Efficient Orchestration Mechanisms for Congestion Mitigation in NFV: Models and Algorithms

Jocelyne Elias*, Fabio Martignon†, Stefano Paris* and Jianping Wang‡

*LIPADE
Université Paris Descartes
{jocelyne.elias, stefano.paris}@parisdescartes.fr

†LRI
University of Paris-Sud
Institut Universitaire de France (IUF)
fabio.martignon@lri.fr

‡Department of Computer Science
City University of Hong Kong
jianwang@cityu.edu.hk

Abstract—Network Functions Virtualization (NFV) has recently gained momentum among network operators as a means to share their physical infrastructure among virtual operators, which can independently compose and configure their communication services. However, the spatio-temporal correlation of traffic demands and computational loads can result in high congestion and low network performance for virtual operators, thus leading to service level agreement breaches.

In this paper, we analyze the congestion resulting from the sharing of the physical infrastructure and propose innovative *orchestration mechanisms* based on both centralized and distributed approaches, aimed at unleashing the potential of the NFV technology. In particular, we first formulate the network functions composition problem as a non-linear optimization model to accurately capture the congestion of physical resources. To further simplify the network management, we also propose a *dynamic pricing strategy* of network resources, proving that the resulting system achieves a stable equilibrium in a completely distributed fashion, even when all virtual operators *independently* select their best network configuration.

Numerical results show that the proposed approaches consistently reduce resource congestion. Furthermore, the distributed solution well approaches the performance that can be achieved using a centralized network orchestration system.

Index Terms—Network Functions Virtualization, Game Theory, Distributed Congestion Control, Non-linear Optimization.

1 Introduction

Nowadays, telecommunication infrastructures are composed of property hardware operated by a single entity to offer communication services to their final users. While this architecture simplifies the design and optimization of the network equipment for specific tasks, its low degree of flexibility represents the main limitation for the evolution of the network infrastructure. For this reason, network operators and equipment manufacturers have started the standardization process of a plethora of virtualization solutions that have been individually developed in recent years for enabling the sharing of general-purpose resources and increasing the flexibility of their network architectures. Such a process has led to the specification of the *Network Functions Virtualization* (NFV) technology [1], which promises to bring about several benefits, such as reduced CAPEX and OPEX (CAPital and OPerational EXpenditure), low time-to-market for new network services, higher flexibility to scale up and down the services according to users' demand, simple and cheap testing of new services.

Nevertheless, the consolidation of the virtualization technology represents one of the main challenging problems for its success and widespread utilization in telecommunication infrastructures, which still consist of a huge set of property hardware appliances and software systems [2]. Indeed, the sharing of the physical infrastructure among multiple virtual operators as well as the simple configuration of network services require the design of complex management mechanisms for the orchestration of the network equipment, with the final goal of dynamically adapting the infrastructure to the resource utilization.

In this paper, we propose novel orchestration mechanisms to optimally control and mitigate the resource congestion of a physical infrastructure based on the NFV paradigm. We first formulate the problem as a non-linear optimization model that accurately captures the congestion of physical network resources, and permits to dynamically control traffic flows and system configurations in order to prevent the congestion of network resources.

While centralized solutions like [3] permit to optimally control the system, the associated costs and responsibilities for satisfying the service (i.e., SLA and corresponding penalties) represent their main obstacle for the operator of the physical infrastructure. Furthermore, the recent debate on Net neutrality in the Comcast versus Netflix dispute has unearthed the economic problems that Over the Top Providers as well as virtual operators might face for an open and fair access to network resources. At the same time, the lack of direct control on the network infrastructure that might affect the quality of service experienced by final customers might discourage virtual operators from using such a technology for offering their services. In particular, this could prevent large-scale diffusion of Infrastructure as a Service (IaaS) with carrier grade QoS requirements, which basically consists in delivering high-level performance services like networking, processing, and storage functions to final customers.

For these reasons, the NFV technology further calls for *distributed* approaches where the best operational point of the system results from individual decisions performed independently by virtual operators according to the network status and customers' requests. In this context, game theory provides the natural framework for both analyzing the evolution of NFV-based systems and designing the rules (e.g., incentives/prices and use policy) to coordinate network allocation decisions of virtual operators [4], [5].

As a second key contribution, we therefore analyze the congestion resulting from the sharing of the physical infrastructure when all virtual operators *independently select* their best network configuration. We formulate the distributed congestion minimization problem as a game, proposing a *dynamic pricing strategy* of network resources to achieve a stable equilibrium in a completely distributed fashion. We demonstrate that the NFV congestion mitigation game admits a unique Nash Equilibrium, under very general conditions, and that efficient solutions can be easily computed in a distributed fashion. We further compare our distributed solution to a centralized approach, using both an optimization model and an efficient heuristic based on the Shortest Path Tree algorithm. Numerical results show that the proposed distributed model significantly decreases network congestion, thus representing a very promising approach for operators to manage network resources in an efficient, fully distributed and dynamic fashion. Furthermore, it well approaches the performance of centralized optimization models, which can hardly be solved to the optimum in real network scenarios.

The paper is structured as follows. Section 2 discusses related work. Section 3 introduces the network model as well as the assumptions considered in our work. Section 4 formulates the centralized, pricing-based congestion control mechanism as a non-linear optimization problem, and provides the conditions for its solution. Section 5 describes the game theoretical approach we design to solve the congestion control problem and achieve a stable operating state in a distributed fashion. Desirable properties, like existence and uniqueness of the Nash equilibrium of the proposed game, are further established in Section 6. Section 7 describes the experimental settings of our performance evaluation campaign, and illustrates numerical results that show the efficiency and validity of our centralized an distributed approaches. Finally, concluding remarks are discussed in Section 8.

2 Related Work

We now discuss the most relevant works that deal with the network functions composition problem, and more in general, with dynamic resource management in virtual networks, which is the focus of our work. The key enabling paradigms that will considerably increase the dynamicity of ICT networks are Software-Defined Networking (SDN) and Network Functions Virtualization (NFV) [1], which are discussed in recent surveys [6], [7], [8], [9]. Indeed, Service Providers and Network Operators are facing increasing problems to design and implement novel network functionalities, following the rapid changes that characterize the current Internet and Telecom operators [10].

Virtualization represents an efficient and cost-effective strategy to exploit, and share, physical network resources. In this context, the network embedding problem has been considered in several recent works [11], [12], [13], [14], [15], [16].

The Virtual Network Embedding problem consists in finding a mapping between a set of requests for virtual network resources and the available underlying physical infrastructure (the *substrate*), ensuring that some given performance requirements (on nodes and links) are guaranteed. Typical node requirements are computational resources (i.e., CPU) or storage space, whereas links have a limited bandwidth and introduce a delay. It has been shown that this problem is NP-hard (it includes as subproblem the multi-way separator problem). For this reason, heuristic approaches have been devised [17].

The problem of virtual resources consolidation is considered in [15], taking into account energy efficiency. The problem is formulated as a mixed integer linear programming (MILP) model, to understand the potential benefits that can be achieved by packing many different virtual tasks on the same physical infrastructure. The observed energy savings are up to 30% for nodes, and up to 25% for link energy consumption. The work [18] presents a solution for the resilient deployment of network functions, using OpenStack for the design and implementation of the proposed service orchestrator mechanism.

An allocation mechanism, based on auction theory, is proposed in [19]. In particular, the scheme selects the most remunerative virtual network requests, while satisfying QoS requirements and physical network constraints. The system is split into two network substrates modeling physical and virtual resources, with the final goal of finding the best mapping of virtual nodes and links into physical resources according to the QoS requirements (i.e., bandwidth, delay, CPU bounds).

A novel network architecture is proposed in [20], to provide efficient, coordinated control of both internal network function state and network forwarding state in order to help operators achieve the following goals: (1) offer and satisfy tight service level agreements (SLAs); (2) accurately monitor and manipulate network traffic;

and (3) minimize operating expenses.

To the best of our knowledge, only few works focus on congestion control in virtual networks. On the contrary, in this work we explicitly address this issue, by studying the effects on the network congestion of services composition in NFV-based infrastructures, in order to derive *numerical bounds* on the congestion reduction that can be achieved by deploying virtualization mechanisms. Furthermore, we formulate the *distributed* congestion minimization problem as a game, proposing a *dynamic pricing strategy* of network resources to achieve a stable equilibrium in a fully distributed fashion. Our proposed distributed approach permits to efficiently compute a solution close to the optimum in a short time, thus enabling the optimal reconfiguration of network resources on the fly.

3 Network Model

This section presents the network model and assumptions we adopt in the design of our mechanisms for controlling and mitigating resource congestion in virtual networks.

We consider a physical network infrastructure managed by a single operator composed of a set $\mathcal{N}$ of general purpose nodes and a set $\mathcal{L}$ of links, as illustrated in Figure 1. Therefore, the topology of the network infrastructure is represented as a weighted directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$. Table 1 summarizes the notation used throughout the paper.

The operator adopts the Network Functions Virtualization (NFV) approach for providing access to its physical resources, since through the virtualization of network functions, virtual operators can share physical resources implementing their own network services independently of each other and the underlying technology. Therefore, each link $(i;j) \in \mathcal{L}$ of the physical infrastructure can be mapped to a set $\mathcal{S}_{ij}$ of different transport services designed and used by VOs that access the network. Such transport services can be implemented using basic network functions provided by the Virtual Operator on top of general purpose network processors like those proposed in [21], [22]. We observe that the bandwidth B_{ij}^s assigned to each transport service $s \in \mathcal{S}_{ij}$ depends on the underlying scheduling mechanism implemented by the operator of the physical infrastructure. As an example, the operator can fairly treat all transport services using a simple round robin scheme, which assigns the same access time of the physical link to all services. While the use of more complex mechanisms is out of the scope of this paper, we emphasize that our analysis on network resources' congestion is general and does not depend on any specific scheduling scheme.

