



AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Dynamic control of NFV forwarding graphs with end-to-end deadline constraints

This is a pre print version of the following article:		
Original Citation:		
Availability:		
This version is available http://hdl.handle.net/2318/1662774	since 2018-03-19T10:06:04Z	
Publisher:		
Institute of Electrical and Electronics Engineers Inc.		
Published version:		
DOI:10.1109/ICC.2017.7996596		
Terms of use:		
Open Access		
Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.		

(Article begins on next page)

Dynamic control of NFV forwarding graphs with end-to-end deadline constraints

Victor Millnert^{*}, Johan Eker^{*†}, Enrico Bini[‡] *Lund University, Sweden [†]Ericsson Research, Sweden [‡]University of Turin, Italy

Abstract-There is a strong industrial drive to use cloud computing technologies and concepts for providing timing sensitive services in the networking domain since it would provide the means to share the physical resources among multiple users and thus increase the elasticity and reduce the costs. In this work, we develop a mathematical model for user-stateless virtual network functions forming a forwarding graph. The model captures uncertainties of the performance of these virtual resources as well as the time-overhead needed to instantiate them. The model is used to derive a service controller for horizontal scaling of the virtual resources as well as an admission controller that guarantees that packets exiting the forwarding graph meet their end-to-end deadline. The Automatic Service and Admission Controller (AutoSAC) developed in this work uses feedback and feedforward making it robust against uncertainties of the underlying infrastructure. Also, it has a fast reaction time to changes in the input.

1. Introduction

Over the last years, cloud computing has swiftly transformed the IT infrastructure landscape, leading to large costsavings for deployment of a wide range of IT applications. Physical resources such as compute nodes, storage nodes, and network fabrics are shared among tenants through the use of virtual resources. This makes it possible to dynamically change the amount of resources allocated to a tenant, for example as a function of workload or cost. Initially the cloud technology was mostly used for IT applications, e.g. web servers, databases, etc., but has now found its way into new domains. One such domain is packages processed by a chain of network functions.

In this work, we are considering a chain of connected network functions through which packets are flowing. Every packet must be processed by each function in the chain within some specific end-to-end deadline. The goal is to ensure that as many packets as possible meet their deadline, while at the same time using as few resources as possible.



Figure 1: Traffic flowing through a switch over 120 hours. The traffic is normalized to have a peak of 10 million packets per second



Figure 2: Several virtual networking functions (VNF) are connected together to provide a set of services. A packet flow is a specific path through the VNFs. Connected VNFs are referred to as virtual forwarding graphs or service chains. The VNFs are mapped onto physical hardware, i.e. compute nodes and network fabrics and this underlying hardware infrastructure is referred to as Network Function Virtualization Infrastructure (NFVI), which is the physical servers and the communication fabric connecting them.

The goal is thus to derive a method for controlling the amount of resources allocated to each network function in the chain. Previously, this has been done by statically allocating some amount of resources to each network function. Since the input is time-varying (see Figure 1 for a trace of traffic flowing through a switch in the Swedish university network, SUNET), such a strategy usually leads to over-allocation of resources for long periods of time (yielding high costs and environmental footprint) as well as overload for shorter periods, when the input is large. To ensure that at least some packets meet their deadlines when the network function is overloaded, one has to use *admission control*, i.e. reject some packets.

Recently a new option has been made available through the advances of virtualization technology for networking services. The standardization body ETSI (European Telecommunications Standards Institute) addresses the standardization of these virtual network services under the name Network Functions Virtualisation (NFV) [1]. These Virtual Network Functions (VNFs) consist of virtual resources, such as virtual machines (VMs), containers, or even processes running in the OS. Using such VNFs it is possible to change the resources allocated to a network function by either vertical scaling (i.e., changing the capacity of the allocated VMs) or horizontal scaling (i.e., changing the number of allocated VMs). Horizontal scaling is considered in this work. These VNFs are connected forming what is commonly called a Forwarding Graph, as illustrated in Figure 2. In this figure, there are two forwarding graphs (corresponding to the blue and the red arrow). The blue forwarding graph consists of VNF_1 , VNF_2 , VNF_3 , and VNF_5 and the red forwarding graph consists of VNF₁, VNF₂, VNF₄, and VNF₅. Each of

the VNF is given a number $m_i \in \mathbb{Z}^+$ of VMs, which are mapped onto the network function virtual infrastructure.

- While the benefit of using NFV technologies is scalability and resource sharing there are two drawbacks:
- a) *starting a new virtual resource takes time*, since it has to be deployed to a physical server and it requires the execution of several initialization scripts and push/pulls before it is ready to serve packets,
- b) *the true performance of the virtual resource differs from the expected performance*, since one does not know what else is running on the physical machines.

In this work we:

- develop a model of a service-chain of network functions and use it to derive a service-controller and admission-controller for the network functions,
- derive a service-controller controlling the number of virtual resources (e.g., VMs or containers) allocated to each network function by using feedback from the true performance of the instances as well as feedforward between the network functions,
- derive an admission-controller that is aware of the actions of the service-controller which it uses in order to reject as few packets as possible,
- evaluate the service- and admission controller using a real-world traffic trace from the Swedish University Network (SUNET).

