

# Getting the Most Out of Your VNFs: Flexible Assignment of Service Priorities in 5G

Francesco Malandrino<sup>\*,†</sup>, Carla-Fabiana Chiasserini<sup>†,\*</sup>

\*: CNR-IEIIT, Torino, Italy

†: Politecnico di Torino, Torino, Italy

**Abstract**—Through their computational and forwarding capabilities, 5G networks can support multiple vertical services. Such services may include several common virtual (network) functions (VNFs), which could be shared to increase resource efficiency. In this paper, we focus on the seldom studied VNF-sharing problem, and decide (i) whether sharing a VNF instance is possible/beneficial or not, (ii) how to *scale* virtual machines hosting the VNFs to share, and (iii) the priorities of the different services sharing the same VNF. These decisions are made with the aim to minimize the mobile operator’s costs while meeting the verticals’ performance requirements. Importantly, we show that the aforementioned priorities should not be determined *a priori* on a per-service basis, rather they should change across VNFs since such additional flexibility allows for more efficient solutions. We then present an effective methodology called FlexShare, enabling near-optimal VNF-sharing decisions in polynomial time. Our performance evaluation, using real-world VNF graphs, confirms the effectiveness of our approach, which consistently outperforms baseline solutions using per-service priorities.

## I. INTRODUCTION

5G networks differ from their previous counterparts in several key features. Among them, two especially relevant aspects are the computational capabilities with which networks are endowed, and the relationship between network operators and vertical industries (e.g., multimedia content providers, automotive industries, smart factories).

Thanks to software-defined networking (SDN) and network function virtualization (NFV), cellular networks have now the ability to *process* data, as well as to forward them. A set of hosts, be them physical servers or virtual machines, run virtual network functions (VNFs), performing network-related (e.g., firewalls), or service-specific (e.g., video transcoding) tasks. Such an arrangement is beneficial for both mobile network operators (MNOs), who can fully utilize their infrastructure, and verticals, who can use it to deploy their services and enjoy lower delays.

The relationship between MNOs and verticals is evolving accordingly. Verticals can enter business relations with MNOs, using their network infrastructure to provide services with the required quality-of-service level. Indeed, vertical services are specified as a set of interconnected VNFs, along with *per-service* target key performance indicators (KPIs), e.g., throughput, delay, or reliability. Supporting services of multiple verticals, with different target KPIs, through the same MNO infrastructure, is possible thanks to *network slicing*. This is a powerful concept enabling multiple logical networks as independent business operations, on a common physical infrastructure [1]. Notably, network slicing also accounts for

composed services, i.e., service VNF graphs including sub-graphs, each representing a child service [2]. Indeed, different network slices may contain one or more common sub-slices [3], [4], a typical example being the evolved packet core (EPC) of cellular networks [2].

When creating a network slice, the MNO is in charge of assigning to it the needed resources (e.g., virtual machines and links connecting them), and of deciding which VNFs each host should run. This problem, known as VNF placement, pursues the twofold objectives of (i) ensuring that the target KPIs are met, and (ii) minimizing the cost for the MNO. Importantly, the latter can be achieved by *sharing* individual VNFs or sub-slices, among multiple services whenever possible.

The vast majority of studies on VNF placement [5]–[7] implicitly assume that (i) all placement decisions are made by one entity, typically the NFV Orchestrator (NFVO) of the Management and Network Orchestration (MANO) framework [4], [8], and (ii) such entity makes fine-grained decisions about how individual hosts and links are used. However, such a behavior is not the only one included in standards, and is not typical of real-world 5G implementations. Indeed, ETSI [9, Sec. 8.3.6] specifies four possible granularity levels for placement decisions, namely, Point of Presence (PoP) – e.g., a data center –, zone<sup>1</sup> group, zone, and individual host, with real-world 5G implementation efforts envisioning the NFVO to make PoP-level decisions [10], [11].

Fine-grained decisions on sharing VNF instances within individual PoPs are made by other, non-MANO entities, as envisioned by IETF [4, Sec. 3], the NGMN alliance [12, Sec. 8.9], and 5G-PPP [3, Sec. 2.2.2]. For sake of concreteness, we take as a reference the latter architecture, which includes an entity called Software-Defined Mobile Network Coordinator (SDM-X). As summarized in Fig. 1, the SDM-X works at a lower abstraction level than the MANO entities, and is in charge of solving the *VNF-sharing* problem, managing the VNFs that are common to multiple services.

This is the problem that, differently from traditional VNF placement studies, we address in this paper. Specifically, we solve the VNF-sharing problem within a PoP, by deciding, for each newly requested service,

- whether it is possible and convenient to re-use the existing VNF instances<sup>2</sup>;
- which priority to give to services sharing the same VNF, in order to meet the target KPIs;

<sup>1</sup>Zone refers to a subset of hosts within a PoP.

<sup>2</sup>For simplicity, and without loss of generality, we will refer to common VNFs only, instead of common VNFs and sub-slices.




Fig. 1. Architectural view of 5G networks according to 5G-PPP. Source: [2].

- whether and to which extent to *scale up* the computational capability of virtual machines (VMs) within the PoP.

A simple instance of the VNF-sharing problem we address is presented in Fig. 2. For each VNF of the newly-requested service  $s_2$ , we have to decide which VM shall run it (this decision is trivial for  $v_4$ , for which no instance exists yet) and the priority to assign to requests, e.g., REST queries, of each service in each shared VNF.

**Contributions.** To effectively tackle the VNF-sharing problem, we make the following main contributions:

- we observe that allowing *flexible* priorities for each VNF and service, makes it possible to meet KPI targets at a lower cost for the MNO;
- based on the above key observation, we present a system model capturing all the relevant aspects of the VNF-sharing problem and the entities it involves, including the capacity-scaling and priority-setting decisions it requires;
- leveraging convex optimization, we devise an efficient, integrated solution methodology called FlexShare, able to make swift, high-quality decisions concerning VM usage, priority assignment, and capability scaling;
- we evaluate the computational complexity of FlexShare and study its performance against real-world VNF graphs.

In the rest of the paper, we begin by presenting an example motivating the need for flexible priorities across VNFs (Sec. II). Then, Sec. III introduces our system model and problem formulation, while Sec. IV describes the FlexShare solution strategy. Our reference scenarios and numerical results are discussed in Sec. V. Finally, after reviewing related work in Sec. VI, Sec. VII concludes the paper.

