Strategies for Network Slicing Negotiation in a Dynamic Resource Market

Alessandro Lieto*†, Eugenio Moro*†, Ilaria Malanchini†, Silvio Mandelli†, Antonio Capone*
*Politecnico di Milano, Italy
Email: {alessandro.lieto, antonio.capone}@polimi.it, eugenio2.moro@mail.polimi.it
†Nokia Bell Labs, Stuttgart, Germany
Email: {ilaria.malanchini, silvio.mandelli}@nokia-bell-labs.com

***Abstract*—One of the disruptive innovations introduced by 5G networks is the opportunity for a new group of stakeholders to be actively involved in the management of network slices with the role of *tenants*. This allows to go beyond the user-centric QoS paradigm of 4G, and to include tools for handling the aggregate performance of multiple services and user groups and to focus on slice resource management, also at the new 5G NR interface. So far, research efforts have privileged a first solution based on the concept of isolation between slices. However, proposed solutions are not particularly efficient due to the loss of pooling gains, and not very reliable due to variable channel conditions that with slice limited resources make performance not easily predictable. We propose a slice management framework where the shared resources are negotiated by tenants in a real-time market based on slice instantaneous demands. Our model, based on game theory, allows tenants to optimize their service strategies acquiring resources when and where it is necessary, according to the level of quality and reliability requested by the specific traffic types they handle. In this paper, we focus on modeling the game theoretical framework and on characterizing its equilibria in a multi-tenant scenario.**

I. Introduction

The availability of digital technologies has characterized the start of what is called the new industrial revolution, which is deeply changing the production of goods and associated processes, as well as the services offered to a large variety of end users in different industrial sectors. So far, however, the connectivity component of the application solutions has been rather critical, since based either on expensive dedicated technologies or on mobile networks mainly optimized for consumer services and thus with limited performance and reliability. The introduction of 5G is expected to radically change this scenario, allowing new players, active in different domains, to benefit of tailored connectivity services that can be directly adapted to needs based on different business strategies and technical solutions.

The new 5G networks and their operators must then manage a large set of use cases with very diverse needs and connectivity requirements, targeting new market segments and vertical industries. *Network slicing* is the key technology that allows network operators to expose their physical infrastructure to support a variety of services with diverse requirements [1]. Different service sets can be associated to logically independent end-to-end networks, i.e. *slices*. Slices are customized and managed by *tenants*, to which operators delegate the control of resource utilization and service performance.

978-1-7281-0270-2/19/$31.0 ©2019 IEEE

In order to guarantee service performance, a relatively easy solution is that of isolating slices on separated resources, so that quality and behavior of one slice cannot be affected by any change on the others. In the end-to-end slicing approach, virtualization technologies, like Software Defined Networking (SDN) and Network Function Virtualization (NFV), can provide flexible and scalable management of virtual network resources for slices based on a proper dimensioning of computation capacity and transport connectivity [2]–[5]. However, the use of isolation and resource partitioning strategies at the radio interface is much more challenging [6]. Indeed, due to the stochastic nature of the wireless medium and the high traffic variability in time and space, it is very hard to provide perfect isolation, unless resources are largely overprovisioned according to worst expected conditions. Obviously, this would lead to an inefficient use of resources for most of the time.

From an economic point of view, the pure slice isolation approach puts the risk of providing performance guarantees on the infrastructure provider, which partitions resources and charges high costs to tenants due to their inefficient use. A different approach can be defined with dynamic resource sharing techniques, where also risks can be shared among the different players allowing more flexible and diverse business models ranging from low-risk high-cost, to high-risk low-cost, according to specific needs of applications and willingness of end users to pay for quality. Indeed, tenants may have different economic evaluation of the type of slice they are willing to deploy. For example, some slices may require high availability and very low flexibility in deviating from target network performance, tolerating the higher costs it will imply; while others may require a trade-off between performance guarantees and its related cost. In this paper, we address this scenario and propose a slicing management framework where tenants renegotiate the shared resources in a real-time market. We follow an approach similar to the one introduced by [7] for smart grid applications in energy markets. As in [7], we structure the problem with a game theoretical model, where the tenants adapt their strategies purchasing resources when and where it is necessary. Although there are previous works modeling the slicing resource allocation problem through game theory [8], they do not consider business implication

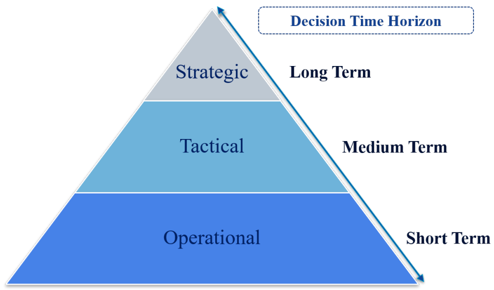

Fig. 1: Decision making process in organization management

on the tenant strategies. To the best of our knowledge, this is one of the first works that model the short time scale resource allocation and slicing renegotiation as a competitive market.

Dynamic sharing of resources has been proved to be an efficient solution when multiple agents coexist on a single physical infrastructure [9], [10]. For instance, [11] proposes different architectural solutions for slicing the Radio Access Network (RAN), that also considers resource management strategies customized for each slice. The works in [12], [13] assume to differentiate the resource allocation strategies for slices through prioritization policies. Although it allows a dynamic share of resources among slices, the slice strategy implementation is static and does not allow tenants to optimize their decisions according to current network conditions.

Our contribution in this paper can be summarized as follows. We propose:

- a general framework that describes the management of radio resource for network slicing;
- a market mechanism that regulates the allocation of radio resources to slices;
- a techno-economic game that allows to study the strategies of the tenants at Nash equilibrium.

We remark that our focus in this paper is on the design of the mechanism that allows tenants to engage in this dynamic market of resources. Therefore, any economic assumption that we take is used only to express the potential interests of the tenants in such an economic system.

The paper is structured as follows. In Sec. II, we introduce the design of our slicing management framework. In Sec. III, we describe the system model for the negotiation of resources among slices. While in Sec. IV the game structure is defined and the properties of the market game are illustrated. The discussion on preliminary results of our approach is given in Sec. V. Finally, Sec. VI concludes the paper.

