Dynamic Pricing for Tenants in an Automated Slicing Marketplace

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

***Abstract*—The paradigm shift from a *one-size-fits-all* architecture to a *service-oriented* network infrastructure promised by network slicing will demand novel technical solutions, as well as new business models. In particular, the role separation between infrastructure providers, i.e. the ones owning the network, and slice tenants, i.e. the ones providing specialized services tailored to their vertical segments, may encourage the definition of a shared platform (or marketplace) where the former can monetize their network infrastructure by leasing network resources at a market price, and the latter can rent on-demand the network resources needed to offer their services at the desired quality. This also enables the flexibility for the slice tenants to optimize the management of their slices by adapting their resource demand to fluctuations of their traffic or variations of the price in the market. In this paper, we extend the market mechanism scheme developed in previous works by including intra-slice radio admission control policies in the utility definition of the tenants in the slicing market game. Moreover, we characterize the mathematical properties of the game with respect to slice configuration, i.e. how diverse strategical behavior of the tenants affects the market operation, in terms of slice resource allocation and performance. Our analysis offers insights to the slice tenants on how they could reconfigure their techno-economic performance indicators in response to the dynamics of network and of the market, namely how to adapt their long-term (and/or real-time) strategies to the fluctuations of the traffic to enhance network performance and increase profits.**

I. Introduction

The key differentiator of the upcoming 5G and beyond systems from previous wireless technologies is the integration of vertical industries in the telcos ecosystem, which paves the way for new market opportunities and innovative business models [1]. In particular, *network slicing* enables an Infrastructure Provider (InP) to support diverse services over a common network infrastructure by offering customized end-to-end (E2E) logical networks, i.e., slices, by sharing the same pool of network resources and functionalities. Once in operation, the slices may need to be dynamically scaled up/down to match any variation of service requirements or adapt their configuration to dynamic changes and/or undesired trend in their monitored Key Performance Indicators (KPIs). Therefore, openness of the network to third parties and flexibility in the management of the slices are key features to encourage vertical players to use existing network infrastructures rather than deploying their own private infrastructure. In [2], the authors define a framework based on a vertical-oriented network slicing design where slice tenants can entirely customize and upgrade their network slices, with zero-touch service and network management. Along the same lines, [3] addresses the key challenges of the life-cycle management of network slices, discussing the tradeoff between the degree of control and customization of network slices (by the vertical tenant) and the operational complexity in the network management (by the InP), specifically for the management of the shared resources in Radio Access Network (RAN) [4], e.g., the spectrum. Due to the random nature of traffic, physical radio resource reservation might not suffice to provide neither resource efficiency nor expected guarantees in terms of Quality of Service (QoS) and Quality of Experience (QoE) [5], [6]. To maintain satisfactory user experience and high profits for slice tenants in a dynamic environment, a slice may need to be reconfigured according to the varying traffic demand and resource availability. However, existing works focus either on static resource allocation, leaving to admission control the decision of accepting a network slice to guarantee the Service Level Agreement (SLA) in varying traffic conditions [7], [8], or defines centralized optimization routines to dynamically allocate the resources to increase resource efficiency or maximize social welfare [9], [10], [11]. However, since these slices will be used by profit business entities, e.g. verticals, the allocation of physical resources to the network slices must consider their private revenue and business models, in addition to the provisioning of desired QoS.

In [12], [13], we introduce a Slicing Management Framework (SMF) that can be applied for the dynamic orchestration of network resources owned by an InP in a multi-tenant shared marketplace. In those works, the market game with its mathematical properties is described, together with an algorithmic implementation to guarantee the convergence to a Nash Equilibrium (NE). In this paper we aim at completing the picture by testing the applicability of the SMF in a realistic scenario and providing additional insights on the impact of each tenant's choice in the evolution of the market due to the configuration of their slices. We show that the system automatically scales and adapts the resource allocation to the dynamics of a real network, being able to react to both short-term effects, like mobility and handovers, and long-term changes in the overall offered traffic at different times of a day, whose variations are taken from traffic traces of an LTE network collected in a urban area. Given the fluctuations of slices' traffic, we extend the analysis of the slicing market game by introducing intra-slice Radio Admission Control (RAC) policies in the strategy evaluation of the tenants, as a mean to control slice congestion that can affect both network performance and costs for the resources, similarly to what is proposed in [14]. Moreover, differently from [12], [13], we also offer insights to the slice tenants for a better customization of their network slice techno-economical configuration. This would allow them to better meet their business intents and desired slice's SLAs.

