Wireless Access Virtualisation: Physical versus Virtual Capacities

Luísa Caeiro[1,2], Filipe D. Cardoso[1,2]
[1]ESTSetúbal
Polytechnic Institute of Setúbal
Setúbal, Portugal
e-mail: {luisa.caeiro, filipe.cardoso}@estsetubal.ips.pt

Luis M. Correia[2]
[2] IST/INOV-INESC
University of Lisbon
Lisbon, Portugal
e-mail: luis.correia@inov.pt

Abstract - **This paper addresses the virtualisation of wireless access in order to provide the required capacity (data rate) to a set of Virtual Base Stations (VBSs). The approach is based on a Virtual Radio Resource Allocation algorithm, OnDemandVRRA, which manages the allocation of the physical radio resources to the VBSs, in order to follow the contract and maintaining isolation among the VBSs, according to the type of guarantees of the VBSs, the amount of contracted capacity, and the VBSs' utilisation. Taking the variability of the wireless medium into account, the algorithm continuously influences RRM mechanisms, namely admission control and MAC scheduling, to be aware of the VBSs' state relative to the service level agreement, in order to compensate for this variability. The algorithm has been tested to evaluate its behaviour, concerning the amount of virtual contracted capacity versus the physical one. From simulation results, it can be concluded that the total capacity contracted for guaranteed VBSs should be limited according to the average capacity provided by the physical set of serving base stations. As an example, although the guaranteed VBS contracted capacity is always achieved, for the case where the guaranteed VBS contracted capacity is about 85% of the average cluster capacity, the cluster serving data rate decreases about 20% relative to the maximum achieved for the *Best Effort Overbooking* use case.**

Keywords: ***Future Internet, Virtual Networks, Heterogeneous Networks, Radio Resource Management, Resource Allocation.***

I. INTRODUCTION

Network virtualisation is an abstraction process, aiming at separating the logical network functionality from the underlying physical network resources. It enables the aggregation and provision of the network by combining different physical networks into a single virtual one, or splitting a physical network into multiple virtual ones, which are isolated from each other. Network virtualisation has been introduced as a tool for large scale experimental networks, e.g., PlanetLab [1] or GENI [2], but it has been also proposed as an approach for a future Internet architecture, [3] and [4].

The virtualisation of wireless resources introduces some new challenges, due to the specific characteristics of the wireless environment. On the one hand, the isolation of traffic cannot be guaranteed due to the scarcity of the radio spectrum, which cannot be over provisioned, while on the other hand, radio signal propagation is a very node-specific property, being difficult to control, and has a significant impact on most Virtual Networks (VNets) [5]. The slicing process in wireless networks has also some specific issues, derived from the characteristics of the medium; the provisioning of slices to multiple VNets with different radio links requires the capability to share radio resources, while at the same time avoiding interference among the different VNets [5].

Wireless virtualisation can be seen as a problem of wireless networks sharing for multi-operator networks. In the current mobile communications marketplace, functionalities that enable various forms of network sharing are becoming more and more important. The topic has already been explored by the research community for the introduction of Mobile Virtual Network Operators (MVNOs) in 3G (Third Generation) systems [6], and is now a hot topic for Long Term Evolution (LTE). LTE network sharing standards can be found in [7], and some proposals have been presented for active RAN sharing in [8]. More recently, the 3GPP group for RAN Sharing Enhancements identified a set of use cases in order to allow a more flexible and efficient RAN sharing [9].

In the context of network virtualisation, the target is to manage radio resource sharing for the VNet's aggregated link, abstracting the involved wireless systems. Wireless virtualisation for specific Wireless and Mobile Networks has been recently addressed in literature, e.g., LTE [10] and WLAN [11]. The majority of these approaches address mainly wireless resources virtualisation, which is not the focus of this work, only a few of them tackling the management of radio resources to be shared among the several VNets. Furthermore, in these approaches, the assignment of radio resources to VNet end-users is handled within one physical resource, in which the virtual resources are instantiated. Still, they do not address the allocation of radio resources based on the capacity required to the virtual resources, but based on a required amount of radio resources, which may perform differently according to the wireless medium conditions, possibly not providing the requested capacity.