II. Slicing Management Framework

In this section we discuss in detail the overall Slicing Management Framework (SMF) that is accounted for the mapping from high level requirements to scheduling resources. For the sake of generality, in this paper we keep the description generic without referring to specific layers of any wireless protocol stack, since the SMF can be applied to any system that requires managing and operating network slices in a multi-tenant scenario. The needs of a new and automated approach

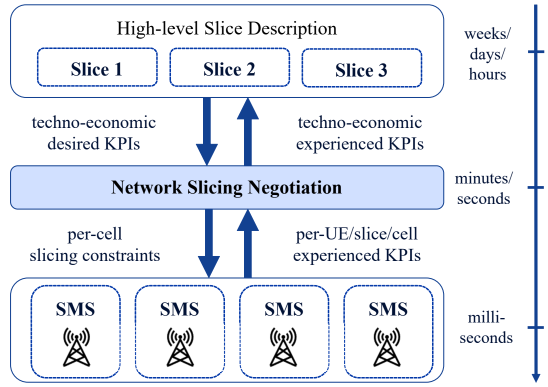

Fig. 2: Slicing Management Framework

comes from the nature of slices and wireless resources, where, on one side, slices are usually defined by tenants in the form of "high-level" Quality-of-Service (QoS) requirements (e.g. required throughput and latency) covering a large geographical area and lasting a variable time duration. On the other side, radio time-frequency resource scheduling must be performed in real time, i.e. each Time Transmission Interval (TTI), to guarantee an efficient usage and quick reactions to both traffic bursts and high priority users. Therefore, the proposed framework aims at managing slice QoS constraints and resource allocation in a distributed and scalable manner. Driven by a top-down approach, the SMF is able to handle the long-term and high-level requirements of the tenants (i.e. expected performance of a slice on large time window and over large macro-areas) as well as the needs of the radio scheduler, which takes care in each TTI of the resource allocation of each user.

As common in organization management [14, Chapter 2], decision making processes follow a pyramidal hierarchy, as sketched in Fig. 1. Namely, decisions are taken at different layers, which correspond to the different management layers of an organization. The strategic management layer deals with the long-term planning of the organization and provides the milestones to the lower layers to lead them in the decisions that need to be taken. Moving down the pyramid, the tactical management layer is responsible to identify a strategic plan for each situation may arise, rather than applying a standard procedure. It monitors and controls the current performance status and identifies the best strategy that benefits the organization. At the lower layer, the operation management takes short term decisions, meaning that predefined routines and rules are applied. By moving from top to down in the pyramid, not only the type of decision structure changes, but also the time horizon over which each decision is effective. According to the pyramidal hierarchy, the effectiveness (i.e. duration) of the decisions taken ranges from long-term horizon, for the strategic management, to short-term scale, for the decisions of the operational management. Clearly, each decision layer is responsible of the overall success of the organization.

By analogy to organization management, we design the SMF as a decision making process starting from the high-level network slice definition, ending up to the allocation of physical resources at the radio scheduler, as shown in Fig. 2. At the top of our hierarchy, the tenants and the infrastructure provider define high-level descriptors for the slices in the form of a set of techno-economical Key Performance Indicators (KPIs), tailored for any specific slice. These KPIs describe the business model of the slice and may include, among others, macro-area/per-cell constraints on aggregate throughput/resource utilization, and/or information about the QoS requested by the users of the slice. In this work, we focus on this last case, leaving the generalization of the SMF to macro-area management for future work. By definition, the aforementioned descriptors are valid for a long-term time window and over the entire spatial coverage area of the slice. Although these descriptors can generalize the expected behavior of the slice, they cannot capture the dynamic evolution of the radio environment. For instance, operators and tenants may agree on the average resource share of a slice, given an average user distribution for a specific area and a user target QoS. However, this does not provide indications on the decisions to be taken in case of bursts of traffic for some slices. To tackle the uncertainty of the wireless channel and network conditions, a baseline approach would over-provision resources to control the risk of performance degradation for slices, leading to highly inefficient solutions in terms of resource utilization for the infrastructure provider, costs of resources for the tenants and QoS for the users. Although we do not exclude the case in which techno-economic KPIs define also reserved resource guarantees to slices in each area, it is well-known that allocating and renegotiating resource shares dynamically (i.e. assigning them only when needed to satisfy the long-term KPIs and user QoS targets) brings gains in terms of efficiency in the resource utilization as well as cost savings. Hereafter we assume that the tenants compete for all the available resources, however our considerations also hold in hybrid scenarios, where a part of the resources is statically assigned to the tenants and the other part is contented among them.

We identify the entity responsible to adapt the slice high-level descriptors into physical requirements/constraints for the MAC scheduler (e.g., resource share for each slice) as the Network Slicing Negotiation (NSN). The NSN allows to accommodate the requirements of the slices according to the actual needs of their users. Renegotiations of resource shares among tenants take place whenever network conditions change, e.g. due to handovers, change of users behavior or position. How these are handled is defined by techno-economic KPIs, assessing the profit of given achieved performance. To model such competition, we propose a resource market mechanism[1] that regulates the negotiation of the resource allocations among slices.

Finally, given the outcome of the NSN in form of low-level KPIs (e.g. slice/user average throughput/resource allocation), radio resources are allocated at every TTI to single users by a Slicing MAC Scheduler (SMS). Previous works in the literature [16], [17] address the SMS problem. Namely, they propose solutions where resources are allocated to users relying on proportional fair metrics, twisted to take into account user/slice constraints received by the higher layers.

In this work, we focus on the modeling and characterization of the NSN. In particular, the NSN must (i) process the high-level KPIs coming from the tenants and input its decision to each cell SMS, (ii) adapt those decisions based on SMS feedback, and (iii) report achieved performance to the topmost layer, where tenants may monitor and correct their long-term policies. According to the scheme in figure 2, the NSN adaptation and decision process takes place in time windows of seconds to minutes. In the next section, we detail the role of the NSN and define the underlying market mechanism.

III. Network Slicing Negotiation

As illustrated in Fig. 2, the NSN takes as input the slice descriptors that are defined individually between each tenant and the infrastructure provider. In particular, those inputs include the slice performance requirements and the economic constraints that describe the profitability of tenants' actions. Since those descriptors are assumed to be valid on a long-term, throughout this work we consider them as fixed input parameters for the proposed NSN, which takes decisions accordingly. However, to make them processable by the NSN, we map them to *utility functions*, as done in [13]. Namely, such functions define how much it is profitable for a slice to buy a certain amount of time-frequency *resources*. In this way, the NSN can validate the strategy of a tenant and evaluate the slice profitability given a specific underlying scenario.

