

RMSC: A Cell Slicing Controller for Virtualized Multi-tenant Mobile Networks

Pablo Caballero Garcés^{*‡}, Xavier Costa Pérez[†], Konstantinos Samdanis[†], Albert Banchs^{*‡}

^{*}IMDEA Networks Institute, Madrid, Spain

[†]NEC Europe Ltd, Kurfursten-Anlage 36, 69115 Heidelberg, Germany

[‡]University Carlos III of Madrid, Spain

Abstract—The traditional model of single ownership of the mobile network infrastructure is being challenged by the forecasted mobile data tsunami and the resulting CAPEX and OPEX costs. In this context, the sharing of network infrastructure among operators has emerged as a way to ensure operators' future cost competitiveness. While the elementary concepts related to passive network sharing are already being exploited today, active network sharing is raising in importance to enable further reduction of network expenses in a substantial and sustainable way. The work presented in this paper addresses this challenge by designing a *RAN Multi-tenant cell Slicing Controller (RMSC)* that allows to flexibly share the RAN resources among multiple virtual operators (tenants). Three different possible designs for the RMSC controller are proposed, ranging from a fully distributed system with no inter-base station communication to a fully centralized solution with all information available. The proposed solutions are benchmarked against a distributed static slicing solution and a centralized load balancing solution, considering realistic scenarios of uneven user location distribution per tenant. The performance results obtained indicate that the proposed RMSC schemes can significantly outperform traditional solutions for multi-tenant mobile networks.

Index Terms—IMT-A, Virtualization, 3GPP LTE, Multi-tenancy, RAN Sharing.

I. INTRODUCTION

Wireless networks are a key element of today's society, enabling communication, access and information sharing. All forecasts agree in predicting that the demand for capacity will grow exponentially over the next years, mainly due to video services. However, as cellular networks move from being voice-centric to data-centric, operators' revenues are not expected to be able to keep pace with the predicted increase in traffic volume. Such pressure on operators' return on investment has pushed research efforts towards achieving more cost-efficient mobile network solutions. In this context, network sharing has emerged as a key business model for reducing deployment and operational costs (CAPEX and OPEX).

Network sharing solutions are already available, standardized [1] 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, etc. Active sharing refers to the reuse of base stations and antenna systems, the reuse of the latter two being labeled as active Radio Access Network (RAN) sharing. According to market surveys, mobile infrastructure

sharing has been already deployed by over 65% of European operators in different ways, and this trend is expected to expand in the future as it can be observed by the virtualization standardization trends [2].

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 [3] it was reported that around 20% of all sites carry about 50% of the total traffic. Estimations regarding the expected savings for operators by implementing active network sharing have been conducted, see e.g., [4]. This study concluded that operators worldwide could reduce combined OPEX and CAPEX costs in up to \$60 Billion over a 5 years period through network sharing and at least 40% of these cost savings are expected to come from active network sharing.

The traditional model of single ownership of all network layers and elements is thus being challenged. While the most basic concepts related to Passive network sharing are easier to implement and have already been partially exploited, Active network sharing is raising in importance to enable substantial and sustainable reduction in network expenses and thus ensure operators' future competitiveness. The work presented in this paper addresses the aforementioned challenges by designing a RAN Multi-tenant Slice Controller (RMSC) that efficiently allocate RAN resources. The proposed solutions consider realistic multi-tenant scenarios of uneven user location distribution per mobile (virtual) network operator.

II. RELATED WORK

Active RAN sharing enables pooling of radio resources enhancing the overall RAN utilization, while reducing infrastructure investments. 3GPP Services Working Group (WG) SA1 studied RAN sharing, analyzing a set of use cases deriving business requirements [5]. The architecture and operations that enable different MNOs to share the RAN are specified by the 3GPP Architecture WG SA2 in [1] detailing two different approaches. One referred to as Multi-Operator Core Network (MOCN), where each operator is sharing eNBs connected to the core network elements of each MNO using a separate S1 interface. In the second one named Gateway Core Network (GWCN), operators share additionally the Mobility Management Entity (MME). GWCN enables extra cost savings compared to MOCN, but at the price of reduced flexibility,

i.e. no mobility among different Radio Access Technologies and no interworking with legacy networks.