The paper is structured as follows. In Section II, we describe the system model of the slicing marketplace and introduce the payoff functions of the game when including RAC. In Section III, we derive some relevant theoretical properties of the game and analyze the impact of the slice configurations on their strategic behavior in the market. Numerical evaluation and simulation setup are discussed in Section IV and, finally, Section V concludes the paper with final remarks and future works.

II. Slicing Market Game

We consider a network slicing setup consisting of a set of network slices, S, each one owned by a single tenant or vertical enterprise[1], that engage in a shared marketplace for purchasing radio network resources, i.e., spectrum, to be assigned to their slices. We further assume that the marketplace is controlled and managed by a single InP, which applies dynamic pricing policies to monetize the utilization of the network resources, e.g., by adapting the price for a unit of resources according to the current demand. Let $x_s \in [0,1]$ be the resource allocation request of a slice s, normalized according to the total available system bandwidth, and let $l = \sum_{s \in S} x_s$ be the current load in the network, namely the total resource demand from all the slices. Then, we define the pricing function as

$$P(l) = 1 + e^{\gamma(l-l_0)}, \tag{1}$$

where parameters γ and l_0 allow one to tune the shape of the pricing function. In particular, when the total load approaches the value l_0, the price increases exponentially to discourage the tenants to purchase extra resources. Accordingly, we compute the costs of the tenants for purchasing x_s resources in the market as

$$C_s(x_s, l) = x_s \cdot P(l), \tag{2}$$

where the term x_s depends solely on the tenant decision, while $P(l)$ accounts for the aggregate decisions of all the tenants (in terms of load). As done in [12], [13], we model the interactions of the tenants in the marketplace as a *slicing market game*. Then, let $\mathcal{G} = \langle S, (X_s)_{s\in S}, (u_s)_{s \in S}\rangle$ be the strategic form of the slicing market game with S as the set of players (the slice tenants), $X_s \in [0,1]$ the strategy space (normalized amount of resources to buy) and u_s the payoff functions of the players. In previous works [12], [13], we analyze the property of the game, by showing that the game $\mathcal{G}$ belongs to the subclass of *aggregative games* [15], given the aggregative property of the cost function $C_s(\cdot)$. We further prove the existence of NE for a class of utility (payoff) functions and provide an algorithmic implementation that guarantees the convergence to at least one of them, also discussing the quality of the achieved equilibria. We assume the payoff function of the players to consist of two terms:

- a private term, $R_s(x_s)$, describing the business model of the tenants that computes the expected returns given the achieved slice performance, which depends solely on the individual action of the slice tenant, and
- a costs term, $C_s(x_s, l)$, that varies according to the law of demand and supply of the market.

Therefore, the payoff function of the tenant can be computed as

$$u_s(x_s, \mathbf{x}_{-s}) = R_s(x_s) - C_s(x_s, \mathbf{x}_{-s}), \tag{3}$$

where we replace the load l with $\mathbf{x}_{-s}$ to highlight the dependence on the other players' strategies. In particular, the revenue function of the tenants is expressed as

$$R_s(x_s) = \sum_{k \in K_s} \chi_k \cdot A_k(x_k), \tag{4}$$

where K_s is the set of users within a slice s, χ_k is a per-user economic value that the tenant assigns to each user k, and $A_k(x_k)$ is an acceptance probability function that quantifies the level of satisfaction of a user for the experienced quality of service. Notice that the value x_k defines the per-user resource allocation, which depends on how the radio scheduler redistributes the x_s resources of a slice among its users[2]. We model the experienced quality of service (in terms of achieved throughput) of a user through the acceptance probability function defined in [18]:

$$A_k(x_k; \mu, r_0) = 1 - q^{\left(\frac{r(x_k)}{r_0}\right)^{\mu}}, \tag{5}$$

where $q \in (0,1)$, μ defines the steepness of the curve, i.e., the sensitivity of the user to performance degradation, r_0 represents the maximal throughput and $r(x_k) = x_k \cdot \eta_k$ denotes the achievable throughput given the allocated resources, x_k, and the experienced spectral efficiency, η_k. As in [12], we assume that the marginal utility shows diminishing return when approaching some reference value (i.e., maximal throughput r_0), namely that the increase in quality of experience of a user vanishes when approaching the requested QoS.