From the bottom layer, i.e. the MAC layer, the NSN can extract periodically (order of seconds) information about the number of users, and average channel conditions, i.e. spectral efficiency. By combining the knowledge about high-level targets with the current context and achieved performance, received, respectively, from the top and bottom layer, the NSN can estimate the amount of resources needed. This is then translated into instructions for the MAC scheduler in the form of QoS constraints to be applied to each slice. In the remaining of this section, we describe the general system model and the market model that represents the behavior of users and tenants, and the parameter settings of the infrastructure provider.

A. System model

We consider a single cell scenario operated by an infrastructure provider where a set of tenants deploy their slices[2]. Let $\mathcal{S}$ be the set of slices and $\mathcal{K}$ the set of active users in the system. We denote by $\mathcal{K}_s \subseteq \mathcal{K}$ the set of users belonging to a generic slice $s \in \mathcal{S}$, and assume that each user belongs to a single slice, i.e. $\{\bigcap_{s\in\mathcal{S}} \mathcal{K}_s = \emptyset\}$ and $\{\bigcup_{s\in\mathcal{S}} \mathcal{K}_s = \mathcal{K}\}$. This last assumption helps us with the notation and does not exclude the

[1]In economics, the market mechanism is a mechanism by which the use of money exchanged by buyers and sellers with an open and understood system of value and time trade-offs in a market tends to optimize distribution of goods and services in at least some ways [15].

[2]For the sake of simplicity, we assume that each tenant deploys a single slice and we will refer to them interchangeably.

case in which the same user has two different traffic profiles, since it would be sufficient to split the two different traffic profiles from the same physical user, and consider them as two separated users.

As mentioned above, we are dealing with a scenario where the shared resources are limited, namely by the total bandwidth of the cell. Therefore, also the requirements of the slices are constrained by resources actually available. Therefore, to model the competition for *scarce resources*, we introduce the definition of a *market*, where tenants compete as business entities with conflicting objectives.

B. Market model

Given the scarce nature of the resources, we propose a mechanism for the market to control the cost of resources according to their scarcity. To model the trend of the price in the market, we refer to the concept of *purchasing power*, well-known in economics. Given the relative normalized cell load $l = \sum_{s \in \mathcal{S}} x_s \in [0, 1]$, where x_s is the portion of resources assigned to slice s, the purchasing power is defined as the amount of resources that can be bought by a unit of money, hence it links the amount of resources that can be bought x with the total cost C. Formally,

$$pp : [0, 1] \to \mathbb{R}^+, \quad C \cdot pp(l) = x. \tag{1}$$

Given that the scarcity of radio resources is the only factor that concurs to the variation of the purchasing power, the infrastructure provider is considering all tenants to be equal in its resource pricing process, without discriminating between them. Nevertheless, considering the case of different functions for tenant is straightforward.

In this work, we model the purchasing power of the market as a decreasing sigmoid function:

$$pp(l) = 1 - \frac{1}{1 + e^{-\gamma(l - l_0)}}. \tag{2}$$

where γ and l_0 are the sigmoid's steepness and midpoint, respectively. As the load in the cell increases, the purchasing power of tenants will decrease accordingly. The price per unit of resources can be derived as the inverse of the purchasing power. This last information will be relevant for the tenants to decide whether to increase or reduce their resource share x_s, based on their own utility functions.

C. Tenants' utility function

As mentioned above, we map the high-level descriptors of a slice in slice-specific utility function, denoted by $u_s(\cdot)$, which depends on the resource share x_s and the total cell load l. Formally, the utility function of a generic slice s is defined as:

$$u_s(x_s, l) = R_s(x_s) - C(x_s, l). \tag{3}$$

where:

- $R_s(x_s)$ is the expected revenue, which depends on the resources allocated to the slice;
- $C(x_s, l)$ is the cost paid for purchasing the resources, calculated according to Eq. (1), which depends not only on the resources assigned to the slice but also on the overall cell load, thus reflecting the cost of competition.

In this work, we assume that the utility functions measure the economic value associated to the decision of the tenants. A possible extension of our proposed model is to map the tenant's revenues and its costs by a generic utility function that models the willingness of the tenant to buy resources

In the following two paragraphs, we elaborate on the two terms of Eq. (3).

1) Revenue model: We assume that slice's revenues are user-centric, meaning that the economic return for a tenant depends on the quality of service provided to the users, thus their satisfaction. To do so, we need to define the portion of resources allocated to user k[3]:

$$\hat{x}_k = \frac{x_s}{|\mathcal{K}_s|}, \; x_s = \sum_{k \in \mathcal{K}_s} \hat{x}_k, \tag{4}$$

where $|\cdot|$ denotes the cardinality of a set. Moreover, we define a per-user *acceptance probability function*, similarly to what is done in [19], as

$$A_k(r(\hat{x}_k)) = 1 - q^{\left(\frac{r(\hat{x}_k)}{r_k^0}\right)^{\mu_k}}, \tag{5}$$

where:

- $r(\hat{x}_k) = \eta_k \hat{x}_k$ is a function measuring the throughput obtained by user k, when $\hat{x}_k$ resources are assigned, where η_k is the k-th user average spectral efficiency measured from the SMS;
- r_k^0 is the maximal performance level for user k;
- $q \in (0, 1)$ defines the level of satisfaction when $r(\hat{x}_k) = r_k^0$;
- μ_k models the user sensitivity to performance variation.

Finally, the revenue function of a slice is defined as

$$R_s(x_s) = \sum_{k \in \mathcal{K}_s} \chi_k A_k(r(\hat{x}_k)), \tag{6}$$

where χ_k is the per-user revenue, that is the monetary value that the tenant s expects from k if it is fully satisfied with the experienced QoS. Since we model the revenues through the users' level of satisfaction, no service provision (no allocated resource) implies no revenues for the slice.

By defining different parameters in Eq. (5) and Eq. (6), one can model different behaviors for users within slices and guarantee service differentiation.

2) Cost model: While the revenue function is customizable for each slice, the cost function is given by the infrastructure provider and common to all slices. Therefore, we define the costs

$$C(x, l) = \frac{x}{pp(l)} \tag{7}$$

for buying x resources with a cell load l. Note that since $C(0, l) = 0$, the market is free to entry for all tenants.