Related works

There are a number of works considering the problem of controlling virtual resources within data centers, and specifically for virtual network functions. However many of them focus on orchestration, i.e. how the virtual resources should be mapped onto the physical hardware. Shen et al [2] develop a management framework, vConductor, for realizing end-toend virtual network services. In [3], Moens and De Turk develop a formal model for resource allocation of virtual network functions. A slightly different approach is taken by Mehraghdam et al [4] where they define a model for formalizing the chaining of forwarding graphs using a context-free language. They solve the mapping of the forwarding graphs onto the hardware by posing it as a MIQCP.

Scaling of virtual network functions is however studied by Mao et al [5] where they develop a mechanism for auto-scaling resources in order to meet some user specified performance goal. Recently, Wang et al [6] developed a fast online algorithm for scaling and provisioning VNFs in a data center. However, they are not considering timingsensitive applications with deadlines for the packets moving through the chain, which is done by Li et al [7] where they present a design and implementation of NFV-RT that aims at controlling NFVs with soft Real-Time guarantees, allowing packets to have deadlines.

The enforcement of an end-to-end deadline for a sequence of jobs is however addressed by by several works, possibly under different terminologies. Di Natale and Stankovic [8] propose to split the E2E deadline proportionally to the local computation time or to divide equally the slack time. Later, Jiang [9] used time slices to decouple the schedulability analysis of each node, reducing the complexity of the analysis. Such an approach improves the robustness of the schedule, and allows to analyse each pipeline in isolation. Serreli et al. [10], [11] proposed to assign local deadlines to minimise a linear upper bound of the resulting local demand bound functions. More recently, Hong et al [12] formulated the local deadline assignment problem as a MILP with the goal of maximizing the slack time.

An alternate analysis was proposed by Jayachandran and Abdelzaher [13], who developed several transformations to reduce the analysis of a distributed system to the single processor case. Or in [14] where Henriksson et al. proposed a feedforward/feedback controller to adjust the processing speed to match a given delay target.

2. Modeling the service-chain



Figure 3: Illustration of the service-chain.

In this section, we present an general model of the forwarding graph and virtual network functions presented in Section 1. We consider a *service-chain* consisting of *n* functions F_1, \ldots, F_n , as illustrated in Figure 3. Packets are flowing through the service-chain and they must be processed by each function in the chain within some end-toend deadline. A *fluid model* is used to approximate the packet flow and at time *t* there are $r_i(t) \in \mathbb{R}^+$ packets per second (pps) entering the *i*'th function. In a recent benchmarking study it was shown that a typical virtual machine can process around 0.1–2.8 million packets per second, [15]. Hence, in this work the number of packets flowing through the functions is assumed to be in the order of millions of packets per second, supporting the use of a fluid model.

A function consists of several parts, as illustrated in Figure 4: an *admission controller*, a *service controller*, $m_i(t)$ *instances*, a *buffer*, and a *load balancer*. It is assumed that all the parts of a function is located on the same location, e.g., the same data center or rack. In [16], Google showed that less than 1% (<1 μ s) of the latency in a data center was due to the propagation in the network fabric. Hence, communication delay within a function is neglected. The functions are also assumed to be user-stateless.



Figure 4: Illustration of the structure and different entities of the function.

2.1. Admission controller

Every packet that enters the service-chain must be processed by all of the functions in the chain within a certain end-to-end (E2E) deadline, denoted D^{\max} . This deadline can be split into local deadlines $D_i(t)$, one for each function in the chain, such that the packet should not spend more than $D_i(t)$ time-units in the *i*'th function. Should a packet miss its E2E deadline it is considered useless. It is thus favorable to use admission control to drop packets that have a high probability of missing their deadline in order to make room for following packets. The goal of the admission controller is to guarantee that the packets that make it through the service-chain do meet their E2E deadline. It is assumed to be possible to do admission control at the entry of every function in the chain.

Packets are admitted into the buffer of the function based on the *admittance flag* $\alpha_i(t) \in \{0, 1\}$. If $\alpha_i(t) = 1$ incoming packets are admitted into the buffer, and if $\alpha_i(t) = 0$ they are rejected. We define the *residual rate* $\rho_i(t)$ to be the rate by which packets are admitted into the buffer:

$$\rho_i(t) = r_i(t) \times \alpha_i(t). \tag{1}$$

2.2. Service controller

At any time instance, function F_i has $m_i(t) \in \mathbb{Z}^+$ instances up and running. Each instance is capable of processing packets and corresponds to a virtual machine, a container, or a process running in the OS. It is possible to control the number of running instances by sending a reference signal $m_i^{\text{ref}}(t) \in \mathbb{Z}^+$ to the service controller. However, as explained in Section 1, it takes some time to start/stop instances since an instantiation of the service is always needed. We denote this as the *time overhead* Δ_i . Hence, the number of instances running in the *i*'th function at time *t* is

$$m_i(t) = m_i^{\text{ref}}(t - \Delta_i).$$
⁽²⁾

The time-overhead is assumed to be symmetric here, but in the real-world it is usually faster to start an instance than it is to stop one. However, for increased readability they are considered equal in this work. It should be noted that it is straight forward to extend the theory to account for an asymmetric time-overhead.

