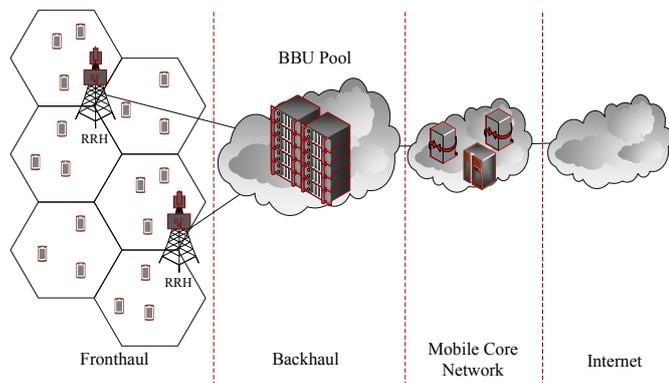


Fig. 1. C-RAN Architecture

The main contribution of this paper is a largely-automatic management tool based on SDN to enable the sharing of a C-RAN. This allows the infrastructure owner to administrate MVNO's traffic, by enforcing different policies based on SLAs. Additionally, the enforcement of different policies is carried out automatically, in order to maximize network utilization and reduce the operational effort. Finally, a prototype for the control plane logic has been implemented using Floodlight [4], and then tested in a C-RAN testbed with 2 MVNOs.

The remainder of the paper is organized as follows. Section II surveys related works in the research community in order to contextualize the contribution of this paper. Section III describes the proposed investigation architecture and the service design, while section IV presents the implementation of the service in a testbed and the model used for the underlying mobile network. Section V presents the results of the service demonstration, and section VI concludes the paper.

## II. BACKGROUND AND RELATED WORK

The feasibility of virtualization and SDN techniques in mobile networks has received substantial interest recently in the research community. Some of the related works take a drastic approach and propose a significantly redesigned mobile network based on the SDN architecture [5, 6, 7 and 8]. These approaches introduce the idea of running the mobile control functions such as the MME in a virtualized fashion on top of the SDNC, which leads to a further flattening of the architecture of LTE-based mobile networks [8]. Additionally, a new protocol is proposed in [6], termed MobileFlow, to cater to the specific needs of a mobile network, arguing that OF is not sufficient. The authors of [5] propose the integration of SDN concepts in the core of the mobile network with SoftCell, a controller that chooses the best path sequence through middleboxes depending on a high-level policy, translated from a service definition. They introduce also the idea of equipping the eNodeBs on the sites with local OF softswitches for fine-grained policy implementation, which is a concept that is also part of our proposed architecture. The applicability of the SDN principles for the radio interface has been studied in SoftRAN [9], which is a framework for decoupling of control and data planes inside the radio nodes. Moreover, a set of geographically-related base stations is seen as a single logical base station, hence resource pool, from where resources are allocated in a centralized fashion along three coordinates: space, time and frequency slots.

With respect to the RAN, centralizing the control plane logic, as SDN promotes, can be readily combined with the C-RAN architecture, which proposes the centralization of radio processing capabilities. The idea of applying SDN in C-RAN setups has been discussed in the literature, and it is a potential solution to the increased flexibility requirements for 5G mobile networks [10] [11] [12]. In [13], the authors investigate the requirements in terms of elasticity and scalability of 5G mobile networks, and propose the use of SDN to ensure such elasticity on the backhaul network. They claim that using SDN in mobile networks is the natural evolution of the control and data plane separation that is being pursued by the 3<sup>rd</sup> Generation Partnership Project (3GPP) with Long Term Evolution (LTE). Their proposed architecture contains several aggregator OF switches between the eNodeBs and the core, having a similar role as the local switch used in [5]. Moreover, Virtual Local Area Network (VLAN) tagging based on IEEE 802.1ah is used to group the UEs in the network, which constitutes another building block in our proposal.