In the wireless networks sharing approach, operators are forced to use similar network functions, as defined by 3GPP specifications, hence, the possibility of having different multiple VNets with its own functions and communication protocols, isolated from each other (the main advantage of network virtualisation), cannot be achieved. Still, without having an integrated perspective relative to multiple radio access technologies, the abstraction of the wireless access is only partially made, avoiding one to take advantage of all available wireless infrastructures. Furthermore, the several models proposed for radio resource sharing are not based on capacity request, the allocation of radio resources being more

5GU 2014, November 26-27, Levi, Finland

DOI 10.4108/icst.5gu.2014.258122

or less fixed and not dynamically adapted to the network state, in order to satisfy the requested capacity. This may lead to situations in which VNets are running out of contract, denying service to end-users, even when radio resources are available.

In this paper, several strategies for the provision of physical capacity to multiple VBSs with different type of requirements are presented and evaluated. The management of radio resource sharing is based on the perspective of VBSs as an aggregated connectivity resource abstracted from a group of Radio resource Units (RUs) of different Radio Access Technologies (RATs), allowing to benefit from Cooperative Radio Resource Management strategies, e.g., [12] and [13]. Instead of looking at wireless virtualisation from the perspective of the instantiation of virtual machines in the wireless nodes, our view is the virtualisation of the wireless access to provide a contracted capacity to the VNet, in order to serve its end-users. Our approach is then agnostic to the point where the virtual node instantiation takes place, being possible to have the virtual nodes in each physical wireless node, or somewhere in the cloud requesting virtual access over a given geographic area covered by a set of wireless nodes. It is worthwhile noting that this capacity can be modified on demand, without manually changing the configuration of the network.

This paper is structured as follows. The Virtual Radio Resource Allocation (VRRA) approach is presented in Section II. Scenarios and results are presented in Sections III and IV, respectively; finally, conclusions are drawn in Section V.

II. VIRTUAL RADIO RESOURCE ALLOCATION

A. Network Architecture

The network architecture to cope with our approach was presented in [14]. It considers the existence of a generic Virtual Network Enabler, which processes the demands for virtual networks and negotiates with the Infrastructure Providers to use their physical networks for the provision of capacity and VNets' creation, Figure 1. A Virtual Resource Allocation function is identified within the Virtual Network Enabler, to include the management of resource allocation to satisfy the VNets' requirements.

From the physical viewpoint, one defines a cluster as a set of BSs from various RATs as the management unit in terms of VRRA. VRRA is implemented at an additional level of abstraction, the virtual RRM level, allowing to follow an approach of integration of both RRM and Cooperative RRM levels. These two levels actively participate in the process to achieve the main target of provision of the contracted level of service for all MVNOs operating over the common infrastructure. VRRA assumes the coordination role of the underlying RRM levels as it is aware of VNets requirements and has the responsibility to satisfy them. Since one is dealing with heterogeneous wireless networks, it coordinates the Cooperative RRM level, managing in a common way the diverse physical networks in the cluster area, and the RRM level, being implemented within the physical networks.

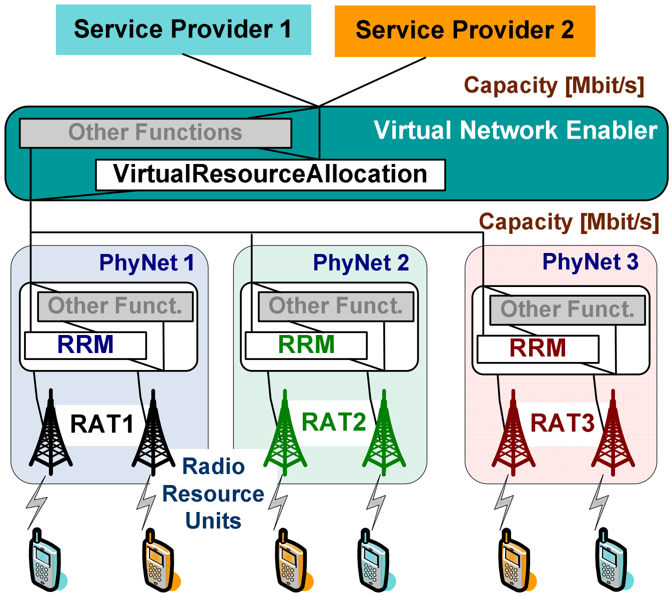

Figure 1. Physical Network Architecture.

MVNOs are the players that manage and operate VNets, including their VBSs, to satisfy Service Providers' requests. They know only the virtual resources that are part of the VNet with their associated capacity, the set of physical resources being hidden from them, Figure 2.