Each virtual operator[1] $u \in \mathcal{U}$ defines its demand b_u by specifying the source and destination nodes as well as the amount of data traffic r^u that is transmitted between them. The sets of all source (or origin) and destination nodes are denoted by two ordered vectors, $\mathcal{O}$ and $\mathcal{D}$, respectively, and we refer to source and destination nodes of VO $u \in \mathcal{U}$ using the notation $\mathcal{O}(u)$ and $\mathcal{D}(u)$. For example, source and destination nodes of VO 1 are denoted by $\mathcal{O}(1)$ and $\mathcal{D}(1)$, respectively.

1. The terms "User" and "Virtual Operator" (VO) are used interchangeably throughout the paper.

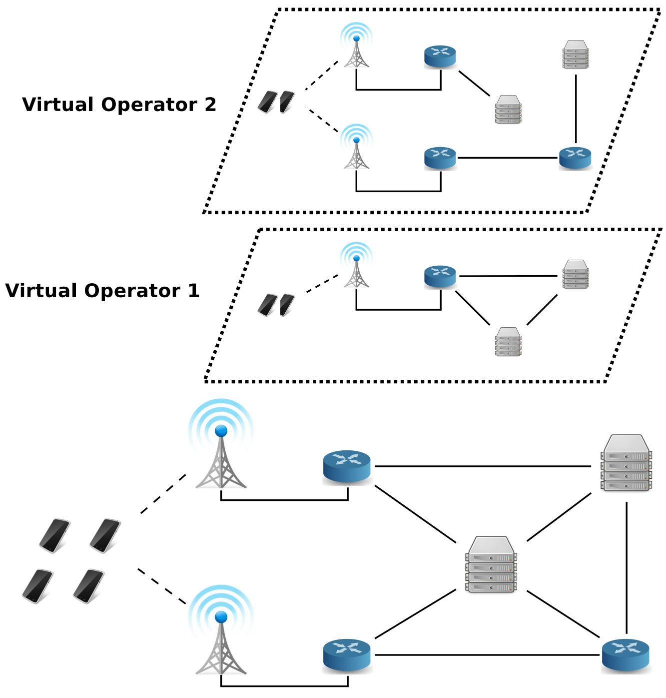

Fig. 1: Network Scenario: A single physical infrastructure (bottom figure) is shared among two Virtual Operators. The network topology, the transmission services of network links, and the services executed by network nodes are selected according to the virtual operator needs in order to minimize the physical network congestion.

In our vision, the virtual operator also provides a list of processing nodes ($\mathcal{P}_u$) through which a fraction of its data traffic $w_j^u \in [0,1], j \in \mathcal{P}_u$ must be routed. These nodes are used to perform intensive computational tasks, like traffic analysis and deep packet inspection (e.g., firewalls, intrusion detection systems and other in-the-cloud middlebox services), and to facilitate key operations like billing, sampling and verification. For example, a VO can verify the service provided by the physical operator based on the performance of packets passing through a specified node in set $\mathcal{P}_u$. Virtual machines for the processing needs of the VO can be allocated dynamically using mechanisms like [23].

The *congestion cost function*, which we propose to use for achieving the best configuration of network functions, depends on the network congestion as defined in [5], [24]. The function will help to fully exploit the physical infrastructure and guarantee, at the same time, a high Quality of Service (QoS). More specifically, for each link $(i;j) \in \mathcal{L}$ and service $s \in \mathcal{S}_{ij}$, we consider an *increasing* and *convex* function *per traffic unit* as follows:

TABLE 1: Basic Notation

Sets and Parameters	
$\mathcal{U}$	Users set (Virtual Operators)
$\mathcal{N}$	Physical node set
$\mathcal{L}$	Physical link set
$\mathcal{S}_{ij}$	Set of transmission/forwarding services for link $(i;j)$
$\mathcal{P}_u \subset \mathcal{N}$	Set of alternative processing nodes selected by user $u \in \mathcal{U}$
B_{ij}^s	Bandwidth assigned to service s on link $(i;j)$

VOs Demands	
$\{\mathcal{O}(u), \mathcal{D}(u), r^u\}$	Source, destination and traffic demand of user u
w_j^u	Fraction of data traffic of user u passing through node j

Variables	
x_{ij}^{us}	Traffic flow of user u passing through link $(i;j)$ and transmitted using service s

$$p_{ij}^s \left(\sum_{u \in \mathcal{U}} x_{ij}^{us} \right) = \left[a_{ij}^s + b_{ij}^s \left(\frac{\sum_{u \in \mathcal{U}} x_{ij}^{us}}{B_{ij}^s} \right)^{\tau} \right], \quad (1)$$

where x_{ij}^{us} is the traffic flow of VO u passing through link $(i;j)$ and transmitted using service s, B_{ij}^s represents the bandwidth assigned to service s on link $(i;j)$, a_{ij}^s and b_{ij}^s are two positive numbers, and τ is a positive integer greater than 1. Expression (1) represents the congestion cost function of a single link $(i;j)$ when using service s, that is, the cost *per traffic unit*, which is a function of the total load $\sum_{u \in \mathcal{U}} x_{ij}^{us}$ on that link.

4 Optimal Congestion Control for Virtual Networks

In this section, we present a *centralized* approach to mitigate the congestion of a NFV-based physical infrastructure operated by a single physical network operator. We first formulate the Congestion Mitigation for Virtual Networks (CMVN) problem as a non-linear integer optimization model to closely capture the congestion of physical resources; specifically, our model aims at minimizing the *total network congestion*. Then, we provide the Karush-Kuhn-Tucker (KKT) conditions using which the physical operator can compute in a short time an optimal operation point for its physical infrastructure.

4.1 Problem Formulation

The total congestion of a transport service implemented over a link $(i;j) \in \mathcal{L}$ experienced by the operator of the physical infrastructure increases steeply with its utilization (or equivalently, with the traffic transported by the link service). Therefore, the congestion of service s of link $(i;j)$ is defined as the product of function (1) and the service utilization as follows:

$$\begin{aligned} J_{ij}^s \left(\sum_{u \in \mathcal{U}} x_{ij}^{us} \right) &= p_{ij}^s \left(\sum_{u \in \mathcal{U}} x_{ij}^{us} \right) \left(\frac{\sum_{u \in \mathcal{U}} x_{ij}^{us}}{B_{ij}^s} \right) = \\ &= \left[a_{ij}^s + b_{ij}^s \left(\frac{\sum_{u \in \mathcal{U}} x_{ij}^{us}}{B_{ij}^s} \right)^{\tau} \right] \left(\frac{\sum_{u \in \mathcal{U}} x_{ij}^{us}}{B_{ij}^s} \right). \end{aligned} \quad (2)$$

We underline that the service congestion (2) is still a convex function that depends on the service utilization. The total network congestion, J, seen by the operator of the physical infrastructure is defined as the sum over all links and transport services of convex function (2), as formulated in Equation (3):

$$\begin{aligned} J &= \sum_{(i;j) \in \mathcal{L}} \sum_{s \in \mathcal{S}_{ij}} J_{ij}^s \left(\sum_{u \in \mathcal{U}} x_{ij}^{us} \right) = \\ &= \sum_{(i;j) \in \mathcal{L}} \sum_{s \in \mathcal{S}_{ij}} \left[a_{ij}^s + b_{ij}^s \left(\frac{\sum_{u \in \mathcal{U}} x_{ij}^{us}}{B_{ij}^s} \right)^{\tau} \right] \cdot \left(\frac{\sum_{u \in \mathcal{U}} x_{ij}^{us}}{B_{ij}^s} \right). \end{aligned} \quad (3)$$

According to such definition, the problem of optimally minimizing the total network congestion can be formulated as follows:

$$\min \quad \sum_{(i;j) \in \mathcal{L}} \sum_{s \in \mathcal{S}_{ij}} J_{ij}^s \left(\sum_{u \in \mathcal{U}} x_{ij}^{us} \right) \quad (4)$$

s.t.

$$\sum_{\substack{i \in \mathcal{N}: \\ (j;i) \in \mathcal{L}}} \sum_{s \in \mathcal{S}_{ji}} x_{ji}^{us} - \sum_{\substack{i \in \mathcal{N}: \\ (i;j) \in \mathcal{L}}} \sum_{s \in \mathcal{S}_{ij}} x_{ij}^{us} = \begin{cases} r^u & \forall u \in \mathcal{U}, \forall j \in \mathcal{O}(u) \\ 0 & \forall u \in \mathcal{U}, \forall j \in \mathcal{N} \\ -r^u & \forall u \in \mathcal{U}, \forall j \in \mathcal{D}(u) \end{cases} \quad (5)$$

$$\sum_{u \in \mathcal{U}} x_{ij}^{us} \le B_{ij}^s \qquad \forall (i;j) \in \mathcal{L}, \forall s \in \mathcal{S}_{ij} \quad (6)$$

$$\sum_{i \in \mathcal{N}} \sum_{s \in \mathcal{S}_{ij}} x_{ij}^{us} = w_j^u r^u \qquad \forall u \in \mathcal{U}, j \in \mathcal{P}_u \quad (7)$$

$$x_{ij}^{us} \ge 0 \qquad \forall (i;j) \in \mathcal{L}, u \in \mathcal{U}, s \in \mathcal{S}_{ij}. \quad (8)$$

Objective function (4) minimizes the overall network congestion. Constraints (5) define the flow balance at node $j \in \mathcal{N}$ for the data traffic demand of user u. Specifically, terms $\sum x_{ij}^{us}$ and $\sum x_{ji}^{us}$ represent the total incoming and outgoing traffic flows, respectively.