## II. THE ROLE OF PRIORITIES

Before addressing the problem of whether it is convenient to share a VNF among multiple services or not, let us highlight the role of priorities while sharing a VNF instance. Three main approaches can be adopted for VNF sharing:

- *per-service* priority, associated with each service and constant across different VNFs;
- *per-VNF* priority, associated with each service and VNF, thus, given a service, it may vary across different VNFs;
- *per-request* priority, associated with individual service requests, e.g., REST queries, it may vary across the different VNFs on a per-request basis.




Fig. 2. Example of the VNF-sharing problem. A PoP is serving service  $s_1$ , with VNFs  $v_1$ – $v_3$  deployed at VMs  $m_1$ – $m_3$ . It is then requested to deploy  $s_2$ , using VNFs  $v_3$  and  $v_4$ . No isolation is requested, so services can share VNFs if convenient. For  $v_4$ , the only option is devoting an unused VM to it,  $m_5$  in the example (pink line). For  $v_3$ , instead, there are three options: re-using the instance of  $v_3$  at  $m_3$ , giving  $s_2$  a lower priority than  $s_1$  (dashed line); doing the same but giving  $s_2$  a higher priority (dotted line); devoting  $m_4$  – currently unused – to  $v_3$  (dash-dotted line), thus having two VMs running  $v_3$ .

In the following example, we focus on the first two steps of the above flexibility ladder and show how higher flexibility in priority assignment is associated with a higher efficiency in handling service traffic.

**Example 1** (The importance of flexible priorities). *Consider the two services,  $s_1$  and  $s_2$ , depicted in Fig. 3, requested by a vertical specialized in video surveillance systems.  $s_2$  includes two VNFs executing transcoding and motion detection, respectively, while  $s_1$  is composed of  $s_2$  and a VNF performing face recognition. Each VNF should run in its own VM, and network transfer times between VMs are neglected. Adopting a well-established and convenient approach [6], [7], let us model VNFs as M/M/1 traffic queues processing service requests, and the services as queuing chains, with arrival rates  $\lambda_1 = 2$  requests/ms and  $\lambda_2 = 1$  requests/ms, respectively. Also, consider service delay as the main performance metric and let the target average delay be  $D_1^{\max} = D_2^{\max} = 1.1$  ms for both services. Then assume that, given the allocated computation resources, the service rate of the transcoding and motion detection is  $\mu_{tc} = \mu_{md} = \mu = 5$  requests/ms, while that of face recognition is  $\mu_{fr} = 9.15$  requests/ms.*

*To meet the delay targets,  $s_2$  requests must traverse the transcoding and motion detection with a combined sojourn time of 1.1 ms, while  $s_1$  requests must do the same in at most 0.96 ms (i.e., the target average delay  $D_1^{\max}$  minus the sojourn time at the face recognition VNF of  $\frac{1}{\mu_{fr} - \lambda_1} = 0.14$  ms). We now show that there is no way of setting per-service priorities that allow this.*

**Case 1: Higher priority to  $s_1$ .** *This choice would make intuitive sense since  $s_1$  requests have to go through more processing stages than  $s_2$  requests, within the same deadline. In this case,  $s_1$  requests incur a sojourn time of  $\frac{1}{\mu - \lambda_1} = \frac{1}{5 - 2} = 0.33$  ms for each of the common VNFs, resulting in a total delay  $D_1 = 0.8$  ms, well within the target. However, the sojourn time of  $s_2$  requests at each of the shared VNFs becomes [13, Sec. 3.2]:  $\frac{1/\mu}{(1 - \frac{\lambda_1}{\mu})(1 - \frac{\lambda_1 + \lambda_2}{\mu})} = \frac{1/5}{(1 - \frac{2}{5})(1 - \frac{1+2}{5})} \approx 0.83$  ms, which results in a total delay of  $D_2 \approx 1.66$  ms  $>$   $D_2^{\max}$ .*

**Case 2: Higher priority to  $s_2$ .** *It is easy to verify that giving*




Fig. 3. Example 1: two video surveillance services,  $s_1$  and  $s_2$ , with  $s_2$  including performing transcoding and motion detection VNFs, and  $s_1$  composed of  $s_2$  and an additional face recognition stage.

higher priority to  $s_2$  implies that  $s_1$  misses its target delay.

Case 3: Equal priority. Giving the same priority to both services results in a sojourn time of  $\frac{1}{\mu - \lambda_1 - \lambda_2} = \frac{1}{5 - 1 - 2} = 0.5$  ms for each of the common VNFs, and in a total delay of  $D_2 = 1$  ms  $< D_2^{\max}$  and  $D_1 = 1$  ms + 0.14 ms  $> D_1^{\max}$ .

Flexible priorities. Assume that  $s_1$  and  $s_2$  have priority in the transcoding and in the motion detection VNF, respectively. Then,  $D_1 = \frac{1}{\mu - \lambda_1} + \frac{1/\mu}{(1 - \frac{\lambda_2}{\mu})(1 - \frac{\lambda_1 + \lambda_2}{\mu})} + 0.14 = 0.33 + 0.625 + 0.14 = 1.095$  ms  $< D_1^{\max}$ , and  $D_2 = \frac{1/\mu}{(1 - \frac{\lambda_1}{\mu})(1 - \frac{\lambda_1 + \lambda_2}{\mu})} + \frac{1}{\mu - \lambda_2} = 0.83 + 0.25 = 1.08$  ms  $< D_2^{\max}$ .

In conclusion, the above example shows that no combination of per-service priorities results in both  $s_1$  and  $s_2$  meeting their target delays. Instead moving one step up the flexibility ladder, i.e., assigning different priorities to service traffic flows across VNFs, allows the MNO to meet the vertical's requirements while increasing efficiency in resource usage, hence lowering the costs. As we will see later, even better performance can be obtained through per-request priorities.

### III. SYSTEM MODEL AND PROBLEM FORMULATION

As mentioned earlier, the SDM-X has to use the VMs under its control to provide the newly-requested services with the required KPIs and at the minimum cost. Specifically, it has to make decisions on (i) *whether* existing VNF instances should be shared; (ii) if so, *how* to assign the priorities to different services sharing the same VNF instance; (iii) *scaling up* the computational capabilities of the VMs if needed and possible.