The current paper leverages some of the concepts mentioned above, and further enhances them by focusing on their applicability in the C-RAN infrastructure-sharing scenario. It is argued that the promise of flexibility carried by SDN is largely needed in C-RAN and that its centralized architecture can ease the integration of such techniques. This however introduces some peculiarities not present in traditional/distributed mobile networks. The most relevant one is that the baseband traffic transported over the fronthaul network has a constant bit rate, as the entire cell bandwidth has to be sampled, regardless of how it is used at a particular moment. Therefore, the physical link capacity cannot be shared as in packet-based networks, unless a packet-based transport is used in the fronthaul, so that it can be shared among the sites and with other access technologies [14]. In a "traditional" C-RAN it is therefore necessary to rethink the sharing at a higher layer, so that the I/Q stream already carries a mix of traffic for the different MVNOs, which reflects the high-level policies in the SLAs between the MVNOs and the infrastructure owner. A new sharing mechanism is proposed herein that is aware of these limitations and overcomes them by using virtualized switches to shape the traffic of each RRH site. This is achieved by installing OF entries in these software switches, which will eventually dictate how the capacity of the site should be used by each MVNO. Moreover, the SLA policies for shaping the traffic span a long time period as they are changed with granularity of hours, therefore reducing the control overhead, as it will be presented in the following.

Besides the considerations on SDN-related works, other QoS management techniques are employed in mobile backhaul networks. Principles such as shaping, fair/weighted round-robin scheduling and weighted fair queueing scheduling are all well-known techniques that can be leveraged and migrated in a SDN-based service design [15]. As the particular focus of this paper is on capacity limitation at the shared Ethernet interface of a virtual switch, techniques related to QoS management with color-based policing are a possible addition to the service in order to deal with excessive traffic demands [15]. As it will be pointed out later – however – the current service definition and implementation is based on sharper capping, while handling of excessive demands is left for future upgrades.

### III. C-RAN SHARING SERVICE MODELLING

In this section, first the C-RAN sharing service will be described, then some considerations on network topology and on modelling assumptions will be addressed, similarly to the work done by the authors of [16].

#### A. Service Description

The core of the service is the sharing of a C-RAN between 2 MVNOs. Their coexistence is ensured by an agreement between them and the infrastructure owner. The SLA is based on a series of values for certain QoS performance metrics offered by the MVNOs to their customers (e.g. bandwidth). The entries in such SLA can consist of guarantees that the metrics will be kept within a range agreed by the two parties, provided that certain network conditions apply. In case an SLA promises very stringent guarantees for a particular performance metric – say a very tight delay budget – the advantage is twofold: on one hand, as it is more challenging to guarantee the metric according to the SLA, then the infrastructure owner can charge more the MVNO for such an agreement. On the other hand, this is also beneficial for the MVNO as it would be able to offer more predictable services. Additionally, an SLA entry could also not provide any guarantee for a certain metric, which gives more flexibility to the infrastructure owner. Depending on the SLA, the performance metrics can address different levels of

granularities, the basic one being a limitation of the overall site capacity that a MVNO can use. This can be extended by policies on a particular traffic type or user group (premium users and similar). More fine-grained policies lay possibly outside the scope of the SLA between the infrastructure owner and the MVNO, as they are more related to how the MVNO manages its own traffic internally. However, it should be possible for the MVNO to map its internal policies to those defined in the SLA. Additionally, the load of a site varies depending on its location and on the time of the day, such as office sites during working hours. In these “busy hours”, the decision and enforcement of the policies is crucial: the SDNC needs to decide which MVNO to prioritize based on the SLA. This should be possibly defined in an intuitive way and should abstract the underlying nature of the C-RAN so that the service logic can be migrated to traditional, distributed mobile networks. A design goal for the service is therefore to keep the high-level definition of the policies as simple as possible, while moving their translation and manipulation to other dedicated components of the SDNC. Some assumptions are thus necessary when modelling the service and the test network, in order to address the challenges of C-RAN, as it will be presented in the next part. In the particular case of the service designed and implemented for the testbed described below, the focus is on a per-site rate limitation per MVNO. In order to exploit load fluctuations throughout the