A preliminary approach of eNB virtualization in LTE that allows multiple MNOs to share the spectrum is analyzed in [6] introducing the notion of the hypervisor, which takes advantage of certain parameters such as channel conditions, sharing contracts, traffic load, etc., or a combination, to perform allocation of resource blocks to each operator. An alternative means for resource sharing is to introduce the notion of reserved resources that are guaranteed for a particular operator. Typically, operators are expected to select a customized mixture of reserved and shared resources. In [7] a partial resource reservation and admission control scheme is proposed to flexibly allocate shared resources to different operators, according to traffic demands and priorities, increasing the spectrum utilization and revenue for the infrastructure owner. A study concentrating on evaluating a range of different resource sharing options for the near term LTE considering a range of variations in capacity and spectrum sharing to advanced virtualized spectrum and resource block sharing is contained in [8], which investigates these models with respect to traffic density and the location of resource availability.

Building on eNB virtualization, [9], [10] introduce the notion of Network Virtualization Substrate (NVS) that operates closely to the MAC scheduler. The NVS separates the scheduling process into two steps: a first level is controlled by each tenant providing slice customization, while the second is managed by a virtualization layer that allocates resources to virtual instances of a base station. A framework called CellSlice [11] adopts a gateway-level solution based on the concept of NVS enabling remotely controlled scheduling decisions ensuring that each MNO receives the appropriate share of the radio resources. In this way CellSlice avoids modifications on MAC schedulers, but instead provides slicing of radio resources from a remote gateway, easing near-term adoption. Authors in [12] presents NetShare, a network-wide radio resource management framework that provides RAN Sharing. NetShare introduces a two-level scheduler to manage and allocate effectively the spectrum resources of a radio access network among multiple different entities that share the network. In [13], the authors present RadioVisor, which proposes an heuristic algorithm to slice time-frequency slots at each base station and share these among network operators while ensuring isolation. [12], [13] come up with similar problem formulation although they target different objective functions for RAN sharing.

Current multi-tenant solutions mostly focus on resource sharing among different slices without considering traditional load balancing means. A relatively simple resource broker based solution that combines multi-tenancy and distributed load balancing is described in [14] with the two mechanisms being performed separately at orthogonal instances. In this paper, we investigate more sophisticated solutions that can integrate multi-tenancy and load balancing considering different flavors of information exchange among neighboring eNBs. Our proposal is based-on the use of a network management

controller, which acts as a resource broker monitoring and controlling the resource allocated per slice and sector.

A network management framework for maintaining virtual network slices and services is described in [15], wherein different tenants abstract their physical resources allowing a capacity broker to enable the creation and management virtual networks. An SDN-based approach for controlling multi-tenant slices is included in [16]. The key feature of SDN is the flexibility in adjusting on-demand the resources allocated to each slice and the fact that it allows application providers to inquire resources and specific QoS associated with certain services. In this paper we adopt such capacity broker architecture and introduce intelligence in allocating resources and assigning user to particular slices considering also load balancing. In particular, we model the system based on Generalized Assignment Problem (GAP) and provide greedy solutions to deal with the inherited complexity of GAP.

III. RAN MULTI-TENANT CELL SLICING CONTROLLER (RMSC)

A. Design Considerations

The proposed RMSC solution introduces a logical controlling entity in a multi-tenant RAN, the *RMSC controller*, which collects information on the mobile virtual network operators radio resource utilization per cell according to the service level agreements (SLAs) as well as the total utilization per cell. This logical entity that can be a *centralized* or *partially distributed* function in the multi-tenant RAN, and may be collocated or integrated with other mobile network management functionalities as for instance a self-organizing-network (SON) server.

The RMSC controller also evaluates the utilization status per cell against the defined service agreements and policies. If one or multiple cells meet the conditions to initiate a corrective action according to the agreed policy, then the controller shall take the corresponding measures, defined by the multi-tenant RAN provider and the mobile virtual network operators involved.

The corrective actions envisioned consist in three major operations which can be performed in a proactive or reactive manner:

- 1) Shift UEs from one particular cell to another within or outside the multi-tenant RAN
- 2) Delay traffic of a specific operator in a cell (e.g., applicable to non real-time traffic)
- 3) Drop traffic of a specific operator in a cell (e.g., applicable to real-time traffic)

As operations 2) and 3) have an impact on the user QoE, the actual actions taken will depend on the configured policies per mobile virtual network operator. Solutions for these two operations can be found in the literature as described for instance in our previous work [10]. The work presented in this paper focuses on operation 1). The RMSC solution proposed first identifies the candidate UEs for which a correction action is needed. Then, based on this pool of candidate UEs, different

policies can be chosen to fulfill different targets for optimizing the resource usage of the multi-tenant RAN and/or optimizing the service experience of the mobile users.