An instance is expected to be able to process packets at an *expected service rate* of \bar{s}_i pps. However, as described in Section 1, the true capacity of the instance will differ from the expected one since there might be other loads running on the infrastructure (i.e. the *physical machine*). Hence, the *true capacity* of the *j*'th instance in the *i*'th function is given by

$$s_{i,j}^{\operatorname{cap}}(t) = \bar{s}_i + \xi_{i,j}(t)$$

where $\xi_{i,j}(t)$ is the *machine uncertainty* for the *j*'th instance in the *i*'th function. It is given by

$$\xi_{i,j}(t) \in [\xi_i^{\mathrm{lb}},\,\xi_i^{\mathrm{ub}}] \text{ pps}, \quad -\bar{s}_i < \xi_i^{\mathrm{lb}} \leq \xi_i^{\mathrm{ub}} < \infty,$$

where ξ_i^{lb} and ξ_i^{ub} are lower and upper bounds of this machine uncertainty, assumed to be known. The machine uncertainty is also assumed to be fairly constant during the lifetime of the instance. Using this, one can express the true capacity of the *i*'th function in the service-chain as

$$s_i^{\rm cap}(t) = \sum_{j=1}^{m_i(t)} \bar{s}_i + \xi_{i,j}(t),$$
(3)

which together with the average machine uncertainty

$$\hat{\xi}_i(t) = \frac{1}{m_i(t)} \sum_{j=1}^{m_i(t)} \xi_{i,j}(t),$$
(4)

can be written as $s_i^{\text{cap}}(t) = m_i(t) \times (\bar{s}_i + \hat{\xi}_i(t))$. Note that it would be natural to allow the time-overhead Δ_i to also have some uncertainty. However, such uncertainty can be translated into a machine uncertainty.

2.3. Processing of packets

The packets in the buffer are stored and processed in a FIFO manner. Once a packet reaches the head of the queue the load balancer will distribute it to one of the instances in the function. Note that this is done continuously due to the fluid approximation. The rate by which the load balancer is distributing packets, and thus by which the function processing packets, is defined as the *service rate*

$$s_i(t) = \begin{cases} \rho_i(t) & \text{if } q_i(t) = 0 \text{ and } \rho_i(t) \le s_i^{\text{cap}}(t) \\ s_i^{\text{cap}}(t) & \text{else} \end{cases}$$
(5)

where $\rho_i(t)$ is residual rate given by (1) and $q_i(t)$ is the number of packets in the buffer:

$$q_i(t) = P_i(t) - S_i(t), \quad q_i(t) \in \mathbb{R}^+, \tag{6}$$

where $P_i(t) = \int_0^t \rho_i(x) dx$ is the total amount of packets that has been admitted into function F_i , and $S_i(t) = \int_0^t s_i(x) dx$ is the total amount of packets that has been served by function F_i . Furthermore, the total amount of packets that has reached the *i*'th function is given by $R_i(t) = \int_0^t r_i(x) dx$.

2.4. Function delay

The time that a packet that exists function F_i at time t has spent inside that function is denoted the *function delay* $d_i(t)$:

$$d_i(t) = \inf\{\tau \ge 0 : P_i(t - \tau) \le S_i(t)\}.$$
 (7)

The expected time that a packet entering the *i*'th function at time t will spend in the function before exiting is defined as the *expected function-delay* $\bar{d}_i(t)$

$$\bar{d}_{i}(t) = \inf\{\tau \ge 0: \\ P_{i}(t) \le S_{i}(t) + \int_{t}^{t+\tau} m_{i}(x) \times (\bar{s}_{i} + \hat{\xi}_{i}(x)) \mathrm{d}x \}.$$
(8)

Equation (8) can be interpreted as finding the minimum time $\tau \ge 0$ such that $S_i(t + \tau) = P_i(t)$, or in other words such that at time $t + \tau$ the function will have processed all the packets that has entered the function at time t. Finally, it is also of interest to find the upper bound of the expected function delay:

$$d_{i}^{ub}(t) = \inf\{\tau \ge 0: \\ P_{i}(t) \le S_{i}(t) + \int_{t}^{t+\tau} m_{i}(x) \times (\bar{s}_{i} + \xi_{i}^{lb}) dx\}.$$
(9)

This can be interpreted as the expected delay in the worst case, i.e. when every instance process packets at the lower bound of the possible service-rate, hence leading to the upper bound on the expected delay.

Computing the expected function-delay $\bar{d}_i(t)$ requires information about $m_i(t)$ and $\hat{\xi}_i(t)$ for the future, whereas computing the expected function delay $d_i^{\text{ub}}(t)$ requires information about $m_i(t)$ for the future. Information about $m_i(t)$ up until time $t+\Delta_i$ is always known since $m_i(t+\Delta_i) = m_i^{\text{ref}}(t)$ and $m_i^{\text{ref}}(x)$ is known for $x \in [0, t]$. It is therefore possible to compute the expected function delay $\bar{d}_i(t)$ whenever it is shorter than the time-overhead Δ_i (which will be used later in Section 3 when deriving the admission controller and the service controller).