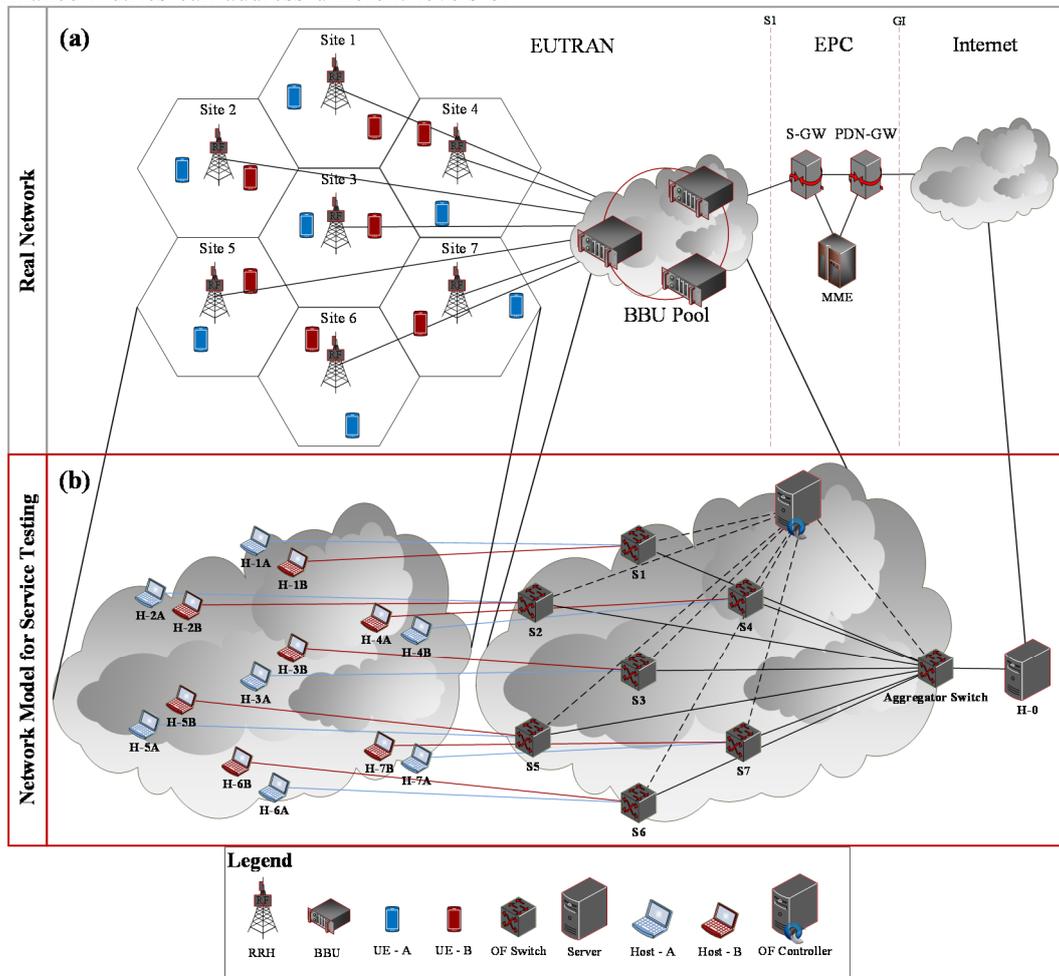


Fig. 2 Mapping of Real Network (a) into Network Model for Service Testing (b)

day and in different site types (e.g. residential, office), the service applies different capping rates automatically, by monitoring the hour of the day, as agreed in the policies between the MVNOs and the infrastructure owner. This has the advantage of isolating the MVNOs' resources and reducing the manual management – and thus the OPEX. Additionally, the service logic in the SDNC deals with rules that are defined to match closely their high-level correspondents in the SLAs, so to allow for a fine-grained control, in case the service is expanded and empowered with further elements from SLAs, such as latency requirements, handling of excessive traffic requests and so forth.