B. Problem formulation and RMSC Optimal

The objective of the RMSC algorithm proposed is defined as follows: given a set of UEs ($N = 1 \dots n$) corresponding to a set of mobile virtual network operators ($O = 1 \dots o$) located within a given area, we want to associate the UEs to the sectors ($M = 1 \dots m$) such that we maximize the total utility value ($p_{UE,sector}$), given a predefined maximum utilization per sector ($T_{Util/Sector}$) and a predefined RAN slice utilization per operator ($Sh = sh_1 \dots sh_o$). This RMSC objective can be expressed as the following optimization problem:

$$\max \sum_{i=1}^m \sum_{j=1}^n p_{i,j} \cdot x_{i,j} \quad (1)$$

subject to

$$\sum_{j=1}^n w_{i,j} \cdot x_{i,j} \leq T_{Util/Sector}, \quad i \in M = 1 \dots m \quad (2)$$

$$\sum_{i=1}^m x_{i,j} \leq 1, \quad j \in N = 1 \dots n \quad (3)$$

$$\sum_{j=1}^n \sum_{i=1}^m Op_{i,k} \cdot x_{i,j} \leq Sh_k, \quad k \in O = 1 \dots o \quad (4)$$

$$x_{i,j} \in \{0, 1\} \quad (5)$$

where ($w_{UE,sector}$) is the weight of a user within a sector, which depends on the users' modulation, Op is a binary matrix that indicates whether a UE belongs to a mobile virtual network operator O or not, and $x_{i,j}$ indicates whether user i is assigned to sector j . The utility value ($p_{UE,sector}$) allows to prioritize among users based on aspects such as the number of cells in range of the user, the user's share or the total cell utilization, among others.

The above problem is an instance of the well-known Generalized Assignment Problem (GAP). Thus, an algorithm that solves GAP would allow us to find the optimal UEs to sector assignment given our constraints. However, GAP is a well-known NP-Hard problem if both weights and utilities values varies along sectors. If only weights or utilities changes, the problem remains APX-Hard, and when they are all identical for all users and sectors, the problem is polynomial time solvable. However, even in the latter case, the solution via bipartite b-matching algorithm has a high complexity cost that might make this algorithm infeasible for a real world implementations.

In the rest of this paper, we will refer to the exact solution of this problem as *RMSC Optimal*, and will use it as an upper bound performance benchmark. Given the complexity of obtaining the exact solution, in the following we design two heuristic algorithms, which we refer to as *Distributed RMSC Heuristic* and *Centralized RMSC Heuristic*. These algorithms aims at achieving a performance close to the *RMSC Optimal* solution but at a lower complexity cost.

C. Distributed RMSC Heuristic

In the distributed solution base stations take their own decisions based on local information available.

The objective of the *Distributed RMSC Heuristic* algorithm is to determine, on a per UE basis, whether connection requests can be directly admitted or should trigger a UE re-associations, given a pre-defined resource utilization target per base station sector ($T_{Util/Sector}$) and the corresponding configured resource share per operator. The rationale here is that as long as the resources of a base station sector are underutilized, no multi-tenant service control is needed since the service provided per tenant will be as agreed. However, as resources become scarce, multi-tenant control is needed to allocate resources according to the SLAs per tenant.

The *Distributed RMSC Heuristic* works as follows. As long as the resource utilization target per sector is not reached, UEs are admitted as usual, following the conventional LTE practices. Once the resource utilization target per sector is reached, the algorithm is triggered for every new UE connection request. In this case, new connection requests are evaluated against currently associated UEs. If a UE requesting association to a specific base station sector contributes a higher utility value than one of the UEs currently associated with this base station sector, a handover shall be initiated for the latter UE in order to keep the resource utilization target.

