

# Resource Efficiency vs Performance Isolation Tradeoff in Network Slicing

Panagiotis Nikolaidis and Asim Zoulikarni and John Baras

Department of Electrical & Computer Engineering and the Institute for Systems Research

University of Maryland, College Park, MD 20742, USA

Email: {nikolaid, asimz, baras}@umd.edu

**Abstract**—We consider the tradeoff between resource efficiency and performance isolation that emerges when multiplexing the resource demands of Network Slices (NSs). On the one hand, multiplexing allows the use of idle resources, which increases resource efficiency. On the other hand, the performance of each NS becomes susceptible to traffic surges in other NSs, which degrades performance isolation. The analysis of this tradeoff enables network operators to determine the effect of performance isolation on the operating cost of each NS.

To study the tradeoff, we solve an optimization problem where we find the multiplexing policy that requires the least provisioned resources to honor the Service Level Agreements (SLAs) of all NSs. The SLA of each NS  $i$  states that its resource demand should be met for  $P_i^H$  fraction of time, and for  $P_i^L \leq P_i^H$  fraction of time, it should be met regardless of the demands of other NSs.

For resource demands that follow ergodic Markov chains, we show that the well-known Max-Weight scheduler is an optimal multiplexing policy. Since the Max-Weight scheduler does not require any knowledge of the statistics of the resource demands, we also propose its use in non-markovian settings. For resource demands obtained in the LTE module of ns-3, we show that the Max-Weight scheduler reduces the provisioned bandwidth by 36.2% when no performance isolation is required. Lastly, for these non-markovian resource demands, the Max-Weight scheduler maintains its optimality since it requires as much provisioned bandwidth as the best non-causal scheduler.

**Index Terms**—network slicing, multiplexing, sharing, overbooking, isolation, Lyapunov optimization, 5G, LTE, ns-3

## I. INTRODUCTION

Future cellular networks need to serve a wide variety of applications. Examples include the traffic control of autonomous vehicles in intersections, the communications between ambulances and hospitals, and augmented reality. Such applications prioritize different types of Quality of Service (QoS) such as high reliability, low packet delays or high bitrates.

These applications may be requested by different companies that may also require strict guarantees regarding the delivery of the desired QoS to their application. Therefore, each company can be viewed as a "tenant" in the cellular network that forms a Service Level Agreement (SLA) with the Network Operator (NO) for the expected QoS delivered to their application.

The above paradigm requires a cellular network that provides varying types of QoS with strict guarantees as stated

This material is based upon work supported by the National Science Foundation award # 2127605 and by the University of Maryland Distinguished University Professor Award.

in SLAs. This requirement has led to the emergence of network slicing, where the physical infrastructure of the cellular network is "sliced" to create multiple virtual networks called Network Slices (NSs). Each NS is an end-to-end network spanning over the Radio Access Network (RAN), the Transport Network (TN), and the Core Network (CN).

To provide various types of QoS, the NO composes each NS using custom network functions that match the needs of its application. For instance, each NS may use its own Medium Access Control (MAC) scheduler in the RAN, routing policies in the TN and CN, and firewalls in the CN. The fast deployment of custom network functions has been enabled by virtualization techniques and software-defined networking.

To ensure that the SLA of each NS is honored, the NO needs to provision in advance enough resources so that the network functions of each NS can deliver the promised QoS. Based on the provisioned resources, the NO can then compute the cost of the SLA and provide a price quote to the tenant. Resource provisioning may involve traffic forecasting for each NS, however forecasting is out of the scope of this paper.

For instance, in the RAN, the NO may need to provision Physical Resource Blocks (PRBs) to the MAC scheduler of a NS to ensure high bitrates, or in the TN and CN, routing paths to ensure high reliability.

Here, we are particularly interested in these provisioning problems. We wish to obtain the amount of resources needed so that the resource demand of a NS  $i$  is met for a high fraction of time  $P_i^H$ . By resource demand, we refer to the amount of resources needed by a network function of a NS to provide the desired QoS given its current traffic state.

An obvious solution approach is to estimate the cumulative distribution function (cdf) of the resource demand of each NS and provision  $P_i^H$ -percentile resources for each of them. However, since the traffic state of a NS is time-varying, its resource demand also varies over time. Thus, provisioning  $P_i^H$ -percentile resources for each NS  $i$  is wasteful since some resources may remain unused for  $P_i^H$  fraction of time.

This observation implies that allowing the unused provisioned resources of a NS  $i$  to be used by some other NS  $j$  may reduce the overall provisioned resources. Thus, it is beneficial to consider a scheduler that multiplexes the resource demands of NSs. Hence, the fulfillment of the SLAs in network slicing involves a joint scheduling and resource provisioning problem.

However, multiplexing may have negative effects on the

performance of a NS  $i$ . For instance, suppose that a NS  $i$  relies solely on the unused  $P_i^H$ -percentile provisions of the other NSs. Then, in case the other NSs experience unexpected traffic surges, e.g., due to Distributed Denial of Service (DDoS) attacks, the performance of NS  $i$  will be severely affected. This is an undesirable outcome since tenants often require that the performance of their NS should remain unaffected by other NSs. This requirement is referred to as performance isolation.

Here, we quantify performance isolation as the fraction of time  $P_i^L$  that the resource demand of NS  $i$  is guaranteed to be met regardless of the traffic state of other NSs. The higher  $P_i^L$  is, the higher the degree of isolation is. A simple solution to provide  $P_i^L$  degree of isolation to NS  $i$  is to allocate  $P_i^L$ -percentile resources that are always available to NS  $i$  if needed.

Given the above, there is a tradeoff between resource efficiency and performance isolation. On one hand, scheduling resource demands reduces the provisioned resources and increases resource efficiency. On the other hand, heavily relying on the unused resources of other NSs degrades performance isolation as indicated in the previous DDoS attack example. The analysis of this tradeoff is the main topic of this paper.

We study this tradeoff by formulating a joint scheduling and resource provisioning optimization problem. The objective is to minimize the provisioned resources s.t. for each NS  $i$ , its resource demand is met for  $P_i^H$  fraction of time, where for  $P_i^L$  fraction of time, it is met regardless of the traffic conditions in other NSs. Adjusting  $P_i^L$  in the above problem and computing the provisioned resources allows us to study the tradeoff.

We primarily focus on the instance of the problem involving the RAN part of a NS, where the network functions to be provisioned are the MAC schedulers of each NS and the resource demands are bandwidth demands measured in PRBs. Nonetheless, our approach can be applied to any network function whose resource provisioning decouples per node.

The main result of our paper is that the Max-Weight scheduler is an optimal multiplexing policy since it requires the least provisioned resources to satisfy the SLAs of all NSs, when their resource demands follow an ergodic Markov Chain (MC). To test its performance in non-markovian settings, we obtain resource demands from the LTE module of ns-3. The results show that the Max-Weight scheduler maintains its optimality and achieved considerable bandwidth savings.