As these decisions are made with reference to a single PoP, processing times are the dominant contribution with respect to network latency, and we therefore neglect the latter. This assumption is also justified by the ongoing work in the datacenter networking community (see, e.g., [14]), where (i) switching is highly optimized, hence network delays tend to be very small, and (ii) network topologies tend to be highly regular, hence network delays between any two VMs tend to be very similar across pairs of VMs. It follows that network delays are often, *within a single PoP*, often negligible with respect to processing times.

Next, Sec. III-A describes how all the entities that are involved in the VNF-sharing problem can be described, while Sec. III-B defines the problem objective and constraints.

#### A. System model

The system model includes VMs  $m \in \mathcal{M}$ , and VNFs  $v \in \mathcal{V}$ <sup>3</sup>. Each VM runs (at most) one VNF, modeled as an M/M/1 queue with FIFO queuing and preemption, as widely assumed in recent works [6], [7], [15]. Also, let  $C(m)$  be the maximum computation capability to which VM  $m$  can be scaled up. We underline that, although in this work we focus on computational capability, memory and storage could be easily accommodated as well. We refer to a VM as active if it hosts a VNF, and we express through binary variables  $y(v, m)$  whether VM  $m$  runs VNF  $v$ .

Different VNFs have different computational requirements, which are modeled through parameter  $l(v)$ , expressing how many units of computational capability are needed to process one request in one second. A VNF with requirement  $l(v) = 1$  running on a VM with capability  $\mu = 1$  takes  $l(v)/\mu = 1$  time unit to process a service request. Using the same VM for a VNF with requirement  $l(v) = 2$  yields a processing time of  $l(v)/\mu = 2$  time units. Notice how  $l(v)$  values do not depend on the service using VNF  $v$ .

Services  $s \in \mathcal{S}$  include one or more VNFs, and service  $s$  requests arrive at VNF  $v$  with rate  $\lambda(s, v)$ ; VNFs that are not used by a certain service have  $\lambda$ -values equal to 0. Through the  $\lambda(s, v)$  parameters, we can account for arbitrarily complex service (VNF) graphs where the number of requests can change between VNFs, and some requests may visit the same VNF more than once. We focus on the maximum average delay  $D^{\max}(s)$  of service  $s$  as target KPI, although our model can be seamlessly extended to account for different (or additional) KPIs, e.g., service request drop probability.

Each VM uses a quantity  $\mu(m)$  of computational capability, which can be scaled up till  $C(m)$ . Then the rate at which a VNF  $v$  deployed at VM  $m$  processes service requests results to be  $\frac{\mu(m)}{l(v)}$ . Finally, binary variables  $x(s, v, m)$  express whether service  $s$  uses the instance of VNF  $v$  at VM  $m$ ; this allows us to account for the fact that multiple instances of the same VNF can be deployed at different VMs. For clarity, the above parameters and variables are summarized in Tab. I.

#### B. Problem formulation

We now discuss the objective of the VNF-sharing problem and the constraints we need to honor.

**Objective.** The high-level goal of the MNO is to minimize the cost, which consists of two components: a fixed cost  $\kappa_f(m)$  paid if VM  $m$  is activated, and a proportional cost  $\kappa_p(m)$  paid for each unit of computational capability used therein. The objective is then given by:

$$\min_{y, \mu} \sum_{m \in \mathcal{M}} \left( \kappa_f(m) \sum_{v \in \mathcal{V}} y(v, m) + \kappa_p(m) \mu(m) \right). \quad (1)$$

<sup>3</sup>VNFs can, in general, include multiple virtual deployment units (VDUs); without loss of generality, we consider that each VNF includes only one VDU. Also, we assume that no VNF requires isolation.

TABLE I  
NOTATION († DENOTES VARIABLES OF THE MODIFIED PROBLEM  
DESCRIBED IN SEC. IV-B)

Symbol	Type	Meaning
$C(m)$	Parameter	Maximum capability to which VM $m$ can be scaled up
$D^{\max}(s)$	Parameter	Target delay for service $s$
$j$	Parameter	Jitter applied when assigning per-request priorities
$l(v)$	Parameter	Computational capability needed to process one request at VNF $v$
$\mathcal{M} = \{m\}$	Set	Set of VMs
$\pi(s, v)$	Random var.	Describes the priority assigned to requests of service $s$ upon entering VNF $v$
$p(s, v)$	Parameter	Per-VNF priority of service $s$ at VNF $v$
$S(s, v)$	Aux. var.	Sojourn time of requests of service $s$ at VNF $v$
$\mathcal{S} = \{s\}$	Set	Set of services
$\mathcal{V} = \{v\}$	Set	Set of VNFs
$x(s, v, m)$	Binary var.	Whether requests of service $s$ use the instance of VNF $v$ running at VM $m$
$y(v, m)$	Binary var.	Whether VM $m$ runs VNF $v$
$\kappa_f(m)$	Parameter	Fixed cost incurred when activating VM $m$
$\kappa_p(m)$	Parameter	Proportional cost incurred when using one unit of computational capability for VM $m$
$\Lambda(s, v)$	Aux. var.	Arrival rate at VNF $v$ of requests that are given priority over requests of service $s$
$\tilde{\Lambda}(s, v)$	Real var.†	Arrival rate at VNF $v$ of requests that are given priority over requests of service $s$
$\lambda(s, v)$	Parameter	Rate at which requests of service $s$ enter VNF $v$
$\mu(m)$	Real var.	Computational capability to use for VM $m$

**Capability and instance limits.** We must account for the maximum value  $C(m)$  to which the capability  $\mu(m)$  of each VM  $m$  can be scaled up:

$$\mu(m) \leq C(m) \quad \forall m \in \mathcal{M}. \quad (2)$$

Also, we can have at most one VNF per VM:

$$\sum_{v \in \mathcal{V}} y(v, m) \leq 1, \quad \forall m \in \mathcal{M}, \quad (3)$$

and only active VMs can be used, i.e.,

$$y(v, m) \geq x(s, v, m), \quad \forall s \in \mathcal{S}, v \in \mathcal{V}, m \in \mathcal{M}.$$

**Service times.** Each service  $s$  has a maximum average service time  $D^{\max}(s)$  that must be honored. Since, as discussed earlier, processing time is the dominant component of service time, this is equivalent to imposing:

$$\sum_{v \in \mathcal{V}} S(s, v) \leq D^{\max}(s), \quad \forall s \in \mathcal{S}, \quad (4)$$

where  $S(s, v)$  is the *sojourn time* (i.e., the time spent waiting or being served) experienced by requests of service  $s$  at VNF  $v$ . By convention, we set  $S(s, v) = 0$  if service  $s$  does not include VNF  $v$ .