As an example of a potential implementation of the *Distributed RMSC Heuristic*, consider a utility function defined as the inverse of the number of cells in range of a UE (defined as the number of cells with a pathloss lower than the lowest pathloss plus a given configurable margin $PL_{threshold}$). With this utility function, incoming UEs trigger the re-association of the UEs with a larger number of options, which increases the probability that re-associated UEs find a suitable cell for handover. Within the UEs that satisfy this, we chose those that belong to virtual operators that use a higher resource share than agreed. Specifically, the proposed algorithm takes the list of UEs with the highest number of cell association options and prioritizes this list in decreasing order with the surplus deviation from the contracted operator's share. In Section IV, we analyze the performance of this specific implementation of the *Distributed RMSC Heuristic* against different algorithmic alternatives.

D. Centralized RMSC Heuristic

In order to evaluate the performance gain that could be obtained by knowing the resource utilization of neighboring base station sectors, we focus next on a centralized version of the *RMSC Heuristic*, which we call *Centralized RMSC Heuristic*. Based on this additional information, when evaluating UEs for re-association purposes, the algorithm can consider not only the utility of the UEs being re-associated, but also the likelihood of suffering no noticeable QoE degradation when handed over to a neighboring cell.

As in the distributed heuristic algorithm case, we provide an example of a potential implementation of the *Centralized RMSC Heuristic* algorithm for illustration purposes. Let us

Parameters	Values
Mobile Network	1 Tier, 19 BSs with 3 Sectors
Virtual Operators	3
IMT-A Scenario	Urban Micro (UMi)
Carrier Frequency	2.5 GHz
Intersite distance	200 m
Number of UEs	50 to 500
Users Mobility	SLAW Model
Pathloss	$36.7 \cdot \log_{10}(d) + 22.7 + 26 \cdot \log_{10}(f_c)$
$PL_{threshold}$	3 dBs
Users Association	Lowest Pathloss / Controller Decision
Simulation Time	1 hour
Warm-up Time	10 minutes

TABLE I
ITU IMT-ADV SIMULATION PARAMETERS CONFIGURATION

consider the same utility function for UE cell association as before (i.e., the inverse of the number of cells that have a pathloss lower than a given configurable margin, $PL_{threshold}$, over the best cell). The difference in this case though is that we will remove those cells with a resource utilization above a predefined target when counting the number of options.

In Section IV, we evaluate the performance of this specific implementation of the Centralized RMSC Heuristic and compare it against different algorithmic alternatives.

IV. PERFORMANCE EVALUATION

In order to evaluate the different algorithms proposed, in the following we describe the simulation setup used for performance benchmarking. For simplicity reasons in our evaluation we consider that all users consume the same number of physical resource blocks (PRB) per UE and thus define $T_{Util}/Sector$ in terms of the number of UEs associated to a base station sector and Sh_k as the total number of UEs in the system per mobile virtual network operator.

A. Simulation Setup

Mobile Network Model: The mobile network scenario considered, based on the IMT-Advanced evaluation guidelines [17], is depicted in Figure 1. It consists of 19 base stations in a hexagonal cell layout with 3 sector antennas. We focus on the 'Urban Micro-cell scenario' to model a 'small cell' deployment. The detailed system configuration parameters are summarized in Table I. The number of UEs in the system starts with 50 and is increased up to 500.

Mobility model: Users mobility follows the SLAW model defined in [18], which is a human walk mobility model based on real GPS-tracking measurements of more than 6 million UEs. This model allows us to synthetically generate UEs movements in our scenario in a realistic way. The configuration parameters used to generate the UEs movement in our experiments are given in Table II, and the resulting distribution depicted in Figure 1.

Benchmark 1 - Distributed Static Slicing Algorithm (Baseline): The first algorithm considered in our study for benchmarking purposes is the *Distributed Static Slicing*. With this approach, UEs from different mobile network virtual operators are associated to a base station sector with the lowest pathloss as long as the resource utilization per sector is below a




Fig. 1. User distribution models for 550 users. The different colours represent users from three different operators

Parameter	
Distance alpha	5
Waypoints(wp)	500
Inverse self-similarity of wp	0.95
Time (hours)	1.25
Clustering range (m)	50

TABLE II
SLAW PARAMETERS CONFIGURATION

predefined threshold. Once the resource utilization threshold is reached, UEs are associated to their next best pathloss option within the $PL_{threshold}$ margin. This models legacy systems with no multi-tenant support for UE re-association where the enforcement of the resource allocation per tenant would be performed at every base station independently at scheduling level (see [10] for an example).

