Radio Access Network Virtualization for Future Mobile Carrier Networks

Xavier Costa-Pérez and Joerg Swetina, NEC Laboratories Europe
Tao Guo, NEC Telecom Modus
Rajesh Mahindra, Sampath Rangarajan, NEC Laboratories America

ABSTRACT

This article presents a survey of cellular network sharing, which is a key building block for virtualizing future mobile carrier networks in order to address the explosive capacity demand of mobile traffic, and reduce the CAPEX and OPEX burden faced by operators to handle this demand. We start by reviewing the 3GPP network sharing standardized functionality followed by a discussion on emerging business models calling for additional features. Then an overview of the RAN sharing enhancements currently being considered by the 3GPP RSE Study Item is presented. Based on the developing network sharing needs, a summary of the state of the art of mobile carrier network virtualization is provided, encompassing RAN sharing as well as a higher level of base station programmability and customization for the sharing entities. As an example of RAN virtualization techniques feasibility, a solution based on spectrum sharing is presented: the network virtualization substrate (NVS), which can be natively implemented in base stations. NVS performance is evaluated in an LTE network by means of simulation, showing that it can meet the needs of future virtualized mobile carrier networks in terms of isolation, utilization, and customization.

INTRODUCTION

Future mobile carrier networks need to address the predicted growth in mobile traffic volume which is expected to have explosive growth in the next years mainly driven by video and web applications [1]. Based on the expected capacity demand increase, mobile network operators (MNOs) are required to increase capital (CAPEX) and operational expenses (OPEX) accordingly. However, as cellular networks move from being voice-centric to data-centric, increase in revenue is not keeping pace with the increase in traffic volume. Forecasted future mobile networks capacity requirements together with the decreasing operators' benefits margin have led to network sharing being considered a key business model for reducing future deployment and operational costs.

Network sharing solutions are already available, standardized [2], and partially used in some mobile carrier networks. These solutions can be divided into *passive* and *active* network sharing. Passive sharing refers to the reuse of components such as physical sites, tower masts, cabling, cabinets, power supply, air-conditioning, and so on. Active sharing refers to the reuse of backhaul, base stations, and antenna systems, the reuse of the latter two labeled as active radio access network (RAN) sharing. According to a market survey [3], mobile infrastructure sharing has already been deployed by over 65 percent of European operators in different ways, and this trend is expected to expand in the future. One of the main motivations for network sharing is that currently, a considerable number of sites consume energy and computational resources, even though they carry a negligible level of traffic. For instance, in [4] it was reported that around 20 percent of all sites carry about 50 percent of total traffic. Estimations regarding the expected savings for operators by implementing active network sharing have been conducted [5]. This study concluded that operators worldwide could reduce combined OPEX and CAPEX costs up to $60 billion over a five-year period, and at least 40 percent of these cost savings is expected to come from active network sharing.

The traditional model of single ownership of all network layers and elements is thus being challenged: since the most basic sharing concepts related to passive network sharing are easier to implement and have already been partially exploited, active network sharing is rising in importance to enable substantial and sustainable reduction in network expenses to ensure operators' future cost competitiveness.

Active RAN sharing involves sharing antennas/base stations across multiple mobile (virtual) network operators with either separate spectrum resources for each entity or shared spectrum resources through spectrum pooling. Current solutions, however, have limitations in terms of separating both data and control planes among operators, flexibility and customizability for accommodating different requirements per operator, and the capability to adapt to new or changing requirements.

We first review 3GPP's network sharing standardized functionality. Emerging business mod-

The research leading to these results has been partially funded by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 257263 (FLAVIA).

As mobile carrier operators look for new ways to reduce costs due to their increasing financial investment burden, new business models for operators and infrastructure owners are emerging, calling for an extension to existing standardized network sharing functionality.

els calling for new sharing capabilities are provided next, followed by the Third Generation Partnership Project's (3GPP's) ongoing efforts in the RAN Sharing Enhancements (RSE) Study Item to address them. Since base station virtualization-based techniques are key enablers for RAN sharing, along with the required management flexibility and independence among the sharing entities, we provide an overview of the state of the art of mobile carrier network virtualization. A concrete implementation of a RAN sharing base station virtualization solution, the network virtualization substrate (NVS), and quantitative performance results in a Long Term Evolution (LTE) system are provided. Finally, we summarize our contributions and conclude the article.

3GPP NETWORK SHARING STANDARDIZATION

3GPP has recognized the importance of supporting network sharing since Release 6, and defined a set of architectural requirements [6] and technical specifications [2] that have been updated, at the time of writing this article, through Releases 10 and 11, respectively.

3GPP NETWORK SHARING FUNCTIONALITY OVERVIEW

In order to derive the network sharing requirements for 3GPP, five main use cases were defined in [6]:

- Scenario 1 corresponds to multiple core networks (CNs) sharing a common RAN in Release 99. For operators that have multiple frequency allocations, it is possible to share the RAN elements, but not to share the radio frequencies. In this case the operators connect directly to their own dedicated carrier layer in the shared radio network controller (RNC) in the shared RAN.
- Scenario 2 corresponds to two or more operators with individual frequency licenses with which their respective RANs cover different parts of a country, but together provide coverage of the entire country.
- Scenario 3 corresponds to one operator deploying coverage in a specific geographical area, and other operators being allowed to use this coverage for their subscribers. Outside this geographical area, coverage is provided by each of the operators independently.
- Scenario 4 corresponds to common spectrum network sharing when:
  - One operator has a frequency license and shares the allocated spectrum with other operators.
  - A number of operators decide to pool their allocated spectra and share the total spectrum.
- Scenario 5 corresponds to multiple RANs sharing a common CN. The multiple RANs can belong to different public land mobile networks (PLMNs) and network operators. Due to operators' deployment choices, different nodes or part of the common CN can be shared.