As detailed below, sojourn times, in turn, depend on:

- the computational capability  $l(v)$  requested by the VNFs;
- the service request arrival rate at the VNFs  $\lambda(s, v)$ ;
- the priority of the service requests at the traversed VNFs;
- the computational capability  $\mu(m)$  assigned to the VM hosting the VNF instance processing the requests.

Using [13, Sec. 3.2] and [16], we can generalize the expression used in Example 1 and write the sojourn time of requests of service  $s$  at VNF  $v$ , as:

$$S(s, v) = \frac{l(v)}{\mu(\bar{m})} \frac{1}{1 - l(v) \frac{\Lambda(s, v)}{\mu(\bar{m})}} \frac{1}{1 - l(v) \frac{\Lambda(s, v) + \lambda(s, v)}{\mu(\bar{m})}}, \quad (5)$$

where  $\bar{m}$  is the VM hosting the instance of VNF  $v$  used by service  $s$ , i.e.,  $\bar{m} = m \in \mathcal{M}: x(s, v, m) = 1$ .

In (5),  $\Lambda(s, v)$  represents the arrival rate of requests (of any service) arriving at VNF  $v$  that are given a priority higher than a generic request of service  $s$ . Let  $\pi(s, v)$  be the random variable describing the priority assigned to requests of service  $s$  at VNF  $v$ , then we have:

$$\Lambda(s, v) = \sum_{t \in \mathcal{S}} \mathbb{P}(\pi(t, v) > \pi(s, v)) \lambda(t, v). \quad (6)$$

The intuitive meaning of (6) is that  $\Lambda(s, v)$  grows as it becomes more likely that requests of other services  $t \neq s$  are given higher priority over requests of service  $s$ .

**Problem complexity.** The actual expression of  $\Lambda(s, v)$  depends on the type of the  $\pi(s, v)$  variables and is not guaranteed to be linear, convex, or even continuous. It follows that, in the general case, the problem of setting the priorities in such a way to optimize (1) has NP-hard complexity and is thus impractical for realistically-sized problem instances. Indeed, as will be more clear from Sec. III-B1 and Sec. III-B2, finding the optimum in the per-request priority case, would require to search over all possible distributions of  $\pi(s, v)$ .

Below, we compute  $\Lambda(s, v)$  for the relevant cases of per-VNF priorities and uniformly-distributed, per-request priorities.

1) *Per-VNF priorities:* We recall that, if per-VNF priorities are supported as in Sec. II, then all requests of each service  $s$  entering VNF  $v$  are given the same, deterministic priority. It follows that, denoted such a priority with  $p(s, v)$ , in the per-VNF case  $\pi(s, v)$  is always distributed according to a Dirac delta function centered in  $p(s, v)$ , i.e.,  $\delta(\pi(s, v) - p(s, v))$ . Hence,  $\Lambda(s, v)$  is discontinuous and given by:

$$\Lambda(s, v) = \sum_{t \in \mathcal{S}} H(p(t, v) - p(s, v)) \lambda(t, v),$$

where  $H(\cdot)$  is the Heaviside step function. Indeed, intuitively a request of service  $s$  will be queued after all requests of services  $t$  with higher priority than  $s$  (since  $H(p(t, v) - p(s, v)) = 1$  if  $p(t, v) > p(s, v)$ ), after half of the requests of services with the same priority as  $s$  (since  $H(p(t, v) - p(s, v)) = 0.5$  if  $p(t, v) = p(s, v)$ ), and before all other requests (since  $H(p(t, v) - p(s, v)) = 0$  if  $p(t, v) < p(s, v)$ ).

2) *Per-request priorities:* This case corresponds to higher flexibility and implies that priorities could follow any distribution. Below, we focus on the simple, yet relevant, case where priorities are distributed uniformly between  $r(s, v) - j$  and  $r(s, v) + j$ . In this case, let us define the quantity  $q(s, t, v) = \mathbb{P}(\pi(t, v) > \pi(s, v))$ , whose value can be computed through the convolution of the pdfs of  $\pi(s, v)$  and  $\pi(t, v)$ . Following the steps in [17], we get:

$$\begin{aligned} q(s, t, v) &= \mathbb{P}(\pi(t, v) > \pi(s, v)) \\ &= \mathbb{P}(\pi(t, v) - \pi(s, v) > 0) \\ &= \begin{cases} 1 & \text{if } r(t, v) - r(s, v) > 2j \\ \frac{1}{2} + \frac{r(t, v) - r(s, v)}{4j} & \text{if } -2j \leq r(t, v) - r(s, v) \leq 2j \\ 0 & \text{if } r(t, v) - r(s, v) < -2j. \end{cases} \end{aligned} \quad (7)$$




Fig. 4. The FlexShare strategy. Step 1 builds a bipartite graph showing which VMs *could* run each VNF. Step 2 runs the Hungarian algorithm on such a graph, obtaining the optimal values for the  $x$ - and  $y$ -variables. Step 3 solves a convex variant of the original problem in Sec. III-B. If feasible, its variables  $(\mu, \bar{\Lambda})$  are used to determine the scaling and the service priorities; otherwise, the bipartite graph is *pruned* (step 4) and the procedure restarts from step 2.

Once the  $q(s, t, v)$  are known, the  $\Lambda(s, v)$  values can be computed by replacing (7) in (6), obtaining:

$$\Lambda(s, v) = \sum_{t \in \mathcal{S}} q(s, t, v) \lambda(t, v). \quad (8)$$

In this case, it is also possible to prove that the choice of parameter  $j$  has no impact on the solution space, hence, on the decisions that are made.

#### IV. THE FLEXSHARE SOLUTION STRATEGY

In light of the problem complexity, we propose a fast, yet highly effective, solution strategy, named FlexShare, which runs iteratively and consists of four main steps, as outlined in Fig. 4. The first step, detailed in Sec. IV-A, consists in building a bipartite graph including the VNFs to deploy, and the VMs that are active or can be activated. The edges of the graph express the possibility of using a VM to provide a VNF, either by sharing an existing instance or by deploying a new one. Edges are labeled with the cost associated with each decision, i.e., the  $\kappa_p$  and/or  $\kappa_f$  terms contributing to the objective (1). In step 2, the Hungarian algorithm [18] is run on the generated bipartite graph, yielding the optimal, minimum-cost assignment of VNFs to VMs, i.e., the  $x$ - and  $y$ -variables of the problem.