A. Admission control policies

In this paper, we extend the analyses of the previous works by including customized RAC policies for each slice. Indeed, the admission of new users in the system for a slice

[1] This assumption holds for simplicity. In general, a single tenant may control multiple network slices and still have different business models for each of them.

[2] In the rest of the work, we assume that the resources are evenly split among the users, i.e. $x_k \simeq \frac{x_s}{|K_s|}$ - with $|\cdot|$ denoting the cardinality of a set - as achieved by the state of the art proportional fair schedulers [16], [17].

may require a renegotiation of the current purchased network resources to avoid degradation of performance of already admitted users. However, depending on the current load of the network, the purchase of extra resources might not necessarily lead to an increase in performance (or profits, due to higher prices). Therefore, the admission of new users can impact the decisions of the tenants in the market. Hence, we replace the payoff functions presented in Eq. (3) with the following formulation

$$u_s(x_s, n_s, \mathbf{x}_{-s}) = R_s(x_s, n_s) - C_s(x_s, \mathbf{x}_{-s}), \qquad (6)$$

where the RAC decision of a tenant translates into the number of users admitted in the system, n_s. Notice that at time of resource negotiation, we assume the tenants to have complete knowledge of their current offered traffic. This assumption holds for tractability reason. However, this approach can be applied and further optimized predicting (e.g. by means of time series) their daily traffic, taking into account that, as a consequence, their RAC policies might be affected by potential prediction errors [8]. Given the complexity of this prediction use case (and the potential impact on its decisions as well as on the decisions of other players in the market game), we focus on a simplified scenario where the amount of traffic is known by the tenants and no prediction errors must be taken into account.

The RAC policy considered hereafter assumes that users arrive in the system at random time and, therefore, the tenants cannot select the users to be accepted according to favorable channel conditions, neither discriminate users based on their location (e.g., reject an user at cell edge). Indeed, the tenant evaluates one user per time and accept a new user only if it generates an increase in its profit function, $u_s(\cdot)$. After a new user is rejected, all following users are also rejected. The optimization variables of the tenants become, therefore, the tuple (x_s, n_s), where the RAC policy implements the following:

$$\forall x_s, \exists n_s^{\text{opt}} : R_s(x_s, n_s^{\text{opt}}) \geq R_s(x_s, n_s), \forall n_s \in [0, |K_s|]. \quad (7)$$

In a such way, the slice tenant will admit a new user only if it generates an increase in revenues, rejecting them otherwise. It is important to remark that

Proposition 1. *The market game $\mathcal{G}$ with the* RAC *policy of Eq.* (7) *always admits at least one NE. Furthermore, the Best Response Dynamics (*BRD*) algorithm always converges to a NE in a finite number of steps.*

Proof. The proposition can be easily verified by construction, by recalling the existence and convergence condition in [12, Theorem 1 and 2]. Indeed, given the formulation in Eq. (7), it can be easily verified that the revenue function $R_s(x_s, n_s^{\text{opt}})$ holds the same property of $R_s(x_s)$ when RAC is not considered, namely being an increasing and differentiable function in x_s, which is a necessary condition for the existence and the convergence to a NE. □

In what follows, we remove the dependency on n_s of $u_s(\cdot)$ for ease of notation.