Benchmark 2 - Centralized Load Balancing Algorithm: The second benchmarking algorithm that we consider a centralized load balancing solution which aims at adjusting equally the number of UEs associated per base station sector. To do so, this algorithm sequentially assigns each UE to the less loaded base station sector. The number of sectors association options per UE is computed dynamically based on the predefined pathloss constraints. The UE to base station assignment is checked regularly by a centralized system based on load information updates.

B. Performance Results

We next present the performance results obtained with the above simulation setup. We start by analyzing the performance from the perspective of the mobile infrastructure provider (e.g., operator with full management access to the physical infrastructure), focusing on the overall network performance. Then, we present the results from the mobile (virtual) operators perspective side, looking at the performance of each operator and showing that proposed algorithms meet the network sharing agreements.

Mobile Infrastructure Provider Perspective

Our analysis of the mobile infrastructure provider perspective focuses on several network performance metrics. Figure 2 shows the deviation of the sector utilization with respect to the predefined target $T_{Util}/Sector$ (which is set to 10 UEs in

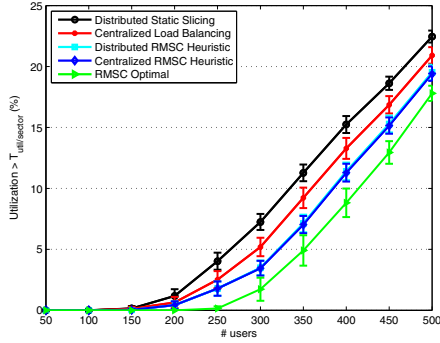


Fig. 2. Deviation in base stations sector utilization with respect to the predefined target $T_{Util/Sector} = 10 UEs$

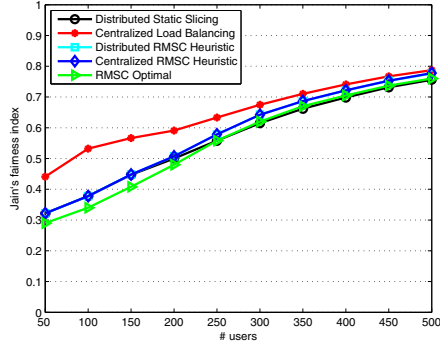


Fig. 3. Jain's fairness index in base stations sector utilization

this experiment according to the IMT-A guidelines for Urban Micro scenarios [17]). The error bars in Figure 2 represent the 95% confidence intervals for the blocking probability in this scenario.¹

As expected, the *Distributed Static Slicing* solution shows the worst performance, since it does not have the capability of handing over UEs according to a predefined base station sector utilization target. *RMSC* solutions clearly outperform the *Load Balancing* one. The reason for this is that, although the *Centralized Load Balancing* solution assigns UEs to base station sectors based on their current resource utilization, it inherently pursues and requires a uniform distribution of UEs in the system, which is not realistic in real networks (as shown in Figure 1) [18]. Among the *RMSC* solutions, as expected *RMSC Optimal* outperforms the others, since it considers all possible assignments. No noticeable differences are observed between *distributed RMSC Heuristic* and *Centralized RMSC Heuristic*, which suggests that knowing the resource utilization in neighboring sectors is useful when taking a handover decision, while the number of handover options of the UEs is critical.

In Figure 3, we analyze the sector load fairness according to Jain's fairness index². The *Centralized Load Balancing* solution shows the best performance, since it is specifically

¹Error bars are not displayed in further figures for readability reasons.

² $fairness(bs = 1...n) = (\sum_{i=1}^n load_i)^2 / (n \sum_{i=1}^n load_i^2)$




Fig. 4. Average number of handovers per minute per UE for $T_{Util/Sector} = 10 UEs$

Algorithm	Execution time
Distributed Static Slicing	1.47 ms
Centralized Load Balancing	19.6 ms
Distributed RMSC Heuristic	18.5 ms
Centralized RMSC Heuristic	40 ms
RMSC Optimal	202.1 s

TABLE III
ALGORITHMS PROFILING TIMES (WORST CASE, 870 UEs)

designed for this purpose. As the load in the network increases, the *RMSC* solutions get closer to the *Centralized Load Balancing* but always keep below, as their target is not to have a similar load in all cells but rather keep the load below the specified limit $T_{Util/Sector}$.