### B. Network Topology and Modelling Assumptions

The reference network topology used is presented in Figure 2a. It consists of a C-RAN with 7 cells. Each cell has 1 RRH and 1 UE/host per each MVNO – referred to as A (blue) and B (red) for simplicity in the following. The RRHs are connected with the BBU pool and subsequently to the EPC towards the internet. The entities in the network acting as UEs can be considered as hosts in the SDN terminology. This is clearly an abstraction and does not consider the LTE-specific lower layers of the protocol stack. These L1 and L2 layers are terminated at the UE on one end and at the eNodeB or BBU at the other, but our investigation focuses on at most the network layer for the UE, so the abstraction is considered reasonable. Furthermore, in a SDN-based network hosts are generally connected directly with an OF-enabled switch, which is not necessarily the case in a mobile network, where the UEs have to get their traffic processed by several nodes before reaching the first switch in the backhaul network or in the core. Taking into account these considerations, the model of the network used for testing the service is presented in Figure 2b – where the EPC has been removed as its functionalities do not affect directly the behaviour of the service. The role of the internet cloud is assumed by a special host acting as a server (H-0). As regards the EUTRAN, the traffic generated by the UEs is modelled with hosts corresponding to the topology of the real network. Finally, in the BBU pool the model does not cover the generation of baseband traffic and LTE's lower layers as said above, instead, only elements for shaping the traffic are considered. These are virtual OF switches (1 per site) that constitute a virtual overlay network on top of the BBUs (Figure 2b) as an additional layer in the protocol stack of the BBUs, above their IP interfaces towards the core. The role of these RRH-related OF switches resembles that of those co-located with the eNodeBs in the works [5,13] presented above. Additionally, an aggregator switch is also needed before the server node. This set of switches is used to rate limit the traffic in both the Downlink (DL) and Uplink (UL) directions, based on the SLA policies enforced by the SDNC that also resides in the BBU pool. This is meant to constrain the maximum available capacity of a site for a certain MVNO and can be seen as a loose slicing of the network resources among the MVNOs, done by the infrastructure owner, directly in the BBU pool. Furthermore, each host is allowed to communicate only with the host used as server, to mimic typical data sessions in a mobile network. This topology eases the definition of the policies in the network so

that it is necessary to limit the traffic only to a single host – that is the one acting as server.

### C. C-RAN Tidal Effect and Traffic Profiles

One of the benefits of using C-RAN solutions can be seen when the same BBU pool serves sites in locations with different traffic profiles throughout the day, by exploiting the so-called *tidal effect* [1]. Sites should then be characterised differently in the model, so that some of them corresponds to areas mainly with offices and the other to areas mainly with residences. The switches corresponding to sites 1-4 are thus labelled as “office areas”, while the switches corresponding to sites 5-7 as “residential areas”. Finally, it is necessary to assume typical traffic profiles for the two site types to define meaningful policies for the MVNOs. Sites that at a certain hour of the day are more loaded require more stringent limitations of the use of the site capacity, while lightly-loaded sites might have more relaxed limitations, or even none – except for the maximum capacity available. A typical trend of the traffic profiles throughout the day for office and residential cells is shown in Figure 3 [1]. It is possible to notice that office cells are more loaded during working hours, when residential cells are lightly loaded, while the vice versa occurs towards the end of the day. This pattern is used to design example policies to demonstrate the service in the testbed.

### D. Service Details

The service has elements of network slicing and QoS management as it primarily limits the bandwidth to the hosts in the network, depending on the site characteristics and on the time of the day. To achieve this, the service uses two policies translated from the SLAs with the MVNOs. Each entry in the policy is specific to a particular host whose bandwidth towards the server is limited. A policy entry is therefore defined by an identifier, source and destination of the flow to be limited, the time interval in which the entry is operative and finally the guaranteed percentage of site capacity for that flow. This last element translates into the maximum bandwidth of an MVNO in that time interval. Currently, the model does not cover the management of excess capacity that can be requested, which is left for future work.