Given these decisions, step 3 aims at assigning the priorities and finding the amount of computational capability to use in every VM. To this end, a simpler (namely, *convex*) variant of the problem defined in Sec. III-B is formulated and solved, as detailed in Sec. IV-B.

If step 3 results in an infeasible problem, we *prune* the bipartite graph (step 4). The underlying intuition is that a cause for infeasibility is overly aggressive sharing of existing VNF instances. Therefore, as detailed in Sec. IV-C, edges that result in an overload of VMs are removed from the bipartite graph. After pruning, the algorithm starts a new iteration with step 2. Moving from one iteration to the next means reducing the likelihood that VNF instances are shared between services, and thus increasing the cost incurred by the MNO, due to the  $\kappa_f$  fixed cost terms. The procedure stops as soon as one feasible solution is found.

##### A. Steps 1–2: Bipartite graph and Hungarian algorithm

**The bipartite graph.** The purpose of the bipartite graph is to represent (i) the possible VNF assignment decisions, i.e., which VNFs can be provided at which VMs and which VNFs

can be shared among services, and (ii) the associated cost incurred by the MNO.

More formally, the bipartite graph is created according to the following rules:

- 1) a vertex is created for each VNF and for each VM;
- 2) an edge is drawn from every VNF to every unused VM;
- 3) an edge is drawn from every VNF to every VM currently running the same VNF, provided that the maximum computational capability of the VM is sufficient to guarantee stability.

Denoted by  $\bar{s}$  is the service now being deployed, VNF  $v$  can be provided at VM  $m$  while (still) ensuring stability if:

$$l(v) \sum_{s \in \mathcal{S}} [(x(s, v, m) + \mathbb{1}_{s=\bar{s}}) \lambda(s, v)] < C(m). \quad (9)$$

The first member of (9) is the total load imposed on VM  $m$  by services that are already being served therein, if any, *and* the new one  $\bar{s}$  (for which the indicator function is one). Through (9), we can then check that this quantity is lower than the maximum VM capability  $C(m)$ .

Note that (9) does *not* imply that  $\bar{s}$ , or existing services, can be served *in time*, i.e., within their deadlines; indeed, this depends on the priority and computational capability assignment decisions, and cannot be checked at the graph generation time. The purpose of step 1 is just to generate a graph accounting for all possible assignment options.

The cost of each edge connecting VM  $m$  with VNF  $v$  is given by the following expression:

$$\left(1 - \sum_v y(v, m)\right) \kappa_f(m) + \kappa_p(m) (l(v) \lambda(\bar{s}, v) + \epsilon). \quad (10)$$

In (10), the first term is the fixed cost associated with activating VM  $m$ , which is incurred only if  $m$  is not already active (the summation can be at most 1, as per (3)). The second term is the proportional cost associated with the additional computation capability needed at VM  $m$  to guarantee stability, with  $\epsilon$  being a positive, arbitrarily small value.

**Hungarian algorithm and assignment decisions.** The Hungarian algorithm [18] is a combinatorial optimization algorithm with polynomial (cubic) time complexity in the number of edges in the graph. When applied to the bipartite graph we generate, it selects a subset of edges such that (i) each VNF is connected to exactly one VM, and (ii) the total cost of the selected edges is minimized.

Selected edges map to assignment decisions. Specifically, for each selected edge connecting VNF  $v$  and VM  $m$ , we set  $y(v, m) \leftarrow 1$  and  $x(\bar{s}, v, m) \leftarrow 1$ , i.e., we activate  $m$  (if not already active) deploying therein an instance of  $v$ , and use it to serve service  $\bar{s}$ . The obtained values for the  $x$  and  $y$ -variables are used in step 3 to decide priorities and computational capability assignment, as set out next.

##### B. Step 3: Priority and scaling decisions

The purpose of step 3 of the FlexShare procedure is to decide the priorities to assign to each VNF and service, as well

as any needed scaling of VM computation capability. Since the complexity of the problem stated in Sec. III-B depends on the presence of the  $\pi(s, v)$  variables, we proceed as follows:

- 1) we formulate a *simplified* problem, which contains no random variables and is guaranteed to be convex;
- 2) we use the variables of the simplified problem to set the  $\mu(q)$  variables of the original problem, as well as the parameters of the distribution of the  $\pi(s, v)$  variables.

**Convex formulation.** To avoid dealing with probability distributions, we replace the  $\Lambda(s, v)$  auxiliary variables of the original problem with independent variables  $\tilde{\Lambda}(s, v)$ , thus dispensing with (6). Given  $x$  and  $y$ , the decision variables of the modified problem are  $\tilde{\Lambda}(s, v)$  and  $\mu(m)$ , while the objective is still given by (1). Having  $\tilde{\Lambda}(s, v)$  as a variable means deciding (intuitively) how many higher-priority service requests each incoming request will find. Such values are later mapped to the parameters of the distributions of  $\pi(s, v)$ .

If we solve the modified problem with no further changes, the optimal solution would always yield  $\tilde{\Lambda}(s, v) = 0, \forall s, v$ , i.e., no request ever encounters higher-priority ones, which is clearly not realistic. To avoid that, we mandate that the average behavior, i.e., the average number of higher-priority requests met, is the same as in the original problem:

$$\sum_{s \in \mathcal{S}} \tilde{\Lambda}(s, v) = \frac{|\mathcal{S}|}{2} \sum_{s \in \mathcal{S}} \lambda(s, v), \quad \forall v \in \mathcal{V}. \quad (11)$$

The intuition behind (11) is that each  $\Lambda(s, v)$ -value (in the original problem) is the sum of several  $\lambda$ -values, i.e., the services arrival rates. The  $\lambda$ -value associated with the highest-priority service will contribute to  $|\mathcal{S}| - 1$   $\Lambda(s, v)$ -values, the one associated with the second-highest-priority service will contribute to  $|\mathcal{S}| - 2$   $\Lambda(s, v)$ -values, and so on. On average, each  $\lambda$ -value contributes to  $\frac{|\mathcal{S}|}{2}$   $\Lambda(s, v)$ -values.