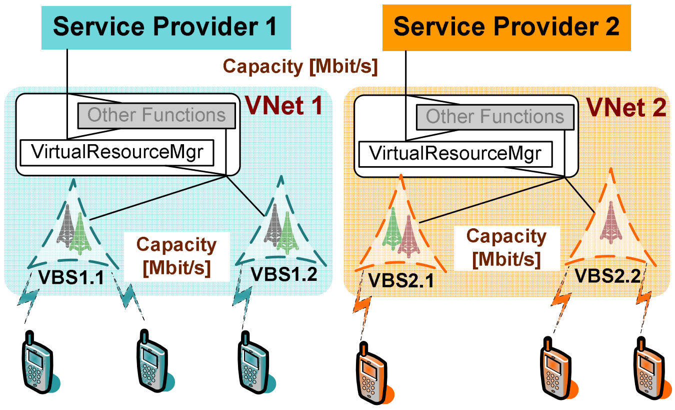

Figure 2. Logical Network Architecture.

In order to make use of a service, the end-user connects physically to the BSs, but the connection to the VNet providing the service is made logically via a VBS, through a virtual link, Figure 2. Concerning the VNets, several VBSs from various VNets may exist in the cluster. The VBS is defined according to the contracted capacity and the capacity assigned to end-users, the *VBS Serving Data Rate,* R_{serv}^{VBS}. Two types of VBSs are considered in this work:

- GRT VBS - characterised by a Minimum Contracted Data Rate, R_{min}^{VBS}, to be guaranteed for all time frames.
- BE VBS - defined by a Reference Contracted Data Rate, R_{ref}^{VBS}, which is indicative and should be followed in a given percentage of the total number of time frames.

B. Strategies and Algorithm

The main target of VRRA is to provide the required capacity to VBSs, optimising radio resources utilisation. VRRA is developed to perform the mapping between virtual and physical links, dynamically adapting the allocation of radio resources to the wireless conditions and VNet utilisation. Supported by the Cooperative RRM, it manages the aggregated capacity provided for the VBSs by sharing the set of available RUs from all RATs.

Applying constraints to the RRM mechanisms according to VBS capacity utilisation, namely, admission control and MAC scheduler, it maps the requested capacity for a particular RAT onto RUs assigned to end-users. It is important to ensure consistency among the decisions taken at the different levels of RRM in order to achieve an overall coherent behaviour. Given that at the VNet level one should have the perspective of the several RATs, the OnDemandVRRA algorithm takes decisions at a time scale that is defined for Cooperative RRM. This time scale is taken as the major common denominator of all RATs, for the sake of simplicity.

The proposed OnDemandVRRA algorithm is a heuristic algorithm, which manages the allocation of RUs among VBSs only when they are requested by VNet end-users. It is responsible for dynamically (re)allocating RUs, satisfying a *Minimum Contracted Data Rate* for GRT VBSs (1), and aiming at a *Reference Contracted Data Rate* for BE VBSs, i.e., minimising its difference to the *VBS Serving Data Rate* (2):

$$R_{serv}^{VBS_i}[\text{bit/s}] \ge R_{min}^{VBS_i}[\text{bit/s}] \quad ,\forall\, VBS_i \in VBS_{GRT} \qquad (1)$$

$$min\left(R_{ref}^{VBS_j} - R_{serv}^{VBS_j}\right)[\text{bit/s}] \quad ,\forall\, VBS_j \in VBS_{BE} \qquad (2)$$

subject to:

$$R_{req}^{VBS_i}[\text{bit/s}] \ge R_{min}^{VBS_i}[\text{bit/s}] \quad ,\forall\, VBS_i \in VBS_{GRT} \qquad (3)$$

$$R_{ref}^{VBS_j}[\text{bit/s}] > R_{serv}^{VBS_j}[\text{bit/s}], \ \forall\, VBS_j \in VBS_{BE} \qquad (4)$$

where, VBS_{GRT} and VBS_{BE} are the set of GRT and BE VBSs in the cluster; $R_{serv}^{VBS_i}$ is the *VBS Serving Data Rate* of VBS i; and $R_{req}^{VBS_i}$ is the *VBS Requested Data Rate*, i.e., the total data rate requested by end-users in VBS i given by:

$$R_{req}^{VBS_i}[\text{bit/s}] = \sum_{n=1}^{N_{EU}^{VBS_i}} R_{req\,n}^{EU}[\text{bit/s}] \qquad (5)$$

where $R_{req\,n}^{EU}$ is the data rate requested by end-user n.

To cope with these objectives, the proposed OnDemandVRRA algorithm is supported by a VNet priority scheme and a data rate reduction strategy, besides the access selection mechanism. Concerning access selection, end-users are connected to the different VBSs according to the requested service and their contract with the MVNO(s). The physical connection is established over one of the existing RATs in the coverage area, according to a list of preferences related to the requested service, end-user's location and available capacity.