III. Theoretical properties of the market game

In this section, we characterize the properties of the market game, based on the customization of the tenants' private revenue function, $R_s(\cdot)$. The main idea behind our investigation comes from the property that in aggregative games the players can analyze the evolution of the game by simply focusing on their individual action and on the value of the aggregate of the other players. In our specific market structure, we can further simplify those assumptions since players only need to know the updated shape of the pricing function, with no need to expose the actual total resource occupation. This is a very interesting property that allows each player to analyze the game independently on the number of players and on the shape of the opponents' utility functions, focusing only on the evolution of the prices in the market. Moreover, we can characterize the effects of changing a parameter of the utility function of any player. For this purpose, we introduce the concept of *positive shock* defined in [19].

Definition 1. *Consider the payoff function of the game $\mathcal{G}$, $u_s(x_s, \mathbf{x}_{-s}, p_s)$ with parameter p_s. An increase in $p_s \in \mathbb{R}$ is a positive shock for player s if the payoff function $u_s(\cdot)$ exhibits increasing differences in x_s and p_s and $u_t(\mathbf{x}, p_s) = u_t(\mathbf{x})$ $\forall t \neq s$.*

From [19], we can also state the following

Proposition 2. *In the game $\mathcal{G}$, a positive shock to any player $s \in S$ determines an increase in his allocation, x_s, and decrease in the aggregate allocations of the remaining players, x_{-s}, at the NE. Moreover, it causes an increase in his own payoff and a decrease in at least one other player's payoff at the NE.*

Specifically, we can prove that, for the slicing market game, holds the following

Proposition 3. *For any slice s, the admission of new users or an increase in the per-user economic value, χ_k, is a positive shock for the player in the market game $\mathcal{G}$.*

Proof. To prove the proposition, one must verify that the utility function $u_s(\cdot)$ exhibits increasing differences in $\left(x_s, n_s^{\text{opt}}\right)$ and (x_s, χ_s), respectively, and that any changing in the value of parameters n_s^{opt} and χ_s for slice s does not affect the utility function of other players $t \in S, t \neq s$. The latter is true by definition, given that the parameters χ_s and n_s^{opt} affect only the revenues, $R_s(x_s)$, which are privately defined by each tenant. The increasing difference property can be verified by checking that the cross-partial derivatives of $u_s(\cdot)$ are nonnegative. This is immediately verified by noticing that $C_s(x_s, \mathbf{x}_{-s})$ does not depend neither on n_s^{opt} nor on χ_s. Therefore, we just need to verify that the cross partials obtained from $R_s(x_s)$ are nonnegative. The analytical derivation is left to the reader. However, intuitively, one can expect that the revenue function satisfies the increasing differences property, being the revenue function $R_s(x_s)$ an increasing function in the purchased resources, x_s, per-user economic value, χ_s, and number of admitted users, n_s^{opt}, as also remarked in Proposition 1. □

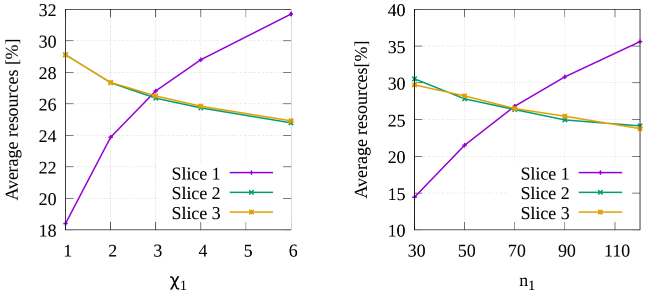

(a) Per-user economic value

(b) Number of users

Fig. 1: The parameter χ and number of users cause a positive shock for Slice 1.

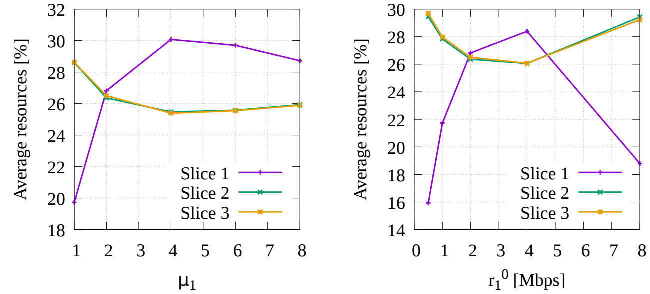

(a) Parameter μ

(b) Parameter r_0

Fig. 2: The parameters μ and r_0 do not cause positive shock for Slice 1.