It can be proved that the modified problem is convex and, thus, solvable in polynomial time (in the problem size, which depends on the number of VNFs and VMs) through off-the-self, commercial solvers:

**Property 1.** *The problem of minimizing (1) subject to constraints (2)–(4) and (11), is convex.*

See [19] for the proof.

**Setting the variables of the original problem.** After solving the convex problem described above, we can use the optimal solution thereof to make scaling decisions, i.e., to set the  $\mu(m)$  variables in the original problem, as well as priority decisions, i.e., the parameters of the distribution of  $\pi(s, v)$ . For  $\mu(m)$ , we can simply use the corresponding variables in the simplified problem, which have the same meaning and are subject to the same constraints. As for priorities, the procedure to follow depends on the type of priority adopted in the system at hand, hence, on the type of the variables  $\pi(s, v)$ .

Specifically, if per-VNF priorities are supported (as in Sec. III-B1), we set the  $p(s, v)$  values in such a way that services associated with a higher  $\Lambda(s, v)$  have lower priority, e.g., by imposing that  $p(s, v) \leftarrow -\tilde{\Lambda}(s, v)$ . If, on the other

hand, we are in the case of Sec. III-B2, i.e., per-requests priorities are uniformly distributed, then we can solve a system of linear equations where the  $\tilde{\Lambda}(s, v)$  from the solution of the simplified problem are known terms, the  $q(s, t, v)$  quantities are the unknowns, and equations have the form (7) and (8).

Regardless the way priorities are assigned, it is important to stress that our approach has general validity and can be combined with *any type* of priority distribution.

### C. Step 4: graph pruning

If the problem we solve in step 3 (priority and scaling decisions) is infeasible, a possible cause lies in the decisions made in step 2 (the Hungarian algorithm), i.e., the  $x$  and  $y$  variables. Therefore, we restart from step 2 considering a different bipartite graph, more likely to result in a feasible problem.

To this end, we consider the *irreducible infeasible set* (IIS) of the problem instance solved in step 3, i.e., the set of constraints therein that, if removed, would yield a feasible problem. Given the IIS, we proceed as follows:

- 1) we identify constraints in the IIS of type (2), thus, a set of VMs that would need more capability;
- 2) among such VMs, we select those that are used by the newly-deployed service  $\bar{s}$ ;
- 3) among them, we identify the one that is the closest to instability, i.e., the VM  $m^*$  minimizing the quantity  $C(m) - \sum_{s \in \mathcal{S}} x(s, v^*, m) l(v^*) \lambda(s, v^*)$ , where  $v^*$  is the VNF deployed at  $m$ ;
- 4) we prune from the bipartite graph the edge between  $v^*$  and  $m^*$ .

The intuitive reason for this procedure is that a cause for delay constraints violations is that the newly-deployed service  $\bar{s}$  is causing one of the VMs it uses to operate too close to instability, and thus with high delays. By removing the corresponding edge from the bipartite graph, we ensure that VM  $m^*$  is not used by service  $\bar{s}$ .

Note that we are guaranteed that the IIS contains at least one constraint of type (2) thanks to the following result, proved in [19]:

**Theorem 1.** *Every infeasible instance of the modified problem presented in Sec. IV-B includes at least one constraint of type (2) in its IIS.*

FlexShare then restarts with step 2, where the Hungarian algorithm takes as an input the pruned bipartite graph.

### D. Computational complexity

The FlexShare strategy has polynomial worst-case computational complexity. Specifically:

- step 1 involves a simple check over at most  $|\mathcal{V}||\mathcal{M}|$  VNF/VM pairs;
- step 2, the Hungarian algorithm, has cubic complexity in the number of nodes in the graph [18];
- step 3 requires solving a convex problem, as proven in Property 1, and the resulting complexity is also cubic;




Fig. 5. VNF graphs in the synthetic scenario.

- step 4 iterates over at most  $|\mathcal{M}|$  constraints of type (2), and thus it has linear complexity;
- the whole procedure is repeated for (at most) as many times as there are edges in the original bipartite graph.

## V. NUMERICAL RESULTS

In this section, we describe the reference scenarios and benchmark solutions we consider, in Sec. V-A; then we present numerical results obtained under the synthetic and realistic scenarios in, respectively, Sec. V-B and Sec. V-C.

### A. Reference scenarios and benchmarks

**Synthetic scenario.** It includes three services  $s_1 \dots s_3$ , sharing five VNFs  $v_1 \dots v_5$ , as depicted in Fig. 5. All VNFs have coefficient  $l(v) = 10^{-3}$  units/request, while the arrival rates associated with each service vary between 1 and 2 requests/ms. Target delays range between 20 ms for  $s_1$  and 5 ms for  $s_3$ . The scenario includes 10 VMs, whose fixed and proportional costs are  $\kappa_f = 8$  and  $\kappa_p = 0.5$  units, respectively, and whose capability is randomly distributed between 5 and 10 units. Such a scenario is small enough to allow a comparison against optimal priority assignments found by brute-force. At the same time, it contains many interesting features, including different combinations of services sharing different VNFs and different cost/capability trade-offs at VMs.

**Realistic scenario.** We consider three services, all connected to the smart-city domain:

- intersection collision avoidance (ICA): vehicles periodically broadcast a message (e.g., CAM) including their position and speed; a collision detector checks if any pair of them are on a collision course and, if so, it issues an alert;
- vehicular see-through (CT): cars display on their on-board screen the video captured by the preceding vehicle, e.g., a large truck obstructing the view;
- urban sensing based on the Internet-of-Things (IoT).

Tab. II, based on [20], [21], reports the VNFs used by each service and the associated arrival rates. All services share the EPC child service, which is itself composed of five VNFs. Furthermore, the car information management (CIM) database can be shared between the ICA and the CT services.