The VNet priority scheme running at cluster level assumes a coordination role and enables to set differentiated end-users according to the type of VNet and the *VBS Serving Data Rate*. VBSs are initialised to be handled with priority, all BSs in the cluster being informed of this. When R_{min}^{VBS} is reached, such priority is deactivated and the end-users who wish to connect to this VBS are handled without differentiation. This priority scheme allows the implementation of a data rate reduction strategy whenever the GRT VBSs have priority, preventing starvation on BE VBSs when the contracted data rate in GRT VBSs is reached.

The data rate reduction strategy, essential to compensate for possible end-user data rate decreases due to degradation of medium conditions or the movement of end-users, is applied whenever the VBS priority is activated. Although another data rate reduction strategy could be applied, the adopted data rate reduction strategy is as follows. Whenever the VBS priority is activated for a GRT VBS, and the end-user tries to connect to a BS in which there are not enough RUs for his/her service, BE end-users connected to the BS are reduced according to: (i) the out of contract rate of the VBS they are connected to; (ii) the QoS priority class of the performed service, end-users performing services with lower priority being the first to be reduced; and (iii) their Signal to Interference plus Noise Ratio (SINR), end-users with lower SINR being reduced first to allow optimising radio resource utilisation. Still, if there are not enough RUs to reach the requested data rate, the evaluation of co-located BSs is performed, in order to select the one with enough RUs available and with the minimum cost to handover end-users.

C. Metrics

The *Average Serving Data Rate* allows evaluating the algorithm ability to allocate the adequate quantity of RUs to the VBS, in order to satisfy the VBS contracted data rate. It is the average of R_{serv}^{VBS} over the total number of time frames in the observation time interval:

$$\overline{R_{serv}^{VBS}}[\text{bit/s}] = \frac{\sum_{n=1}^{N_{TF}} R_{serv\,n}^{VBS}[\text{bit/s}]}{N_{TF}} \qquad (6)$$

where $R_{serv\,n}^{VBS}$ is the VBS serving data rate in time frame n and N_{TF} the total number of time frames.

The *Average Cluster Serving Data Rate* is the metric that evaluates the performance of the overall cluster, allowing one to observe the impact of using VRRA algorithms for different use cases. It is defined as the average of the *Cluster Serving Data Rate* over the total number of time frames in the observation time interval:

$$\overline{R_{serv}^{Cl}}[\text{bit/s}] = \frac{\sum_{n=1}^{N_{TF}} R_{serv\,n}^{Cl}[\text{bit/s}]}{N_{TF}} \qquad (7)$$

where $R_{serv_n}^{Cl}$ is the *Cluster Serving Data Rate* for *n*–th time frame.

The *Average Cluster Utilisation* is a measure of the RUs utilisation within the cluster. This metric should be analysed together with the *Average Cluster Data Rate*, since the efficiency of the use of the RUs is as important as maximising their use. It is defined as the average of the *Cluster Utilisation* over the total number of time frames in the observation time interval:

$$\overline{\eta_{Cl}} = \frac{\sum_{n=1}^{N_{TF}} \eta_{Cl_n}}{N_{TF}} \qquad (8)$$

where η_{Cl_n} is the ratio between the maximum data rate corresponding to the RUs occupied by end-users and the maximum data rate the cluster can provide in time frame n, given by:

$$\eta_{Cl_n} = \frac{\sum_{i=1}^{N_{RAT}} \left(N_{RU_{occ\,n}}^{RAT} \cdot R_{RU_{max}\,[\text{bit/s}]}^{RAT_i} \right)}{\sum_{i=1}^{N_{RAT}} \left(N_{RU}^{RAT} \cdot R_{RU_{max}\,[\text{bit/s}]}^{RAT_i} \right)} \qquad (9)$$

where $N_{RU_{occ\,n}}^{RAT}$ is the number of RUs occupied by end-users in each RAT in time frame n, being subjected to:

$$N_{RU_{occ}}^{RAT_i} \le N_{RU_{max}}^{RAT_i} \quad , \forall\, RAT_i \in RAT^{Cl} \qquad (10)$$

where RAT^{Cl} is the number of RATs in the cluster.

The *Average Cluster Utilisation* should be analysed together with the *Average Cluster Data Rate*, since the efficiency of the use of the RUs is as important as maximising their use.