Note that the (expected) function delay does not distinguish between queueing delay and processing delay. In [16], Google profiled where the latency in a data center occurred and showed that 99% of the latency ($\approx 85\mu s$) occurred somewhere in the kernel, the switches, the memory, or the application. It is very difficult to say exactly which of this 99% is due to processing or queueing, hence they are considered together as the function delay.

2.5. Concatenation of functions

The *n* functions in the service-chain are concatenated with the assumption of no loss of packets in the communication channel between them. Therefore the input of function F_i is exactly the output of function F_{i-1} :

$$r_i(t) = s_{i-1}(t), \quad \forall i = 2, 3, \dots, n.$$

Finally, no communication latency between functions are assumed. However, it is possible to account for it, and would be necessary should the different functions reside in different locations, i.e. different data centers. However, adding a communication latency is straightforward, and if such communication latency (say C) were to be constant between the functions one could easily account for it by properly decrementing the end-to-end deadline: $\tilde{D}^{\max} = D^{\max} - C$, and then use the framework developed in this paper.

2.6. Problem formulation

The goal of this paper is to derive a service-controller and an admission-controller that guarantees that packets that pass through the service-chain meet their E2E deadline. This should be done using as few resources as possible while still achieving as high throughput as possible. This is captured in a simple, yet intuitive *utility function* $u_i(t)$. Later in Section 3, the utility function is used to derive an automatic service- and admission controller, denoted AutoSAC.

Utility function. The utility function measures the *availability* $a_i(t)$ and the *efficiency* $e_i(t)$ of the each function in the service chain. The availability is defined as the ratio between the service-rate and the input-rate of the function, and the efficiency is defined as the ratio between service-rate and the capacity of the function:

$$a_i(t) = \frac{\text{service}}{\text{demand}} = \frac{s_i(t)}{r_i(t - d_i(t))} \in [0, 1 + \epsilon], \quad (10)$$

$$e_i(t) = \frac{\text{service}}{\text{capacity}} = \frac{s_i(t)}{s_i^{\text{cap}}(t)} \in [0, 1].$$
(11)

The reason why $a_i(t)$ can grow greater than 1 is due to the buffer—it is possible to store packets for a short interval and then process them at a rate greater than what they arrived with. However, it is not possible to have $a_i(t) > 1$ for an infinite amount of time. In practice, ϵ is very small, and it is not possible to achieve a $a_i(t) > 1$ for any significant period of time.

A low availability corresponds to a large percentage of the incoming load being rejected by the admission controller, since there is not enough capacity to serve them. A low efficiency, on the other hand, correspond to over-provisioning of resources. It is therefore difficult to achieve both a high availability and a high efficiency. The utility function $u_i(t)$ captures this, by combining the availability and efficiency

$$u_i(t) = a_i(t) \times e_i(t) = \frac{s_i^2(t)}{s_i^{\text{cap}}(t) \times r_i(t - d_i(t))}.$$
 (12)

Note that the utility function, as well as the availability and efficiency function, all have the good property of being normalized making it easy to compare the performance of service-chains having different input load.

To evaluate the performance between service-chains of different lengths and over different time-horizons the *average utility* U(t) is defined:

$$u(t) = \frac{1}{n} \sum_{i=1}^{n} u_i(t), \quad U(t) = \frac{1}{t} \int_0^t u(x) dx.$$
(13)

3. Controller design

In this section an automatic service- and admissioncontroller (AutoSAC) is derived. The admission controller is then derived in Section 3.1 and the service controller in Section 3.2. In Section 3.3 a short discussion of the properties of AutoSAC is presented.

The difficulty when deriving AutoSAC lies in the different time-scales for starting/stopping instances, the E2E deadlines, and the rate-of-change of the input. They are all assumed to be of different orders of magnitudes, given by Table 1. However, these timing assumptions will be exploited when deriving AutoSAC later.

Parameter	timing assumption	
long-term trend change of the input	1min – 1h	
service-rate change overhead Δ_i	1s – 1min	
request end-to-end deadline D^{\max}	1ms - 100 ms	

TABLE 1: Timing assumptions for the the end-to-end deadline, the changeof-rate of the input, and the overhead for changing the service-rate. These timing assumptions are used when deriving the automatic service- and admissioncontroller.

3.1. Admission controller

Every request that enters the service chain has an *end-toend deadline* D^{\max} . It has to pass through every function in the chain within this time. Furthermore, each function can impose a *local deadline* $D_i(t)$ for the packet entering the *i*'th function at time *t*.