Based on these scenarios, 3GPP has defined a network sharing architecture with the objective of allowing different CN operators to connect to a shared RAN [2]. Operators may share not only network elements, but radio resources as well. Two possible architectural network sharing configurations have been specified: the gateway CN (GWCN) and the multi-operator CN (MOCN).

- GWCN: In a GWCN network sharing configuration, CN operators share CN in addition to RAN nodes.
- MOCN: In an MOCN network sharing configuration, multiple CN nodes are connected to the same RNC, and the CN nodes are operated by different operators.

In both configurations, the network sharing agreement between operators is transparent to the users. Base stations broadcast multiple PLMN-ids according to the number of mobile (virtual) network operators. Mobile stations, user equipment (UE) in 3GPP terminology, that support network sharing functionality are able to discriminate among operators in a spectrum network sharing configuration based on the PLMN-ids. Non-supporting UE devices ignore the broadcast system information, and the shared network selects a CN operator from the available ones. PLMN-ids are used for routing and handover purposes within the network.

DEVELOPING BUSINESS MODELS FOR NETWORK SHARING

As mobile carrier operators look for new ways to reduce costs due to their increasing financial investment burden, new business models for operators and infrastructure owners are emerging, calling for an extension to existing standardized network sharing functionality. See [7] for use cases under discussion in 3GPP and [8] for an extensive industry survey on mobile operators' views. In the following, we provide several examples of these disruptive developing business models.

On-Demand Capacity — There are several business cases where different service providers could be interested in buying capacity from MNOs for a specific period of time and/or purpose. Examples include specialized mobile virtual network operators (MVNOs) offering specific services (e.g., VoIP, video telephony, live streaming), mobile operators requiring additional capacity for major events (e.g., sports, concerts, fairs), and emerging machine-to-machine (M2M) service providers (e.g., meter measurements, security surveillance). Dynamic and flexible sharing of the infrastructure would be required in this case on timescales smaller than the current times needed for contract agreement and implementation.

Wholesale-Only Network — In this model MNOs do not offer services to end customers but sell network capacity to businesses who do not have their own wireless network or have limited geographic coverage/spectrum. This represents an extreme case for network sharing requirements where full-fledged flexibility in terms of dynamic resource allocation, traffic isolation, and customizability is needed to maximize its potential.

Over-The-Top Service Providers — In current mobile networks the biggest capacity share is consumed by the content of third party service providers (over-the-top, OTT, services). Such OTT service providers may wish to pay wireless network operators in the future to ensure a satisfactory quality of experience for their users. In such cases no additional M(V)NO infrastructure might be involved in the RAN sharing process.

These new business models together with additional ones under discussion [7, 8], omitted here for the sake of brevity, have triggered the creation of a RAN Sharing Enhancements Study Item in 3GPP, described next.

3GPP RAN SHARING ENHANCEMENTS STUDY ITEM OVERVIEW

The 3GPP RAN Sharing Enhancements (RSE) Study Item [7] of the System Architecture Working Group 1 (SA1) is defining new scenarios of multiple mobile carrier operators sharing radio network resources. Some of the authors of this article are actively contributing to this effort and have been appointed as the official 3GPP RSE Study Item Rapporteurs.

The objective of the RSE Study Item is to create new requirements that complement existing system capabilities for sharing common RAN resources. The new scenarios aim to:

- Provide means to be able to verify that the shared network elements provide allocated RAN resources according to the sharing agreements and/or policies
- Provide means to efficiently share common RAN resources according to identified RAN sharing scenarios (e.g., pooling of unallocated radio resources)
- Provide means to flexibly and dynamically allocate RAN resources on demand at smaller timescales than those supported today
- Provide means to act on overload situations considering sharing agreements and/or policies

Based on these objectives, we describe below some of the main use cases being used for defining new 3GPP requirements within the SA1 RSE Study Item.[1]

RSE Monitoring — This use case describes the situation of a *hosting RAN provider* (primary RAN operator) that allows *participating operators* (e.g., MVNOs) to use a share of the RAN capacity where the amount of resources might be different for each participating operator. In such a scenario, the hosting RAN shall allow participating operators to retrieve operation, administration, and management (OAM) status information at the same level of detail as in a non-shared RAN for the share of their resources.

RSE Dynamic Capacity — This use case describes the situation of participating operators requiring varying network capacities during different time periods (e.g., within a day or a week). Participating operators might request various allocations of a portion of the shared RAN in order to meet the projected variation in network usage requirements. An example of this use case could be an MVNO requiring significant RAN capacity during business hours but lower capacity during the night or weekends.

RSE On-Demand Automated Capacity Brokering — This use case describes the situation of a hosting RAN provider that supports on-demand automated capacity requests from participating operators. A typical situation for this use case would be a hosting RAN provider with excess capacity during the night that participating operators could request on demand (e.g., M2M services such as smart meter measurements. Another example could be a major event (e.g., sports, concerts, trade fairs, etc.) requiring additional capacity from a participating operator for that event. On-demand requests for capacity shall indicate the time period in which the capacity is needed, the amount of capacity required, and the specific service(s) treatment desired, for example, based on standardized quality of service (QoS) class identifiers (QCIs). The hosting RAN provider then verifies automatically whether the RAN sharing request can be fulfilled; if so, the shared RAN is re-configured accordingly.