The policy time intervals are in the range 0 to 23, as it is assumed for simplicity that rules last at least 1 hour. This is considered a reasonable design assumption as the service focuses on long-term changes happening on a daily basis.

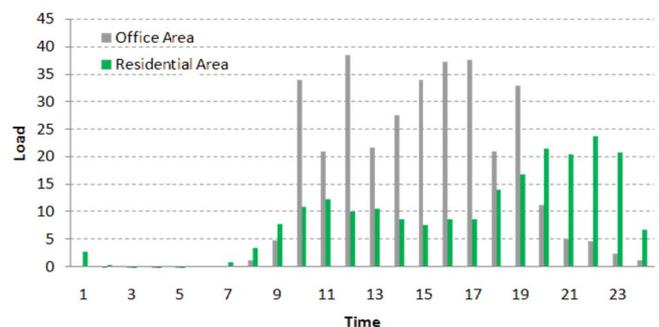


Fig. 3. Typical Traffic Load Variations Throughout the Day

Finally, the rules in the policies should be consistent among MVNOs. This means that for a particular site and time of the day, the aggregated guaranteed capacity of overlapping rules should not exceed the maximum capacity of the site. This is a simplifying assumption as in some cases it could be possible to “over-sell” capacity if additional knowledge about the site and the traffic patterns of the MVNOs is known to the infrastructure owner. In the case modelled herein, it is assumed that no “over-selling” is done, and the policies are designed with this assumption in mind. This can be seen in the following Table I. Here, the values of the guaranteed bandwidth do not sum up to more than 100% when entries for the two MVNOs refer to the same site in the same time interval. This except in the case where there is no limitation on the fraction of site capacity that a MVNO can use (e.g. first entry in Table I). The cases where there is no capacity limitation are particularly interesting for the service, as it will be pointed out in Section V. These cases, in fact, correspond to a network where the service is not implemented and the traffic flows belonging to each MVNO compete to access the site capacity and eventually share it in a best-effort fashion.

TABLE I. EXAMPLES OF POLICY DEFINITIONS

Rule ID	Site ID	Interval	Guaranteed Capacity (%)	
			MVNO A	MVNO B
1	1 – 4	00 – 09	N.L. <sup>a</sup>	N.L. <sup>a</sup>
2	1 – 4	09 – 13	80	20
3	1 – 4	13 – 17	80	20
4	1 – 4	17 – 00	40	40
5	5 – 7	00 – 10	N.L. <sup>a</sup>	N.L. <sup>a</sup>
6	5 – 7	10 – 17	20	N.L. <sup>a</sup>
7	5 – 7	17 – 21	40	40
8	5 – 7	21 – 00	20	80

<sup>a</sup> Not Limited

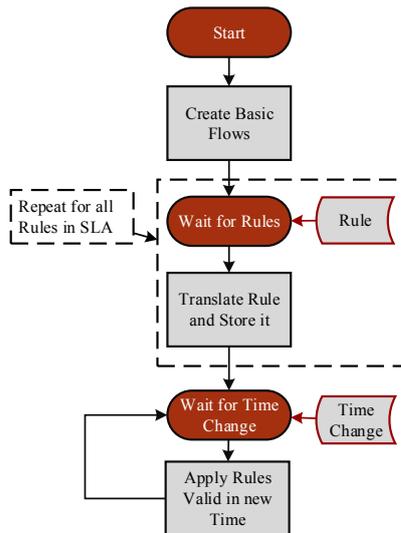


Fig. 4. Service Behavior Flowchart

For testing and emulation purposes, the service should allow manual changes of the time of the day, as well as enforcing, editing, and deleting policy rules. The possibility of deleting/editing rules is relevant when the service is used to apply updates to the SLA. As a result, Figure 4 presents a basic flowchart of the service behavior without the editing and deletion of rules.