[3]We assume resources of a slice to be equally split among their users for ease of notation. Note that this corresponds to the allocation of a proportional fair scheduler [18] with static user channel conditions. However, this assumption is not necessary and any other mapping between slice resource and users' allocation vector, i.e. $\hat{\mathbf{x}}_s = g(x_s)$, can be applied

Algorithm 1 Best Response Dynamics

procedure NE($\mathcal{S}, X$)
  initialize $\mathbf{x}^{old} \leftarrow \mathbf{x}^{(0)}$
  while $stop\_bucket < |\mathcal{S}|$ **do**
    for $s \in \mathcal{S}$ **do**
      $u_s^{old} \leftarrow u_s(x_s^{old}, \mathbf{x}_{-s}^{old})$
      $x_s^* \leftarrow \arg\max_{x_s} u_s(x_s, \mathbf{x}_{-s}^{old})$
      if $u_s(x_s^*, \mathbf{x}_{-s}^{old}) \leq u_s^{old}$ **then**
        $stop\_bucket \leftarrow stop\_bucket + 1$
      else
        $x_s^{old} \leftarrow x_s^*$ ▷ update the allocation
        $stop\_bucket \leftarrow 0$ ▷ reset the token

IV. Slicing Market Game

We define a game to model the competitive behaviour of the tenants in the market.

Let $\Gamma = \langle \mathcal{S}, (X_s)_{s\in\mathcal{S}}, (u_s)_{s\in\mathcal{S}} \rangle$ be the game in strategic form, where:

- $\mathcal{S} \doteq \{1, 2, ..., s\}$ is the set of players (i.e. the tenants);
- $X_s = [0, 1]$ is the strategy space (i.e. resources to buy) of the generic player s;
- $\mathbf{X} \doteq X_1 \times X_2 \times \cdots \times X_S$ is the cartesian product of all the strategy spaces, and $\mathbf{x} \in \mathbf{X}$ is a strategy profile of the game;
- $u_s : \mathbf{X} \to \mathbb{R}$ is the utility function of player s, as defined in Eq. (3).

Recalling the cost function defined in Eq. (7), one can replace the dependency on the total load of the cell l by $\mathbf{x}_{-s}$, i.e. $C(x_s, \mathbf{x}_{-s})$, where $\mathbf{x}_{-s}$ denotes the vector of strategies of the players but s. Indeed, the total load of the cell is simply the sum of the allocation strategies of the players, i.e.:

$$l = x_s + l_{-s} = x_s + \sum_{t\in\mathcal{S}\setminus\{s\}} x_t, \tag{8}$$

where l_{-s} denotes the load caused by the players except s. Similarly, we denote the utility function of a player, defined in Eq. (3), as $u_s(x_s, \mathbf{x}_{-s})$.

Let x_s^* be the strategy maximizing the payoff of the tenant s, then the strategy profile $\mathbf{x}^* = (x_s^*, \mathbf{x}_{-s}^*)$ is a Nash Equilibrium (NE) [20] if and only if

$$u_s(x_s^*, \mathbf{x}_{-s}^*) \geq u_s(x_s, \mathbf{x}_{-s}^*), \quad \forall s \in \mathcal{S}, x_s \in X_s.$$

To find the NE of the proposed game we apply a Best Response Dynamics (BRD) algorithm, described in Algorithm 1. As common in game theory, the *Best Response (BR)* strategy of a player is the strategy that maximizes the payoff given the other players' strategies. Formally, the best response correspondence is given by:

$$BR_s(\mathbf{x}_{-s}) = \arg\max_{x_s} u_s(x_s, \mathbf{x}_{-s}), \quad \forall s \in \mathcal{S}. \tag{9}$$

A. Existence of Nash equilibria

In literature, fixed point theorems are extensively used to obtain results on the existence, quantity and quality of the Nash equilibria. Indeed, any strategy profile is a Nash equilibrium if and only if it is a fixed point of the joint best response. More rigorously, let $\mathbf{x} \in \mathbf{X}$ be a generic strategy profile and $BR : \mathbf{X} \to \mathbf{X}$ the joint best response correspondence. Then $\mathbf{x} \in \mathbf{X}$ is a Nash equilibrium if and only if $\mathbf{x} \in BR(\mathbf{x})$.

In order to prove the existence of a Nash equilibrium for our game model, we use the following reformulation of the fixed point theorem in [21].

Theorem 1. *A game $\langle \mathcal{S}, (X_s)_{s\in\mathcal{S}}, (u_s)_{s\in\mathcal{S}} \rangle$ with the following properties:*

(i) the players are finite and each strategy space X_s is a compact of $\mathbb{R}$,
(ii) each utility function u_s is bounded, continuous and at least twice differentiable in X_s,
(iii) $\frac{\partial}{\partial x_s \partial x_t}[u_s(x_s, \mathbf{x}_{-s})] \leq 0 \; \forall s, t \in \mathcal{S}, s \neq t$,
(iv) each utility function u_s depends only on the player's strategy and on the sum of the opponents' strategies,

always admits at least one Nash equilibrium in pure strategies.

The proof of the theorem can be found in the Appendix. Concisely, the argument of the theorem is that any class of games verifying the properties of Theorem 1 belongs to a subclass of games that proves the hypotheses of the theorem in [21].

By means of Theorem 1, we can state the following proposition.

Proposition 1. *The game Γ always admits at least one Nash equilibrium in pure strategies.*

Proof. We proceed by showing that Γ satisfies properties (*i*)-(*iv*) of Theorem 1. (*i*) is verified by definition of strategy spaces of the game Γ. (*ii*) is implied by definition of the utility function in Eq. (3). Indeed, it is a sum of two continuous, bounded and C^∞ functions. Moreover, these two functions depend only on the player's own strategy (Eq. (6)) and the sum of the opponents' strategies, as shown in Eq. (8), which also verifies (*iv*). Finally, (*iii*) can be analytically verified from $\frac{\partial}{\partial x_s \partial x_t}[u_s(x_s, \mathbf{x}_{-s})] = -\frac{\partial}{\partial x_s \partial x_t}[C(x_s, \mathbf{x}_{-s})] \leq 0$. ☐

B. Quantity and quality of the Nash equilibria.

While Theorem 1 is instrumental to prove that game Γ admits at least one Nash equilibrium, it gives no information about the uniqueness of the equilibrium. Nonetheless, based on our simulations experiments (whose results are presented in greater details in Section V), we can claim that there exist some game configurations in which the equilibrium is not unique. We have also empirically observed that the Nash equilibria can be generally classified in two categories:

(i) equilibria whose strategies are all greater than zero;
(ii) equilibria where at least one strategy of a player is zero.