RSE Load Balancing — This use case describes the situation of overlapping cells in a shared RAN. The hosting RAN shall be able to support load balancing within a shared RAN while respecting the agreed shares of RAN resources based on the whole cell load level and the load level for each participating operator. If the load levels of individual cells are exceeded, the hosting RAN provider shall enforce agreed maximal usage limits of each participating operator by handing over UE devices to neighboring cells.

Based on the described use cases and additional ones under discussion, a set of new consolidated requirements has been derived. In Table 1 we summarize some of these main new requirements, which illustrate the new functionality expected to be specified based on the RSE Study Item contributions.

These set of new requirements illustrate the paradigm change being defined by the RSE Study Item work. The 3GPP acceptance of automated means (no human intervention) to allocate RAN resources dynamically and on demand deserves special mention. This enhanced functionality opens the door to new industry players as well as to an evolution of operators' business models which could potentially have a deep impact on the industry. Additionally, the enhanced sharing flexibility along the capacity, time, and sector dimensions enable higher efficiency of RAN infrastructure exploitation.

The RSE work in TR 22.852 [7] is expected to be completed and kick off a new Work Item producing normative text this year. In the next section, we discuss the current state of the art of mobile carrier networks' virtualization technologies that could potentially address the requirements of the previously described emerging RAN sharing use cases.

In current mobile networks the biggest capacity share is consumed by the content of third party service providers (over-the-top, OTT, services). Such OTT service providers may wish to pay wireless network operators in the future to ensure a satisfactory quality of experience for their users.

---

[1] *Note that at the time of writing this article the Study Item is close to completion, but changes might still be made in the final document release.*

Hosting RAN resources allocation
• A hosting RAN provider shall be able to allocate the share of RAN capacity per participating operator as: – Fixed, minimum allocation guaranteed – Fixed for a period of time and/or sectors – First come first served (i.e., on demand) • A shared RAN shall be capable of providing differentiated traffic treatment per participating operator • A shared RAN shall conduct admission control based on the proportion of assigned RAN usage per participating operator
On-demand capacity negotiation
• The hosting RAN shall be able to offer by automatic means shareable RAN resources as on-demand capacity to participating operators • Participating operators shall be able to automatically request hosting-RAN-offered on-demand resources • The hosting RAN provider shall allow a participating operator to request the cancellation of granted on-demand requests • The hosting RAN provider shall be able to withdraw a granted request (within SLA limits)
OAM access to the hosting RAN
• The hosting RAN shall be able to provide and control selective OAM access to participating operators • The hosting RAN shall be able to allow participating operators to retrieve selective OAM status information at the same level of detail as would be available from a non-shared RAN
Handovers due to RAN sharing agreements
• Participating operators shall be able to direct UE toward the hosting RAN at the beginning of a RAN sharing period • A hosting RAN provider shall be able to direct UE away from the hosting RAN at the end of a RAN sharing period • Participating operators shall be involved in the handover decisions at the end of a RAN sharing period

Table 1. *3GPP SA1 RSE Study Item, consolidated requirements sample, TR 22.852 [7].*

MOBILE CARRIER NETWORKS VIRTUALIZATION

Although network virtualization in the wired domain has received significant attention in the past years [9], mobile carrier network virtualization is still in its nascent stage. The main focus of wired network virtualization has been on the design of an adaptive network substrate that can support multiple virtual networks running customized services over a physical network. To enable end-to-end cellular network virtualization, the mobile CN and RAN have to be virtualized. While solutions from the wired network and server domains can be used to virtualize the mobile CN, RAN or base station virtualization has to deal with problems specific to the characteristics of wireless access links. Wireless access links usually have more dynamic sets of users, user mobility, and varying channel conditions that make it harder to virtualize the wireless resources across multiple entities.

Base station virtualization may be performed at either the hardware level (dedicated spectrum) or the flow level (shared spectrum). Hardware virtualization solutions exist commercially today for traditional MNOs to cut operating costs. The virtual base transceiver station (vBTS) [10] is one such virtualized base station solution that enables sharing radio components at the hardware level and running multiple base station protocol stacks in software. To support spectrum-sharing-based models, virtualization at higher levels such as the flow level at the base station is more appropriate and leads to better multiplexing of resources. Furthermore, it can support multiple deployment scenarios such as MVNOs where the virtual networks do not own spectrum. Multiple MNOs can also pool their spectrum to save costs and accelerate network rollouts.

Several GENI design documents describe proposals and issues of wireless network virtualization [11]. The aim of GENI is to create a unified nationwide wired-wireless virtualized testbed that can support multiple concurrent experiments. In this context, recent work [12] focuses on virtualizing the wireless resources of WiMAX base stations remotely from an access service network (ASN) gateway. Other efforts [13] propose to virtualize the LTE network by implementing a hypervisor in the eNodeB. Each entity runs its LTE stack in a virtual machine. The hypervisor allocates spectrum to the different entities in accordance with a specified guarantee. An entity can request either a fixed or dynamic guarantee based on its current load up to a maximum amount of resources. However, all of the above mentioned schemes only ensure isolation of resources and do not guarantee bandwidth across the entities. Moreover, these schemes fail to provide flexibility to the slices to customize the resource allocation across their users.

The European FP7 project Flexible Architecture for Virtualizable Future Wireless Internet Access (FLAVIA) [14] is defining and prototyping a new base station architecture with the objective of enabling a higher level of programmability. FLAVIA promotes the concept of a wireless medium access control (MAC) processor, a programmable device that:

- Provides a set of stateless MAC commands
- Embeds a MAC protocol engine in charge of executing a finite state machine able to exploit and compose the sequence of commands forming a desired protocol

This enhanced base station programmability functionality is being explored in the context of wireless access virtualization, and project participants contribute to the 3GPP RSE Study Item efforts.