## II. RELATED WORK

We are not the first to identify the resource efficiency vs performance isolation tradeoff. In [1], the authors studied the effect of overbooking strategies on resource allocation and service violations. Specifically, the authors considered the same  $P^H$  and  $P^L$  values for all NSs. Two scenarios where investigated; perfect sharing and network slicing. In perfect sharing, the base station (BS) simply sums the resource demands of all NSs and provisions  $P^H$ -percentile resources. In network slicing, performance isolation is considered by allocating dedicated  $P^L$ -percentile resources to each NS. The remaining resources needed to achieve  $P^H - P^L$  fraction of time acceptance for each NS are computed based on past data.

Although their method provides significant insight, it requires the same  $P^H$  and  $P^L$  value for all NSs. Also, there is no insight on how to find the optimal multiplexing strategy.

In [2], the authors consider that the fixed provisioned resources are more than the sum of the dedicated resources to each NS. The remaining resources that are not dedicated to any NS are viewed as auxiliary resources that are provided to each NS with probabilistic guarantees. The main focus of the paper is on overbooking strategies for the auxiliary resources. The authors highlight the importance of forecasting and propose interesting pricing models for this scenario. However, to analyze the tradeoff we are interested in, the joint consideration of overbooking and resource provisioning is needed.

The authors in [3] study the effect of  $P_i^H$  on the provisioned resources by solving at each timeslot  $t$  a bandwidth minimization problem. However, multiplexing is not considered and the tradeoff cannot be fully studied. Also, solving a separate optimization problem each timeslot  $t$  may lead to SLA violations over time if the resource demands are not iid.

Several other works address SLA fulfillment in various contexts. In [4], the authors present an overbooking strategy to efficiently satisfy SLAs in end-to-end NSs and implement it on a testbed composed of small-scale LTE BSs and OpenFlow switches. In [5], a prediction model for each NS's PRB usage is developed to limit overprovisioning. Also, past literature on cloud computing has addressed SLA fulfillment [6]–[10]. However, a clear problem formulation and solution approach regarding the tradeoff we are interested in is still missing.

Lastly, in [11], the authors discuss the concept of isolation in terms of performance, security, and dependability for various network functions in the RAN, TN, and CN. In this paper, we primarily focus on performance isolation as the title suggests.

## III. SYSTEM ARCHITECTURE

We consider  $N$  NSs served by a single BS in a RAN. Each NS may have different type of QoS requirements. For instance, a NS may wish to upper bound the average packet delay of its users, whereas another NS may wish to provide constant bitrates to its users. For this reason, we also consider that NSs may use different MAC schedulers.

Let vector  $\mathbf{W}(t) = [W_1(t), W_2(t), \dots, W_N(t)]$  contain the bandwidth demands at timeslot  $t$  of all  $N$  NSs. By bandwidth demand  $W_i(t)$ , we refer to the number of PRBs that the MAC scheduler of NS  $i$  needs in order to provide the desired QoS of NS  $i$  throughout timeslot  $t$ .

The determination of the bandwidth demand  $W_i(t)$  given the current traffic state  $s_i(t)$  of NS  $i$  is not trivial, especially for complex QoS requirements. However, for simpler QoS metrics, it may be straightforward. For instance, suppose NS  $i$  needs to deliver constant bitrates to each of its users. Then, the bandwidth demand  $W_i(t)$  is computed using the Modulation and Coding Scheme (MCS) of each user at timeslot  $t$ .

In any case, we consider that the bandwidth demand vector  $\mathbf{W}(t)$  is estimated at each timeslot  $t$  by a network function called bandwidth demand estimator whose operation is outside of the scope of this paper. Next, we consider a NS-level

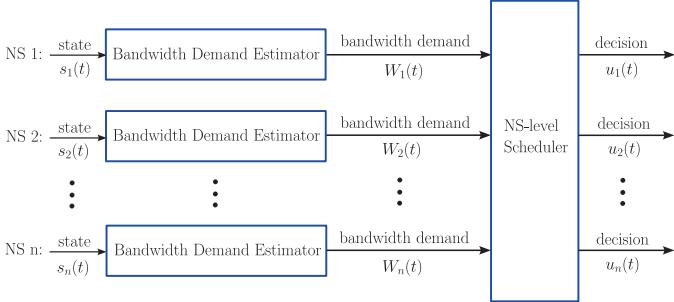


Fig. 1. In our system architecture, we consider two network functions. First, the Bandwidth Demand Estimator observes the traffic state  $s_i(t)$  of a NS and estimates the number of PRBs  $W_i(t)$  needed to meet its desired QoS. Second, the NS-level Scheduler receives all bandwidth demands  $W_i(t)$  and decides which ones to satisfy given the limited bandwidth at the BS. The timeslot length depends on the timescales supported by these two network functions.

scheduler that at each timeslot  $t$  decides whether to allocate the  $W_i(t)$  PRBs to the MAC scheduler of NS  $i$  given the overall demand vector  $\mathbf{W}(t)$  and the limited bandwidth at the BS. Thus, the scheduler's output is a binary decision vector  $\mathbf{u}(t)$ . Figure 1 depicts the system architecture.

#### IV. PROBLEM FORMULATION

The overall objective of our system is to satisfy a high percentage of the bandwidth demands  $\mathbf{W}(t)$  over a long period of time. Considering that  $u_i(t) = 1$  if and only if the demand  $W_i(t)$  is met, then we wish to satisfy the following constraint:

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(t) \geq P_i^H, \forall i. \quad (1)$$

Constraint (1) states that NS  $i$  needs to receive its desired QoS for a high percentage of time  $P_i^H$ . This type of constraints are known as availability constraints and are widely used in real networks for bandwidth provisioning purposes, as in Google's software-defined network B4 [12].

Next, we consider that the scheduler provisions  $W_i^r$  bandwidth at the BS for each NS  $i$ . Therefore, regardless of the traffic in the other NSs, the demand of NS  $i$  is always met if it is less than  $W_i^r$ . Thus, bandwidth  $W_i^r$  affects the degree of performance isolation that NS  $i$  enjoys. Here, we consider the following performance isolation constraint:

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{W_i(t) \leq W_i^r} \geq P_i^L, \forall i, \text{ where } P_i^L < P_i^H. \quad (2)$$

Note that  $P_i^L$  is the guaranteed percentage of time that NS  $i$  receives its desired QoS regardless of unexpected traffic surges in other NSs. Thus,  $P_i^L$  can be viewed as the degree of performance isolation that NS  $i$  enjoys. The values of  $P_i^H$  and  $P_i^L$  are specified in the SLA between tenant  $i$  and the NO.

Next, note that since  $u_i(t) = 1$  if  $W_i(t) \leq W_i^r$ , constraint (1) is equivalently reformulated as follows:

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{W_i(t) \leq W_i^r} + u_i(t) \mathbb{1}_{W_i(t) > W_i^r} \geq P_i^H, \forall i. \quad (3)$$

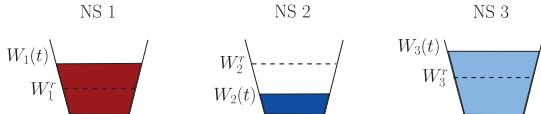


Fig. 2. Whenever  $W_i(t) \leq W_i^r$ , the demand of NS  $i$  is met regardless of traffic in other NSs. In case  $W_i(t) > W_i^r$ , the scheduler may allocate the unused resources of the other NSs to NS  $i$ . Here, the unused resources of NS 2 do not suffice for the excess demands of both NS 1 and NS 3. Thus, the scheduler needs to decide which one to satisfy. Note that the larger  $W_i^r$  is, the less frequently NS  $i$  needs to rely on the unused resources of other NSs. Thus, increasing  $W_i^r$  increases performance isolation. However, increasing  $W_i^r$  also increases the provisioned bandwidth, hence the tradeoff.