The number of handovers per minute per UE, which is another relevant performance metric, is shown in Figure 4. In terms of this metric, *RMSC Optimal* is the worst performing solution, followed by the *Centralized Load Balancing*; the reason is that both approaches do not consider the previous UE to sector assignment. Both *RMSC Heuristic* approach present similar handover numbers to the *Distributed Static Slicing* when the load is low, and slightly higher as the load increases. This the price to pay in order to achieve the utilization gains shown in Figure 2.

Finally, we compare the different solutions in terms of their execution time³. Table III shows the running time of the different solutions. The *Centralized Load Balancing* and *RMSC Heuristic* solutions have an execution time that is one order of magnitude higher than the *Distributed Static Slicing*. Additionally, the *RMSC Optimal* approach has an execution time in several orders of magnitude larger, and thus is not feasible in practice.

Mobile (Virtual) Operators Perspective

In following we show the results from the mobile (virtual) operators perspective, by focusing on how the proposed solutions affect the individual performance of virtual network operators. In order to do so, we consider three different network sharing agreements: triangular, gaussian and irregular sharing distributions (see Table IV). Given the performance gains observed for the *RMSC* solutions in the previous sections,

³Time consumption measured in a regular server with 4 cores. Note that the algorithms implemented are not optimized, so their execution time could be lowered.

Distribution	Operator 1, Sh_1	Operator 2, Sh_2	Operator 3, Sh_3
Triangular	0.2	0.6	0.2
Gaussian	0.158	0.684	0.158
Irregular	0.1	0.6	0.3

TABLE IV
CAPACITY SLICING DISTRIBUTION PER OPERATOR $\in \{0, 1\}$

and the fact that the *RMSC Optimal* approach is unfeasible in realistic scenarios, in what follows we focus on the *Distributed RMSC Heuristic* and *Centralized RMSC Heuristic* approaches.

Figure 5 shows the results in terms of utilization of base station sectors. As it can be seen in the figure, both *RMSC Heuristic* solutions successfully manage to distribute UEs per operator and sector according to their agreed capacity share. In the case of the *triangular* distribution, the percentage of UEs for operator 2 that are located in sectors with a utilization above $T_{Util}/Sector$ increases proportionally to their contracted share as the number of UEs in the system increases. Similar results are observed for the gaussian and irregular distributions. This behavior is achieved by prioritizing the re-association of UEs from operators that are above their configured share.

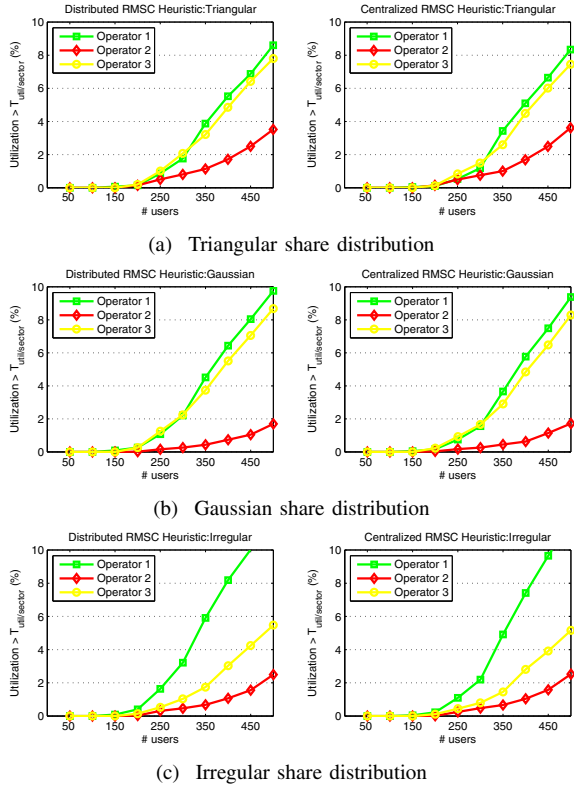


Fig. 5. Share distribution agreements.

V. SUMMARY & CONCLUSIONS

The traditional model of single ownership of the mobile network infrastructure is being challenged by the forecasted mobile data tsunami and the resulting CAPEX and OPEX costs. In this context, active network sharing is raising as a means to reduce network costs and enable the sustainable deployment and operation of the mobile network.