In parallel, the emerging software-defined networking (SDN) paradigm is a key enabler to simplify the network provisioning, management, (re)configuration, and control of such virtualized and shared mobile carrier networks. Wireless SDN requires the identification of abstractions for wireless primitives and functions, which compromise between flexibility and ability of the abstractions to permit developing innovative wireless functions and mechanisms. An overview of ongoing SDN standardization efforts can be found in [15].

With respect to commercial deployments, there are several MVNOs that lease spectrum

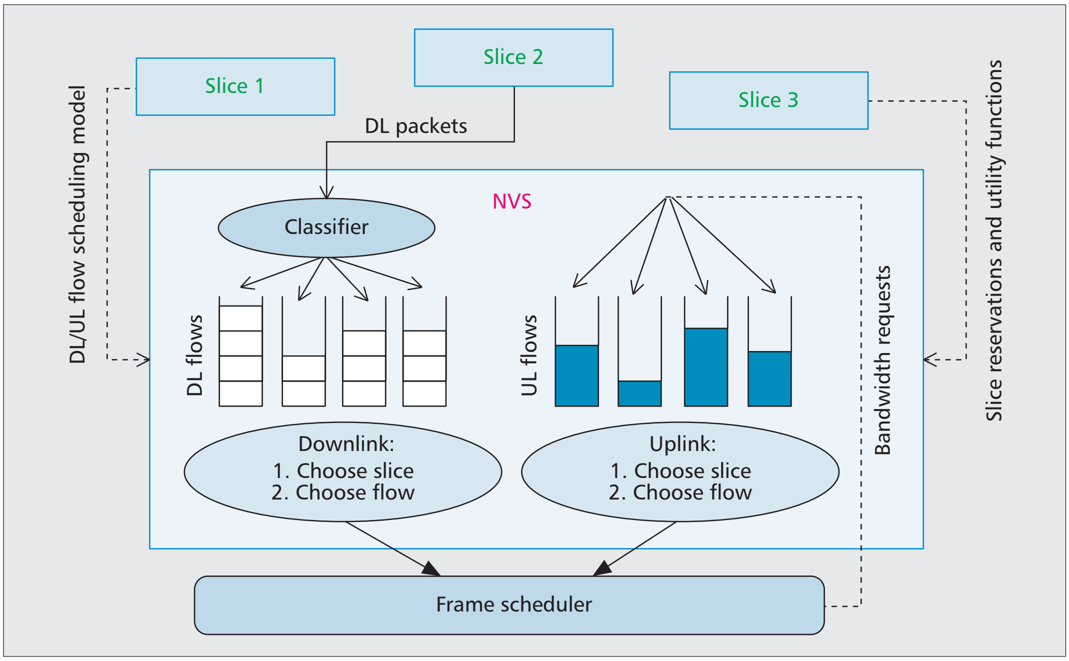

Figure 1. *Network virtualization substrate architecture illustration.*

NVS operates at fine time-scales in order to fill up the resources of a physical layer frame. At every scheduling instant, NVS selects a slice that maximize the overall utility or revenue. Once a slice is selected, NVS schedules flows within the slice according to a custom slice policy.

from existing MNOs to provide cellular services. There are also a few instances of MNO-based RAN sharing in different parts of the world [8].

However, there is very little publicly available information on the base station virtualization aspects of such sharing solutions.

BASE STATION VIRTUALIZATION FEASIBILITY

So far we have considered mobile carrier networks virtualization from a conceptual perspective. In this section, we take the next step and focus on describing a specific solution to implement it in real networks. In order to do so, we build on our work on base station virtualization, which allows multiple entities to share the same spectrum, achieving effective virtualization of its wireless resources: the network virtualization substrate (NVS) [16]. NVS can be readily deployed to virtualize orthogonal frequency-division multiple access (OFDMA)-based 4G base stations such as LTE and WiMAX, and aims to provide key virtualization features as defined by the 3GPP SA1 RSE consolidated requirements:

- *Isolation* ensures that traffic, mobility, and fluctuations in channel conditions of users of one entity do not affect the reserved resource allocations of other entities sharing the same base station: RSE hosting RAN resources allocation.
- *Customization* provides flexibility to the different entities to program the base station to optimize their service delivery: RSE OAM access to the hosting RAN.
- *Utilization* of base stations is maximized by allowing usage of resources unused by one entity by other entities: RSE hosting RAN resources allocation.

From this point onward we call virtual networks *slices*. A slice is defined as a virtual network consisting of a group of flows that share a physical base station with multiple other virtual networks.

BASE STATION VIRTUALIZATION: DEFINITIONS

To meet the three aforementioned requirements for virtualization, NVS is designed as a hierarchical scheduler, as shown in Fig. 1, with two components:

- A *slice scheduler* ensuring resource isolation across the slices
- A *flow scheduling framework* enabling flexibility of flow scheduling to the different slices

Separating the slice scheduler from the flow scheduler(s) enables greater flexibility in customizing the slices in both the uplink and downlink directions, and also makes service level agreements (SLAs) between the slice owners and the physical network owner easier to define and manage.

NVS operates in fine timescales in order to fill up the resources of a physical layer frame. At every scheduling instant, NVS selects a slice that maximizes the overall utility or revenue. Once a slice is selected, NVS schedules flows within the slice according to a custom slice policy.