In case  $W_i(t) \leq W_i^r$ , we allow the scheduler to allocate the unused bandwidth provisions  $W_i^r - W_i(t)$  of NS  $i$  to some other NS  $j$  that needs it, where  $W_j(t) > W_j^r$ . Thus, we allow the scheduler to multiplex the bandwidth demands of NSs.

Note that NS  $i$  may reduce its  $W_i^r$  provisions by relying on the unused provisions  $\sum_{j \neq i} (W_j^r - W_j(t))$  of the other NSs. Thus, multiplexing increases resource efficiency since the overall provisioned bandwidth is reduced. On the other hand, if a NS relies heavily on the unused resources of other NSs, unexpected traffic surges in the other NSs will affect its performance, which degrades performance isolation.

In the multiplexing paradigm, the scheduler decides at each time  $t$  which excess bandwidth demands to satisfy given the available bandwidth. Thus, the scheduler is constrained by<sup>1</sup>:

$$\mathbf{u}(t)^\top [\mathbf{W}(t) - \mathbf{W}^r]^+ \leq W^c + \mathbf{1}^\top [\mathbf{W}^r - \mathbf{W}(t)]^+, \forall t, \quad (4)$$

where  $W^c$  is an auxiliary bandwidth that is under the full control of the scheduler. Although the introduction of  $W^c$  is not necessary, it helps with the analysis later on. The multiplexing paradigm is shown in Fig. 2.

Next, note that  $\mathbf{W}(t)$  may be a random vector given by some traffic forecasting model. Thus, to find the scheduler that needs the least provisioned bandwidth to meet each availability (3) and isolation (2) constraint, we need to solve:

$$\begin{aligned} & \underset{W^c, \mathbf{W}^r, \{\mathbf{u}(t)\}_{t \in \mathbb{N}}}{\text{minimize}} \quad W^c + \mathbf{1}^\top \mathbf{W}^r \\ & \text{s.t.: } \Pr \left( \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{W_i(t) \leq W_i^r} + u_i(t) \mathbb{1}_{W_i(t) > W_i^r} \geq P_i^H \right. \\ & \quad \left. \cap \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{W_i(t) \leq W_i^r} \geq P_i^L \right) = 1, \forall i, \\ & \quad \mathbf{u}(t)^\top [\mathbf{W}(t) - \mathbf{W}^r]^+ \leq W^c + \mathbf{1}^\top [\mathbf{W}^r - \mathbf{W}(t)]^+, \forall t, \\ & \quad \mathbf{u}(t) \in \{0, 1\}^N, \forall t. \end{aligned} \quad (5)$$

The probabilistic constraint allows us to neglect extreme cases that occur w.p.0 that may result in over-provisioning in practice. Next, to simplify the first constraint, we assume the existence of a function  $F_{W_i}$  that satisfies:

$$\Pr \left( \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{W_i(t) \leq W_i^r} = F_{W_i}(W_i^r) \right) = 1, \forall W_i^r, \forall i. \quad (6)$$

<sup>1</sup>We use the shorthand notation  $[\cdot]^+$  for  $\max\{\cdot, 0\}$ .

Assumption (6) states that the time-average converges w.p.1 to a constant that depends on  $W_i^r$ . It implies that a single realization provides information for all time-averages over all sample paths. All ergodic processes  $\mathbf{W}(t)$  satisfy (6). We note that although  $F_{W_i}$  is not readily available to the NO, it must be an increasing function. Next, consider the following quantity:

$$W_i^L \triangleq \min\{W_i^r \in \mathbb{R} : F_{W_i}(W_i^r) \geq P_i^L\}. \quad (7)$$

Bandwidth  $W_i^L$  can be obtained through binary search by computing the time-average in (6) through simulations due to the monotonicity of  $F_{W_i}$ . Next, note that the probability in the first constraint does not change if we intersect its events with the event in (6). Thus, it can be seen that (5) is equivalent to:

$$\begin{aligned} & \underset{W^c, \mathbf{W}^r, \{\mathbf{u}(t)\}_{t \in \mathbb{N}}}{\text{minimize}} \quad W^c + \mathbf{1}^\top \mathbf{W}^r \\ & \text{s.t.: } \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(t) \mathbf{1}_{W_i(t) > W_i^r} \geq P_i^H - F_{W_i}(W_i^r), \text{ w.p.1, } \forall i, \\ & \quad W_i^r \geq W_i^L, \quad \forall i, \\ & \quad \mathbf{u}(t)^\top [\mathbf{W}(t) - \mathbf{W}^r]^+ \leq W^c + \mathbf{1}^\top [\mathbf{W}^r - \mathbf{W}(t)]^+, \quad \forall t, \\ & \quad \mathbf{u}(t) \in \{0, 1\}^N, \quad \forall t. \end{aligned} \quad (8)$$

Problem (8) is a joint scheduling and bandwidth provisioning problem with probabilistic time-average constraints. Its solution involves the determination of a scheduler  $\{\mathbf{u}(t)\}_{t \in \mathbb{N}}$  that satisfies the constraints with the least total bandwidth  $W^c + \mathbf{1}^\top \mathbf{W}^r$  and the derivation of bandwidths  $W^c$  and  $\mathbf{W}^r$ .

Lastly, note that by adjusting the degree of performance isolation  $P_i^L$ , we affect  $W_i^L$  in (7) which in turn affects the optimal value of (8), i.e., the provisioned bandwidth. Thus, problem (8) allows us to study the resource efficiency vs performance isolation tradeoff. In the next section, we transform (8) to simpler forms until it becomes tractable.

## V. EQUIVALENT TRANSFORMATIONS

From now on, we use the shorthand notation  $\mathbf{u}$  for  $\{\mathbf{u}(t)\}_{t \in \mathbb{N}}$ . Due to the second constraint in (8), all feasible solutions can be written as  $(W^c, \mathbf{W}^L + \mathbf{e}, \mathbf{u})$ , where  $\mathbf{e}$  is a vector with non-negative components. To simplify (8), we observe that the existence of  $W^c$  allows us to set  $\mathbf{W}^r = \mathbf{W}^L$  and proceed without loss of optimality. To prove this, we first show the following proposition.

**Proposition 1.** *Let  $\mathbf{e} \in \mathcal{R}_+^N$ . If  $(W^c, \mathbf{W}^L + \mathbf{e}, \mathbf{u})$  optimally solves (8), then  $(W^c + \mathbf{1}^\top \mathbf{e}, \mathbf{W}^L, \mathbf{v})$  also optimally solves (8), where  $v_i(t) = 1$  if  $u_i(t) \mathbf{1}_{W_i(t) > W_i^L + e_i} = 1$  or  $W_i^L \leq W_i(t) \leq W_i^L + e_i$ , otherwise  $v_i(t) = 0$ .*

*Proof:* See Appendix A.