The set of constraints (6) ensures that the total traffic routed on a link established between two devices i and j using transmission service s does not exceed the bandwidth assigned by the operator to such service, which is denoted by B_{ij}^s. The set of constraints (7) forces the fraction of data traffic that must be processed or analyzed, w_j^u, to pass through the processing nodes ($\mathcal{P}_u$) selected by the virtual operator. Therefore, the physical operator can select the set of physical nodes that minimize the congestion to perform the computational tasks requested and developed by the virtual operator u for its data traffic. Finally, constraints (8) ensure the positiveness of the flow variables.

$$
\begin{aligned}
\mathcal{L}\left(x,\alpha,\beta,\gamma,\eta,\theta,\mu\right) = &- \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} J_{ij}^{s} + \sum_{u\in\mathcal{U}} \alpha^{\mathcal{O}(u)} \left(- \sum_{\substack{i\in\mathcal{N}:\\ (\mathcal{O}(u);i)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{\mathcal{O}(u)i}} x_{\mathcal{O}(u)i}^{us} + \sum_{\substack{i\in\mathcal{N}:\\ (i;\mathcal{O}(u))\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{i\mathcal{O}(u)}} x_{i\mathcal{O}(u)}^{us} + r^{u} \right) + \\
&+ \sum_{u\in\mathcal{U}} \sum_{\substack{j\in\mathcal{N}:\\ j\neq s_u \wedge j\neq d_u}} \beta_{j}^{u} \left(- \sum_{\substack{i\in\mathcal{N}:\\ (i;j)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} + \sum_{\substack{i\in\mathcal{N}:\\ (j;i)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ji}} x_{ji}^{us} \right) + \sum_{u\in\mathcal{U}} \gamma^{\mathcal{D}(u)} \left(- \sum_{\substack{i\in\mathcal{N}:\\ (\mathcal{D}(u);i)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{\mathcal{D}(u)i}} x_{\mathcal{D}(u)i}^{us} + \sum_{\substack{i\in\mathcal{N}:\\ (i;\mathcal{D}(u))\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{i\mathcal{D}(u)}} x_{i\mathcal{D}(u)}^{us} - r^{u} \right) + \\
&+ \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} \eta_{ij}^{s} \left(B_{ij}^{s} - \sum_{u\in\mathcal{U}} x_{ij}^{us} \right) + \sum_{u\in\mathcal{U}} \sum_{\substack{j\in\mathcal{N}:\\ j\in\mathcal{P}_u}} \theta_{j}^{u} \left(w_{j}^{u} r^{u} - \sum_{i\in\mathcal{N}} \sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} \right) + \sum_{u\in\mathcal{U}} \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} \mu_{ij}^{us} x_{ij}^{us} \qquad (9)
\end{aligned}
$$

Comments: Similarly to classical traffic engineering techniques proposed for wired networks [25], our work is based on the idea of using a non-linear increasing and convex function to strongly penalize network configurations that intensively use only few links. However, our work is unique with respect to the underlying traffic and network models, which accurately capture the flexibility and reconfigurability features of infrastructures based on the NFV technology. In particular, differently from virtual embedding problems like [11], [26], our model considers different transmission services for each link (e.g., MAC protocols and scheduling policies) and the nonlinear effect on the link's congestion and capacity degradation caused by scheduling mechanisms when the contention level increases. Such an effect is typical in communication systems based on resource sharing, which show an exponential response time. Finally, besides the accurate modeling of NFV transmission services, our proposed orchestration mechanism provides a certain degree of flexibility for the placement of network services (like billing, caching, traffic sampling and verification) that require the execution of complex functions on physical machines.

4.2 Optimal Congestion Control Solution

Since the congestion control problem (4)-(8) is a nonlinear convex optimization problem, we propose to use the Karush-Kuhn-Tucker (KKT) conditions to speed-up the computation of the optimal solution. In particular, the objective function $J = \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} J_{ij}^{s} \left(\sum_{u\in\mathcal{U}} x_{ij}^{us}\right)$ in (3) is differentiable in the x_{ij}^{us} variables and convex, and constraints (5)-(8) are linear in x_{ij}^{us}, hence the solution given by the KKT conditions is indeed an optimal solution of our problem. The Lagrangian function for the centralized congestion control problem (4)-(8) of the physical infrastructure can be written as indicated in (9), where $\alpha^{\mathcal{O}(u)}$, β_{j}^{u}, $\gamma^{\mathcal{D}(u)}$, η_{ij}^{s}, θ_{j}^{u} and μ_{ij}^{us} are the Lagrangian multipliers (nonnegative real numbers). Therefore, the following KKT conditions provide the optimal solution for the centralized congestion control problem:

$$
\begin{aligned}
\frac{\partial\mathcal{L}}{\partial x_{ij}^{us}} = &- \frac{a_{ij}^{s}}{B_{ij}^{s}} - \frac{b_{ij}^{s}}{B_{ij}^{s}}(1+\tau)\left(\frac{\sum_{u\in\mathcal{U}} x_{ij}^{us}}{B_{ij}^{s}}\right)^{\tau} \\
&- \alpha^{\mathcal{O}(u)} - \beta_{j}^{u} - \gamma^{\mathcal{D}(u)} - \eta_{ij}^{s} - \theta_{j}^{u} + \mu_{ij}^{us} = 0 \\
&\forall u\in\mathcal{U}, \forall s\in\mathcal{S}_{ij}, \forall (i;j)\in\mathcal{L}
\end{aligned}
$$

$$
\sum_{\substack{i\in\mathcal{N}:\\ (j;i)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ji}} x_{ji}^{us} - \sum_{\substack{i\in\mathcal{N}:\\ (i;j)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} = \begin{cases} r^{u} & \forall u\in\mathcal{U}, \forall j\in\mathcal{O}(u) \\ 0 & \forall u\in\mathcal{U}, \forall j\in\mathcal{N} \\ -r^{u} & \forall u\in\mathcal{U}, \forall j\in\mathcal{D}(u) \end{cases}
$$

$$
\begin{aligned}
&\sum_{u\in\mathcal{U}} x_{ij}^{us} - B_{ij}^{s} \leq 0 && \forall (i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij} \\
&\eta_{ij}^{s}\left(\sum_{u\in\mathcal{U}} x_{ij}^{us} - B_{ij}^{s}\right) = 0 && \forall (i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij} \\
&\eta_{ij}^{s} \geq 0 && \forall (i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij} \\
&\sum_{i\in\mathcal{N}} \sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} = w_{j}^{u} r^{u} && \forall u\in\mathcal{U}, j\in\mathcal{P}_u\subseteq\mathcal{N} \\
&x_{ij}^{us} \geq 0 && \forall u\in\mathcal{U}, (i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij} \\
&\mu_{ij}^{us} x_{ij}^{us} = 0 && \forall u\in\mathcal{U}, (i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij} \\
&\mu_{ij}^{us} \geq 0 && \forall u\in\mathcal{U}, (i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij}.
\end{aligned}
$$

5 Distributed Congestion Control for Virtual Networks

The centralized control policy that we presented in the previous section might discourage VOs from using the physical infrastructure, since they might perceive the network infrastructure as non-neutral. Indeed, a central entity could treat differently VOs' services by leveraging pricing policies that introduce hidden fees to fairly serve their resource requests. To this end, we present in this section a *fully distributed approach* to control the congestion of physical resources, proposing a dynamic pricing mechanism that drives the system to a stable and efficient operating point without any central coordination. Note that the dynamic pricing mechanism we propose is different from the hidden fees that the physical operator can charge on virtual operators for using "normally" the network. Indeed, these hidden costs act as a barrier for the VOs, since their services are treated unfairly as long as they do not pay for restoring or improving the resources' share they obtain from the physical infrastructure. On the contrary, our pricing mechanism simply discourages the use of heavily loaded resources. A virtual operator might be willing to slightly increase the latency of its data connections by selecting a longer yet unloaded path that provides higher bandwidth. Hereafter, we illustrate our proposed *distributed* congestion mitigation approach tailored for virtual networks.

5.1 Problem Formulation

Before describing the game that models the distributed congestion control mechanism, let us introduce some additional notation, which is used throughout this section. For each VO u, $\boldsymbol{X^u}$ represents its network flow, that is the portion of traffic x_{ij}^{us} sent over any link $(i;j)$ using any transport service s. In contrast, we use $\boldsymbol{X^{-u}}$ to denote the network flow of all other VOs, namely the portion of traffic x_{ij}^{vs}, $\forall v \in \mathcal{U} : v \neq u$.