We leverage the real-world mobility trace [22], to assign user density and request rates. Focusing on a downtown intersection, we consider that (i) all vehicles within 50 m from the intersection are users of the ICA service, and send a CAM every 0.1 s; (ii) all vehicles within 100 m from the intersection are users of the CT service, and send a request (i.e., refresh their video) every 200 ms; (iii) a total of 200 sensors are deployed in the area, each generating, according to the traffic model described in the 3GPP standard [23], one request

TABLE II  
REALISTIC SCENARIO: REQUEST RATE AND COMPUTATIONAL LOAD ASSOCIATED WITH EVERY VNF

VNF	Rate $\lambda(s)$	Requirement $l(v)$
<i>Intersection collision avoidance (ICA)</i>		
eNB	117.69	$10^{-4}$
EPC PGW	117.69	$10^{-4}$
EPC SGW	117.69	$10^{-4}$
EPC HSS	11.77	$10^{-4}$
EPC MME	11.77	$10^{-3}$
Car information management (CIM)	117.69	$10^{-3}$
Collision detector	117.69	$10^{-3}$
Car manufacturer database	117.69	$10^{-4}$
Alarm generator	11.77	$10^{-4}$
<i>See through (CT)</i>		
eNB	179.82	$10^{-4}$
EPC PGW	179.82	$10^{-4}$
EPC SGW	179.82	$10^{-4}$
EPC HSS	17.98	$10^{-4}$
EPC MME	17.98	$10^{-3}$
Car information management (CIM)	179.82	$10^{-3}$
CT server	179.82	$5 \cdot 10^{-3}$
CT database	17.98	$10^{-4}$
<i>Sensing (IoT)</i>		
eNB	50	$10^{-4}$
EPC PGW	50	$10^{-4}$
EPC SGW	50	$10^{-4}$
EPC HSS	5	$10^{-4}$
EPC MME	5	$10^{-3}$
IoT authentication	20	$10^{-4}$
IoT application server	20	$10^{-3}$

every 0.1 s. Tab. II reports the resulting request rates and the requirement  $l(v)$  associated with each VNF. As discussed in Sec. III, the  $\lambda(s, v)$  values also incorporate the fact that not all requests visit all VNFs of a service, e.g., all ICA requests visit the collision detector but only one in ten needs the alarm generator.

Finally, we assume that the PoP contains 10 VMs, each of which can be scaled up to at most  $C(m) = 1000$  units, and each associated with fixed and proportional costs of  $\kappa_f = 1000$  units and  $\kappa_p = 1$  unit, respectively.

**Benchmark solutions.** We study the performance of the following strategies, in increasing order of flexibility:

- **service**: per-service priorities, with lowest-delay services having the highest priority;
- **VNF/FS**: per-VNF priorities, assigned with FlexShare;
- **VNF/brute**: optimal per-VNF priorities (with brute-force);
- **req/FS**: per-request, uniformly distributed priorities, assigned with FlexShare.

### B. Results: synthetic scenario

We start by considering the synthetic scenario and, in order to study different traffic conditions, multiply the arrival rates by a factor  $n$  ranging between 1 and 2.

Fig. 6(left) focuses on the main metric we consider, namely, the total cost incurred by the MNO. We can observe that, as one might expect, higher traffic translates into higher cost. More importantly, more flexibility in priority assignment results in substantial cost savings. As for per-VNF priorities, they exhibit an intermediate behavior between per-service and per-request ones, with virtually no difference between the case where FlexShare is used to determine the priorities




Fig. 6. Synthetic scenario: total cost (left); average number of services sharing a VNF instance (center); used and maximum VM capability (right).




Fig. 7. Synthetic scenario,  $n = 1.8$ : priorities assigned to each service with per-service (left), per-VNF (center), and per-request priorities (right).

(“VNF/FS”) and that where all possible options are tried out in a brute-force fashion (“VNF/brute”). This highlights the effectiveness of the FlexShare strategy, which can make optimal decisions in almost all cases with low complexity.

Fig. 6(center) shows the average number of services sharing a VNF instance. It is clear that a higher flexibility in priority assignment results in more sharing, hence fewer VNF instances deployed. As  $n$  increases, the number of services per instance decreases: scaling up (i.e., increasing the capability of VMs) is insufficient, and scaling out (i.e., increasing the number of instances) becomes necessary.

This is confirmed by Fig. 6(right), depicting the total used VM capability (i.e.,  $\sum_{m \in \mathcal{M}} \mu(m)$ ) as well as the sum of the maximum values to which the capability of active VMs can be scaled up (i.e.,  $\sum_{m \in \mathcal{M}} C(m)y(m)$ ), denoted by solid and dotted lines, respectively. Both quantities grow with  $n$  and decrease as flexibility becomes higher. This makes intuitive sense for the maximum capability: Fig. 6(left) shows that under higher-flexibility strategies fewer VNF instances, hence fewer active VMs, are needed. Importantly, *used* capability values, i.e.,  $\mu(m)$ , also decrease with flexibility. Indeed, higher flexibility makes it easier to match the computational capability obtained by each service within each VNF, with its needs.

Fig. 7 provides a qualitative view of how priorities are assigned to different services across different VNF instances. When priorities are assigned on a per-service basis (Fig. 7(left)), services with lower target delay invariably have higher priority. If priorities are assigned on a per-VNF basis, as in Fig. 7(center), the priorities of different services can change across VNF instances, e.g.,  $s_2$  has priority over  $s_1$  in the  $v_1$  instance deployed at VM  $m_2$ , but the opposite happens in the  $v_3$  instance deployed at VM  $m_1$ . Fig. 7(right) shows that if per-request priorities are possible, services can be combined in any way at each VNF instance.

### C. Results: realistic scenario

We now move to the realistic scenario, again multiplying the arrival rates reported in Tab. II by a factor  $n$  varying between 1 and 2. Recall that, owing to the larger scenario size, no comparison with the brute-force strategy is possible.

Fig. 8(left) shows how the total cost yielded by the different strategies has the same behavior as in the synthetic scenario (Fig. 6(left)): the higher the flexibility, the lower the cost. Furthermore, for very high values of  $n$ , all strategies yield the same cost; in those cases, few or no VNF instances can be shared, regardless of how priorities are assigned.

Fig. 8(center) shows that VNF instances are shared among services; by comparing it to Fig. 6(center) we can observe how the behavior of per-VNF priorities tends to be closer to per-request priorities than to per-service ones. This suggests that, even in large and/or complex scenarios, per-VNF priorities can be a good compromise between performance and implementation complexity.

Fig. 8(right) shows a much larger difference between used and maximum capabilities compared to Fig. 6(right). This is due to the fact that, as can be seen from Tab. II, there are fewer VNFs that are common among different services, and thus fewer opportunities for sharing.