We now focus our attention on the latter. An equilibrium belonging to category (ii) is such that the player or the players with a zero strategy cause no load to the system, thus having no influence on the other players' strategies. It means that these zero-strategy players are actually excluded from the game. For this reason, from now on we will address to category (ii) as Nash equilibria with *exclusion effect* (NE-EE). The reason why a player's best response might be zero is to be found in the very nature of the model: if the cost of any amount of resource is higher than the revenue it could generate, the player is able to avoid an economic loss only if its strategy is zero. This allows us to define, for each player $s \in \mathcal{S}$, a *maximum bearable external load*[4], $\bar{l}_s$, such that $l_{-s} \geq \bar{l}_s \implies BR_s(l_{-s}) = 0$ and $l_{-s} < \bar{l}_s \implies BR_s(l_{-s}) > 0$. With this definition, we can state the following necessary condition for the existence of a NE with *exclusion effect*:

Proposition 2. *Given an instance of the game* Γ*, if there exists a player* $t \in \mathcal{S}$ *such that the game* $\hat{\Gamma} = \left\langle \hat{\mathcal{S}} = \mathcal{S} \setminus \{t\}, (X_s)_{s\in\hat{\mathcal{S}}}, (u_s)_{s\in\hat{\mathcal{S}}} \right\rangle$ *admits an equilibrium in which the sum of the strategies is larger or equal than* $\bar{l}_t$*, then there exists an equilibrium for* Γ *where* $x_t = 0$*.*

Proof. Consider the game Γ, and assume there exist a player $t \in \mathcal{S}$ such that $BR_t(l_{-t}) = 0,\ l_{-t} \geq \bar{l}_t$. Now, consider an equilibrium profile $\hat{\mathbf{x}}$ of the game $\hat{\Gamma}$ that causes on overall load equal to $\hat{l} \geq \bar{l}_t$. Then, a strategy profile $\mathbf{x} = (x_t, \mathbf{x}_{-t})$, where $x_t = 0$ and $\mathbf{x}_{-t} = \hat{\mathbf{x}}$, is a Nash equilibrium for Γ, being $x_t \in BR_t(\mathbf{x}_{-t})$ and causing no effect on the strategies of the other players.

□

As we have already mentioned, there exist some particular game configurations for which the Nash equilibrium is not unique. In these cases, the equilibria can be either all NE-EE, or they can be of mixed types, but we have never encountered any instance for which there are multiple equilibria of type (i) only.

When the game admits more than one equilibrium and at least one of them is not a NE-EE, one might want to bootstrap the Best Response Dynamics to reach this type of equilibrium. In this paper, we do not investigate initialization strategies of the BRD algorithm to appropriately steer the BR path. However, for completeness, we can state that there are at least two degrees of freedom that determine the selection of the equilibria. One is the starting point of the BRD, i.e. the initial value of the strategy of each player, defined in Algorithm 1 as the strategy profile $\mathbf{x}^{(0)}$. The second is the order in which tenants play their BRs. An effect of how these two degrees of freedom can influence the result of the algorithm is shown in Sec. V.

V. Numerical Results

A. Simulation setup

The simulation scenario adopted for our numerical analysis considers three different types of tenants competing for the resources of a single cell, where the total bandwidth of the cell is equal to 20 MHz. Since we focus on the competition among tenants, we assume that within a slice there is no service differentiation among users. This allows us to simplify the acceptance probability function in Eq. (5) and the reward function in Eq. (6) by setting the same $r_s^0 = r_k^0$, $\mu_s = \mu_k$ and $\chi_s = \chi_k$, $\forall\ k \in \mathcal{K}_s$. Furthermore, the parameters used for the purchasing power sigmoid function in Eq. (1) are set to $\gamma = 10$ and $l_0 = 0.75$.

Type	$\lvert\mathcal{K}_s\rvert$	r_s^0	μ_s	χ_s	η_s	q
LP	5	5 Mbps	2	$2 * 10^6$	5	0.001
MP	5	5 Mbps	4	$2 * 10^6$	5	0.001
HP	5	5 Mbps	6	$2 * 10^6$	5	0.001

TABLE I: Tenants simulation parameters

We assume not to have any instantaneous information about the channel gain and spectral efficiency of each user, but to collect, from the SMS scheduler, only the average spectral efficiency of the slice. This allows us to define a slice-wide spectral efficiency η_s, by means of which we further simplify the relation $r(\hat{x}_k)$ as follows:

$$r(\hat{x}_k) = \hat{x}_k \cdot \eta_s \ \ \forall\ k \in \mathcal{K}_s, \forall\ s \in \mathcal{S} \tag{10}$$

In Table I, we report the parameters that characterize the three types of tenants considered, unless differently stated. As already mentioned in Sec. III, these parameters describe the tenant's behavior in the game and, consequently, determines the allocation of resources at the equilibrium. In more details, we consider a low performance (LP) type of tenant (and corresponding slice) with a relatively lower (compared to other tenants) μ_s. A low value of μ_s allows the tenant to be willing to accept lower resource allocations, since the acceptance probability is susceptible to performance degradation. On the contrary, we consider a high performance (HP) tenant with a comparatively higher μ_s that guarantees only marginal quality degradation. Finally, the middle performance (MP) type has a behavior that is intended to be somewhere in the middle between the previous two. Note that for all types, the other slice parameters are fixed to the same value.

The numerical results reported hereafter have been obtained using a MATLAB implementation of the Best Response Dynamics algorithm.

B. Slice performance at Nash Equilibrium

At first, we consider a scenario where only two tenants compete for the radio resources in the cell. Namely, one tenant (that we call Tenant 1) of type LP, and on tenant (Tenant 2) that can be either of type LP or type HP. We start from a symmetric scenario in which both tenants are of type LP. Graphically, one can see the Nash Equilibrium as the intersection point of the two players' best responses. This is presented in Fig. 3 as a black cross located at the intersection between $BR_1(x_2)$ (solid red curve) and BR_2^{LP} (solid blue curve). From Fig. 3, one can easily notice that there exists a unique equilibrium for this specific game configuration. The corresponding allocation

[4]Note that this definition is well posed, since we have already shown that the best responses are not increasing, thus once the best response goes to 0, it will never become strictly positive again.