Using the local deadline each function can evaluate the upper bound of the expected delay, $d_i^{\rm ub}(t)$ for a new packet. If this is longer than the local deadline the admission controller should drop the packet. The admittance-flag $\alpha_i(t)$ is thus controlled as

$$\alpha_i(t) = \begin{cases} 1 & \text{if } D_i(t) \ge d_i^{\text{ub}}(t) \\ 0 & \text{if } D_i(t) < d_i^{\text{ub}}(t) \end{cases}$$
(14)

3.2. Service controller

The goal for the service-controller is to find $m_i^{\text{ref}}(t)$ such that the utility function is maximized once the reference signal is realized in Δ_i time-units, i.e. such that $u_i(t + \Delta_i)$ is maximized. Recall that the utility function is given by

$$u_i(t) = a_i(t) \times e_i(t) = \frac{s_i^2(t)}{s_i^{\text{cap}}(t) \times r_i(t - d_i(t))}.$$

As explained in the introduction of this section, the input load is assumed to change relatively slowly over a time interval of a few milliseconds. Hence, one can approximate

$$r_i(t - d_i(t)) \approx r_i(t), \tag{15}$$

since the goal of both the admission controller and the service controller is to keep $d_i(t)$ in the order of milliseconds or less. Therefore it is possible to approximate the utility function with

$$u_i(t) \approx rac{s_i^2(t)}{s_i^{\operatorname{cap}}(t) \times r_i(t)}.$$

Furthermore, the service rate $s_i(t)$ can be approximated to be either at the capacity of the function, $s_i^{cap}(t)$, or at the input rate $r_i(t)$

$$s_i(t) \approx \min\{s_i^{\operatorname{cap}}(t), r_i(t)\}.$$
(16)

where the min is used since the function cannot process packets at a faster rate than what they are entering the function for a prolonged period of time. Likewise, it cannot process packets at a rate higher than the capacity of the function when the input were to be higher than this. This leads to the utility function being approximated as

$$u_i(t) \approx \begin{cases} \frac{(s_i^{\operatorname{cap}}(t))^2}{s_i^{\operatorname{cap}}(t) \times r_i(t)} = \frac{s_i^{\operatorname{cap}}(t)}{r_i(t)}, & \text{if } s_i^{\operatorname{cap}}(t) \le r_i(t) \\ \frac{r_i^2(t)}{s_i^{\operatorname{cap}}(t) \times r_i(t)} = \frac{r_i(t)}{s_i^{\operatorname{cap}}(t)}, & \text{else} \end{cases}$$

With $s_i^{\rm cap}(t)$ given by (3) and the average machine uncertainty $\xi_i(t)$ given by (4) the utility function can finally be approximated as

$$u_{i}(t) \approx \begin{cases} \frac{m_{i}(t)(\bar{s}_{i} + \xi_{i}(t))}{r_{i}(t)}, & \text{if } m_{i}(t)(\bar{s}_{i} + \hat{\xi}_{i}(t)) \leq r_{i}(t) \\ \frac{r_{i}(t)}{m_{i}(t)(\bar{s}_{i} + \hat{\xi}_{i}(t))}, & \text{else} \end{cases}$$
(17)

Since the goal is to find $m_i^{\text{ref}}(t)$ in order to maximize $u_i(t+\Delta_i)$ one needs knowledge of $\hat{\xi}_i(t+\Delta_i)$ and $r_i(t+\Delta_i)$ which is not available. However, one can assume that the machine uncertainty will be fairly constant during Δ_i timeunits such that $\hat{\xi}_i(t+\Delta_i) \approx \hat{\xi}_i(t)$. Furthermore, one has to estimate the future input-rate to the function. For the first function, F_1 , this can be done by using the derivative of the (preferably low-pass filtered) input-rate:

$$\hat{r}_1(t) = r_1(t) + \Delta_1 \frac{\mathrm{d}r_1(t)}{\mathrm{d}t}$$

For the succeeding functions, i = 2, ..., n the input-rate will change in a step-wise fashion and can therefore not approximate it with the expression above. However, since $r_i(t) = s_{i-1}(t)$ and $m_{i-1}(x)$ is known for $x \in [0, t + \Delta_{i-1}]$ (with t being the current time) one could estimate the future input-rate $\hat{r}_i(t)$ with

$$\hat{r}_i(t) \approx \min\left(s_{i-1}^{\text{cap}}(t + \Delta_{i-1}), \hat{r}_{i-1}(t)\right), \quad i = 2, \dots, n.$$

Note that $s_{i-1}^{\operatorname{cap}}(t+\Delta_{i-1})$ is used here, instead of $s_{i-1}^{\operatorname{cap}}(t+\Delta_i)$. The reason is that if $\Delta_i > \Delta_{i-1}$ one does not have enough information to compute $s_{i-1}^{\operatorname{cap}}(t+\Delta_{i-1})$. However, one can use the assumption that $\Delta_i \approx \Delta_{i-1}$. Furthermore, since

$$s_{i-1}^{\text{cap}}(t + \Delta_{i-1}) \approx m_{i-1}^{\text{ref}}(t) \times (\bar{s}_{i-1} + \hat{\xi}_{i-1}(t))$$

one can summarize the predicted input $\hat{r}_i(t)$ as

$$\hat{r}_{i}(t) = \begin{cases} r_{i}(t) + \Delta_{i} \frac{\mathrm{d}r_{i}(t)}{\mathrm{d}t}, & i = 1\\ \min\left(m_{i-1}^{\mathrm{ref}}(t) \times (\bar{s}_{i-1} + \hat{\xi}_{i-1}(t)), \ \hat{r}_{i-1}(t)\right), & \text{else} \end{cases}$$