## VI. RELATED WORK

5G networks based on network slicing have attracted substantial attention, with several works focusing on 5G architecture [2], [24], associated decision-making issues [25], and security [26].

As one of the most important decisions to make in 5G environments, VNF placement has been the focus of several studies. One popular approach is optimizing a network-centric metric, e.g., load balancing [27] or network utilization [28].






Fig. 8. Realistic scenario: total cost (left); average number of services sharing a VNF instance (center); used and maximum VM capability (right).

Other papers use cost functions, e.g., [29], possibly including energy-efficiency considerations [30].

The aforementioned works typically result in mixed-integer linear programming (MILP) models. Other works cast VNF placement into a generalized assignment [31] or a set cover problem [32].

**Novelty.** A first novel aspect of our work is the problem we consider, i.e., VNF-sharing within one PoP as opposed to traditional VNF placement. From the modeling viewpoint, we depart from existing works in three main ways: (i) priorities are used as a decision variable rather than as an input; (ii) different priority-assignment schemes with different flexibility are accounted for and compared; (iii) the relationship between the amount of computational resources assigned to VNFs and their performance is modeled and studied; (iv) VM capacity scaling is properly accounted for as a necessary, complementary aspect of VNF sharing.

## VII. CONCLUSION

We considered the problem of sharing VNFs among 5G services using the same PoP, and identified in service priority management one of its key aspects. We then proposed a solution strategy called FlexShare, able to efficiently make high-quality priority and VM scaling decisions. Our performance evaluation has shown that higher flexibility in priority assignment yields lower costs, and that FlexShare is able to provide near-optimal performance in all scenarios.

Future work will focus on formally characterizing the performance gap between FlexShare and the optimal solution. Furthermore, with reference to the per-request priority case, we will study additional distributions and assess their impact on the resulting performance.

## REFERENCES

- [1] NGMN Alliance, "Description of network slicing concept," 2016.
- [2] P. Rost *et al.*, "Network slicing to enable scalability and flexibility in 5G mobile networks," *IEEE Comm. Mag.*, 2017.
- [3] 5G PPP Architecture Working Group. (2017) View on 5G Architecture.
- [4] IETF. (2017) Network Slicing Management and Orchestration.
- [5] J. Cao *et al.*, "VNF placement in hybrid NFV environment: Modeling and genetic algorithms," in *IEEE ICPADS*.
- [6] R. Cohen *et al.*, "Near optimal placement of virtual network functions," in *IEEE INFOCOM*, 2015.
- [7] S. Agarwal *et al.*, "Joint VNF Placement and CPU Allocation in 5G," in *IEEE INFOCOM*, 2018.
- [8] ETSI. (2017) Network Functions Virtualisation (NFV); Management and Orchestration.

- [9] ——. IFA 007: Network Functions Virtualisation (NFV); Management and Orchestration; Or-Vnfm reference point – Interface and Information Model Specification.
- [10] K. Antevski *et al.*, "Resource orchestration of 5G transport networks for vertical industries," in *IEEE PIMRC*, 2018, <http://arxiv.org/abs/1807.10430>.
- [11] B. Sayadi *et al.*, "SDN for 5G Mobile Networks: NORMA perspective," in *Springer CROWNCOM*, 2016.
- [12] NGMN Alliance. (2017) 5G Network and Service Management including Orchestration.
- [13] L. Kleinrock, *Queueing systems: Computer applications*. John Wiley & Sons, 1976.
- [14] W. Xia *et al.*, "A survey on data center networking (DCN): Infrastructure and operations," *IEEE Comm. surveys & tutorials*, 2017.
- [15] D. Bhamare *et al.*, "Optimal virtual network function placement in multi-cloud service function chaining architecture," *Computer Communications*, 2017.
- [16] Malathi Veeraraghavan. (2014) Priority queueing. <http://www.ece.virginia.edu/mv/edu/715/lectures/PQ.pdf>.
- [17] Dimitrios Milios. (2009) Probability Distributions as Program Variables. <http://www.inf.ed.ac.uk/publications/thesis/online/IM090722.pdf>.
- [18] H. W. Kuhn, "The hungarian method for the assignment problem," *Wiley Naval Research Logistics*, 1955.
- [19] Proofs. <https://dl.dropbox.com/s/9ylmq71iyyuqjrk/proofs.pdf>.
- [20] C. Casetti *et al.*, "Arbitration among vertical services," in *IEEE PIMRC*, 2018, <http://arxiv.org/abs/1807.11196>.
- [21] T. Taleb, A. Ksentini, and A. Kobbane, "Lightweight mobile core networks for machine type communications," *IEEE Access*, 2014.
- [22] L. Codeca, R. Frank, and T. Engel, "Luxembourg SUMO Traffic (LuST) Scenario: 24 hours of mobility for vehicular networking research," in *IEEE VNC*, 2015.
- [23] 3rd Generation Partnership Project. "3GPP specification: 37.868; RAN improvements for machine-type communications," Tech. Rep.
- [24] H. Zhang *et al.*, "Network slicing based 5G and future mobile networks: mobility, resource management, and challenges," *IEEE Comm. Mag.*, 2017.
- [25] K. Samdanis *et al.*, "5G Network Slicing – Part 2: Algorithms and practice," *IEEE Comm. Mag.*, 2017.
- [26] M. A. S. Santos *et al.*, *Security Requirements for Multi-operator Virtualized Network and Service Orchestration for 5G*. Springer, 2017.
- [27] A. Hirwe and K. Kataoka, "LightChain: A lightweight optimization of VNF placement for service chaining in NFV," in *IEEE NetSoft*, 2016.
- [28] T. W. Kuo *et al.*, "Deploying chains of virtual network functions: On the relation between link and server usage," in *IEEE INFOCOM*, 2016.
- [29] M. Mechtri, C. Ghribi, and D. Zeghlache, "A scalable algorithm for the placement of service function chains," *IEEE Trans. on Network and Service Management*, 2016.
- [30] A. Marotta and A. Kassler, "A power efficient and robust virtual network functions placement problem," in *IEEE ITC*, 2016.
- [31] R. Cohen *et al.*, "Near optimal placement of virtual network functions," in *IEEE INFOCOM*, 2015.
- [32] A. Tomassilli *et al.*, "Provably efficient algorithms for placement of service function chains with ordering constraints," in *IEEE INFOCOM*, 2018.