Proposition 1 states that if the provisioned bandwidth  $W_i^r$  is bigger than  $W_i^L$  by some  $e_i$ , this  $e_i$  can be instead added to the auxiliary bandwidth  $W^c$  and still there exists a scheduler  $\mathbf{v}$  that satisfies the constraints. Let  $W_i^e(t) \triangleq W_i(t) - W_i^L$  denote the excess demand of NS  $i$  when  $W_i^r = W_i^L$ , and let  $P_i^M \triangleq F_{W_i}(W_i^L)$ . We mention that if  $F_{W_i}$  is continuous, then

$P_i^M = P_i^L$  due to (7). By adding the constraint  $\mathbf{W}^r = \mathbf{W}^L$  in (8), we obtain the following problem:

$$\begin{aligned} & \underset{W^c, \mathbf{u}}{\text{minimize}} \quad W^c \\ & \text{s.t.: } \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(t) \mathbf{1}_{W_i^e(t) > 0} \geq P_i^H - P_i^M, \text{ w.p.1, } \forall i, \\ & \quad \mathbf{u}(t)^\top [\mathbf{W}^e(t)]^+ \leq W^c + \mathbf{1}^\top [-\mathbf{W}^e(t)]^+, \quad \forall t, \\ & \quad \mathbf{u}(t) \in \{0, 1\}^N, \quad \forall t. \end{aligned} \quad (9)$$

**Proposition 2.** *If  $(W^c, \mathbf{W}^r, \mathbf{u})$  optimally solves (8), then  $(W^c + \mathbf{1}^\top \mathbf{e}, \mathbf{v})$  optimally solves (9), where  $\mathbf{e} = \mathbf{W}^r - \mathbf{W}^L$  and  $\mathbf{v}$  as defined in Proposition 1. If  $(W^c, \mathbf{u})$  optimally solves (9), then  $(W^c, \mathbf{W}^L, \mathbf{u})$  optimally solves (8).*

*Proof:* See Appendix B.

Thus, problems (8) and (9) are equivalent. However, problem (9) has  $N$  less optimization variables than (8). To solve (9), we need to find a scheduler  $\mathbf{u}$  that satisfies the constraints using the smallest possible amount of bandwidth  $W^c$  and compute this bandwidth  $W^c$ . Let  $\mathcal{G}(W^c, \mathbf{P}^H, \mathbf{P}^M, \{\mathbf{W}^e(t)\}_{t \in \mathbb{N}})$  denote the set of schedulers  $\mathbf{u}$  that satisfy the constraints of (9) for a fixed value of  $W^c$ . We use the shorthand notation  $\mathcal{G}_{W^c}$  for  $\mathcal{G}(W^c, \mathbf{P}^H, \mathbf{P}^M, \{\mathbf{W}^e(t)\}_{t \in \mathbb{N}})$ . Then, problem (9) can be rewritten more compactly as:

$$\underset{W^c, \mathbf{u}}{\text{minimize}} \quad W^c \quad \text{s.t. } \mathbf{u} \in \mathcal{G}_{W^c}. \quad (10)$$

As mentioned before, problem (10) is a joint scheduling and bandwidth provisioning problem. A possible solution approach is to first find a feasible scheduler that requires the least bandwidth possible, and then solve the resulting provisioning problem only for this scheduler. This observation leads us to the following proposition.

**Proposition 3.** *Consider the optimization problem:*

$$\underset{W^c, \mathbf{u}}{\text{minimize}} \quad W^c \quad \text{s.t. } \mathbf{u} \in \mathcal{F}_{W^c}. \quad (11)$$

*Suppose  $\exists \{\mathbf{u}_{W^c}\}_{W^c \in \mathbb{R}}$  s.t.  $\forall W^c \in \mathbb{R}$ , if  $\mathbf{u}_{W^c} \notin \mathcal{F}_{W^c}$ , then  $\mathcal{F}_{W^c} = \emptyset$ . Next, consider optimization problem:*

$$\underset{W^c}{\text{minimize}} \quad W^c \quad \text{s.t. } \mathbf{u}_{W^c} \in \mathcal{F}_{W^c}. \quad (12)$$

*If  $W^{c*}$  optimally solves (12), then  $(W^{c*}, \mathbf{u}_{W^{c*}})$  optimally solves (11). Also, if  $(W^{c*}, \mathbf{v}^*)$  optimally solves (11), then  $W^{c*}$  optimally solves (12).*

*Proof:* See Appendix C.

Proposition 3 implies that if there exists a set of schedulers parameterized by  $W^c$  s.t. for every bandwidth  $W^c$ , either its scheduler corresponding to  $W^c$  can satisfy the constraints or no other scheduler can, then it suffices to consider only that set of schedulers and solve only for  $W^c$  as shown in (12).

Thus, our goal now is to identify a scheduler that satisfies the premise of Proposition 3 for  $\mathcal{F}_{W^c} = \mathcal{G}_{W^c}$ , and then solve provisioning problem (12) for that scheduler for  $\mathcal{F}_{W^c} = \mathcal{G}_{W^c}$ . Therefore, the original joint scheduling and bandwidth provisioning problem can be divided into a scheduling problem and a bandwidth provisioning problem.

In Sec. VI, we show that the well-known Max-Weight scheduler [13] solves the scheduling problem, i.e., it satisfies the premise of Proposition 3 for  $\mathcal{F}_{W^c} = \mathcal{G}_{W^c}$ , under the assumption that bandwidth demands  $\mathbf{W}(t)$  follow an ergodic MC. Although in practice the markovian assumption may be violated, it still motivates us to use the Max-Weight scheduler in general settings. Later in Sec. VII, we address the bandwidth provisioning problem for the Max-Weight Scheduler.

## VI. SCHEDULING PROBLEM

To identify a set of schedulers that satisfies the premise of Proposition 3, note that the second and third constraint in (9) specify a set of allowed control actions  $\mathcal{U}(t)$  that depends on  $\mathbf{W}^e(t)$  and  $W^c$ . Thus, problem (9) can be rewritten as:

$$\begin{aligned} & \underset{W^c, \mathbf{u}}{\text{minimize}} \quad W^c \\ \text{s.t.:} & \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(t) \mathbf{1}_{W_i^e(t) > 0} \geq P_i^H - P_i^M, \text{ w.p.1, } \forall i, \\ & \mathbf{u}(t) \in \mathcal{U}(\mathbf{W}^e(t), W^c), \forall t. \end{aligned} \quad (13)$$

Given that schedulers operate online, we consider only causal schedulers, i.e., schedulers that at time  $t$  do not know the future bandwidth demands  $W_i(\tau), \tau > t$ . Therefore at time  $t$ , the control decisions  $\mathbf{u}(t)$  may depend only on past stored information such as  $W_i^e(\tau)$  and  $\mathbf{u}(\tau)$  where  $\tau < t$ . Here, we consider that the scheduler computes  $\mathbf{u}(t)$  based on a state  $\mathbf{x}(t)$  that summarizes the system's history at time  $t$ .