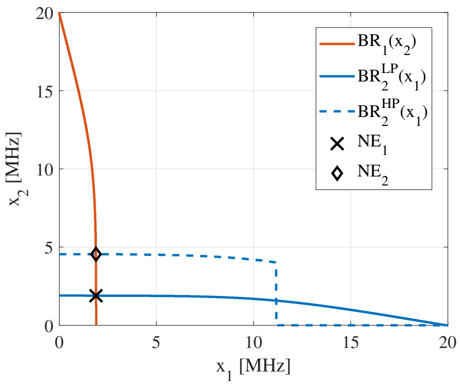

Fig. 3: Best response intersections when Tenant 1 is LP type and Tenant 2 is LP (solid blue) and HP (dashed blue)

Game	LP vs LP		LP vs HP	
Tenant	1	2	1	2
Type	LP	LP	LP	HP
x_s	9.4%	9.4%	9.4%	22.7%
$r(\hat{x}_k)$	1.89 Mbps	1.89 Mbps	1.89 Mbps	4.55 Mbps
$A_k(r(\hat{x}_k))$	0.93	0.93	0.93	0.98

TABLE II: Achieved resource allocation x_s, user rate $r(\hat{x}_k)$, and acceptance probability $A_k(r(\hat{x}_k))$ per tenant for the cases LP vs LP and LP vs HP

at the equilibrium is reported in Table II, where one can verify that the outcome of the game is also symmetric. Furthermore, we can observe that the obtained bit-rate x_s is lower than the maximal performance r_s^0. This is due to the low value of μ for the LP tenant. However, we can also comment that generally this is caused by the nature of the game model itself: the best strategy for a tenant is not to blindly give users the requested maximal performance, but rather buying resources as long as it is profitable.

Now, we analyze the scenario in which Tenant 2 belongs to type HP. As before, we show the Nash Equilibrium as a black diamond in Fig. 3, being it the intersection between the Tenant 1 best response (solid red) and the Tenant 2 type HP best response (dashed blue). The corresponding allocation at the equilibrium is also reported in Table II. In this case, the rate obtained by the users of Tenant 2 $r(\hat{x}_k)$ is much higher than the opponent's, even if the maximal performance r_s^0 is the same. This effect is due to the different values of μ_s. It is worth noticing that there is no difference in the allocation for type LP between the two scenarios. This means that the increase in load due to the HP slice does not affect the equilibrium strategy for the LP slice. From Fig. 3, we can also appreciate the different behaviors of the two types of tenants as the opponent's strategy increases by looking at the best response of Tenant 2 in the two different configurations. Indeed, $BR_2^{LP}(x_1)$ indicate a certain slice flexibility with respect to load variations, since it is able to reduce the required allocation in response to a higher adversary strategy. On the other hand, $BR_2^{HP}(x_1)$ shows a stronger rigidity to the external load, as the tenant best strategy will slightly decrease as the external load increases, up to a point in which the best response suddenly goes to 0. From this point on (in this case around $x_1 \simeq 11$ MHz), the adversary strategy causes a load in the cell so high that the cost of buying any amount of resource is higher than the corresponding generated revenue. In this case, we can say that the tenant loses any interests in taking part in the game.

Tenant	1	2	3
Type	LP	MP	HP
x_s	8.9%	20.6%	22.4%
$r(\hat{x}_k)$	1.79 Mbps	4.13 Mbps	4.49 Mbps
$A_k(r(\hat{x}_k))$	0.92	0.96	0.97

TABLE III: Achieved resource allocation x_s, user rate $r(\hat{x}_k)$, and acceptance probability $A_k(r(\hat{x}_k))$ per tenant for the case LP vs MP vs HP

In Table III, we report the numerical results at the equilibrium for a scenario in which all the three different types of tenant compete against each other in the cell. Differently from previous cases, adding one slice to the system determines an increase in the load in the cell. Consequently, it implies an increase in the price of the resources. This results in a smaller (but still noticeable) difference in the allocation for both tenant types LP and HP with respect to the previous scenario, in which the tenant of type MP was not part of the game. Finally, as expected, the percentage of resources allocated to each tenant increases with μ_s.

C. Sensitivity analysis on the tenant parameters

While a thorough sensitivity analysis of all the parameters involved in our game model is out of the scope of this paper, in this section we present some preliminary results to give an intuitive understanding of the behavior of some of them.

Consider two tenants of type MP (namely T1 and T2) competing for the resources in the cell. We study the resulting allocation and average user acceptance probability of both tenants at the equilibrium when one slice (i.e. the one of T1) either experiences different intra-cell condition (i.e. in number of users $|\mathcal{K}_1|$ or average spectral efficiency η_1) or changes the settings of the slice (i.e. different μ_1 or r_1^0). Results are reported in Fig. 4. At first, one can easily noticed that, while the behavior of T1 changes while varying its own parameters, allocation strategy and acceptance probability for T2 (blue curves) are generally slightly affected.

Let us then focus on how the equilibrium changes for T1. Fig. 4(a) shows that the obtained resources linearly increase with the number of users, although the average amount per user decreases. This also causes a decrease in the acceptance probability (dashed squared line). As a matter of fact, while the revenues grow linearly with the number of users, the costs have a non-linear behavior (cf. Eq. (6) and Eq. (7)), which leads to decrease the amount of requested resources. Similar behavior can also be observed in Fig. 4(b), when varying the maximal performance r_1^0 of T1.