Slice Scheduler — NVS defines two possible ways in which slices can reserve a part of the base station resources:

- *Resource-based provisioning*: Defines resource allocation to a slice in terms of a fraction of the total base station resources. This implies that the slice is guaranteed to receive at least a fraction of the total spectrum resources. For instance, if two MNOs share RAN resources, MNO A can reserve 60 percent of the total resources while MNO B gets 40 percent of them.
- *Bandwidth-based provisioning*: Defines resource allocation to a slice in terms of

Although NVS introduces a hierarchical scheduler in the base station, the implementation is lightweight to ensure it is readily deployable on today's base station hardware. The NVS scheduler is integrated into the current base station flow scheduling framework.

aggregate bandwidth (in megabits per second) that will be obtained by the flows of that slice.

For instance, a content provider like youtube can reserve 2 Mb/s rate on an MNO's network so that its users receive good quality video even during periods of congestion. Note that while MNOs and MVNOs may prefer to share resources based on resource provisioning, content or application service providers may prefer to reserve bandwidth on the cellular network.

To satisfy both bandwidth- and resource-based provisioning on the same base station, NVS defines different utility functions for the two provisioning schemes while ensuring that they are comparable.

Assuming that G_1 and G_2 are the sets of slices with bandwidth-based and resource-based provisioning, respectively, NVS defines the following utility functions:

$$\mathcal{U}_g(r_g) = \frac{r_g^{\text{rsv}}}{r_g^{\text{eff}}}\log(r_g), g \in G_1 \tag{1}$$

$$\mathcal{V}_g(t_g) = t_g^{\text{rsv}}\log(t_g), g \in G_2 \tag{2}$$

where r_g^{rsv} and r_g denote the minimum reserved cumulative bandwidth and the achieved cumulative bandwidth (units of megabits per second), while t_g^{rsv} and t_g denote the minimum reserved cumulative resources and the achieved cumulative resources (units of resource blocks). r_g^{eff} denotes the average effective transmission rate of a slice g in terms of the resource blocks allocated to it.

The slice utilities $U_g(r_g)$ and $V_g(t_g)$ are based on SLAs between the slice owners and the physical network owner. NVS uses concave utility functions to model the revenue generated from the different slices. Specifically, NVS uses a logarithmic function of resources or bandwidth allocated to the slice. This choice of utility function assumes that the marginal utility of a slice decreases as its achieved resources or bandwidth increases. By employing weighted log-utility functions, NVS effectively provides fairness among slices that is similar to proportional fairness after meeting the reserved bandwidth or resources for each slice.

The slice scheduler is designed to maximize the overall utility of the base station, while ensuring the requirements of individual slices. To achieve maximum utility, the slice scheduler employs the following weight function and picks the slice with the largest weight at every time instant j.

$$w_{g,j} = \begin{cases} \dfrac{r_g^{\text{rsv}}}{r_{g,j}^{\exp}} & \text{if } g \in G_1 \\[2ex] \dfrac{t_g^{\text{rsv}}}{t_{g,j}^{\exp}} & \text{if } g \in G_2. \end{cases} \tag{3}$$

where $r_{g,j}^{\exp}$ and $t_{g,j}^{\exp}$ are exponential moving averages of bandwidth and resource blocks that are allocated to slice g at time j, respectively.

If there are unused resource blocks in the system after the flow scheduling within the selected slice, the slice with the second largest weight at this time instant may be selected for flow scheduling. This process continues until all the resources are utilized or no slice remains. In this way, the utilization of the base station can be maximized.

In addition to the data traffic generated by users, note that NVS also accounts for the control signaling generated by each user of a slice when updating the cumulative resources t_g or the cumulative bandwidth r_g of a slice. This ensures that NVS provides isolation of resources for a slice against slices that generate high amounts of signaling (e.g., due to high user mobility and/or increased transitions to and from idle mode).

Flow Scheduling Framework — NVS builds a generic flow scheduling framework that facilitates slices to emulate a wide variety of flow schedulers. NVS defines three modes of performing flow scheduling for slices. Each mode is designed with trade-offs between complexity and flexibility.

Scheduler Selection — In this mode, NVS provides several common flow schedulers already implemented in the base station. A slice can choose from one of these schedulers. While this mode is simple, it is not suitable for evaluating or employing more innovative schedulers.

Model Specification — In this mode, NVS provides an interface for slices to specify a weight distribution function to schedule the flows of that slice. The weight distribution can be a function of several parameters like average rate, and modulation and coding scheme (MCS). When a slice is scheduled, NVS chooses a flow within this slice with the smallest weight and then updates the weights of all flows. This mode enables arbitrary weight-based schedulers.

Virtual-Time Tagging — In this mode, NVS provides real-time flow-level feedback to a slice to perform flow scheduling itself. It exposes higher flexibility to the slice to implement its own arbitrary scheduler, but it is also more complicated as slices need to implement their own scheduler.

Although NVS introduces a hierarchical scheduler in the base station, the implementation is lightweight to ensure it is readily deployable on today's base station hardware. The NVS scheduler is integrated into the current base station flow scheduling framework. The implementation of NVS simply requires each flow to be tagged by the slice ID to which it belongs, and at every scheduling epoch, only those flows that belong to the slice selected as per Eq. 3 are scheduled.