To determine the information that the state  $\mathbf{x}(t)$  should include, note that the knowledge of  $\mathbf{W}^e(t)$  and  $W^c$  is needed, otherwise the set of feasible control actions cannot be found. Thus,  $(\mathbf{W}^e(t), W^c) \in \mathbf{x}(t)$ . Next, note that the first constraint (13) is a time-average constraint. Thus, Lyapunov optimization is applicable [13]. For this reason, we introduce the following:

$$d_i(t+1) = [d_i(t) - u_i(t) \mathbf{1}_{W_i^e(t) > 0}]^+ + P_i^H - P_i^M, \forall i. \quad (14)$$

Quantities  $d_i(t)$  can be viewed as deficits; in each timeslot the deficit to NS  $i$  increases by  $P_i^H - P_i^M$  unless there is a positive excess demand and the scheduler satisfies it, i.e.,  $W_i^e(t) > 0$  and  $u_i(t) = 1$ , in which case the deficit decreases. Deficits have been widely used for scheduling packets with deadlines where the goal is to meet delivery ratios [14].

Deficits  $d_i(t)$  can also be viewed as virtual queue lengths where in each timeslot, the number of arrivals is  $P_i^H - P_i^M$  and the amount of service received is  $u_i(t) \mathbf{1}_{W_i^e(t) > 0}$  for each queue. The introduction of deficits  $d_i(t)$  is useful due to the following implication [13, Theorem 2.8b]:

$$\begin{aligned} & \limsup_{T \rightarrow \infty} \sum_{t=1}^T \mathbb{E}[d_i(t)] < \infty \\ \Rightarrow & \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(t) \mathbf{1}_{W_i^e(t) > 0} \geq P_i^H - P_i^M, \text{ w.p.1.} \end{aligned} \quad (15)$$

Implication (15) states that if deficits  $d_i(t)$  are strongly stable, then the time-average constraints in (13) are satisfied. Consequently, the knowledge of deficits  $\mathbf{d}(t)$  may be useful

to the scheduler. Thus, we consider schedulers that produce  $\mathbf{u}(t)$  at each timeslot  $t$  based on  $\mathbf{x}(t) = (\mathbf{W}^e(t), \mathbf{d}(t), W^c)$ .

Now, consider a scheduler that observes  $\mathbf{x}(t)$  and chooses to satisfy the combination of NSs with the largest sum of deficits given the available bandwidth. Equivalently, consider a scheduler whose decisions  $\mathbf{u}_{W^c}^*(t)$  are obtained by solving:

$$\begin{aligned} & \underset{\mathbf{u}(t)}{\text{maximize}} \quad \mathbf{u}(t)^\top \mathbf{d}(t) \\ \text{s.t.:} & \sum_{i: W_i^e(t) > 0} u_i(t) W_i^e(t) \leq W^c + \sum_{i: W_i^e(t) < 0} |W_i^e(t)|, \\ & \mathbf{u}(t) \in \{0, 1\}^N. \end{aligned} \quad (16)$$

The scheduler obtained by solving (16) at each timeslot  $t$  is called the Max-Weight scheduler. Note that (16) is a binary knapsack problem in each timeslot  $t$ . Although, the binary knapsack problem is known to be NP-Hard, there exist fully polynomial time approximation schemes. Thus, in practice, close to optimal solutions can be obtained fast.

Intuitively, we expect that if the Max-Weight scheduler  $\mathbf{u}_{W^c}^*$  cannot stabilize the deficits for bandwidth  $W^c$ , then no other scheduler can, since the Max-Weight scheduler reduces the sum of deficits as much as possible at each timeslot. Indeed, it can be shown that the Max-Weight scheduler satisfies the premise in Proposition 3 when the bandwidth demands  $\mathbf{W}(t)$  follow an ergodic MC [13, Chapter 4.9.2]. Thus, it holds:

**Proposition 4.** *Let  $\mathbf{u}_{W^c}^*$  be the Max-Weight scheduler for bandwidth  $W^c$  obtained by solving (16) for  $\forall t$ . Suppose  $\mathbf{W}(t)$  follows an ergodic MC. If  $\mathbf{u}_{W^c}^* \notin \mathcal{G}_{W^c}$ , then  $\mathcal{G}_{W^c} = \emptyset$ .*

*Proof:* See Appendix D.

We note that the Max-Weight scheduler not only satisfies the premise of Proposition 3 for  $\mathcal{F}_{W^c} = \mathcal{G}_{W^c}$  as stated in Proposition 4, but also it does not require any knowledge of the statistics of bandwidth demands  $\mathbf{W}(t)$ . This is a major advantage since in case of incorrect estimation of the statistics of  $\mathbf{W}(t)$ , the scheduler does not need to be re-adjusted.

Another benefit of the Max-Weight scheduler is its simplicity. When used as the NS-level scheduler in Fig. 1, it computes the excess demands  $\mathbf{W}^e(t) = \mathbf{W}(t) - \mathbf{W}^L(t)$  and decides which ones to satisfy by solving (16). Clearly, all negative excess demands are satisfied. Then, based on these decisions  $\mathbf{u}_{W^c}^*(t)$ , it updates the deficit  $d_i(t)$  owed to each NS  $i$  using (14). This process is repeated online at each timeslot  $t$ .

## VII. BANDWIDTH PROVISIONING PROBLEM

Due to Proposition 3 and Proposition 4, the original joint scheduling and bandwidth provisioning problem is equivalent to a bandwidth provisioning problem for the Max-Weight scheduler, if the bandwidth demands  $\mathbf{W}(t)$  are markovian. Consequently, in that case, it suffices to solve:

$$\begin{aligned} & \underset{W^c}{\text{minimize}} \quad W^c \quad \text{s.t. } \forall i: \\ & \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_{W^c, i}^*(t) \mathbf{1}_{W_i^e(t) > 0} \geq P_i^H - P_i^M, \text{ w.p.1.} \end{aligned} \quad (17)$$

Although we proved the optimality of the Max-Weight scheduler under the assumption of ergodic markovian demands  $\mathbf{W}(t)$ , it still motivates us to consider the Max-Weight scheduler in general settings, where the markovian assumption may not hold. Thus, in all cases, we are interested in solving (17).

However, the derivative of the probability in the constraint w.r.t.  $W^c$  is not readily available. Thus, derivative-free methods are needed. Since  $u_{W^c,i}^*(t)$  is increasing w.r.t.  $W^c$ , then (17) is a monotonic optimization problem. Thus, binary search or polyblock outer approximation [15] can be used to solve it.

### VIII. OVERALL SOLUTION APPROACH

Here, we summarize the overall solution approach for the original joint scheduling and bandwidth provisioning problem (8). Regarding scheduling, we use the Max-Weight scheduler motivated by its optimality under ergodic markovian demands  $\mathbf{W}(t)$ . The scheduler's decisions are found by solving (16) online at each timeslot  $t$ . Since (16) is a binary knapsack problem, we can solve it fast using approximation algorithms.