In Fig. 4(c), we show how the system responds to variations of the spectral efficiency η_1. First, one can note that for $\eta_1 = 2$, T1 plays a zero strategy at the equilibrium, effectively excluding itself from the game. In this case, the spectral efficiency is

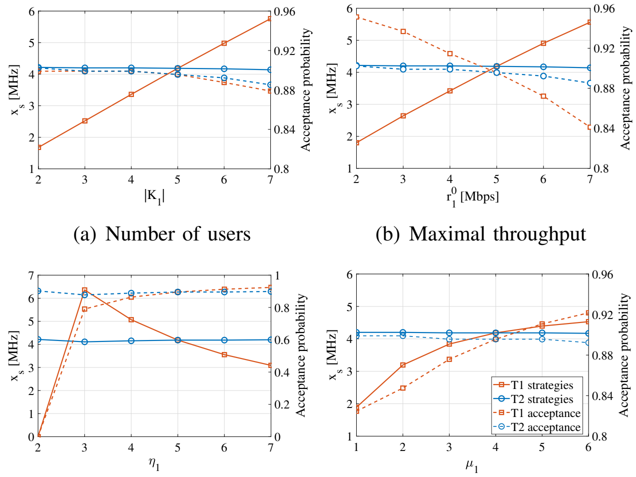

(a) Number of users (b) Maximal throughput

(c) Spectral efficiency (d) Quality sensitivity

Fig. 4: Allocation strategy and acceptance probability while varying slice parameters of T1.

Tenant	$\|\mathcal{K}_s\|$	r_s^0	μ_s	χ_s	η_s	q
T1	5	3 Mbps	6	$1.4 * 10^6$	3	0.001
T2	5	11.8 Mbps	2	$4 * 10^6$	2.5	0.001

TABLE IV: NE-EE scenario: tenants simulation parameters

so low that any strategy will lead to a negative profit, causing T1 to lose any interests in taking part in the game. For values of $\eta_1 > 2$, the tenant has actually interest in taking part to the game and therefore buys resources to provide the required service to its slice. Moreover, for increasing values of η_1, the amount of resources needed to satisfy the users' requests gets incrementally lower (due to the better channel conditions) and thus also the tenant's strategies decrease accordingly. Finally, we show the effect of variations of μ_1 in Fig. 4(d). In this case, as the quality sensitivity parameter increases, the strategy of T1 also increases. This is expected, since as discussed above, increasing μ_1 corresponds to incentivizing T1 to provide a more reliable service to its own users.

D. Nash equilibria with exclusion effect and equilibrium selection

Certain combinations of cell and tenant parameters lead to game configurations for which multiple equilibria exist. Moreover, we conjecture that in case of multiple equilibria, at least one of them involves one or more players that do not to take part to the game, leading to a NE-EE. Consider the two-tenants game, whose parameters are reported in Table IV. In this particular scenario, the game admits two different Nash equilibria in pure strategies, as shown in Fig. 5. The numerical values of the allocated resources and acceptance probability at the equilibria are reported in Table V. We notice that Tenant 1 receives a null allocation in one of the equilibria, causing no load in the system and effectively leaving Tenant 2 to play the highest strategy its best response permits. In contrast, NE_2, is a Nash equilibrium where both players play strategies that are far above zero.

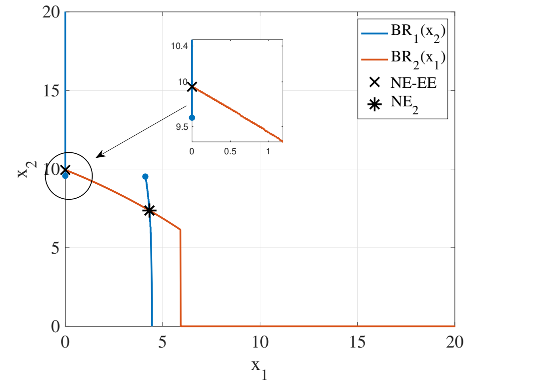

Fig. 5: HP type vs LP type tenants: best response intersection

Equilibrium	T1 allocation	T2 allocation	T1 acceptance	T2 acceptance
NE-EE	0%	49.7%	0	0.71
NE_2	21.6%	36.8%	0.94	0.49

TABLE V: Achieved resource allocation and acceptance probability per tenant for the NE-EE scenario

As mentioned at the end of Sec. IV, the Best Response Dynamics algorithm may converge to different Nash equilibria (when multiple) depending on the initial starting conditions of the algorithm and the order in which tenants play. In this case, NE-EE of Fig. 5, is reached by letting T2 start playing its strategy in empty cell condition. This cause an overall load that penalizes T1, whose best response will never move from zero (cf. Proposition 2). However, there exist also initial conditions that allow the BRD algorithm to converge to different equilibria, i.e. NE_2, in which both tenants receive a positive resource allocations.

VI. Conclusion

A novel slicing management framework is presented to handle short term renegotiations of radio resources among slices. Our proposed market mechanism is defined through a game theoretical model that dynamically regulates negotiations among diverse slice tenants. The existence of at least one Nash equilibrium in pure strategies in such game is proven and the game's properties are analyzed.

We evaluated our framework in multiple scenarios, always reaching a Nash equilibrium at the end of the Best Response Dynamics algorithm implemented. Still, for few critical instances, multiple equilibria may be reached depending on the algorithm initialization, making the need of a proper procedure to bootstrap the Best Response Dynamics. This will be object of future work, together with an extensive study with a more detailed system simulator and performance comparison with alternative approaches based on a static allocation paradigm.

References

[1] NGMN, "5G White Paper," 2015. [Online]. Available: http://www.ngmn.de/5gwhite-paper.html

[2] P. Rost, A. Banchs, I. Berberana, M. Breitbach, M. Doll, H. Droste, C. Mannweiler, M. A. Puente, K. Samdanis, and B. Sayadi, "Mobile network architecture evolution toward 5G," *IEEE Communications Magazine*, vol. 54, no. 5, pp. 84–91, 2016.

[3] C. Liang, F. R. Yu, and X. Zhang, "Information-centric network function virtualization over 5G mobile wireless networks," *IEEE network*, vol. 29, no. 3, pp. 68–74, 2015.

[4] X. Zhou, R. Li, T. Chen, and H. Zhang, "Network slicing as a service: enabling enterprises' own software-defined cellular networks," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 146–153, 2016.

[5] C. Liang and F. R. Yu, "Wireless network virtualization: A survey, some research issues and challenges," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 358–380, 2015.

[6] M. Richart, J. Baliosian, J. Serrat, and J.-L. Gorricho, "Resource slicing in virtual wireless networks: A survey," *IEEE Transactions on Network and Service Management*, vol. 13, no. 3, pp. 462–476, 2016.

[7] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. W. Wong, "Advanced demand side management for the future smart grid using mechanism design," *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1170–1180, 2012.