The congestion of a transport service s implemented over link $(i;j) \in \mathcal{L}$ experienced by VO u (J_{ij}^{us}) depends both on its own flow and the traffic transmitted by other VOs, as follows:

$$J_{ij}^{us}\left(\boldsymbol{X^u}, \boldsymbol{X^{-u}}\right) = p_{ij}^{s}\left(\sum_{u\in\mathcal{U}} x_{ij}^{us}\right)\left(\frac{x_{ij}^{us}}{B_{ij}^{s}}\right) = \\ = \left[a_{ij}^{s} + b_{ij}^{s}\left(\frac{\sum_{u\in\mathcal{U}} x_{ij}^{us}}{B_{ij}^{s}}\right)^{\tau}\right]\left(\frac{x_{ij}^{us}}{B_{ij}^{s}}\right), \tag{10}$$

where the first factor represents the congestion cost function (per traffic unit) of a link as seen by the operator of the physical infrastructure and defined in (1).

Each VO u naturally wants to select the network configuration that minimizes the overall network congestion it experiences. To this end, we adopt the cost function J_u^{VO} (Equation (11)) that closely models the network congestion experienced only by the traffic of VO u due to the overall utilization of network resources requested by all virtual operators.

$$J_u^{VO}\left(\boldsymbol{X^u}, \boldsymbol{X^{-u}}\right) = \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} J_{ij}^{us} = \\ = \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} \left[a_{ij}^{s} + b_{ij}^{s}\left(\frac{\sum_{v\in\mathcal{U}} x_{ij}^{vs}}{B_{ij}^{s}}\right)^{\tau}\right]\left(\frac{x_{ij}^{us}}{B_{ij}^{s}}\right). \tag{11}$$

We can now formulate our distributed congestion control framework as a game. Formally, each VO $u \in \mathcal{U}$ solves the following MILP to find the best strategy that minimizes its traffic congestion:

$$\min \quad J_u^{VO}\left(\boldsymbol{X^u}, \boldsymbol{X^{-u}}\right) \tag{12}$$

s.t.

$$\sum_{\substack{i\in\mathcal{N}:\\ (j;i)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ji}} x_{ji}^{us} - \sum_{\substack{i\in\mathcal{N}:\\ (i;j)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} = \begin{cases} r^u & \forall j \in \mathcal{O}(u) \\ 0 & \forall j \in \mathcal{N} \\ -r^u & \forall j \in \mathcal{D}(u) \end{cases} \tag{13}$$

$$\sum_{vin\mathcal{U}} x_{ij}^{vs} \leq B_{ij}^{s} \qquad \forall (i;j) \in \mathcal{L}, \forall s \in \mathcal{S}_{ij} \tag{14}$$

$$\sum_{i\in\mathcal{N}} \sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} = w_j^u r^u \qquad \forall j \in \mathcal{P}_u \tag{15}$$

$$x_{ij}^{us} \geq 0 \qquad \forall (i;j) \in \mathcal{L}, s \in \mathcal{S}_{ij}. \tag{16}$$

Flow balance at the source, at each node $j \in \mathcal{N}$, and at the destination are imposed by constraints (13) for the traffic demand of user u. In particular, $\sum\sum x_{ij}^{us}$ and $\sum\sum x_{ji}^{us}$ are the total incoming and outgoing traffic flows, respectively.

Constraints (14) ensure that the VO's traffic routed on a link established between devices i and j, using transport service s, does not exceed the bandwidth assigned by the operator to such service (denoted by B_{ij}^{s}).

The set of constraints (15) forces a fraction w_j^u of the traffic to pass through the processing nodes ($\mathcal{P}_u$) selected by the VO u for traffic analysis and processing purposes. Therefore, the physical operator can select the set of physical nodes that minimize the congestion to perform the computational tasks requested and developed by the VO u for its data traffic. Finally, constraints (16) ensure the positiveness of the flow variables.

5.2 Distributed Congestion Control Solution

We now turn to the computation of the equilibrium solutions of our distributed congestion mitigation game. In particular, each VO u aims at minimizing his cost function J_u^{VO}. By definition, a Nash equilibrium is the solution to the individual utility optimization problem for each VO given all other virtual operators' actions. Since the distributed congestion control problem (12)-(16) is a nonlinear convex optimization problem, we propose to use the Karush-Kuhn-Tucker (KKT) conditions to determine the optimal solution. Indeed, the objective function $\sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} J_{ij}^{us}$ in (11) is differentiable in the x_{ij}^{us} variables and convex, and constraints (13)-(16) are linear in x_{ij}^{us}, hence the solution given by the KKT conditions is indeed an optimal solution of our problem. As in Section 4.2, we derive the best solution for VO u by detailing the Lagrangian function $\mathcal{L}^{VO}$, which can be written as expressed in Equation (17), where $\alpha^{\mathcal{O}(u)}$, β_j^u, $\gamma^{\mathcal{D}(u)}$, η_{ij}^{s}, θ_j^u and μ_{ij}^{us} are the Lagrangian multipliers.

Therefore, based on nonlinear convex programming theory [27], the following KKT conditions are necessary and sufficient for a solution $\boldsymbol{x}$ to be a Nash equilibrium:

$$\frac{\partial\mathcal{L}^{VO}}{\partial x_{ij}^{us}} = -\frac{a_{ij}^{s}}{B_{ij}^{s}} - \frac{b_{ij}^{s}}{B_{ij}^{s}}\left(\frac{\sum_{u\in\mathcal{U}} x_{ij}^{us}}{B_{ij}^{s}}\right)^{\tau} - \frac{b_{ij}^{s}\tau x_{ij}^{us}}{(B_{ij}^{s})^2}\left(\frac{\sum_{u\in\mathcal{U}} x_{ij}^{us}}{B_{ij}^{s}}\right)^{\tau-1} \\ - \alpha^{\mathcal{O}(u)} - \beta_j^u - \gamma^{\mathcal{D}(u)} - \eta_{ij}^{s} - \theta_j^u + \mu_{ij}^{us} = 0 \\ \forall u \in \mathcal{U}, \forall s \in \mathcal{S}_{ij}, \forall (i;j) \in \mathcal{L}$$

$$\sum_{\substack{i\in\mathcal{N}:\\ (j;i)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ji}} x_{ji}^{us} - \sum_{\substack{i\in\mathcal{N}:\\ (i;j)\in\mathcal{L}}} \sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} = \begin{cases} r^u & \forall u \in \mathcal{U}, \forall j \in \mathcal{O}(u) \\ 0 & \forall u \in \mathcal{U}, \forall j \in \mathcal{N} \\ -r^u & \forall u \in \mathcal{U}, \forall j \in \mathcal{D}(u) \end{cases}$$

$$\mathcal{L}^{VO}(\overrightarrow{x},\alpha,\beta,\gamma,\eta,\theta,\mu) = -\sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}} J_{ij}^{s,VO} + \alpha^{\mathcal{O}(u)}\left(-\sum_{\substack{i\in\mathcal{N}:\\(\mathcal{O}(u);i)\in\mathcal{L}}}\sum_{s\in\mathcal{S}_{\mathcal{O}(u);i}} x_{\mathcal{O}(u)i}^{us} + \sum_{\substack{i\in\mathcal{N}:\\(i;\mathcal{O}(u))\in\mathcal{L}}}\sum_{s\in\mathcal{S}_{i\mathcal{O}(u)}} x_{i\mathcal{O}(u)}^{us} + r^u\right) +$$
$$\sum_{\substack{j\in\mathcal{N}:\\ j\neq s_u\wedge j\neq d_u}}\beta_j^u\left(-\sum_{\substack{i\in\mathcal{N}:\\(i;j)\in\mathcal{L}}}\sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} + \sum_{\substack{i\in\mathcal{N}:\\(j;i)\in\mathcal{L}}}\sum_{s\in\mathcal{S}_{ji}} x_{ji}^{us}\right) + \gamma^{\mathcal{D}(u)}\left(-\sum_{\substack{i\in\mathcal{N}:\\(\mathcal{D}(u);i)\in\mathcal{L}}}\sum_{s\in\mathcal{S}_{\mathcal{D}(u);i}} x_{\mathcal{D}(u)i}^{us} + \sum_{\substack{i\in\mathcal{N}:\\(i;\mathcal{D}(u))\in\mathcal{L}}}\sum_{s\in\mathcal{S}_{i;\mathcal{D}(u)}} x_{i\mathcal{D}(u)}^{us} - r^u\right) +$$
$$\sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}}\eta_{ij}^s\left(B_{ij}^s - \sum_{u\in\mathcal{U}} x_{ij}^{us}\right) + \sum_{\substack{j\in\mathcal{N}:\\ j\in\mathcal{P}_u}}\theta_j^u\left(w_j^u r^u - \sum_{i\in\mathcal{N}}\sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us}\right) + \sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}}\mu_{ij}^{us}x_{ij}^{us} \tag{17}$$

$$\begin{aligned}
&\sum_{v\in\mathcal{U}} x_{ij}^{vs} - B_{ij}^s \le 0 && \forall(i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij}\\
&\eta_{ij}^s\left(\sum_{u\in\mathcal{U}} x_{ij}^{us} - B_{ij}^s\right) = 0 && \forall(i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij}\\
&\eta_{ij}^s \ge 0 && \forall(i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij}\\
&\sum_{i\in\mathcal{N}}\sum_{s\in\mathcal{S}_{ij}} x_{ij}^{vs} = w_j^v r^u && \forall v\in\mathcal{U}, j\in\mathcal{P}_u\subseteq\mathcal{N}\\
&x_{ij}^{vs}\ge 0 && \forall v\in\mathcal{U},(i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij}\\
&\mu_{ij}^{vs}x_{ij}^{vs} = 0 && \forall v\in\mathcal{U},(i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij}\\
&\mu_{ij}^{vs}\ge 0 && \forall v\in\mathcal{U},(i;j)\in\mathcal{L}, s\in\mathcal{S}_{ij}.
\end{aligned}$$

5.3 Socially-Aware Pricing Function

While the game theoretic control approach we illustrated achieves a stable operation point for the physical infrastructure in a distributed fashion, the total network congestion might result larger than the value obtained by the centralized control scheme. In order to reduce the gap between the congestion levels obtained with the centralized and distributed mechanisms, we propose to introduce additional individual VO pricing factors as a means to drive the output of our distributed framework (i.e., the VOs choices) towards the optimal configuration computed using the centralized scheme. To this end, we compare the social welfare (i.e., the overall network congestion) and the overall VOs utility (i.e., the sum over all VOs of their congestion costs) obtained using the centralized and distributed mechanisms, respectively.