With these simplifications and approximations, $m_i^{\text{ref}}(t)$ can be found by solving

$$\begin{split} m_i^{\text{cr}}(t) &= \\ \begin{cases} \arg\max_{x\in\mathbb{Z}^+} \Big\{ \frac{x\times(\bar{s}_i + \hat{\xi}_i(t))}{\hat{r}_i(t)} \Big\}, & \text{if } x\times(\bar{s}_i + \hat{\xi}_i(t)) \leq \hat{r}_i(t) \\ \arg\max_{x\in\mathbb{Z}^+} \Big\{ \frac{\hat{r}_i(t)}{x\times(\bar{s}_i + \hat{\xi}_i(t))} \Big\}, & \text{else} \end{cases} \end{split}$$

where $x \in \mathbb{Z}^+$ is the number of instances. Here one can see that the first case of the above equation is maximized when x is as large as possible, but since this case is only valid when $x \leq \hat{r}_i(t)/(\bar{s}_i + \hat{\xi}_i(t))$ it leads to $x = \lfloor \frac{r_i(t)}{\bar{s}_i + \hat{\xi}_i(t)} \rfloor$. Similarly, the second case is maximized when x is as small as possible, and since this case is valid for $x \geq \hat{r}_i(t)/(\bar{s}_i + \hat{\xi}_i(t))$ it leads to $x = \lceil \frac{r_i(t)}{\bar{s}_i + \hat{\xi}_i(t)} \rceil$, leading to the final control-law:

$$m_{i}^{\text{ref}}(t) = \begin{cases} \left\lfloor \frac{\hat{r}_{i}(t)}{\bar{s}_{i} + \hat{\xi}_{i}(t)} \right\rfloor, & \text{if } \left\lfloor \frac{\hat{r}_{i}(t)}{\bar{s}_{i} + \hat{\xi}_{i}(t)} \right\rfloor \frac{\bar{s}_{i} + \hat{\xi}_{i}(t)}{\hat{r}_{i}(t)} \geq \\ & \frac{\hat{r}_{i}(t)}{\bar{s}_{i} + \hat{\xi}_{i}(t)} \times \frac{1}{\left\lceil \frac{\hat{r}_{i}(t)}{\bar{s}_{i} + \hat{\xi}_{i}(t)} \right\rceil} \\ & \left\lceil \frac{\hat{r}_{i}(t)}{\bar{s}_{i} + \hat{\xi}_{i}(t)} \right\rceil, & \text{else} \end{cases}$$

$$(18)$$

3.3. Properties of AutoSAC

There are several interesting properties captured by the admission controller and service controller presented in this section. First of all, the admission controller (14) ensures, by design, that every packet that is admitted into a function, and thus exits the function, meets its deadline. Therefore, no packets that exit the service-chain will miss their end-to-end deadline.

The service-controller given by equation (18) captures both the feedback used from the true performance of the instances (when computing $\hat{\xi}_i(t)$) as well as feedforward information about future input coming from functions earlier in the service-chain (when computing $\hat{r}_i(t)$). This makes it robust against machine uncertainties but also ensures that it reacts fast to sudden changes in the input. For instance, given a service-chain of 6 functions, function F_5 will know that in Δ_4 time-units, F_4 will have $m_4^{\text{ref}}(t)$ instances running and can thus start as many instances as needed to process this new load.

4. Evaluation

In this section, we evaluate the automatic service- and admission-controller (AutoSAC) developed in Section 3 through a Monte Carlo simulation with $15 \cdot 10^4$ runs. It is compared against two state-of-the-art methods for autoscaling VMs in industry; *dynamic auto-scaling* (DAS) and *dynamic over-provisioning* (DOP). However, since these two methods do not use any admission control they are also augmented with the admission controller presented in Section 3.1. The two augmented methods are denoted by "DAS with AC" and "DOP with AC". Hence, in total the method presented in Section 3 is compared with four other methods.



Figure 5: Average utilities for *one* of the runs in the Monte Carlo simulation. AutoSAC- $U_a(t)$, DAS- $U_b(t)$, DOP- $U_c(t)$, DAS with AC- $U_d(t)$, DOP with AC- $U_e(t)$. One can see that AutoSAC performs better than DAS and DOP, as well as that the utility for DAS and DOP is the same as for DAS with AC and DOP with AC except for some short time-intervals. During these time-intervals these functions were overloaded, and the lack of admission control in DAS and DOP lead to packets missing their deadlines, thus leading to a drop in the utility. However, even when augmenting DAS and DOP with admission control they do not perform as well as AutoSAC, the main reason being that AutoSAC uses feedback from the true performance of the instances as well as feedforward between the functions in the service-chain to obtain a faster reaction time.

A real-world trace of traffic data is used as input to the simulation. The trace was gathered over 120 hours from a port in the Swedish University NETwork (SUNET) and then normalized to have a peak of 10,000,000 packets per second as shown in Figure 1. The simulation was written in the open-source language Julia [17]. The code used for this simulation is provided on GitHub¹.