III. SCENARIOS FOR SIMULATION

The scenario for simulation was defined in order to assess the OnDemandVRRA algorithm regarding situations in which the total amount of capacity contracted by MVNOs is on average, under and over booking. It is considered that an under booking situation, *Under* use case, occurs when the amount of contracted data rate by all the VBSs instantiated in the cluster is lower than the average cluster capacity, i.e., the data rate the cluster can provide when the modulation and coding schemes applied to all the radio resource units within the cluster is between the second and third higher data rates. Two over booking situations were considered, *GRTOver* and *BEOver* use cases, in which the total contracted data rate is greater than the average cluster capacity. Finally, an *Average* use case is considered to depict the situation when the contracted capacity is near the average cluster capacity. TABLE I presents a summary of the parameters for the four use cases under study.

The physical cluster is composed of 2 TDMA, 1 CDMA, 8 OFDM and 4 OFDMA BSs, all RATs overlapping. End-users are uniformly distributed within the cluster, increasing from 1 000 to 15 000, in order to depict low and high loaded situations. Concerning services, the GRT VBS provides VoIP and Video, and the BE VBS provides Web and File Sharing.

TABLE I. CONTRACTED VERSUS CLUSTER CAPACITY RELATED USE CASE'S PARAMETERS.

Use Cases	Physical Capacity		Contracted Data Rate	
	Max [Gbit/s]	Average [Gbit/s]	GRT VBS [Gbit/s]	BE VBS [Gbit/s]
Average	6170	2800	1.25	1.5
Under			1.5	0.05
BEOver			1.0	2.2
GRTOver			2.2	1.0

IV. SIMULATION RESULTS

Simulation results are obtained considering an observation time interval of 1 hour of network operation. It should be referred that end-users' SINR changes over time to reflect unpredictable BS distance and channel variability, which affects differently RUs performance according to each RAT.

From Figure 3, it can be observed that $\overline{R_{serv}^{VBS}}$ for GRT VBS reaches the data rate contracted for all the use cases as soon as the data rate requested by end-users exceeds that value. Furthermore, the GRT VBS can achieve $\overline{R_{serv}^{VBS}} \ge R_{min}^{VBS}$ when the BE VBS is operating within the contract and there is some remaining capacity in the cluster, e.g., for *BEOver* when $1 < R_{req}^{VBS} < 1.5$ Gbit/s ($\overline{R_{serv}^{VBS}} \ge 1\,\text{Gbit/s} = R_{min}^{VBS}$).

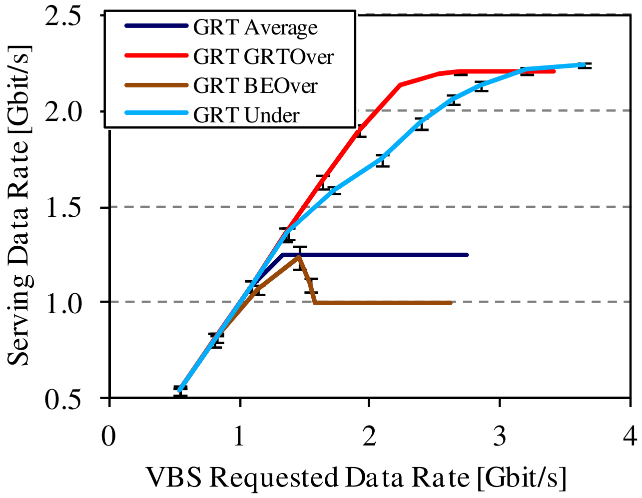

Figure 3. GRT VBS *Average Serving Data Rate*.

Figure 4 shows that the BE VBS follows R_{ref}^{VBS} for $R_{req}^{VBS} \ge 6$ Gbit/s, corresponding to the GRT VBS running with at least the minimum contracted capacity. When $R_{req}^{VBS} \approx 9$ Gbit/s, $\overline{R_{serv}^{VBS}} \approx R_{ref}^{VBS}$ on BE VBS for both *Average*, $R_{ref}^{VBS} = 1.5$ Gbit/s, and *Under*, $R_{ref}^{VBS} = 50$ Mbit/s, use cases. It is also seen that for *BEOver*, although the BE VBS is close to R_{ref}^{VBS}, it cannot reach that value because the total contracted capacity is above the average capacity of the cluster. The same reason is underlying the *GRTOver* use case, in which the BE VBS cannot follow R_{ref}^{VBS} , because OnDemandVRRA is allocating RUs to GRT VBS to satisfy its minimum contracted data rate, $R_{min}^{VBS} = 2.2$ Gbit/s, near the cluster average capacity.