As indicated in Equation (3), the total cost J incurred by the physical operator due to the overall network congestion caused by VOs requests can be computed summing the cost in (2) over all links and transport services. To better highlight the contribution of a generic VO u to the total congestion cost J and simplify the analysis, we reformulate Equation (3) as follows:

$$J = \sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}} J_{ij}^s = \sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}}\left[\frac{a_{ij}^s}{B_{ij}^s}x_{ij}^{us} + \frac{a_{ij}^s}{B_{ij}^s}\sum_{v\in\mathcal{U}\setminus\{u\}} x_{ij}^{vs} + \right.$$
$$\left. + \frac{b_{ij}^s}{(B_{ij}^s)^{(\tau+1)}}\left(\sum_{u\in\mathcal{U}} x_{ij}^{us}\right)^{\tau}\left(x_{ij}^{us} + \sum_{v\in\mathcal{U}\setminus\{u\}} x_{ij}^{vs}\right)\right]. \tag{18}$$

Similarly, by expanding the product of $J_u^{VO}\left(\boldsymbol{X^u},\boldsymbol{X^{-u}}\right)$ in Equation (11), the VO cost function can be equivalently expressed as:

$$J_u^{VO}\left(\boldsymbol{X^u},\boldsymbol{X^{-u}}\right) = \sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}} J_{ij}^{s,VO} =$$
$$= \sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}}\frac{a_{ij}^s}{B_{ij}^s}x_{ij}^{us} + \frac{b_{ij}^s}{(B_{ij}^s)^{(\tau+1)}}x_{ij}^{us}\left(\sum_{u\in\mathcal{U}} x_{ij}^{us}\right)^{\tau}. \tag{19}$$

At this stage, if we compare equations (18) and (19), which represent respectively the cost incurred by the physical infrastructure and by Virtual Operators (under a centralized and distributed framework, respectively), we can observe that a possible solution to guide the output of the distributed approach to an efficient one is to define *additional prices* to be charged to virtual operators (or, more in general, to users of the physical infrastructure). Such an approach forces VOs to cooperate, since it aligns user and system objectives, thus reducing the Price of Anarchy (PoA) of the distributed solution. To compute the additional price charged to each VO u, let us first evaluate the gap between the system and user costs of a link $(i;j)\in\mathcal{L}$ and transport service $s\in\mathcal{S}_{ij}$:

$$\begin{aligned}
\Delta_{ij}^{us} &= J_{ij}^s\left(\sum_{v\in\mathcal{U}} x_{ij}^{vs}\right) - J_{ij}^{us}\left(\boldsymbol{X^u},\boldsymbol{X^{-u}}\right) =\\
&= \frac{1}{B_{ij}^s}p_{ij}^s\left(\sum_{v\in\mathcal{U}} x_{ij}^{vs}\right)\left(\sum_{v\in\mathcal{U}} x_{ij}^{vs} - x_{ij}^{us}\right)\\
&= \frac{a_{ij}^s}{B_{ij}^s}\sum_{v\in\mathcal{U}\setminus\{u\}} x_{ij}^{vs} + \frac{b_{ij}^s}{(B_{ij}^s)^{(\tau+1)}}\left(\sum_{v\in\mathcal{U}\setminus\{u\}} x_{ij}^{vs}\right)\left(\sum_{v\in\mathcal{U}} x_{ij}^{vs}\right)^{\tau}.
\end{aligned}$$

Hence, each VO u incurs the following additional cost for transmitting x_{ij}^{us} units of traffic on link $(i;j)$, and using service s:

$$c_{ij}^{us} = \frac{\partial\Delta_{ij}^{us}}{\partial x_{ij}^{us}} = \frac{\tau b_{ij}^s}{(B_{ij}^s)^{(\tau+1)}}\left(\sum_{v\in\mathcal{U}\setminus\{u\}} x_{ij}^{vs}\right)\left(\sum_{u\in\mathcal{U}} x_{ij}^{us}\right)^{(\tau-1)} \tag{20}$$

and therefore the total additional cost that the physical operator can charge to VO u is:

$$c_u^{VO} = \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} c_{ij}^{su} =$$

$$= \sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} \frac{\tau b_{ij}^s}{(B_{ij}^s)^{(\tau+1)}} \left(\sum_{v\in\mathcal{U}\setminus\{u\}} x_{ij}^{vs} \right) \left(\sum_{v\in\mathcal{U}} x_{ij}^{vs} \right)^{(\tau-1)}. \tag{21}$$

The cost c_u^{VO} can be directly added to the first constraint of the KKT conditions in the distributed formulation of the congestion control problem described in Section 5.2. In the numerical results section we will compare the distributed approaches with the centralized optimization framework (both solved using KKT conditions) to demonstrate that the utilization of our proposed socially-aware pricing functions consistently reduces the performance loss of the fully distributed scheme.

6 Existence and Uniqueness of the Nash equilibrium

In this section, we demonstrate that our proposed distributed congestion mitigation game admits a unique Nash equilibrium. Since the vector of all virtual operators' strategies (denoted as $\boldsymbol{x}$) is selected from a convex, closed and bounded set, and the cost function $\sum_{(i;j)\in\mathcal{L}} \sum_{s\in\mathcal{S}_{ij}} J_{ij}^{us}$ is continuous and differentiable in all variables x_{ij}^{us}, according to Theorem 1 of [28] there exists at least one pure Nash equilibrium.

While the *existence* of a Nash equilibrium guarantees that the NFV system can operate in a stable configuration, it does not permit to analyze in depth the performance of our distributed orchestration mechanism. Indeed, network configurations corresponding to Nash equilibria might result in very different congestion levels, which might be far from the optimal solution. In order to evaluate the gap between our distributed approach and a centralized solution computed by the physical operator, we demonstrate that our game admits a *unique* Nash equilibrium.

To this aim, let us first slightly simplify the notations for ease of clarity. We denote by l (instead of $(i;j)$) a link chosen from the set $\mathcal{L}$, and by $\mathcal{S}_l$ (instead of $\mathcal{S}_{ij}$) the set of services on link l. Let J_l^{us} represent the cost perceived by VO u on link l using service s. The objective function of u can be rewritten as follows:

$$J_u^{VO}\left(\boldsymbol{X^u}, \boldsymbol{X^{-u}}\right) = \sum_{l\in\mathcal{L}} \sum_{s\in\mathcal{S}_l} J_l^{us}\left(\boldsymbol{X^u}, \boldsymbol{X^{-u}}\right)$$

$$= \sum_{l\in\mathcal{L}} \sum_{s\in\mathcal{S}_l} \left\{ \left[a_l^s + b_l^s \left(\frac{\sum_{v\in\mathcal{U}} x_l^{vs}}{B_l^s} \right)^{\tau} \right] \left(\frac{x_l^{us}}{B_l^s} \right) \right\} \tag{22}$$

The cost function J_l^{us} defined in (10) can be reformulated as follows:

$$J_l^{us}\left(\boldsymbol{X^u}, \boldsymbol{X^{-u}}\right) = \left[\frac{a_l^s}{B_l^s} x_l^{us} + \frac{b_l^s}{(B_l^s)^{(\tau+1)}} x_l^{us} \left(\sum_{v\in\mathcal{U}} x_l^{vs} \right)^{\tau} \right].$$

We now introduce the gradient column vector $g_l^s(\boldsymbol{x}_l^s)$:

$$g_l^s(\boldsymbol{x}_l^s) = \left[\nabla_1 J_l^{1s}(x_l^{1s}, \hat{x}_l^s), \ldots, \nabla_N J_l^{Ns}(x_l^{Ns}, \hat{x}_l^s) \right]^T, \tag{23}$$

where $N = |\mathcal{U}|$ is the total number of VOs, the vector $\boldsymbol{x}_l^s = (x_l^{1s}, x_l^{2s}, \ldots, x_l^{Ns})$ represents the VOs flows on link $(i;j)$ using service s, $\hat{x}_l^s = \sum_{v\in\mathcal{U}} x_l^{vs}$ is the overall flow, while ∇_u is the gradient operator with respect to the variable x_l^{us}, and T is the transpose operator.