Dynamic auto-scaling (DAS). This method is currently being offered to customers using Amazon Web Services [18]. It allows the user to monitor different metrics (e.g., CPU utilization) of their VMs using CloudWatch. One can then use it together with their auto-scaling solution to achieve *dynamic auto-scaling*. This allows the user to scale the number of VMs as a function of these metrics. One should note that the CPU utilization can be considered the same as the efficiency metric $e_i(t)$ defined in (11). For the Monte Carlo simulation the following rules were used:

- add a VM if the efficiency is above 99%,
- remove a VM if the efficiency is below 95%,

which might seem as a high and tight interval, but it is necessary in order to achieve a high utility.

Dynamic over-provisioning (DOP). A downside with DAS is that it reacts slow to sudden changes in the input. A natural alternative to would therefore be to instead do *dynamic over-provisioning*, where one measures the input to each function and allocate virtual resources such that there is an expected over-provision by 10%.

Monte Carlo Simulation. The five methods are compared using a Monte Carlo simulation with $15 \cdot 10^4$ runs. For every run, one hour of input data was randomly selected from the total of 120 hours shown in Figure 1. In every run a new service-chain with 5 functions was generated. The E2E deadline for the service-chain was set to 50ms, which in turn was split into local deadlines of 10ms for each function. The other parameters (i.e., \bar{s}_i , Δ_i , $\xi_i^{\rm lb}$, and $\xi_i^{\rm ub}$) for every function in the service-chain was generated randomly. The expected service-rate \bar{s}_i was chosen uniformly at random from the interval [100,000, 200,000] pps. The time-overhead Δ_i was drawn uniformly at random from

the interval [30, 120] seconds. The machine uncertainty was chosen to be in the range of $\pm 30\%$ of the expected servicerate \bar{s}_i . The lower bound of the machine uncertainty was drawn from the interval $[-0.3\bar{s}_i, 0]$ pps and likewise, the upper bound was drawn from $[0, 0.3\bar{s}_i]$ pps.

The evaluation of the Monte Carlo simulation is based on the average utility $U(t) = \frac{1}{t} \int_0^t \sum_{i=1}^n u_i(x) dx$. Since a packet that misses its deadline (which is possible when using DAS or DOP) is considered useless it is evaluated as a dropped packet when exiting the function. It therefore impacts the availability metric and in turn the utility. Should all packets miss their deadlines in function F_i for a time interval τ , then $a_i(t) = 0 \quad \forall t \in \tau$, i.e. the availability would be evaluated as 0 during this time-interval since the output of the function is considered useless.

Results. The mean of the average utility U(t) for all the simulation runs is presented in Table 2 for each of the five methods. One can see that AutoSAC achieves a utility that is 30–40% better than that of DAS and DOP. The main reason for this is that they are lacking admission control leading to packets missing their deadlines, which eventually results in a low utility. This can be seen in Figure 5, illustrating a "typical" run in the Monte Carlo simulation. There, the sudden drops in the utility function for DAS $(U_b(t))$ as well as the step-wise increase of DAS and DOP $(U_c(t))$ in the beginning, indicate that they are overloaded and packets are missing their deadlines in some of the functions in the service-chain.

Method	Mean(U(t))	$\operatorname{Var}(U(t))$	Figure 5
AutoSAC	0.99	$3.0 \cdot 10^{-5}$	$U_a(t)$
DAS	0.67	$2.7 \cdot 10^{-2}$	$U_b(t)$
DOP	0.75	$1.6 \cdot 10^{-2}$	$U_c(t)$
DAS with AC	0.94	$2.8 \cdot 10^{-3}$	$U_d(t)$
DOP with AC	0.91	$4.9 \cdot 10^{-4}$	$U_e(t)$

TABLE 2: Results from the Monte Carlo simulation. AutoSAC performs 30–40% better than DAS and DOP. The main reason is the admission controller used in AutoSAC. When augmenting DAS and DOP with this admission controller, their performance is increased by more than 20%. However, AutoSAC still outperforms the augmented methods by 5–10% since it uses feedforward, making it faster to react to input changes, as well as feedback making it more robust to machine uncertainties.

When augmenting DAS and DOP with the admission controller derived in Section 3.1 the performance is increased by 20-40%, purely as a result of not having these sudden

^{1.} https://github.com/vmillnert/ICC17simulation

drops in performance. However, AutoSAC still performs 5– 10% better, due to the feedforward property of AutoSAC giving it having a faster reaction time to changes in the input as well as the feedback property leading to better prediction and robustness against the machine uncertainties.

5. Summary

In this work we have developed a mathematical model for a NFV Forwarding Graphs residing in a Cloud environment. The model captures, among other things, the time needed to start/stop virtual resources (e.g., virtual machines or containers), and the uncertainty of the performance of the virtual resources which can deviate from the expected performance due to other tenants running loads on the physical infrastructure. The packets that flow through the forwarding graph must be processed by each of the virtual network functions (VNFs) within some end-to-end deadline.

A utility function is defined to evaluate performance between different methods for controlling NFV Forwarding Graphs. The utility function is also used to derive an automatic service- and admission-controller (AutoSAC) in Section 3. It ensures that packets that exit the forwarding graph meet their end-to-end deadline. The service-controller use feedback from the actual performance of the virtual resources making it robust against uncertainties and deviations from the expected performance. Furthermore, it uses feedforward between the VNFs making it fast to react to changes in the input load.