The work presented in this paper addresses this challenge by designing a *RAN Multicell Share Controller (RMSC)* that

allows to flexibly sharing the RAN resources in realistic scenarios. Three different possible designs for the *RMSC* controller are designed, ranging from a fully distributed system with no inter-base station communication to a fully centralized solution with all information available.

The proposed solutions have been benchmarked against a *Distributed Static Slicing* and a *Centralized Load Balancing* approach. The results show that: (i) the *RMSC* are effective in keeping the utilization per sector below a predetermined limit, (ii) the gain resulting from exchanging information among neighbours through the x2 interface is very small, (iii) very high gains result from choosing the candidate UEs for handover based on their utility value, and (iv) while the user distribution patterns have an impact on the gains observed, the *RMSC* solutions always outperform the other approaches.

REFERENCES

- [1] 3GPP, "Network Sharing; Architecture and Functional Description." TS 23.251, v12.1.0, Jun. 2014.
- [2] X. Costa-Perez, A. Festag et al., "Latest Trends in Telecommunication Standards," *ACM Computer and Communications Review*, 2013.
- [3] H. Guan, T. Kolding, and P. Merz, "Discovery of cloud-ran," in *Cloud-RAN Workshop*, 2010.
- [4] "Active RAN Sharing Could Save \$60 Billion for Operators." <http://www.cellular-news.com/story/36831.php>, 2009.
- [5] 3GPP, "Study on Radio Access Network (RAN) Sharing Enhancements," TR 22.852, v12.0.0, Jun. 2013.
- [6] Y. Zaki, L. Zhao, C. Goerg, and A. Timm-Giel, "LTE Wireless Virtualization and Spectrum Management," in *Wireless and Mobile Networking Conference (WMNC), 2010 Third Joint IFIP*, 2010, pp. 1–6.
- [7] T. Guo and R. Arnott, "Active LTE RAN Sharing with Partial Resource Reservation," in *VTC Fall 2013.*, 2013, pp. 1–5.
- [8] J. Panchal, R. Yates, and M. Buddhikot, "Mobile Network Resource Sharing Options: Performance Comparisons," *Wireless Communications, IEEE Transactions on*, vol. 12, no. 9, pp. 4470–4482, 2013.
- [9] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan, "NVS: A Substrate for Virtualizing Wireless Resources in Cellular Networks," *Networking, IEEE/ACM Transactions on*, vol. 20, no. 5, pp. 1333–1346, 2012.
- [10] X. Costa-Perez, J. Swetina, T. Guo, R. Mahindra, and S. Rangarajan, "Radio Access Network Virtualization for Future Mobile Carrier Networks," *Communications Magazine, IEEE*, vol. 51, no. 7, 2013.
- [11] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan, "CellSlice: Cellular wireless resource slicing for active RAN sharing," in *Communication Systems and Networks (COMSNETS), 2013 Fifth International Conference on*, 2013, pp. 1–10.
- [12] R. Mahindra, M. Khojastepour, H. Zhang, and S. Rangarajan, "Radio access network sharing in cellular networks," in *Network Protocols (ICNP), 2013 21st IEEE International Conference on*, 2013, pp. 1–10.
- [13] A. Gudipati, L. E. Li, and S. Katti, "Radiovisor: A slicing plane for radio access networks," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 237–238.
- [14] M. Li, L. Zhao, X. Li, X. Li, Y. Zaki, A. Timm-Giel, and C. Gorg, "Investigation of network virtualization and load balancing techniques in lte networks," in *Vehicular Technology Conference (VTC Spring), 2012 IEEE 75th*, 2012, pp. 1–5.
- [15] ITU-T, "Report ITU-T Y.3011, Framework of Network Virtualization for Future Networks," Next Generation Networks Future networks., Jan. 2012.
- [16] C. Bernardos, A. De La Oliva, P. Serrano, A. Banchs, L. Contreras, H. Jin, and J. Zuniga, "An Architecture for Software Defined Wireless Networking," *Wireless Communications, IEEE*, vol. 21, 2014.
- [17] ITU-R, "Report ITU-R M.2135-1, Guidelines for evaluation of radio interface technologies for IMT-Advanced," Tech. Rep., 2009.
- [18] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, "Slaw: self-similar least-action human walk," *IEEE/ACM Transactions on Networking (TON)*, vol. 20, no. 2, pp. 515–529, 2012.