The gradient $\nabla_u J_l^{us}(x_l^{us}, x_l^s), \forall u \in \mathcal{U}$ can be written as:

$$\nabla_u J_l^{us}(x_l^{us}, x_l^s) = \frac{a_l^s}{B_l^s} + \frac{b_l^s}{(B_l^s)^{(\tau+1)}} \left[(\hat{x}_l^s)^{\tau} + \tau x_l^{us} (\hat{x}_l^s)^{(\tau-1)} \right].$$

It follows from Theorem 2 in [28] and Corollary 2 in [29] that the Nash equilibrium is unique if for any set of vectors $\boldsymbol{x}_l^s$ and $\tilde{\boldsymbol{x}}_l^s, l \in \mathcal{L}, s \in \mathcal{S}_l$ with $\boldsymbol{x}_l^s \neq \tilde{\boldsymbol{x}}_l^s$ (in the vector sense) satisfying flow constraints (13), we have:

$$\sum_{l\in\mathcal{L}} \sum_{s\in\mathcal{S}_l} (\boldsymbol{x}_l^s - \tilde{\boldsymbol{x}}_l^s)^T (g_l^s(\boldsymbol{x}_l^s) - g_l^s(\boldsymbol{x}_l^s)) > 0,$$

or equivalently each term in the above summation is positive:

$$(\boldsymbol{x}_l^s - \tilde{\boldsymbol{x}}_l^s)^T (g_l^s(\boldsymbol{x}_l^s) - g_l^s(\boldsymbol{x}_l^s)) > 0 \qquad \forall l \in \mathcal{L}, s \in \mathcal{S}_l.$$

We reasonably assume that the traffic load is such that all capacity constraints in (6) are implicitly satisfied. In fact, admission control mechanisms can be adopted to ensure that the total traffic admitted in the network does not exceed the available physical resources. On the other hand, if the network operates in a high-traffic regime (i.e., capacity constraints are active), we observe that we have a game with *coupled constraints*. In this case, the solution concept that can be adopted is a special case of Nash equilibrium, called *normalized Nash equilibrium* [28], which is out of the scope of the present study.

We first consider the special case of $\tau = 1$. The u^{th} element of the gradient column vector $g_l^s(\boldsymbol{x}_l^s)[u]$ corresponding to VO u is equal to $\frac{a_l^s}{B_l^s} + \frac{b_l^s}{(B_l^s)^2} \times (\hat{x}_l^s + x_l^{us})$. At this point, we need to demonstrate that the Jacobian of $g_l^s(\boldsymbol{x}_l^s)$, denoted as $G_l^s(\boldsymbol{x}_l^s)$, is positive definite:

$$G_l^s(\boldsymbol{x}_l^s) = \left\{ \frac{\partial^2 J_l^{us}(x_l^{us}, \hat{x}_l^s)}{\partial x_l^{us} \partial x_l^{vs}} \right\}_{u,v\in\mathcal{U}}$$

$$= \left(\boldsymbol{1}\boldsymbol{1}^T + \boldsymbol{I} \right) \times \frac{b_l^s}{(B_l^s)^2}, \tag{24}$$

with $\boldsymbol{1}$ an N-dimensional column vector of ones and $\boldsymbol{I}$ the identity matrix. For the reader's convenience, we have:

$$\frac{\partial^2 J_l^{us}(x_l^{us}, \hat{x}_l^s)}{\partial x_l^{us} \partial x_l^{vs}} = \frac{b_l^s}{(B_l^s)^2}, \tag{25}$$

and

$$\frac{\partial^2 J_l^{us}(x_l^{us}, \hat{x}_l^s)}{\partial^2 x_l^{us}} = \frac{2 \times b_l^s}{(B_l^s)^2}. \tag{26}$$

It is easy to verify that the Jacobian G_l^s is positive definite. Therefore, the Nash equilibrium is unique in the case of $\tau = 1$.

Let us now demonstrate that the NE is unique for general τ values ($\tau > 1$). In this case,

$$\frac{\partial^2 J_l^{us}(x_l^{us}, \hat{x}_l^s)}{\partial x_l^{us} \partial x_l^{vs}} = \\ = \frac{b_l^s}{(B_l^s)^{(\tau+1)}} \left[\tau(\hat{x}_l^s)^{(\tau-1)} + \tau(\tau-1) x_l^{us} (\hat{x}_l^s)^{(\tau-2)} \right],$$

$$\frac{\partial^2 J_l^{us}(x_l^{us}, \hat{x}_l^s)}{\partial^2 x_l^{us}} = \\ = \frac{b_l^s}{(B_l^s)^{(\tau+1)}} \left[2\tau(\hat{x}_l^s)^{(\tau-1)} + \tau(\tau-1) x_l^{us} (\hat{x}_l^s)^{(\tau-2)} \right].$$

Therefore, the Jacobian becomes

$$G_l^s(\boldsymbol{x}_l^s) = \frac{b_l^s \tau (\hat{x}_l^s)^{(\tau-2)}}{(B_l^s)^{(\tau+1)}} [\boldsymbol{\mathcal{G}}_l + \hat{x}_l^s \boldsymbol{I}], \tag{27}$$

where $\boldsymbol{\mathcal{G}}_l = \boldsymbol{q}_l \times \boldsymbol{1}^T$, and $\boldsymbol{q}_l$ is an N-dimensional column vector whose i-th element, $\boldsymbol{q}_l(i)$, is equal to $\hat{x}_l^s + (\tau-1) x_l^{is}$.

According to the proof in Section 4.2 of [30], the Nash equilibrium is unique for general values of τ if all flow variables are strictly non-negative, which is a very general assumption easily verified in real network operation scenarios.

7 Numerical Results

This section presents and discusses numerical results obtained using the proposed approaches to mitigate congestion in virtual networks. Specifically, we test the sensitivity of the proposed models to key parameters like (i) the number of virtual operators, and (ii) their requests. We formalize the distributed and centralized congestion mitigation problems in AMPL (A Mathematical Programming Language) and solve the instances, which are generated by a standalone program, using the SNOPT 7.2 (Sparse Nonlinear OPTimizer) solver.

In the following, we first describe the settings of our simulations. Then, we analyze the performance achieved by the proposed congestion mitigation schemes.

7.1 Experimental Methodology

In our analysis, we consider real network topologies obtained from the on-line archive maintained by the Internet Topology Zoo project[2]. More specifically, we extend the *Abilene* (Fig. 2) and *Geant* topologies (Fig. 3), which contain 11 and 40 nodes connected through 28 and 122 directed links, respectively.

The capacity of all links has been normalized to simplify the analysis of the network congestion. Furthermore, we vary the number of transmission services for each link in the range $[2, 5]$ to quantify the control overhead that may be introduced using multiple virtual services. Note that the link capacity has been evenly divided among all transmission services, since we assume the implementation of a round robin scheme for scheduling multiple virtual services.

For each Virtual Operator $u \in \mathcal{U}$, we randomly select the source ($\mathcal{O}(u)$) and the destination ($\mathcal{D}(u)$) of its data connection, which represent the ingress and egress points of the VO. The bandwidth demand of each VO is drawn according to a uniform distribution in $[0, 1]$. Note that the granularity of our representation for VOs' requests is highly flexible: in fact, we can easily represent VOs' demands with multiple connections by simply defining a different VO for each pair of ingress/egress points.

As for *processing nodes*, which can implement key network functionalities such as traffic analysis, filtering, caching/storage, security and billing, we vary their number in the $[2, 4]$ range, randomly selecting the network nodes that implement the processing functions for each VO. We further assume that all network functions are replicated on all processing nodes, for both reliability and load balancing purposes. Therefore, we split equally the portion of data traffic that must flow through them, i.e., $w_j^u = \frac{1}{|\mathcal{P}_u|}, j \in \mathcal{P}_u$.

In addition to testing our distributed approaches (with and without the socially-aware pricing scheme introduced in Section 5.3, identified by the labels "*S.D.A.*" and "*D.A.*", respectively), and the Centralized Approach ("*C.A.*"), we further consider a heuristic algorithm based on the Shortest Path Tree (*SPT*) to compute the resource allocation that minimizes the congestion costs experienced by VOs. Indeed, routes in physical networks

2. Available at http://www.topology-zoo.org/

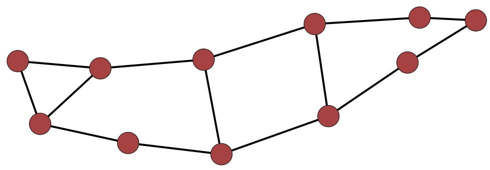

Fig. 2: The *Abilene* topology (11 nodes, 28 directed links).

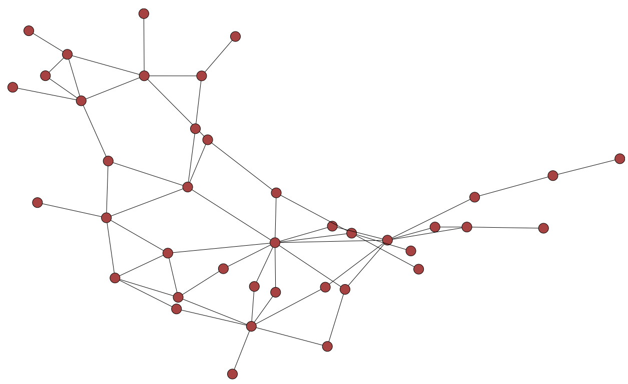

Fig. 3: The *Geant* topology (40 nodes, 122 directed links).