EVALUATION

In our previous work [16], NVS was implemented and evaluated on a picochip-based WiMAX testbed as part of the MAC layer in software. Several experiments were conducted that demonstrated the isolation, customization, and utilization efficiency of NVS in a real testbed, with a special focus on bandwidth reservation. Due to

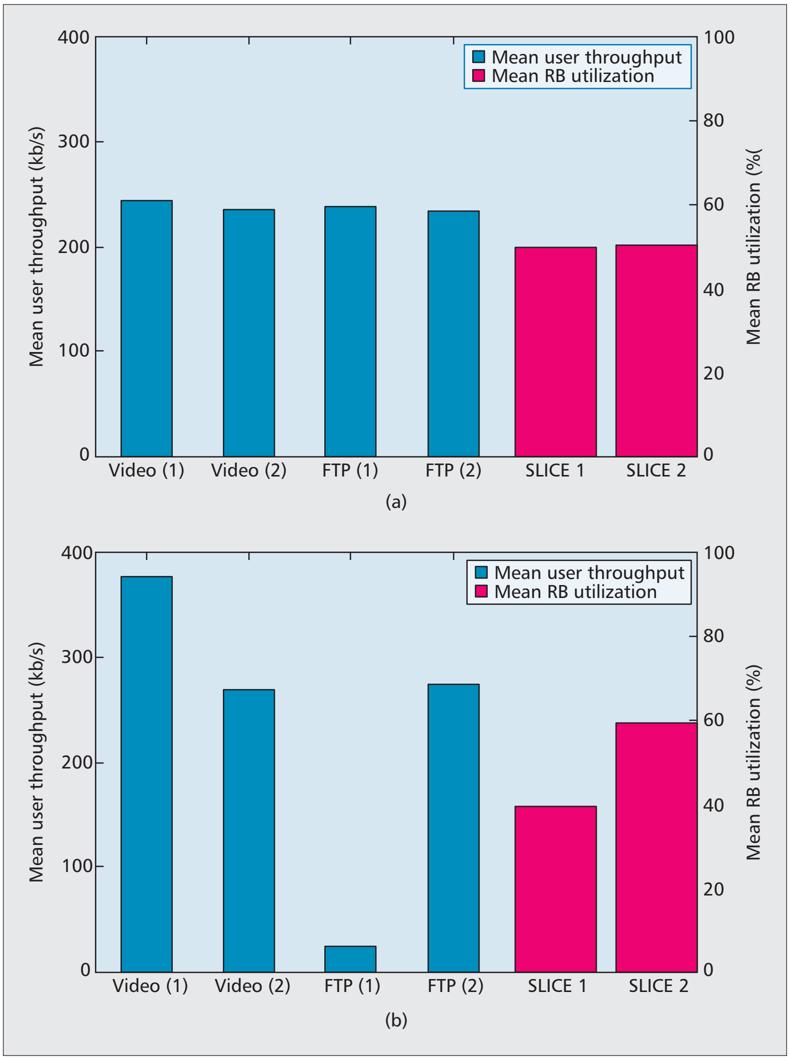

Figure 2. *NVS isolation and customization features: a) full sharing (FS); b) network virtualization substrate (NVS).*

hardware limitations, the experiments included only a limited number of users, and no admission control mechanism was considered.

In this article, we evaluate an LTE system, and implement and analyze NVS on a C++ system-level simulator in a typical cellular network scenario. We focus on the resource reservation approach with an LTE resource block (RB) as the minimum allocable resource unit. A hexagonal cell layout with seven sites and three sectors per site is considered. The system bandwidth is 10 MHz consisting of 50 RBs. NVS is compared with two baseline schemes. The interaction between NVS and admission control is also evaluated.

Full Sharing — In the first set of simulations, we compare NVS with a full spectrum sharing (FS) scheme to illustrate the isolation and customization capabilities of NVS. We set up each base station with two slices. In the FS case there is no resource reservation, and the whole bandwidth is shared by all users, whereas in the NVS case, slices 1 and 2 reserve 40 and 60 percent of the total RBs in the system bandwidth, respectively. Ten video users and 20 FTP users are placed per sector in each slice. Video users continuously stream an average rate of 384 kb/s in the downlink. FTP users have backlogged traffic also in the downlink. In the FS case, a MAC scheduler based on conventional proportional fairness without traffic type prioritization is used. In the NVS case, slice 1 runs a proportional fair MAC scheduler with video users having strict priority over FTP users, while for slice 2 the same scheduler as in the FS case is used.

Figure 2a shows the mean throughput received by the end users and the mean RB utilization of both slices in the FS case. As expected, all users, regardless of their traffic type and associated slices, receive a similar mean throughput. Both slices also experience 50 percent RB utilization on average. As a result, the video rate requirements are not met for all users in both slices. On the other hand, Fig. 2b shows that with NVS, slices 1 and 2 receive 40 and 60 percent of the total RBs, as configured, although they have the same traffic load. Slice 1 users are prioritized based on their traffic type according to the customized slice scheduler satisfying the video requirements, while in slice 2 the video and FTP users receive similar throughput, although the value is larger than the one in the FS case since slice 2 has 60 percent of the total RBs allocated.

Static Reservation — In the second set of simulations, we compare NVS with a static reservation (SR) scheme to illustrate the benefits of NVS in terms of resource utilization. In SR each slice is allowed to use only its reserved RBs. RBs 0–14 (30 percent) are permanently reserved for slice 1, while RBs 15-49 (70 percent) are permanently reserved for slice 2. In the NVS case, we set up the base station with two slices such that slices 1 and 2 are allocated 30 and 70 percent of the total RBs in the system, respectively. In this set of experiments, all users in slice 1 download FTP files with a size of mean 250 kbytes, while all users in slice 2 stream video at an average rate of 128 kb/s in the downlink that last for a duration of 10 s on average. Following common practice in commercial cellular networks, FTP requests are always admitted regardless of the load conditions, while RB-usage-based admission control (AC) is used in slice 2 to support the QoS of the video users. For illustration purposes, the AC threshold is set to 80 percent of the RBs. In both slices a proportional fair MAC scheduler is used.