An example of a positive shock for a player is plotted in Figure 1, for the two aforementioned values. We consider a scenario with two symmetric slices and one slice which spans different parameter configurations. In particular, in Fig. 1(a), we show that an increase in the per-user economic value, χ_1, allows the tenant to get an higher amount of resources, causing a decrease in his opponents' strategies (and, therefore, to the aggregate, as stated in Proposition 2). The same applies when a slice experiences an increase in his traffic demand, expressed as number of accepted users n_1, where the admission of new users induces the same effect (cf. Fig. 1(b)). Conversely, if we look at the parameters modeling the quality of experience of the users, they do not induce a positive shock in the allocation of the network slices. Indeed, as shown in Figure 2, an increase in the sensitivity parameter μ and in the maximal throughput r_0 - for example due to a change in the user behavior - may encourage the tenants to purchase more resources in the market (to improve the user performance), but it may also result in the opposite behavior due to non sustainable costs, causing, e.g., an higher user rejection rate and/or lower revenues.

IV. Numerical Evaluation

In this section, we simulate the online trading of radio resources in the slicing marketplace. We consider the same slicing setting presented in [12], with $S = \{$critical IoT (cIoT), eMBB Premium (eMBB Pr.), eMBB Basic (eMBB Bs.)$\}$, to span different service characteristics and user behaviors, i.e.,

- critical applications with low-rate requirements but high QoS guarantees (cIoT),
- non-critical applications with high-rate requirements and high user economic value (eMBB Pr.),
- non-critical applications with adaptive QoS and low user economic value (eMBB Bs.),

that we model by means of different values of the tuple (χ_s, μ_s, r_s^0). In our experiments, the users of each slice are uniformly distributed throughout the coverage area of 3 cells and their position over time varies according to their random movements. In particular, we assume that the users of the cIoT slice have fixed locations, while the users of eMBB slices move at constant speed of 3 km/h. We further consider, without

Tenant	cIoT	eMBB Pr.	eMBB Bs.
μ_s	8	4	2
r_s^0	0.5 Mbps	4 Mbps	2 Mbps
q_s	0.001	0.001	0.001
χ_s	3	8	3

TABLE I: Slice parameters setting

loss of generality, that all the users within a slice have the same QoS requirements, which are defined in terms of throughput, r_s^0, and the same economic value, i.e. $\chi_k = \chi_s, \forall k \in K_s$. In Table I we resume the parameter settings of the 3 slices.

The simulations reproduce a realistic scenario, where the incoming traffic of the slices vary over time. The simulated traffic traces are taken from real measurements collected from an LTE network in an urban area and they are artificially re-distributed among the 3 slices. In our simulations, we consider the traffic patterns of two consecutive weekdays to leverage the network utilization and the resource demand during different time of the day (e.g., showing different behavior during peak hours or during non-peak hours - for example at night). We assume that cIoT end-users have a deterministic network behavior, with constant data transmission during daily hours and only background traffic transmitted during night hours. Contrarly, the eMBB slices exhibit more randomic behavior, due to sudden activation of users requesting data. The traffic traces are taken from an original dataset that consists of time series collecting average downlink datarate observations (measured in bps) with 15-minute granularity from 3 cells of an LTE network for a duration of two days. The subdivision of the traffic traces among the 3 slices during the two days of measurements is shown in Fig. 3, for the whole 48 hours with 4 samples per hour (for a total of 192 simulated scenarios). For every traffic sample of 15 minutes, each user random movement is simulated for 3 out of the 15 minutes to reduce the simulation time, determining the variable serving cell and channel during that time span.