Next, we need to find quantities  $W_i^L$  and  $P_i^M$  to solve provisioning problem (17). To do so, we consider that whenever a new NS needs to be deployed, a data collection period begins during which we observe previous bandwidth demands and generate future bandwidth demands  $\{\mathbf{W}(t)\}_{t \in \mathbb{N}}$  based on some traffic forecasting model. Then, since  $F_{W_i}$  is an increasing function, we find  $W_i^L$  by applying binary search on this sequence  $\{\mathbf{W}(t)\}_{t \in \mathbb{N}}$ . Note that during the process of finding  $W_i^L$ , we also obtain  $P_i^M$  as follows from (7).

Once  $W_i^L$  and  $P_i^M$  are obtained, problem (17) can be solved using binary search. Once this is done, both the scheduler  $\mathbf{u}$  and the required bandwidths  $\mathbf{W}^r = \mathbf{W}^L, W^c$  are specified. Thus, a feasible solution to (8) has been found. In case demands the sequence  $\{\mathbf{W}(t)\}_{t \in \mathbb{N}}$  was generated by an ergodic MC, this feasible solution is the optimal solution.

This overall process is repeated each billing cycle to update the forecasting model and charge each tenant accordingly. Some forecasting models for RANs can be found in [16], [17].

### IX. EXPERIMENTATION IN NS-3

Here, we test our overall solution approach on bandwidth demands obtained using the LTE module of ns-3 [18, Chapter 19]. All simulations were run on a home computer with an Intel i7-10700K processor using 16 GB of RAM running on Windows 10. Our simulation setup is depicted in Fig. 3 and the associated Radio Environment Map (REM) in Fig. 4. We deploy one voice, one video streaming and two web-browsing NSs. The simulation time in all scenarios is 10 minutes due to time complexity. The User Equipment (UE) activity changes every few seconds instead of minutes to observe various traffic patterns within this short simulation time.

For the voice NS, we create 100 UEs receiving traffic from the remote host with 1kBps data rate and 20B packet size to simulate the G.729A codec. Each UE's call duration is modeled by a Pareto random variable with a mean of 2s. The idle duration of each UE follows the exponential distribution with a mean of 5s. Hence, each UE is active 29% of the time.

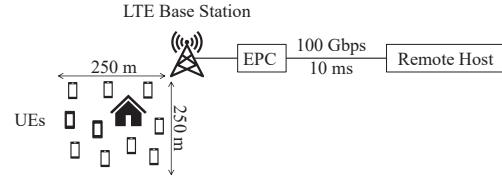


Fig. 3. In our ns-3 simulation setup, we consider a sector of an LTE BS covering a square area. A building was placed to simulate channel conditions in urban areas. The UEs move with 1.5m/s speed within this square and their propagation losses are computed by the "OhBuildings" model [18, p. 40]. The system bandwidth is 100 PRBs. We collect the bandwidth demands for each NS by executing an instance of this setup with a different number of UEs and a different application on the remote host generating downlink traffic.

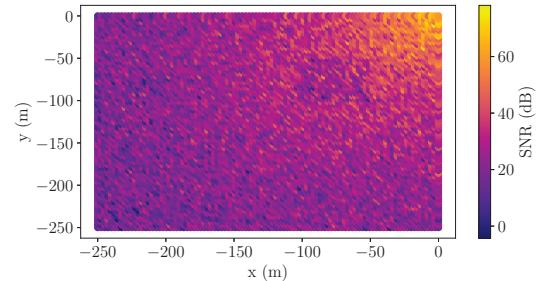


Fig. 4. The resulting REM of the considered network topology at 1m above the ground. The UEs move within this area and their downlink SNR changes accordingly. The BS is placed at (0, 0) and 20m above the ground. The building's dimensions are 50m x 50m x 10m and placed at (-100, -100).

We model the traffic of the video streaming NS as previously, however we consider 15 UEs with average idle duration of 10s, 8Mbps data rate and 1kB packet length.

The two web-browsing NSs were simulated by the default model in the 3GPP HTTP applications module of ns-3 [18, Chapter 5]. This model considers random byte sizes required to load objects of a website, random parsing times of these objects, and random reading times during which the UE remains idle. For the first web-browsing NS, we consider 10 UEs and for the second one 20 UEs.

Since we have not yet developed a bandwidth demand estimator as in Fig. 1, we approximate the bandwidth demand of a NS by the number of PRBs that the MAC scheduler allocated to the NS. Here, we encounter a challenge. Most MAC schedulers in ns-3 always allocate all 100 PRBs since it is wasteful not to do so [18, Chapter 19.1.7]. For instance, even if there is only one UE in the voice NS that requires just 1kBps to receive good QoS, all PRBs are allocated to that UE. Then, the bandwidth demand of each NS is a horizontal line at 100 PRBs which nullifies the benefits of multiplexing.

To tackle this issue, we need a MAC scheduler that does not always use all 100 PRBs. We found that the Token Bank Fair Queueing (TBFQ) MAC scheduler possesses this desired behavior [18, p. 217]. Specifically, this MAC scheduler stops allocating extra PRBs if the already allocated ones are enough to empty the packet buffers of all UEs. Thus, using the TBFQ MAC scheduler, we can approximate the resource demand of a NS when the desired QoS is to transmit all packets of all UEs within 1 ms. Although this is an extreme QoS requirement, it

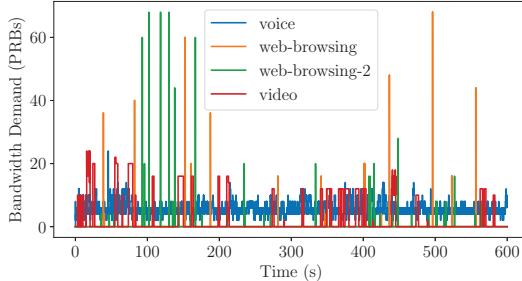


Fig. 5. The bandwidth demands of all four NSs for a 10 minute simulation. Since their peaks are not synchronized, multiplexing can be highly beneficial.

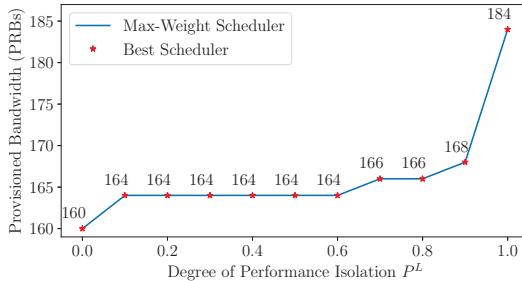


Fig. 6. Here, we consider  $P^L = P^H = 1$  for all NSs except from the voice NS whose degree of isolation  $P^L$  we vary. The figure also shows that the Max-Weight scheduler requires as much bandwidth as the best scheduler, even though the bandwidth demands are not generated by a MC.

allows us to obtain bandwidth demands that vary over time.

Since the simulation time is 10 minutes, we obtain 600K bandwidth demands for each NS. Then, we find the maximum every 100 ms and obtain the bandwidth demands in Fig. 5.

Given the obtained sequence of bandwidth demands, we can study the tradeoff of interest by following Sec. VIII for various  $P_i^H$  and  $P_i^L$  values. Some such results are depicted in Fig. 6. More extensive results are provided in Fig. 7 where we consider the voice, web-browsing and video NSs.