#### IV. SERVICE IMPLEMENTATION

In the testbed, network emulation is done with Mininet [17], while the controller logic is implemented using Floodlight [4]. In the remaining of this section, the different parts of the implementation will be presented.

##### A. Network Model

For the scope of the service presented herein, it is possible to consider the UEs only from the IP layer and up, and so model them as hosts connected through Ethernet interfaces to the OF switch corresponding to the RRH site. The site capacity is set to 10 Mbps, which corresponds to the maximum rate allowed on the Ethernet link between each RRH switch and the aggregator switch (Figure 2.b). This is a fraction of a typical estimation for the maximum capacity achievable by a single LTE cell with 20 MHz of available bandwidth (~100 Mbps). For the sake of simplicity and to avoid performances issues in the virtualized environment (i.e. Mininet), the testbed is downsized to 10% of the estimated capacity of a real LTE cell. To implement the QoS classes, rate-limited queues are added to the switches’ ports, so that the host-server flows are capped according to the values defined in the policies. Such policies have to be checked for consistency by the infrastructure provider, which is currently done at the definition stage, prior to their installation. For completeness, in section V a case of “over selling” will also be presented as it is an interesting scenario to test the service, as mentioned earlier. This case consists of a combination of more than 100% for the guaranteed capacity.

##### B. Controller Logic

The main control and management plane functionalities needed are the topology discovery, rule consistency check, flows installation and modification, depending on the hour of the day and on the site type, and finally interaction with the infrastructure owner to receive rules. Node identification for hosts and switches is done by the SDNC with the IP address in the former case or with the Data Path ID (DPID) – an identifier for switches used in Mininet – in the latter.

#### V. TESTING AND DEMONSTRATION

Testing and demonstration of the service in the Mininet network model is achieved by installing the two lists of policies presented above through an interface for the infrastructure owner. As stated above, the capacity of each site (and so the bandwidth of each link between a RRH switch and the aggregator switch) is limited to 10 Mbps in this example. This eases the rate limitation as for example a policy rule stating a guaranteed capacity of 80% in Table I, corresponds to a limitation of the maximum rate to 8 Mbps for that MVNO’s traffic. Figure 5 presents a summary of the measurements taken

for the average Transmission Control Protocol (TCP) throughput, using the tool iPerf [18] between the hosts and the server with MVNO A in blue, and B in dotted red. The aggregated throughput does not reach 100% of the site capacity, due to the congestion-control in TCP and to the fact that average values are plotted. Figure 5a shows that the network capacity is shared among the hosts of the 2 MVNOs, when no policies are implemented in the network. As can be seen in the resulting graph, there is no strict guarantee on how the capacity is going to be used and the connections tend to adapt in a best-effort way, both at around 50% of the total available, though with notable fluctuations. On the other hand, when policies are in place and therefore the C-RAN sharing service is being used, the mixture of traffic in the network can be controlled, depending on the contents of the SLAs, as can be seen in Figure 5b. This graph illustrates a combination of three different cases of capacity sharing. In the sites 1 – 3, MVNO A is limited to 2 Mbps, while MVNO B to 8 Mbps. Site 4 has 4 Mbps per MVNO, while in sites 5 – 7 the configuration is specular to the one used in the sites 1 – 3.