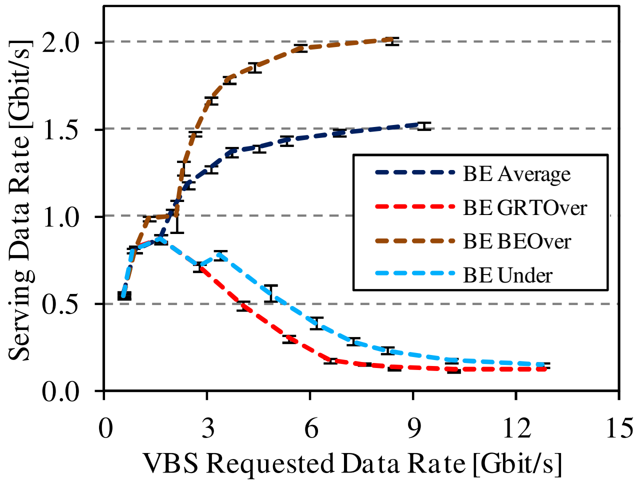

Figure 4. BE VBS *Average Serving Data Rate.*

Concerning the *Cluster Serving Data Rate*, Figure 5, the worst performance is obtained for *GRTOver* and *Under* use cases, $\overline{R_{serv}^{Cl}} \approx 2.3$ Gbit/s. In fact, for both use cases a large number of end-users is accepted in GRT VBS, though for different reasons: for *GRTOver* due to a high minimum contracted data rate of GRT VBS, $R_{min}^{VBS} = 2.2$ Gbit/s, and for *Under* due to a low limit of the BE VBS use case, $R_{ref}^{VBS} = 50$ Mbit/s.

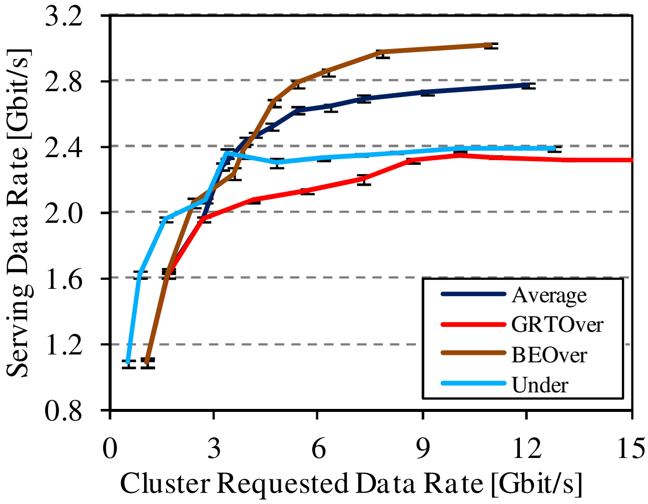

Figure 5. *Average Cluster Serving Data Rate.*

Given that some of the GRT end-users are receiving service in bad performance conditions, the number of RUs providing low data rate increases, and consequently $\overline{R_{serv}^{Cl}}$ decreases.

It is worth to note that the maximum $\overline{R_{serv}^{Cl}}$ is achieved for *BEOver,* $\overline{R_{serv}^{Cl}} \approx 3.0$ Gbit/s, since BE end-users use all the remaining capacity as soon as the GRT VBS minimum contracted capacity, $R_{min}^{VBS} = 1.0$ Gbit/s, is satisfied. The value obtained for *Average* use case, $\overline{R_{serv}^{Cl}} = 2.7$ Gbit/s, is also interesting, because in this situation the total contracted data rate, by both VBSs, is approximately the average cluster capacity, being the traffic in each VBS shaped to fit this value.

The *Average Cluster Utilisation*, Figure 6, increases with the number of end-users in all uses cases, reaching 100% when $R_{req}^{CL} >$ 4 Gbit/s for *Under* and *GRTOver*, and $R_{req}^{CL} > 9$ Gbit/s for *BEOver* and *Average* use cases. It should be highlighted that for *BEOver* and *Average*, η_{Cl} may decrease due to the need to assign RUs to all end-users requesting GRT services, since the reduction of the RUs allocated to the end-users on BE services is made primarily to those in poor performance conditions. This is the case when the number of end-users in the cluster corresponds to $R_{req}^{VBS} \approx R_{min}^{VBS}$ for GRT VBS and $R_{req}^{VBS} \geq R_{ref}^{VBS}$ in BE VBS. As an example, for the *Average*, the decrease of η_{Cl} is verified for $R_{req}^{Cl} = 3.25$ Gbit/s, when $R_{req}^{VBS} = 1.35$ Gbit/s in GRT VBS and $R_{req}^{VBS} = 1.9$ Gbit/s in BE VBS, the contracted capacity in each VBS being $R_{min}^{VBS} =$ 1.25 Gbit/s and $R_{ref}^{VBS} = 1.5$ Gbit/s.