In the Fig. 6, we also compare the Max-Weight scheduler to the best scheduler. The best scheduler is found by solving (9) for the obtained sequence of bandwidth demands which is a mixed integer linear program. To solve it, we use the `scipy.optimize.milp` function in python which uses the HiGHS optimization software [19]. Note that the best scheduler depends on the perfect forecasting of all future bandwidth demands and thus is non-causal. Clearly, its computation and performance is susceptible to errors in forecasting.

Lastly, we measure the execution time of the Max-Weight scheduler, which requires the solution of a binary knapsack problem. To do so, we use the branch and bound algorithm from google OR-tools that obtains an optimal solution [20]. The average execution time over 10K tests for a varying number of NSs and arbitrary demands is shown in Fig. 8.

## X. CONCLUSION

In this paper, we studied the tradeoff between resource efficiency and performance isolation in network slicing. We showed that for bandwidth demands following an ergodic MC, the optimal multiplexing policy is the Max-Weight sched-

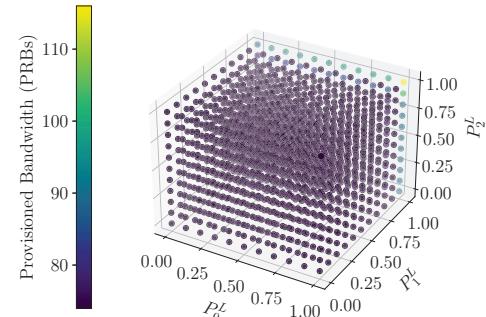


Fig. 7. The effect of the performance isolation vector  $\mathbf{P}^L$  on the provisioned bandwidth. Point  $(0, 0, 0)$  corresponds to the case where no isolation is required and all NSs fully rely on multiplexing. Point  $(1, 1, 1)$  corresponds to full isolation for each NS where multiplexing is not allowed. The first case requires 74 PRBs and the second case requires 116 PRBs. Hence, multiplexing in this scenario reduces the provisioned bandwidth by 7.56 MHz or by 36.2%.

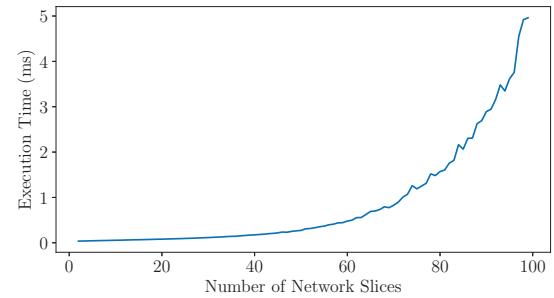


Fig. 8. The Max-Weight scheduler supports 1 ms time-granularity for a large number of NSs deployed at the same BS. Thus, the timeslot length in Fig. 1 is affected by the timescale that the bandwidth demand estimator can support.

uler. We then tested the scheduler's performance in a non-markovian setting and observed that it is still optimal. Also, when no performance isolation was required, the Max-Weight scheduler reduced the provisioned bandwidth by 36.2%.

We note that multiplexing in general requires the operation of a bandwidth demand estimator as in Fig. 1. The bandwidth savings for the simple scenarios studied in this paper motivates its development. For future work, we wish to develop such a bandwidth demand estimator using online learning methods.

## APPENDIX A PROOF OF PROPOSITION 1

First, note that the first point  $(W^c, \mathbf{W}^L + \mathbf{e}, \mathbf{u})$  and the second point  $(W^c + \mathbf{1}^\top \mathbf{e}, \mathbf{W}^L, \mathbf{v})$  have the same objective function value. Thus, it suffices to show that if the first point satisfies the constraints, then the second point satisfies them as well. Clearly, if the first point satisfies the second and fourth constraints of (9), then so does the second point.

Next, we consider the third constraint. Suppose the first point satisfies it. We need to show that  $(W^c + \mathbf{1}^\top \mathbf{e}, \mathbf{W}^L, \mathbf{v})$  also satisfies it. Equivalently, by re-arranging terms and using the definition of  $\mathbf{v}$ , we need to show:

$$\sum_{i: u_i(t)=1, W_i(t) \geq W_i^L + e_i} W_i(t) - W_i^L \leq W^c + \sum_i e_i + \sum_{i: W_i(t) \leq W_i^L + e_i} W_i^L - W_i(t). \quad (18)$$

Note however that if the first point  $(W^c, \mathbf{W}^L + \mathbf{e}, \mathbf{u})$  satisfies the third constraint, it holds:

$$\begin{aligned} & \sum_{i:u_i(t)=1, W_i(t) \geq W_i^L + e_i} W_i(t) - W_i^L - e_i \\ & \leq W^c + \sum_{i:W_i(t) < W_i^L + e_i} W_i^L + e_i - W_i(t). \end{aligned} \quad (19)$$

We prove (18) by moving the sum of  $e_i$  in the left-side of (19) and by increasing its right-side by  $\sum_{i:u_i(t)=0, W_i(t) \geq W_i^L + e_i} e_i$ .

Lastly, we need to show that if the first point satisfies the first constraint, so does the second one. From (6), we need to show:

$$\begin{aligned} & \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{W_i(t) \leq W_i^L} + v_i(t) \mathbb{1}_{W_i(t) > W_i^L} \\ & \geq \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{W_i(t) \leq W_i^L + e_i} + u_i(t) \mathbb{1}_{W_i(t) > W_i^L + e_i}. \end{aligned} \quad (20)$$

By definition of  $\mathbf{v}$ , inequality (20) holds with equality.

## APPENDIX B PROOF OF PROPOSITION 2

We start with the first statement. Since  $(W^c, \mathbf{W}^r, \mathbf{u})$  feasible, then  $\mathbf{W}^r = \mathbf{W}^L + \mathbf{e}$ , where  $\mathbf{e} \in \mathcal{R}_+^N$ . Due to Proposition 1, it follows that  $(W^c + \mathbf{1}^\top \mathbf{e}, \mathbf{W}^L, \mathbf{v})$  is also an optimal solution to (8). Next, note that the feasibility region of (9) is a subset of the feasibility region of (8). Thus, since  $(W^c + \mathbf{1}^\top \mathbf{e}, \mathbf{W}^L, \mathbf{v})$  optimally solves (8), then  $(W^c + \mathbf{1}^\top \mathbf{e}, \mathbf{v})$  must be an optimal solution to (9).

We prove the second statement by contradiction. Suppose  $\exists (W^{c'}, \mathbf{W}^{r'}, \mathbf{u}')$  s.t.  $W^{c'} + \mathbf{1}^\top \mathbf{W}^{r'} < W^c + \mathbf{1}^\top \mathbf{W}^L$ . If  $\mathbf{W}^{r'} = \mathbf{W}^L$ , it follows that  $(W^{c'}, \mathbf{u}')$  is a feasible solution to (9) with  $W^{c'} < W^c$ . Thus,  $(W^c, \mathbf{u})$  does not optimally solve (9), which is a contradiction. If  $\mathbf{W}^{r'} = \mathbf{e} + \mathbf{W}^L$ , where  $\mathbf{e} \in \mathcal{R}_+^N$ , due to Proposition 1,  $\exists (W^{c'} + \mathbf{1}^\top \mathbf{e}, \mathbf{W}^L, \mathbf{v})$  that optimally solves (8). Thus,  $(W^{c'} + \mathbf{1}^\top \mathbf{e}, \mathbf{v})$  feasible solution to (9) with  $W^{c'} + \mathbf{1}^\top \mathbf{e} \leq W^c$ , which is a contradiction.