[8] P. Caballero, A. Banchs, G. de Veciana, and X. Costa-Pérez, "Network slicing games: Enabling customization in multi-tenant networks," in *INFOCOM 2017*. IEEE.

[9] I. Malanchini, S. Valentin, and O. Aydin, "Generalized resource sharing for multiple operators in cellular wireless networks," in *Wireless Communications and Mobile Computing Conference (IWCMC), 2014 International*. IEEE, pp. 803–808.

[10] K. Samdanis, X. Costa-Perez, and V. Sciancalepore, "From network sharing to multi-tenancy: The 5G network slice broker," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 32–39, 2016.

[11] O. Sallent, J. Perez-Romero, R. Ferrus, and R. Agusti, "On radio access network slicing from a radio resource management perspective," *IEEE Wireless Communications*, vol. 24, no. 5, pp. 166–174, 2017.

[12] M. Jiang, M. Condoluci, and T. Mahmoodi, "Network slicing management & prioritization in 5G mobile systems," in *European wireless*, 2016.

[13] A. Lieto, I. Malanchini, and A. Capone, "Enabling dynamic resource sharing for slice customization in 5G networks," in *GLOBECOM 2018*. IEEE.

[14] R. Stair and G. Reynolds, *Principles of information systems*. Cengage Learning, 2013.

[15] M. Friedman, *The optimum quantity of money*. Transaction Publishers, 2005.

[16] S. Mandelli, M. Andrews, S. Borst, and S. Klein, "Satisfying network slicing constraints via 5G MAC scheduling," *INFOCOM 2019*.

[17] B. Khodapanah, A. Awada, I. Viering, D. Oehmann, M. Simsek, and G. P. Fettweis, "Fulfillment of service level agreements via slice-aware radio resource management in 5G networks," in *IEEE 87th Vehicular Technology Conference (VTC Spring)*. IEEE, 2018.

[18] F. P. Kelly, A. K. Maulloo, and D. K. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research society*, vol. 49, no. 3, pp. 237–252, 1998.

[19] L. Badia, M. Lindstrom, J. Zander, and M. Zorzi, "Demand and pricing effects on the radio resource allocation of multimedia communication systems," in *GLOBECOM 2003*, vol. 7. IEEE, pp. 4116–4121.

[20] J. Nash, "Non-cooperative games," *Annals of mathematics*, pp. 286–295, 1951.

[21] N. S. Kukushkin, "A fixed-point theorem for decreasing mappings," *Economics Letters*, vol. 46, no. 1, pp. 23–26, 1994.

Appendix

Before showing the proof of Theorem 1, we need the following two lemmas.

Lemma 1. *The best response correspondences of a game* $\langle \mathcal{S}, (X_s)_{s\in\mathcal{S}}, (u_s)_{s\in\mathcal{S}} \rangle$ *verifying properties (i) and (ii) of Theorem 1 are upper hemi-continuous.*

Proof. Given the compactness of the strategy spaces (*i*) and the fact that the utility functions are continuous and bounded (*ii*), then the range of the best response correspondence $BR_s(\mathbf{x}_{-s})$ is compact. Under this condition, the properties of upper hemi-continuity and closed-graphness are equivalent. Therefore, we need to show that $BR_s(\mathbf{x}_{-s})$ has a closed graph. Consider the sequence $t_n \in X_s$ and the tuple of sequences $v_n \in \mathbf{X}_{-s}$ whose limits are t and v, respectively. Suppose these sequences are such that $t_n \in BR_s(v_n)$ for each n. Then it must be $u_s(t_n, v_n) \geq u_s(z, v_n)\ \forall z \in X_s$ and for each n. Exploiting the continuity of u_s we consider the limits and write $u_s(t, v) \geq u_s(z, v)\ \forall z \in X_s$. The latter is equivalent to $t \in BR_s(v)$ and proves the closed-graphness of BR_s. □

Lemma 2. *The best response correspondences of a game* $\langle \mathcal{S}, (X_s)_{s\in\mathcal{S}}, (u_s)_{s\in\mathcal{S}} \rangle$ *verifying properties (ii), (iii) and (iv) of Theorem 1 allow decreasing single-valued selections.*

Proof. Property (*iv*) allows us to rewrite the generic best response $BR_s(\mathbf{x}_{-s})$ as $BR_s(\sum_{t\in\mathcal{S}\setminus\{s\}} x_t)$. We consider now the following single-valued selection: $\phi(\sum_{t\in\mathcal{S}\setminus\{s\}} x_t) = \sup\{BR_s(\sum_{t\in\mathcal{S}\setminus\{s\}} x_t)\}$. We say that ϕ is decreasing if any $x, y \in \mathbb{R}$ such that $x \geq y$ imply $\phi(x) \leq \phi(y)$. Then

$$\frac{\partial}{\partial x_s \partial x_t} [u_s(x_s, \mathbf{x}_{-s})] \leq 0\ \forall s, t \in \mathcal{S}\ , s \neq t \tag{11}$$

is sufficient to claim that ϕ is decreasing with respect to any increment of $\sum_{t\in\mathcal{S}\setminus\{s\}} x_t$. Namely, Eq. (11) implies that when any of the opponents' strategies increase, an increase of the own strategy necessarily leads to a smaller utility. Thus, the best response to an increase of opponents' strategies can never be a larger strategy, proving that the single-valued selection can only decrease. □

Proof of Theorem 1

Consider a game $\langle \mathcal{S}, (X_s)_{s\in\mathcal{S}}, (u_s)_{s\in\mathcal{S}} \rangle$ verifying the hypotheses of Theorem 1. Then by property (*iv*) we can rewrite the generic best response $BR_s(\mathbf{x}_{-s})$ as $BR_s(\sum_{t\in\mathcal{S}\setminus\{s\}} x_t)$. Moreover, Lemmas 1 and 2 imply that this correspondence is upper hemi-continuous and allows a single-valued selection. Under the previous conditions and because of property (*i*), one can apply the fixed point theorem proposed in [21], which ensures the existence of a point $\mathbf{x}^0 \in X = \prod_{s\in\mathcal{S}} X_s$ such that $x_s^0 \in BR_s(\sum_{t\in\mathcal{S}\setminus\{s\}} x_t^0)$ for each $s \in \mathcal{S}$. □