In Fig. 3a the mean RB utilization of slices 1 and 2 is shown with respect to an increasing offered traffic load for both NVS and SR schemes. As the load increases, slice 1 with SR reaches saturation first as its mean RB utilization is upper bounded to 30 percent of the system capacity. In contrast to this, since NVS can opportunistically allocate the unused resources in slice 2 to slice 1, slice 1 reaches an RB utilization above 30 percent as long as there are unused resources in slice 2. Once slice 2 reaches saturation at 56 percent, both SR and NVS provide resources per slice as configured. NVS reaches higher system throughput than SR due to its capability to allocate the free resources in slice 2

(20 percent blocked for video traffic due to the AC mechanism) to FTP traffic from slice 1.

SUMMARY AND CONCLUSIONS

The virtualization of mobile carrier networks is an appealing opportunity for operators to handle the forecasted "mobile data apocalypse" while at the same time increasing the return on investment in CAPEX and OPEX infrastructure costs. In this article, we first review 3GPP's network sharing standardized functionality followed by an overview of the RAN sharing enhancements currently being discussed in the 3GPP RSE Study Item based on emerging business models. Then a summary of the state of the art of mobile carrier network virtualization research is provided with a special focus on RAN sharing, base station programmability, and customization.

The feasibility of mobile network virtualization is analyzed by presenting a RAN sharing technique meeting the required isolation, network utilization, and customization virtualization needs, NVS, which can be natively implemented in base stations. LTE performance results are presented and benefits of the proposed approach discussed.

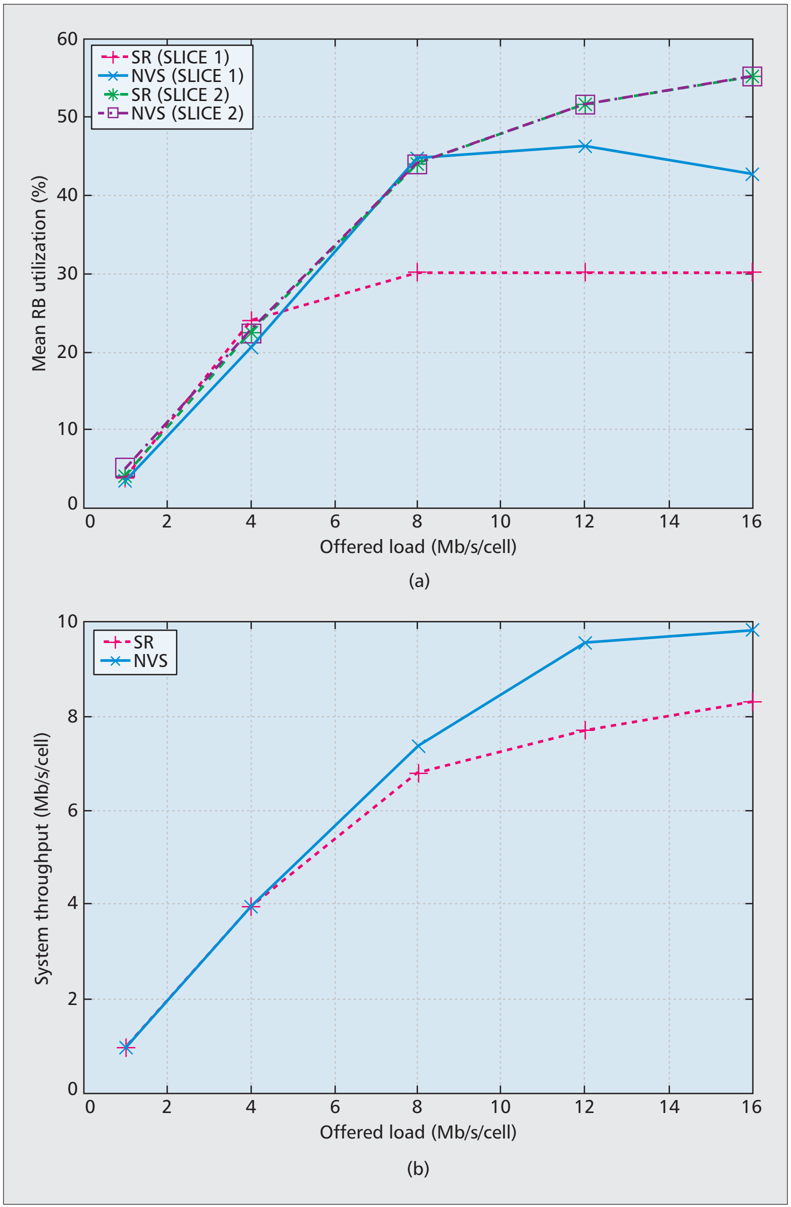

Figure 3. *NVS resource utilization features: a) mean RB utilization; b) system throughput.*

Based on the business trends reported in this article, the corresponding 3GPP RAN sharing enhancements efforts and research results in the area, we conclude that future mobile networks will increasingly include advanced virtualization solutions, opening the market to a wide range of new business models building on the enhanced capabilities of such virtualized mobile carrier networks and likely to be exploited by emerging wireless SDN frameworks.

REFERENCES

[1] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012–2017, http://www.cisco.com/.

[2] 3GPP TS 23.251, "Network Sharing; Architecture and Functional Description," v. 11.3.0, 2012.