The experiments are performed in a downlink system level simulator which is 3GPP-calibrated [20] for a 3D Urban Macro (3D-UMa) scenario and abstracts the physical-layer effects through a link-to-system level interface. The interface

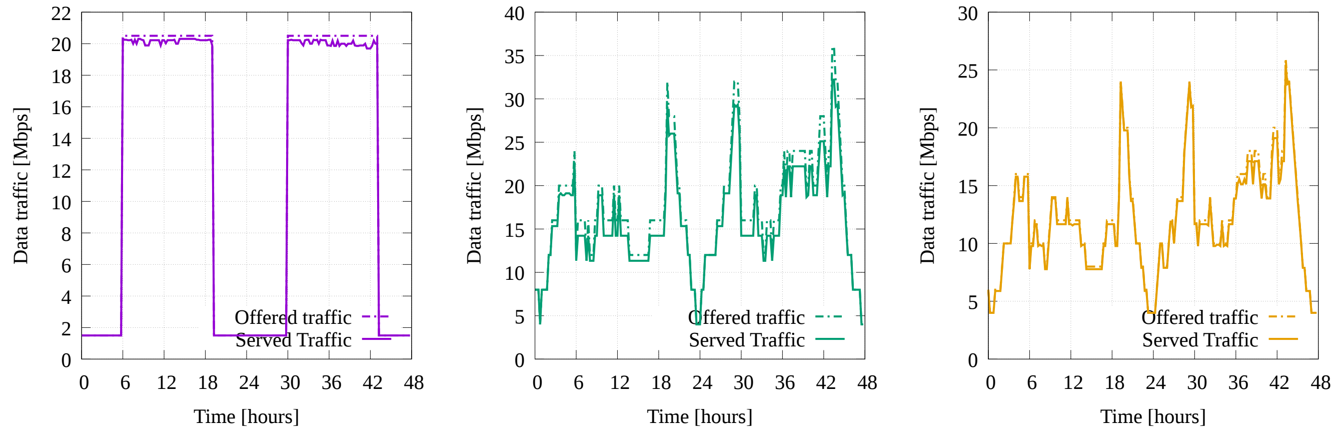

(a) cIoT traffic request

(b) eMBB Pr. traffic request

(c) eMBB Bs. traffic request

Fig. 3: Time evolution of traffic request for each slice. In dashed lines the incoming offered traffic; in solid lines the accepted served traffic after RAC.

applies an equivalent Signal-to-Interference-and-Noise Ratio (SINR), computed given the cell topology, the active user transmissions, and a vertically polarized antenna configuration. The radio environment and other relevant simulation parameters are taken from [17]. We assume that the strategy step size of each slice is equal to $\Delta x = 15$ kHz, that is the subcarrier spacing and a total bandwidth of 10 MHz. The traffic traces are then generated in the 3GPP-calibrated system level simulator, where each user data transmission is modeled as Constant Bit Rate (CBR) traffic, and the traffic variation in Fig. 3 is obtained via the activation and/or deactivation of new users in each slice. We assume that the tenants have perfect knowledge of their incoming traffic at time of resource negotiation and they can implement their own independent admission control policies to shape the traffic they are going to accept. The market game is executed at each cell on a timescale of 30 seconds to let slice tenants adjust the per-cell required resources due to user mobility (e.g., after handovers) and/or variation of channel quality of their users, in addition to traffic fluctuations.