In Section 4 AutoSAC is evaluated and compared against 4 other methods in a Monte Carlo simulation with $15 \cdot 10^4$ runs. The input load for the simulation is a real-world trace of traffic data gathered over 120 hours. The traffic is normalized to have a peak of 10,000,000 packets per second. AutoSAC is shown to have better performance than what is offered in the cloud industry today. We also show that when augmenting the industry-methods with the admission controller derived in Section 3 they have a significant increase in performance.

It would be interesting to extend this work by investigating how to derive a controller when the true performance is unknown or when the time-overhead needed to start virtual resources is unknown. Moreover, it would be interesting to investigate how as well as by developing a model for a forwarding graph that has forks and joins, i.e. a graph structure.

Acknowledgements The authors would like to thank Karl-Erik Årzén and Joao Monteiro Soares for the useful comments on early versions of this paper.

Source code The source code for the simulation in Section 4 can be found on Github at https://github.com/vmillnert/ICC17simulation.

References

- [1] ETSI, "Network Functions Virtualization (NFV)," https://portal.etsi.org/nfv/nfv_white_paper.pdf, October 2012.
- [2] W. Shen, M. Yoshida, T. Kawabata, K. Minato, and W. Imajuku, "vconductor: An nfv management solution for realizing end-to-end virtual network services," in *Network Operations and Management Symposium (APNOMS)*, 2014 16th Asia-Pacific. IEEE, 2014, pp. 1–6.

- [3] H. Moens and F. De Turck, "Vnf-p: A model for efficient placement of virtualized network functions," in 10th International Conference on Network and Service Management (CNSM) and Workshop. IEEE, 2014, pp. 418–423.
- [4] S. Mehraghdam, M. Keller, and H. Karl, "Specifying and placing chains of virtual network functions," in *Cloud Networking (CloudNet)*, 2014 IEEE 3rd International Conference on. IEEE, 2014, pp. 7–13.
- [5] M. Mao, J. Li, and M. Humphrey, "Cloud auto-scaling with deadline and budget constraints," in 2010 11th IEEE/ACM International Conference on Grid Computing. IEEE, 2010, pp. 41–48.
- [6] X. Wang, C. Wu, F. Le, A. Liu, Z. Li, and F. Lau, "Online vnf scaling in datacenters," arXiv preprint arXiv:1604.01136, 2016.
- [7] Y. Li, L. Phan, and B. T. Loo, "Network functions virtualization with soft real-time guarantees," in *IEEE International Conference on Computer Communications (INFOCOM)*, 2016.
- [8] M. Di Natale and J. A. Stankovic, "Dynamic end-to-end guarantees in distributed real time systems," in *Proceedings of the 15-th IEEE Real-Time Systems Symposium*, Dec. 1994, pp. 215–227.
- [9] S. Jiang, "A decoupled scheduling approach for distributed real-time embedded automotive systems," in *Proceedings of the 12th IEEE Real-Time and Embedded Technology and Applications Symposium*, 2006, pp. 191–198.
- [10] N. Serreli, G. Lipari, and E. Bini, "Deadline assignment for component-based analysis of real-time transactions," in 2nd Workshop on Compositional Real-Time Systems, Washington, DC, USA, Dec. 2009.
- [11] —, "The demand bound function interface of distributed sporadic pipelines of tasks scheduled by EDF," in *Proceedings of the 22-nd Euromicro Conference on Real-Time Systems*, Bruxelles, Belgium, Jul. 2010.
- [12] S. Hong, T. Chantem, and X. S. Hu, "Local-deadline assignment for distributed real-time systems," *IEEE Transactions on Computers*, vol. 64, no. 7, pp. 1983–1997, Jul. 2015.
- [13] P. Jayachandran and T. Abdelzaher, "Delay composition algebra: A reduction-based schedulability algebra for distributed real-time systems," in *Proceedings of the 29-th IEEE Real-Time Systems Sympo*sium, Barcelona, Spain, Dec. 2008, pp. 259–269.
- [14] D. Henriksson, Y. Lu, and T. Abdelzaher, "Improved prediction for web server delay control," in *Proceedings of the 16th Euromicro Conference on Real-Time Systems*, Jun. 2004, pp. 61–68.
- [15] R. Bonafiglia, I. Cerrato, F. Ciaccia, M. Nemirovsky, and F. Risso, "Assessing the performance of virtualization technologies for nfv: a preliminary benchmarking," in 2015 Fourth European Workshop on Software Defined Networks. IEEE, 2015, pp. 67–72.
- [16] R. Kapoor, G. Porter, M. Tewari, G. M. Voelker, and A. Vahdat, "Chronos: Predictable low latency for data center applications," in *Proceedings of the Third ACM Symposium on Cloud Computing*, ser. SoCC '12. New York, NY, USA: ACM, 2012, pp. 9:1–9:14. [Online]. Available: http://doi.acm.org/10.1145/2391229.2391238
- [17] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," arXiv preprint arXiv:1411.1607, 2014.
- [18] (2016, 10). [Online]. Available: https://aws.amazon.com/