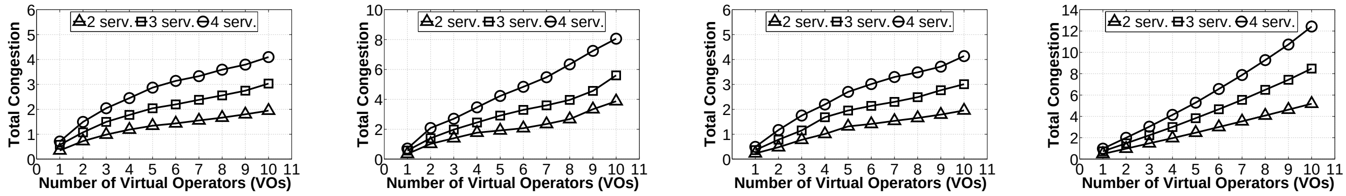

(a) C.A. (b) D.A. (c) S.D.A. (d) SPT

Fig. 4: Average total congestion using the centralized, distributed, socially-aware, and heuristic approaches as a function of the number of Virtual Operators (VOs) for different numbers of services in the *Abilene* network.

are usually computed by routing protocols using the SPT algorithm. Therefore, a physical operator can use the routing protocol that is running on its network to map VOs requests over the physical infrastructure, as suggested in [31]. Specifically, the physical operator can merge shortest paths from the source $\mathcal{O}(u)$ to the processing nodes $\mathcal{P}_u$, and from the processing nodes to the destination $\mathcal{D}(u)$ to satisfy the request of a VO u, provided that the residual capacity on all links is enough to satisfy the VO bandwidth demand r^u.

In order to evaluate the performance achieved by different congestion minimization techniques (distributed, centralized and heuristic) which we designed to compose virtual network functions, we consider the *total congestion* caused to the entire network infrastructure and the *number of links* (NL) selected to satisfy the VOs requests, defined according to Equations (4) and (28), respectively:

$$NL = \sum_{(i;j)\in\mathcal{L}} \mathbb{I}\left(\sum_{u\in\mathcal{U}}\sum_{s\in\mathcal{S}_{ij}} x_{ij}^{us} > 0\right) \tag{28}$$

where $\mathbb{I}(\cdot)$ is the indicator function, which is equal to 1 if the condition is satisfied, 0 otherwise. Furthermore, we compare the performance of the distributed and centralized approaches computing the *Price of Anarchy* (PoA). The PoA was originally defined in [32] to capture the worst case selfish performance of a simple game of N players that compete for the utilization of M parallel links. In particular, the PoA can be defined in our context as the ratio between the worst distributed (Nash equilibrium) cost and the best optimal centralized solution cost, as follows:

$$PoA = \frac{\sum_{u\in\mathcal{U}} J_u^{VO}}{\sum_{(i;j)\in\mathcal{L}}\sum_{s\in\mathcal{S}_{ij}} J_{ij}^{s}}. \tag{29}$$

For each network scenario, the results we obtained represent the average of the performance metric measured over 500 network instances.

7.2 Abilene Network Scenario

In this subsection we consider the *Abilene* network topology, and vary the number of VOs and services in the ranges $[1, 10]$ and $[2, 4]$, respectively. Each VO selects randomly the ingress and egress points (i.e., source and destination) of its virtual network and 4 processing devices among the remaining nodes (i.e., $|\mathcal{P}_u| = 4$) in order to perform intensive computational tasks and run virtual machines for management operations. The normalized traffic flowing on each virtual network is fixed to 0.05. Furthermore, parameters a_{ij}^s and b_{ij}^s in the congestion cost function (1) are set to 1 for all links $(i; j) \in \mathcal{L}$ and services $s \in \mathcal{S}$, while $\tau = 2$, in order to simulate the *nonlinear* overhead caused by the access coordination mechanism of the transmission link and the increasing congestion.

Figures 4(a), 4(b), 4(c), and 4(d) show the *total* network congestion as a function of the number of virtual operators for different numbers of services (viz., 2, 3 and 4) using the centralized, distributed, socially-aware, and heuristic approaches, respectively. It can be observed that the centralized approach (referred to as "C.A", Fig. 4(a)) computes the network configuration (i.e., traffic flows and resource allocations) that achieves the lowest congestion. However, the dynamic pricing scheme that we have designed for our distributed approach drives the VO strategies towards efficient solutions that well approach the optimal configuration. Indeed, the maximum ratio of the total network congestion between the distributed and centralized approaches is lower than 2, and always well below the congestion level caused by the centralized heuristic algorithm. In fact, by using the SPT algorithm to select network paths and services for each virtual network, the total congestion grows up to three times with respect to the centralized approach.

Furthermore, we underline that the distributed approach that uses the socially-aware pricing scheme we proposed in this work (referred to as "S.D.A" in Figure 4(c)) shows very good performance, even when compared to the centralized approach. In fact, the total network congestion obtained with S.D.A. is always very close to that achieved by C.A. (in particular, the gap is always less than 0.09 with 4 transmission services, and even lower for 2 and 3 services). This indeed confirms the key feature of the socially-aware pricing function, which guides the distributed system towards efficient resource allocation solutions that are close to the optimum.

The efficiency of the network configuration achieved by VOs in a fully distributed fashion using our pricing rule is highlighted in more detail in Figure 5, which illustrates the empirical Cumulative Distribution Function (CDF) of the PoA, as defined in Equation (29). The line labeled "1 VO (C.A.)" corresponds to the centralized

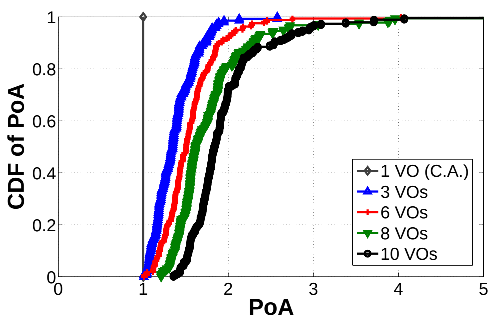

Fig. 5: The CDF of the PoA as a function of the number of VOs in the *Abilene* network.

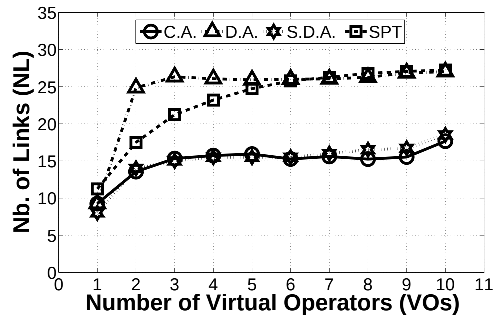

Fig. 6: Average number of used links (with 3 transmission services and 4 processing nodes) for the centralized, distributed, socially-aware, and heuristic approaches as a function of the number of Virtual Operators (VOs) in the *Abilene* network.

solution, whereas the other curves show the CDF of the PoA computed in scenarios with increasing VOs requests, each composed of 500 network scenarios. The figure shows that in more than 75% of the scenarios, and even for a large number of VOs, the network configuration obtained by individual VOs choices results in a total network congestion at most twice the best congestion level, which can be achieved only using the optimal, centralized approach, which is computationally cumbersome. The figure further shows a small variance for the PoA (e.g., $PoA < 3$ in 95% of network scenarios).

As for the S.D.A. mechanism, we measured experimentally that the total congestion achieved in the worst Nash equilibrium is even lower, and in particular its ratio with respect to the social optimum is less than 1.3 in more than 95% of all considered scenarios. This means that the distributed solution coupled with our dynamic pricing rule is very effective, and is only slightly affected by fluctuations in the traffic distribution, thus producing more robust and stable network configurations.

We observe that the efficiency obtained using the distributed scheme comes at the cost of a higher resource utilization, as illustrated in Figure 6 (we report the results with 3 transmission services, since others are overlapped). Indeed, VOs tend to select unloaded (or lightly loaded) network paths to avoid congestion. As a consequence, when VOs can decide autonomously, they fully use the network capacity, thus increasing the number of used links with respect to the centralized solution. However, under the socially-aware pricing-based approach, VOs tend to share common links to reduce their pricing function that depends on the total flow that passes through the links and the choices of other VOs. As a consequence, the number of used links under S.D.A. practically overlaps with that of the C.A. curve. Note, finally, that by using the centralized heuristic approach (SPT), the number of used links increases with the number of VOs requests, thus reducing the gain of the centralized scheme on the fully distributed approach.

7.3 Geant Network Scenario

We now consider the *Geant* network topology (40 nodes, 122 links), and evaluate the impact of different parameters on the proposed approaches. All parameters are the same as in the previous scenario.

Figures 7(a), 7(b), 7(c), and 7(d) illustrate the total network congestion achieved in such topology for the considered approaches (optimal centralized, distributed, socially-aware, and SPT), with 2 processing nodes, different transmission services (viz., 2, 3 and 4), a normalized rate equal to 0.05 for each VO request, and an increasing number of VOs (from 1 to 10). The trends confirm the findings we highlighted for the Abilene topology, which can be observed also in the larger Geant network. Specifically, our proposed distributed approach obtains solutions that are at most double the cost of the optimal approach, but operating in a fully distributed way.

A closer look at the empirical CDF of the PoA, which is depicted in Figure 8 for different numbers of VOs (the processing nodes and services are set to 2 and 4, respectively) shows that even for a large number of VOs, in more than 60% of the scenarios the ratio between their performance is always less than 2. Conversely, for fewer VOs the performance of our proposed distributed approach is even closer to the optimal one, achieved in a centralized way. However, in larger network scenarios like the *Geant* topology, by using our distributed scheme, VOs select more links than any centralized solution, thus increasing the total network congestion as illustrated in Figure 9. Indeed, the higher is the network size, the higher is the number of available paths that are used to build virtual topologies, and the higher is the congestion level of the entire network. In contrast, either optimal or heuristic centralized solutions can exploit the complete knowledge of the status of the physical infrastructure and VOs requests to select a network configuration that leads to lower congestion than autonomous decisions.