In Fig. 3, we show the offered traffic (in dashed lines) and the served traffic (in solid lines) of each slice, after the RAC policies are implemented. One can notice that both the cIoT and the eMBB Bs. slices (Fig. 3(a) and Fig. 3(c)) can almost serve 100% of their traffic, rejecting only few users during peak hours, while the eMBB Pr. slice shows an higher amount of dropped traffic, specifically in proximity of a burst of traffic for its slice (cf. Fig. 3(b)). This behavior depends on multiple factors, hidden in the strategic behavior of the players in the marketplace, but they can be partly addressed by analyzing the results given in Fig. 4, where we show the achieved performance of the slices with and without RAC. One can observe in Fig. 4(a) the variation in costs for each slice, which follows the increase in price due to their daily traffic oscillation, while in Fig. 4(b) the average acceptance values of the slices, which we use as a measure to estimate the monitored slice KPIs. If we compare the results for the cIoT and the eMBB Bs. when RAC policies are implemented, we see that there is no significant variation in their KPIs. Indeed, given the nature of the two slices, the cIoT slice can always afford higher costs (to accept more users) to keep the desired level of performance, due to the criticality of its service, while the eMBB Bs. slice can always accept users in the system while keeping low costs by offering lower QoS. Slightly variations can be observed only during high traffic peaks, where an optimized utilization of the resources can benefit the slices in terms of lower costs and more affordable performance. Conversely, the eMBB Pr. slice can effectively improve his KPIs only by rejecting part of the traffic, in order to get lower prices and higher performance for the accepted users. However, as shown in the previous section, the slice tenants can manipulate their configuration parameters to enhance some of their KPIS, being reducing costs or improving performance. In particular, we consider the case when the eMBB Pr. slice intends to improve his performance indicators, namely increasing the percentage of accepted users in the system. For example, by raising the per-user economic value from $\chi_s = 8$ to $\chi_s = 20$, we can observe a radical behavioral change in the marketplace. Indeed, from Fig. 5, one can see the positive shock effect on resource allocation (Fig. 5(a)) and on the served traffic (Fig. 5(b)), plotted as differences with respect to the same values illustrated in previous configuration. In particular, during non-peak hours, where the traffic demand is not high, there are no evident variations in the strategic behavior of all the tenants (due to the lower price for resources); contrarily, the eMBB Pr. can accept more traffic for his slice (by increasing the resource demand) during peak hours. As expected, we can appreciate also the reduction in the resource demand of the other players, with consequences also on their accepted traffic. Although this behavior might look undesirable, it exactly describes how diverse can be the strategic behavior of the slice tenants in this market and

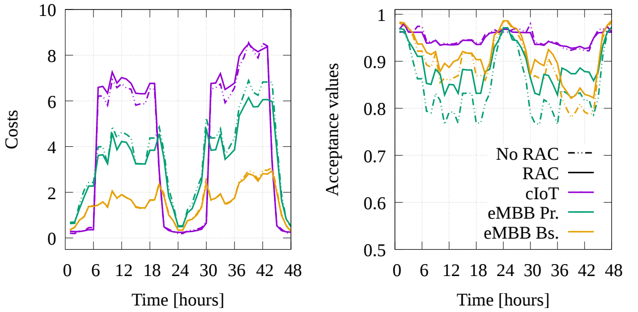

(a) Costs variation over time (b) User acceptance values over time

Fig. 4: Real-time KPIs monitoring during marketplace evolution

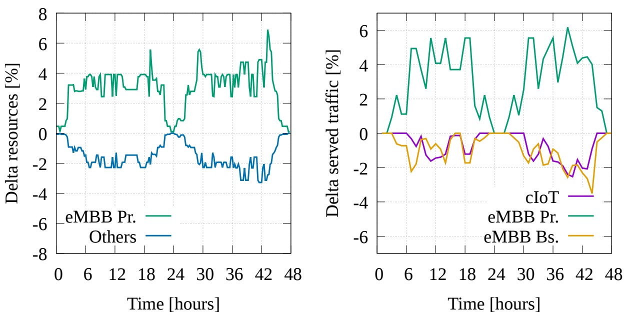

(a) Variation in resource allocation (b) Variation in served traffic

Fig. 5: Positive shock and performance enhancement for eMBB Pr. slice

how they might influence each other in the decisions to be taken. One can imagine that the optimal setup of the techno-economic KPIs (i.e., the ones reflecting the strategic behavior of the players in the market) can either define the long-term or real-time slice configuration, depending on their private business model. By our approach, we offer slice tenants a mean to optimize the management of their slices, by keeping simultaneously an eye to technical performance indicators, like achieved throughput or latency, and business intents, like cost reductions or profit maximization.

V. Conclusion

In this paper, we discuss a slicing marketplace that enables a dynamic negotiation of radio resources among network slice tenants. We model the interaction of the slice tenants in the market through game theory and describe the theoretical properties of the market game, by means of theoretically demonstrations and numerical simulations. We introduce intra-slice radio admission control policies in the utility function of the players and describe how the tenants behavior in the market is affected by the parameter configuration of the slices. We test the proposed scheme on a dynamic environment, showing that the system automatically scales and adapts the slice resource allocation to the fluctuations of the traffic demand of slice users. Moreover, we show how slice tenants can modify their techno-economic parameter configuration in response to the evolution of the market, e.g., for performance enhancement and/or costs reduction. The automated optimization of those parameters is left to future works.