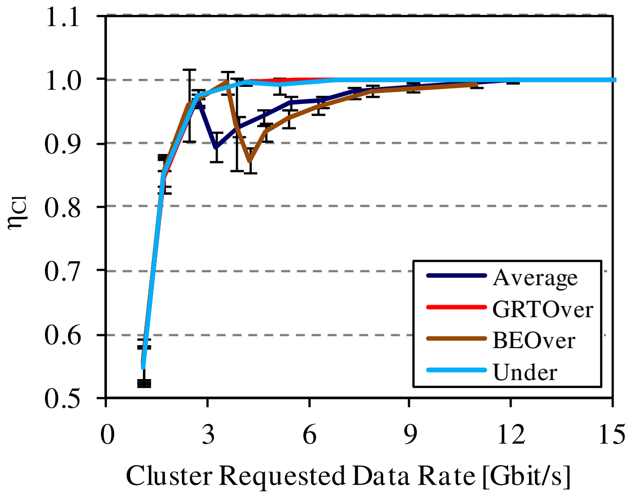

Figure 6. *Average Cluster Utilisation.*

Analysing the *Average Cluster Serving Data Rate* and the *Cluster Utilisation* simultaneously, Figures 5 and 6, one can say that the best RU efficiency is achieved when the strategy for the overall capacity provision is to limit the capacity contracted by GRT VNets, overbooking the capacity contracted by BE VNets, i.e., the *BEOver* use case. The relative inefficiency for both *Under* and *GRTOver* use cases is related to the quantity of end-users in the GRT VBS. Due to the fact that GRT services have a minimum data rate to be performed, the RUs may be assigned to end-users in poor performance conditions. For the *Under* use case, the problem is originated by the priority in handling end-users of GRT services whenever all the VBSs in the cluster have their contracted data rate satisfied. For the *GRTOver* use case, the inefficiency is related to the value for contracted capacity of the GRT VBS, which is about 85% of the average cluster capacity, causing most of the connected end-users to be in GRT VBS.

V. CONCLUSIONS

This paper addresses the virtualisation of the wireless access in order to provide required capacity to a set of VBSs instantiated in a geographical area. Instead of looking at the wireless virtualisation in the perspective of the instantiation of virtual machines in the wireless nodes, the approach used here takes the perspective of the virtual resource as an aggregated connectivity resource abstracted from a group of radio resource units (RUs) from different radio access technologies, being agnostic to the point where the virtual node instantiation takes place.

An algorithm is proposed to manage, in a coordinated manner, a pool of RUs from all the heterogeneous wireless networks serving a given area, for allocation on demand to several virtual base stations (VBSs). This allocation is done indirectly by enforcing the decisions taken from the cluster point of view, to be considered by RRM algorithms when to schedule the RUs to be assigned to end-users. All available RUs in the cluster may be allocated to any VBS if they have been requested, as soon as all the other VBSs in the cluster have their contracted capacity being satisfied. By handling differently the VBSs according to their type of requirements, it supports the deployment of VNets with minimum guaranteed capacity, GRT VNets, and with a reference capacity to be provided whenever possible, BE VNets.

Results from simulations were presented to evaluate the performance of the OnDemandVRRA algorithm when the amount of capacity contracted by VBSs is over, on average and under the physical capacity of a cluster of BSs from different RATs serving a given geographic area. It is verified that the *Cluster Serving Data Rate* may increase by approximately 20% in *BEOver*, in which the BE VBS contracted capacity is 85% of the average capacity of the cluster, compared to the *GRTOver* use case, where the capacity contracted by the GRT VBS is the one with 85% of the cluster average data rate. Furthermore, the *Serving Data Rate* of the GRT VBSs is always greater than the minimum contracted being constrained by the defined BE VBSs reference data rate, which tends to be followed. It can be also noted that, independently of the use case, both GRT and BE VBSs can use the remaining capacity of the cluster even when the contracted capacity has been achieved, since the other VBSs are not requesting for its use.