[3] Mobile Network Sharing Report 2010-2015, Development, Analysis & Forecasts, Market Study, Visiongain, 2010.

[4] H. Guan *et al.*, "Discovery of Cloud-RAN," *Cloud-RAN Workshop*, 2010.

[5] "Active RAN Sharing Could Save $60 Billion for Operators," http://www.cellular-news.com/story/36831.php.

[6] 3GPP TS 22.951, "Service Aspects and Requirements for Network Sharing," v. 11.0.0, 2012.

[7] 3GPP TR 22.852, 3GPP System Architecture Working Group 1 (SA1) RAN Sharing Enhancements Study Item.

[8] Telecoms.com Intelligence, Industry Survey, 2013.

[9] I. Stoica *et al.*, "A Hierarchical Fair Service Curve Algorithm for Link-Sharing, Real-Time, and Priority Services," *IEEE/ACM Trans. Net. J.*, 2000.

[10] Vanu Networks, http://www.vanu.com/.

[11] S. Paul and S. Seshan, "Virtualization and Slicing of Wireless Networks," GENI Design Doc. 06-17, 2006.

[12] G. Bhanage *et al.*, "Virtual Base Station: Architecture for an Open Shared WiMAX Framework," *Proc. ACM SIGCOMM VISA*, 2010.

[13] Y. Zaki *et al.*, "LTE Wireless Virtualization and Spectrum Management," *Proc. IFIP WMNC Conf.*, Oct 2010.

[14] EU FP7 FLAVIA Project, http://www.ict-flavia.eu/.

[15] X. Costa-Pérez *et al.*, "Latest Trends in Telecommunication Standards," *ACM Comp. and Commun. Review*, Apr. 2013.

[16] R. Kokku *et al.*, "NVS: A Substrate for Virtualizing Wireless Resources in Cellular Networks," *IEEE/ACM Trans. Net.*, vol. 20, issue 5, 2012.

BIOGRAPHIES

XAVIER COSTA-PÉREZ [M] (xavier.costa@ieee.org) is chief researcher at NEC Laboratories, where he has managed several projects related to mobile networks. In the wireless LAN area, he led a project contributing to 3G/WiFi mobile phones evolution and received NEC's R&D Award for his work on N900iL, NEC's first 3G/WiFi phone. In the 4G area he managed a team researching base station enhancements and recently received NEC's R&D Award for successful technology transfers. In 3GPP he is contributing to the SA1 RSE Study Item where new requirements for future systems are being defined. He has served on the Program Committees of several conferences, including IEEE Greencom, WCNC, and INFOCOM, and holds over 20 patents. He received both his M.Sc. and Ph.D. degrees in telecommunications from Universitat Politècnica de Catalunya (UPC) and was the recipient of the national award for the best Ph.D. thesis on multimedia convergence in telecommunications.

JOERG SWETINA (joerg.swetina@neclab.eu) studied chemistry and mathematics at the University of Vienna, Austria, and later conducted research in theoretical chemistry. Moving from academia to industry, he led a development team dealing with GSM call processing and software testing at Siemens Austria. Since the early days of 3GPP he represented Siemens and later Nokia Siemens Networks in standardization bodies like ETSI SMG, 3GPP, and OMA. In 2008 he moved to NEC Europe, Heidelberg, Germany, continuing his work in standards. In addition to 3GPP, his current field of interest covers machine-to-machine communication. He is active in ETSI TC M2M and is currently acting as Vice Chair of the requirements group of the "oneM2M Global Initiative" organization.

TAO GUO (tao.guo@emea.nec.com) received his B.Sc. degree in information engineering from Xi'an Jiaotong University, China, in 2004, his M.Sc. degree with distinction in communications engineering from the University of Birmingham, United Kingdom, in 2005, and his Ph.D. degree in wireless networking from Newcastle University, United Kingdom, in 2009. In May 2009 he joined the University of Surrey, United Kingdom, as a research fellow working on various EU projects and the Mobile VCE project. Since June 2012 he has been with NEC Telecom Modus Ltd as a senior systems engineer working on the development of the 3GPP LTE system. His current research interests focus on radio resource and mobility management, self-organizing networks, and network virtualization.

RAJESH MAHINDRA (rajesh@nec-labs.com) received his M.S. degree from the School of Electrical and Computer Engineering, Rutgers University, New Jersey, and his B. Tech. degree from the Department of Electronics and Communications Engineering, M. S. Ramaiah Institute of Technology, Bangalore, India, in 2005. He is currently working as a senior associate research staff member at NEC Labs America. His research interests include video streaming over wireless, wireless network virtualization, and wireless resource management with focus on next-generation cellular networks.

SAMPATH RANGARAJAN [SM] (sampath@nec-labs.com) received his M.S. degree in electrical and computer engineering and Ph.D. degree in computer sciences from the University of Texas at Austin. He heads the Mobile Communications and Networking Research Department at NEC Laboratories America in Princeton, New Jersey. Previously, he was a researcher at Bell Laboratories in New Jersey. He was also a co-founder and vice president of technology at Ranch Networks, Morganville, New Jersey, a venture-funded startup in the IP networking space. Before joining Bell Laboratories, he was an assistant professor with the Electrical and Computer Engineering Department, Northeastern University, Boston, Massachusetts. His research interests span the areas of mobile communications, mobile networks, and distributed systems. He has been on the Editorial Boards of *IEEE Transactions on Computers*, *IEEE Transactions on Parallel and Distributed Systems*, and ACM Mobile Computing and Communications Review.