References

[1] 5GPPP Architecture Working Group, "5G empowering vertical industries," 2016. [Online]. Available: https://5g-ppp.eu/wp-content/uploads/2016/02/BROCHURE%205PPP%20BAT2%20PL.pdf

[2] Q. Wang, J. Alcaraz-Calero *et al.*, "Slicenet: End-to-end cognitive network slicing and slice management framework in virtualised multi-domain, multi-tenant 5G networks," in *IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, 2018.

[3] S. E. Elayoubi, S. B. Jemaa, Z. Altman, and A. Galindo-Serrano, "5G RAN slicing for verticals: Enablers and challenges," *IEEE Communications Magazine*, vol. 57, no. 1, pp. 28–34, 2019.

[4] S. Vassilaras, L. Gkatzikis *et al.*, "The algorithmic aspects of network slicing," *IEEE Communications Magazine*, pp. 112–119, 2017.

[5] I. Malanchini, S. Valentin, and O. Aydin, "Wireless resource sharing for multiple operators: Generalization, fairness, and the value of prediction," *Computer Networks*, pp. 110–123, 2016.

[6] J. M. Peha, "Sharing spectrum through spectrum policy reform and cognitive radio," *Proceedings of the IEEE*, pp. 708–719, 2009.

[7] M. R. Raza, A. Rostami, L. Wosinska, and P. Monti, "A slice admission policy based on big data analytics for multi-tenant 5G networks," *Journal of Lightwave Technology*, vol. 37, no. 7, pp. 1690–1697, 2019.

[8] D. Bega, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Perez, "Deepcog: Cognitive network management in sliced 5G networks with deep learning," in *IEEE Conference on Computer Communications (INFOCOM)*, 2019.

[9] G. Wang, G. Feng, W. Tan, S. Qin, R. Wen, and S. Sun, "Resource allocation for network slices in 5G with network resource pricing," in *IEEE Global Communications Conference (GLOBECOM)*, 2017.

[10] L. Liang, Y. Wu, G. Feng, X. Jian, and Y. Jia, "Online auction-based resource allocation for service-oriented network slicing," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 8, pp. 8063–8074, 2019.

[11] M. Jiang, M. Condoluci, and T. Mahmoodi, "Network slicing in 5G: An auction-based model," in *IEEE International Conference on Communications (ICC)*, 2017.

[12] A. Lieto, I. Malanchini, S. Mandelli, E. Moro, and A. Capone, "Strategic network slicing management in radio access networks," *IEEE Transactions on Mobile Computing*, 2020.

[13] A. Lieto, E. Moro, I. Malanchini, S. Mandelli, and A. Capone, "Strategies for network slicing negotiation in a dynamic resource market," in *IEEE 20th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*, 2019.

[14] J. Zheng, P. Caballero, G. de Veciana, S. J. Baek, and A. Banchs, "Statistical multiplexing and traffic shaping games for network slicing," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2528–2541, 2018.

[15] M. Jensen, "Aggregative games and best-reply potentials," *Economic Theory*, pp. 45–66, 2010.

[16] 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.

[17] M. Andrews, S. Borst, S. Klein, H. Kroener, and S. Mandelli, "The effect of additive and multiplicative scheduler weight adjustments on 5G slicing dynamics," in *IEEE 2nd 5G World Forum (5GWF)*, 2019.

[18] L. Badia, M. Lindstrom, J. Zander, and M. Zorzi, "Demand and pricing effects on the radio resource allocation of multimedia communication systems," in *IEEE Global Communications Conference (GLOBECOM)*, 2003.

[19] D. Acemoglu and M. K. Jensen, "Aggregate comparative statics," *Games and Economic Behavior*, vol. 81, pp. 27–49, 2013.

[20] 3GPP TR 38.901, "Study on channel model for frequencies from 0.5 to 100 GHz," 2019.