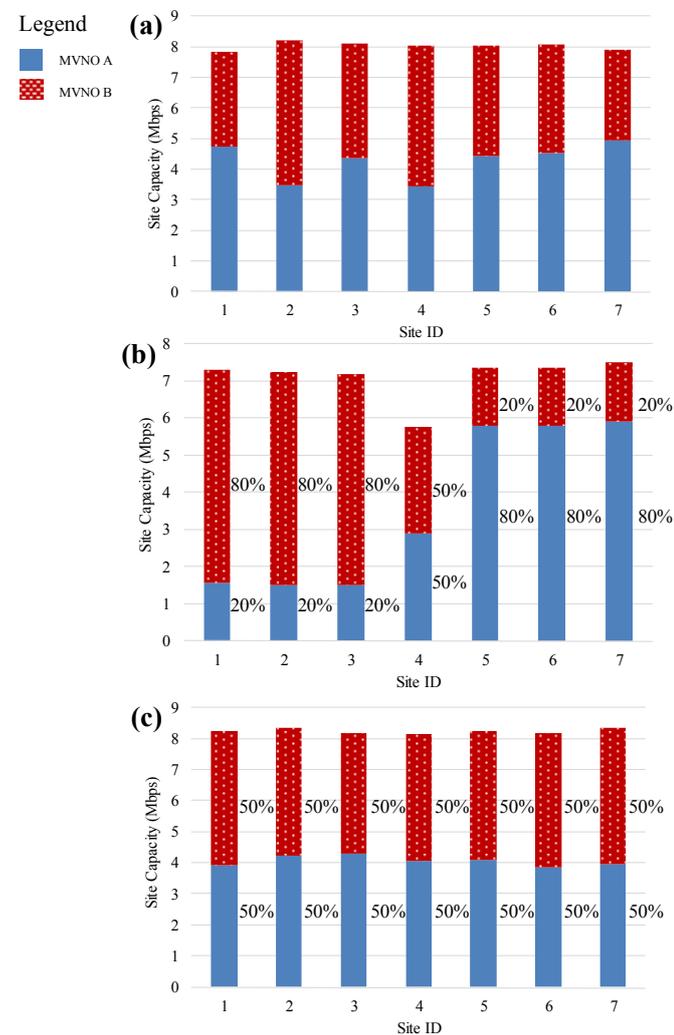


Fig. 5. TCP Throughput Measurements: a) Without Policies; b) With 3 Different Policy Rules; c) With 1 Common Policy Rule

As can be seen, the average throughput reflects the configuration enforced by the controller in the network, still keeping into consideration what said above for the maximum values of the TCP bandwidth. This allows more control in the network usage and it is therefore possible to avoid part of the variability noticed in Figure 5a. Finally, the case of site 4 in Figure 5b gives an interesting insight in what happens when a 50/50 share is enforced. In Figure 5c, results are presented for measurements taken when all the sites implement the policy of up to 8 Mbps for both MVNOs. This case could correspond to the one in Figure 6a without rules so that the MVNOs share the available bandwidth. However, in this case the result shows a more consistent sharing of 50/50 throughout the network where both MVNOs achieve ~4 Mbps. This also means that the design of the policies should be done carefully as presented in Section III, as the maximum of a policy might not be guaranteed when the aggregated guaranteed capacity exceeds the total available.

## VI. CONCLUSIONS

In this paper we presented the design of a sharing service for a C-RAN based mobile network where an infrastructure owner shares the network with other operators (MVNOs). This is based on policies that limit the amount of capacity in each site that a MVNO can use. The limitation broadly depends on two factors: the nature of the site (located in office or residential area), and on the time of the day. Additionally, the infrastructure owner is able to ensure a good level of isolation between the resources allocated to each MVNO by using the proposed policy service. The purpose of the service was also to provide an easy way for the management of the network and of the policies, which is achieved by means of an interface for the infrastructure owner in the implementation presented. The format of the rules has been designed to adhere closely to high-level policies that could be defined in SLAs between the MVNOs and the infrastructure owner.

The service has been implemented in a testbed based on the SDNC Floodlight and a network model in Mininet to demonstrate its role for infrastructure sharing. TCP throughput measurements show that the usage of the service ensures that the network capacity is shared according to the SLAs and not to a best-effort behavior, giving more control to the infrastructure owner to offer guarantees to its MVNOs.

Finally, the service allows for automation of policies enforcement as the rate limitations are changed automatically throughout the day to match their high-level description in the SLAs.

## ACKNOWLEDGMENT

This work was partially sponsored by the 7<sup>th</sup> Framework Programme for Research of the European Commission COSIGN project, under grant number 619572 – COSIGN.