From this analysis, one can say that a limit for the data rate contracted by GRT VBSs should be established in order to allow an efficient use of RUs among all the VBSs deployed within the cluster. From these results, one can say that this limit should be less than 50% of the average capacity of the physical cluster, while the BE VBSs contracted capacity can be overbooked.

ACKNOWLEDGMENT

The support of the European Commission by partially funding this work via the FP7-ICT-2009-5-SAIL project, Grant Agreement Number 257448, is acknowledged.

REFERENCES

[1] Peterson,L., Muir,S., Roscoe,T. and Klingaman,A., *PlanetLab Architecture: An Overview*, Technical Report PDN–06–031, PlanetLab Consortium, May 2006.

[2] Sanjoy,P. and Srini,S., *GENI: Global Environment for Network Innovations*, Technical Document on Wireless Virtualisation, GENI project, Wireless Working Group, Document GDD-06-17, Sep. 2006 (http://www.geni.net).

[3] Schaffrath,G., Werle,C., Papadimitriou,P., Feldmann,A., Bless,R., Greenhalgh,A., Wundsam,A., Kind,M., Maennel,O. and Mathy,L., "Network Virtualization Architecture: Proposal and Initial Prototype", in *Proc. of VISA'09 – 1st ACM SIGCOMM Workshop on Virtualized Infrastructure Systems and Architectures*, Barcelona, Spain, Aug. 2009.

[4] Zhu,Y., Zhang-Shen,R., Rangarajan,S. and Rexford,J., "Cabernet: Connectivity Architecture for Better Network Services.", in *Proc. of ReArch'08 - ACM Workshop on Re-Architecting the Internet*, Madrid, Spain, Dec. 2008.

[5] Sachs,J. and Baucke,S., "Virtual radio: a framework for configurable radio networks", in *Proc. of WICON'08 - 4th Annual International Conference on Wireless Internet*, Maui, HI, USA, Nov. 2008.

[6] AlQahtani,S.A., Mahmoud,A.S., Sheltami,T.R. and El-Tarhuni,M., "Adaptive Radio Resource Management for Multi-Operator WCDMA Based Cellular Wireless Networks with Heterogeneous Traffic", in *Proc. of PIMRC'06 - 17th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, Helsinki, Finland, Sep. 2006.

[7] 3GPP, *3GPP Technical Specification Universal Mobile Telecommunications System (UMTS); LTE; Network Sharing; Architecture and functional description (Release 11)*, Technical Specification TS 23.251 V11.4.0, Jan. 2013 (http://www.3gpp.org).

[8] Kokku,R., Mahindra,R., Zhang,H., and Rangarajan,S., "Cellular Wireless Resource Slicing For Active RAN Sharing", in *Proc. of COMSNETS 2013 – Fifth International Conference on Communication Systems and Networks*, Bangalore, India, Jan. 2013.

[9] 3GPP, Study on Radio Access Network (RAN) sharing enhancements (Release 12), RAN, Technical Report TR 22.852 v12.0.0, Jun. 2013 (http://www.3gpp.org).

[10] Zaki,Y., Liang,Z., Goerg,C. and Timm-Giel,A., "LTE wireless virtualisation and spectrum management", in *Proc. of WMNC'10 - Wireless and Mobile Networking Conference*, Budapest, Hungary, Oct. 2010.

[11] Xia,L., Kumar,S., Yang,X., Gopalakrishnan,P., Liu,Y., Schoenberg,S. and Guo,X., "Virtual WiFi: Bring Virtualisation from Wired to Wireless", in *Proc. of VEE'11 – International Conference on Virtual Execution Environments*, Newport Beach, CA, USA, Mar. 2011.

[12] 3GPP, *Improvement of RRM across RNS and RNS/BSS (Release 5)*, RAN, Technical Report TR 25.881 v5.0.0, Dec. 2001 (http://www.3gpp.org).

[13] Sachs,J., Aguero,R., Berg,M., Gebert,J., Jorguseski,L., Karla,I., Karlsson,P., Koudouridis,G.P., Lundsjo,J., Prytz,M. and Strandberg,O., "Migration of Existing Access Networks Towards Multi-Radio Access", in *Proc. of VTC'06 Fall - 64th IEEE Vehicular Technology Conference*, Montreal, Canada, Sep. 2006.

[14] Caeiro,L., Cardoso,F.D., and Correia,L.M., "OConS Supported On Demand Radio Resource Allocation for Virtual Connectivity", in *Proc. of MONAMI - 4th International Conference on Mobile Networks and Management*, Hamburg, Germany, Sep. 2012.