## APPENDIX C PROOF OF PROPOSITION 3

Suppose the first statement does not hold, i.e.,  $\exists (W^{c'}, \mathbf{u}')$  that optimally solves (11), where  $W^{c'} < W^{c*}$ . Then,  $\mathbf{u}' \in \mathcal{F}_{W^{c'}}$ , thus  $\Rightarrow \mathcal{F}_{W^{c'}} \neq \emptyset$ . Due to the premise, it follows that  $u_{W^{c'}} \in \mathcal{F}_{W^{c'}}$ . Thus,  $W^{c'}$  and  $W^{c*}$  optimally solve (12) which is a contradiction. For the second statement,  $(W^{c*}, \mathbf{v}^*)$  optimally solves (11), thus  $\mathcal{F}_{W^{c*}} \neq \emptyset$ . Due to the premise, it follows that  $\mathbf{u}_{W^{c*}} \in \mathcal{F}_{W^{c*}}$ . Thus,  $(W^{c*}, \mathbf{u}_{W^{c*}})$  also optimally solves (11). Therefore,  $W^{c*}$  optimally solves (12).

## APPENDIX D PROOF OF PROPOSITION 4

Due to page limitations, the complete proof is available in [21, Appendix D]. Here, we provide a sketch of the proof. Our proof follows the same methodology described over several chapters in [13]. By noticing that  $\mathbf{W}^e(t)$  and  $\mathbf{d}(t)$  correspond to quantities  $\omega(t)$  and  $\mathbf{Q}(t)$  in [13] respectively, it

is straightforward to apply the methods in [13] to our problem. Our proof consists of two parts.

First, we consider a class of schedulers that make decisions based only on bandwidth demands  $\mathbf{W}^e(t)$  called  $\mathbf{W}^e(t)$ -only schedulers. Using [13, Theorem 4.5], we show that there exists a set of schedulers of this class that satisfies the premise of Proposition 3 for ergodic markovian demands.

Second, by following the steps in [13, p.34-p.36], we show that the Max-Weight scheduler strongly stabilizes the deficits  $d_i(t)$  whenever that set of  $\mathbf{W}^e(t)$ -only schedulers stabilizes them. Thus, the Max-Weight scheduler also satisfies the premise of Proposition 3 for ergodic markovian demands.

## REFERENCES

- [1] C. Marquez, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Pérez, “Resource sharing efficiency in network slicing,” *IEEE Trans. Netw. Service Manag.*, vol. 16, no. 3, pp. 909–923, 2019.
- [2] C. Sexton, N. Marchetti, and L. A. DaSilva, “On provisioning slices and overbooking resources in service tailored networks of the future,” *IEEE/ACM Trans. Ntw.*, vol. 28, no. 5, pp. 2106–2119, 2020.
- [3] J. Li, J. Liu, T. Huang, and Y. Liu, “Dra-ig: The balance of performance isolation and resource utilization efficiency in network slicing,” in *IEEE ICC 2020*, 2020, pp. 1–6.
- [4] J. X. Salvat, L. Zanzi, A. Garcia-Saavedra, V. Sciancalepore, and X. Costa-Perez, “Overbooking network slices through yield-driven end-to-end orchestration,” in *ACM CoNEXT ’18*, 2018, p. 353–365.
- [5] C. Gutierrez, E. Grinshpun, S. Sharma, and G. Zussman, “Ran resource usage prediction for a 5g slice broker,” in *ACM MobiHoc ’19*, New York, NY, USA, 2019, p. 231–240.
- [6] S. A. Baset, L. Wang, and C. Tang, “Towards an understanding of oversubscription in cloud,” in *USENIX Hot-ICE ’12*, 2012.
- [7] F. Caglar and A. Gokhale, “ioverbook: Intelligent resource-overbooking to support soft real-time applications in the cloud,” in *2014 IEEE CLOUD*, 2014, pp. 538–545.
- [8] D. Shue, M. J. Freedman, and A. Shaikh, “Performance isolation and fairness for Multi-Tenant cloud storage,” in *USENIX OSDI ’12*. Hollywood, CA: USENIX Association, Oct. 2012, pp. 349–362.
- [9] S. Walraven, W. De Borger, B. Vanbrabant, B. Lagaisse, D. Van Landuyt, and W. Joosen, “Adaptive performance isolation middleware for multi-tenant saas,” in *2015 IEEE/ACM UCC*, 2015, pp. 112–121.
- [10] S. A. Baset, “Cloud slas: Present and future,” *ACM SIGOPS Oper. Syst. Rev.*, vol. 46, no. 2, pp. 57–66, jul 2012.
- [11] A. J. Gonzalez *et al.*, “The isolation concept in the 5g network slicing,” in *2020 EuCNC*, 2020, pp. 12–16.
- [12] C.-Y. Hong *et al.*, “B4 and after: Managing hierarchy, partitioning, and asymmetry for availability and scale in google’s software-defined wan,” in *SIGCOMM*. Budapest, Hungary: ACM, 2018, p. 74–87.
- [13] M. J. Neely, “Stochastic network optimization with application to communication and queueing systems,” *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1–211, 2010.
- [14] C. Tsanikidis and J. Ghaderi, “Randomized scheduling of real-time traffic in wireless networks over fading channels,” *IEEE/ACM Trans. Ntw.*, pp. 1–14, 2022.
- [15] Y. J. Zhang, L. Qian, and J. Huang, *Monotonic Optimization in Communication and Networking Systems*. Now Publishers, 2013.
- [16] V. Sciancalepore, K. Samdanis, X. Costa-Perez, D. Bega, M. Gramaglia, and A. Banchs, “Mobile traffic forecasting for maximizing 5g network slicing resource utilization,” in *IEEE INFOCOM 2017*, 2017, pp. 1–9.
- [17] R. Li, Z. Zhao, X. Zhou, J. Palicot, and H. Zhang, “The prediction analysis of cellular radio access network traffic: From entropy theory to networking practice,” *IEEE Commun. Mag.*, vol. 52, no. 6, 2014.
- [18] NS-3, “ns-3.38 model library.” 2023. [Online]. Available: <https://www.nsnam.org/docs/release/3.38/models/ns-3-model-library.pdf>
- [19] Q. Huangfu and J. A. J. Hall, “Parallelizing the dual revised simplex method,” *Math. Program. Comput.*, vol. 10, pp. 119–142, 2015.
- [20] Google-OR-tools, “The knapsack problem.” 2023. [Online]. Available: <https://developers.google.com/optimization/pack/knapsack>
- [21] P. Nikolaidis, A. Zoukarni, and J. Baras, “Resource efficiency vs performance isolation tradeoff in network slicing,” 2023, *arXiv:2303.14634*.