## REFERENCES

- [1] A. Checko, H. L. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M. S. Berger and L. Dittmann, "Cloud RAN for Mobile Networks - a Technology Overview", *IEEE Communications Survey & Tutorials*, vol. 17(1), pp. 405-426, 2015.
- [2] A. Pizzinat, P. Chanclou, T. Diallo and F. Saliou, "Things you should know about fronthaul," in *European Conference on Optical Communication (ECOC)*, Cannes, 2014.
- [3] H. Ali-Ahmad, C. Cicconetti, A. De La Oliva, V. Mancuso, M. Reddy Sama, P. Seite and S. Shanmugalingam, "An SDN-Based Network Architecture for Extremely Dense Wireless Networks," in *IEEE SDN for Future Networks and Services (SDN4FNS)*, Trento, 2013.
- [4] Big Switch Networks, "Project Floodlight," [Online]. Available: <http://www.projectfloodlight.org/floodlight/>. [Accessed 21 August 2015].
- [5] X. Jin, L. E. Li, L. Vanbever, J. Rexford, "SoftCell: Taking Control of Cellular Core Networks," [Online]. Available: <http://arxiv.org/pdf/1305.3568.pdf>. [Accessed 19 August 2015].
- [6] K. Pentikousis, Y. Wang and W. Hu, "MobileFlow: Toward software-defined mobile networks", *IEEE Communications Magazine*, vol. 51, no. 7, pp. 44-53, 2013.
- [7] C. J. Bernardos, A. De La Oliva, P. Serrano, A. Banchs, L. M. Contreras, H. Jin, and J. C. Zúñiga, "An architecture for software defined wireless networking", *IEEE Wireless Communications Magazine*, vol. 21, no. 3, pp. 52-61, 2014.
- [8] J. Costa-Requena, "SDN integration in LTE mobile backhaul networks," *International Conference on Information Networking (ICOIN)*, Phuket, 2014.
- [9] A. Gudipati, D. Perry, L. Erran and S. Katti, "SoftRAN: Software Defined Radio Access Network", *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, 2013.
- [10] P. Rost, J. C. Bernardos, A. D. Domenico, M. D. Girolamo, M. Lalam, A. Maeder, D. Sabella and D. Wübben, "Cloud Technologies for Flexible 5G Radio Access Networks," *IEEE Communications Magazine*, vol. 52, no. 5, pp. 68-76, 2014.
- [11] M. Draxler and H. Karl, "Dynamic Backhaul Network Configuration in SDN-based Cloud RANs", [Online]. Available: <http://arxiv.org/pdf/1503.03309>. [Accessed 21 August 2015].
- [12] M. Arslan, K. Sundaresan and S. Rangarajan, "Software-defined networking in cellular radio access networks: potential and challenges", *IEEE Communications Magazine*, vol. 53, no. 1, pp. 150-156, 2015.
- [13] J. Costa-Requena, R. Kantola, J. Llorente, V. Ferrer, J. Manner, A.Y. Ding, L. Yanhe and S. Tarkoma, "Software defined 5G mobile backhaul," *1<sup>st</sup> International Conference on 5G for Ubiquitous Connectivity (5GU)*, Levi, 2014.
- [14] M. Artuso, A. Marcano, H. L. Christiansen, "Cloudification of mmWave- and Packet-based Fronthaul for Future Heterogeneous Mobile Networks," *IEEE Wireless Communications*, vol. 22, no. 4, 2015.
- [15] J. Salmelin and E. Metsälä, "Mobile Backhaul", John Wiley and Sons, Inc. 2012.
- [16] M. Fiorani, A. Rostami, L. Wosinska and P. Monti, "Transport Abstraction Models for an SDN-Controlled Centralized RAN," *IEEE Communications Letters*, vol. 19, no. 8, 2015.
- [17] Mininet Team, "Mininet," [Online]. Available: <http://mininet.org/>. [Accessed 21 August 2015].
- [18] Iperf Team, "iPerf," [Online]. Available <https://iperf.fr/> [Accessed 21 August 2015].