As for the distributed approach with a social pricing scheme, both the total network congestion and the number of used links exhibit a trend which is very close to the one obtained by the centralized approach. The gap between the total congestion of S.D.A and C.A. is always less than 0.92 and 1.125 for 2 and 4 transmission services, respectively. Indeed, this is expected since the proposed pricing function drives the VOs towards efficient decisions, close to the optimal ones.

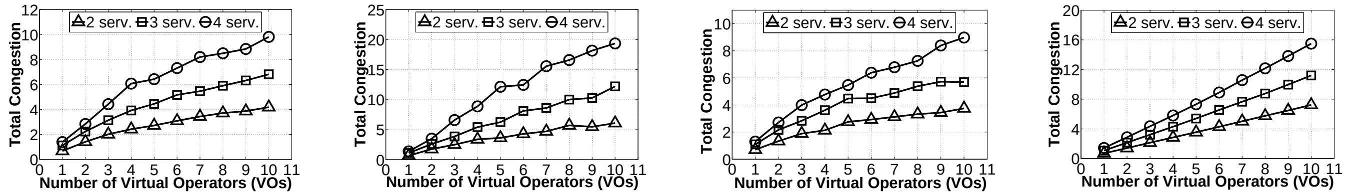

(a) C.A. (b) D.A. (c) S.D.A. (d) SPT

Fig. 7: Average total congestion using the centralized, distributed, socially-aware, and heuristic approaches as a function of the number of Virtual Operators (VOs) for different numbers of services in the *Geant* network ($|\mathcal{P}_u| = 2$, $|\mathcal{S}| = 4$).

Finally, in all considered scenarios we measured the *average computing time* necessary to obtain optimal centralized as well as distributed solutions, using the SNOPT 7.2 solver on a Dual Intel Xeon E5-2630 v2 @ 2.60GHz machine with 64 GByte of RAM. In the Abilene topology, our approaches took, in the worst case (the centralized solution), up to $10.5s$, and up to $50s$ in the Geant topology. This timing allows the physical operator to reconfigure network resources on-the-fly. We underline that, in the considered scenarios, our proposed distributed approaches proved to be at least 10 times faster than the centralized ones, while guaranteeing very good solutions thanks to the socially-aware pricing function we implemented.

8 Conclusion

This paper provided a novel, holystic approach to the congestion mitigation problem in virtual networks by proposing and comparing both centralized and fully distributed solutions.

In particular, we first proposed an optimization approach to design novel orchestration mechanisms to optimally control and reduce the resource congestion of a physical infrastructure based on the NFV paradigm. We then illustrated a fully distributed approach, based on a *dynamic pricing strategy* of network resources, to mitigate congestion in virtual networks where the physical infrastructure is shared among multiple Virtual Operators. In our approach, which we model and study using Game Theory, VOs independently select their best configurations, and achieve stable and efficient network allocations due to the existence and uniqueness of the network equilibrium which our game exhibits.

Numerical results, which have been obtained in both medium and large-size realistic ISP topologies, show that the proposed distributed solution significantly reduces resource congestion, well approaching the performance that can be achieved using a centralized network orchestration scheme. In particular, in large networks with 40 nodes (i.e., the Geant topology) and several transmission services activated for each link (up to 4), our distributed scheme well approaches the total congestion experienced using the optimal solution. The maximum measured ratio with respect to a network optimum (which can be obtained only through computationally intensive operations) is indeed always lower than 2, and it is significantly lower (and close to the optimum) when we adopt our proposed socially-aware pricing function. Furthermore, our proposed distributed approach outperforms heuristics commonly adopted for solving virtual embedding problems in small-to-medium size network scenarios.

For these reasons, our proposed congestion mitigation framework represents a very promising approach for operators to manage network resources in NFV-based systems in an efficient, fully distributed and dynamic fashion.

References

[1] ETSI. Network Functions Virtualisation - Introductory White Paper (v.2). http://portal.etsi.org/nfv/nfv_white_paper2.pdf.

[2] Raj Jain and Subharthi Paul. Network Virtualization and Software Defined Networking for Cloud Computing - A Survey. *IEEE Communications Magazine*, pages 24–31, vol. 51(11), November 2013.

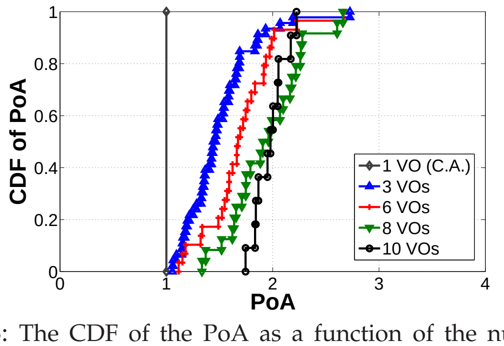

Fig. 8: The CDF of the PoA as a function of the number of VOs in the *Geant* network ($|\mathcal{P}_u| = 4$, $|\mathcal{S}| = 2$).

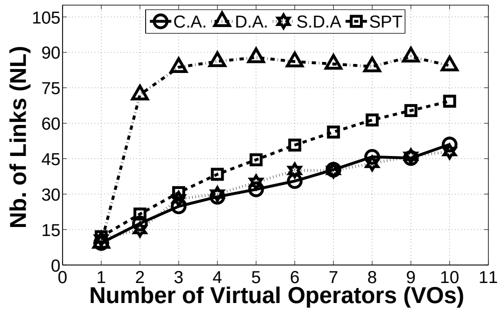

Fig. 9: Average number of used links for the centralized, distributed, socially-aware distributed, and heuristic approaches as a function of the number of VOs in the *Geant* network ($|\mathcal{P}_u| = |\mathcal{S}| = 2$).

[3] J. Elias, F. Martignon, S. Paris, and J. Wang. Optimization models for congestion mitigation in virtual networks. In *Proc. of the 22nd IEEE International Conference on Network Protocols (ICNP 2014)*, The Research Triangle, NC, October 2014.

[4] D. Ardagna, B. Panicucci, and M. Passacantando. Generalized nash equilibria for the service provisioning problem in cloud systems. *IEEE Trans. on Services Computing*, 6(4):429–442, 2013.

[5] E. Altman, T. Basar, T. Jimenez, and Nahum N. Shimkin. Competitive routing in networks with polynomial costs. *IEEE Trans. on Automatic Control*, 47(1):92–96, 2002.

[6] A. Khan, A. Zugenmaier, D. Jurca, and W. Kellerer. Network virtualization: a hypervisor for the Internet? *IEEE Communications Magazine*, pages 136–143, vol. 50(1), January 2012.

[7] Q. Duan, Y. Yan, and A.V. Vasilakos. A Survey on Service-Oriented Network Virtualization Toward Convergence of Networking and Cloud Computing. *IEEE Trans. on Network and Service Management*, pages 373–392, vol. 9(4), December 2012.

[8] G.M. Roy, S.K. Saurabh, N.M. Upadhyay, and P.K. Gupta. Creation of virtual node, virtual link and managing them in network virtualization. *IEEE WICT*, pages 738–742, 2011.

[9] X. Costa-Perez, J. Swetina, Tao Guo, R. Mahindra, and S. Rangarajan. Radio access network virtualization for future mobile carrier networks. *IEEE Communications Magazine*, pages 27–35, vol. 51(7), July 2013.

[10] A. Manzalini et al. Software-defined networks for future networks and services. main technical challenges and business implications. *White Paper based on the IEEE Workshop SDN4FNS*, 2014.

[11] M. Yu, Y. Yi, J. Rexford, and M. Chiang. Rethinking Virtual Network Embedding: Substrate Support for Path Splitting and Migration. *ACM SIGCOMM Comput. Commun. Rev.*, 38(2):17–29, Apr. 2008.

[12] A. Jarray and A. Karmouch. Decomposition Approaches for Virtual Network Embedding With One-Shot Node and Link Mapping. *IEEE/ACM Trans. Networking*, 2014.

[13] N.M.M.K. Chowdhury, M.R. Rahman, and R. Boutaba. Virtual network embedding with coordinated node and link mapping. In *IEEE INFOCOM*, pages 783–791, 2009.

[14] X. Cheng, S. Su, Z. Zhang, H. Wang, F. Yang, Y. Luo, and J. Wang. Virtual Network Embedding Through Topology-Aware Node Ranking. *ACM SIGCOMM Comp. Commun. Rev.*, 41(2):38–47, 2011.

[15] J. F. Botero, X. Hesselbach, M. Duelli, D. Schlosser, A. Fischer, and H. De Meer. Energy Efficient Virtual Network Embedding. *IEEE Communications Letters*, 16(5):756–759, 2012.

[16] A. Leivadeas, C. Papagianni, and S. Papavassiliou. Efficient resource mapping framework over networked clouds via iterated local search-based request partitioning. *IEEE Trans. on Parallel and Distributed Systems*, 24(6):1077–1086, 2013.

[17] A. Fischer, J.F. Botero, M. Till Beck, H. de Meer, and X. Hesselbach. Virtual Network Embedding: A Survey. *IEEE Communications Surveys Tutorials*, 15(4):1888–1906, 2013.

[18] M. Schöller, M. Stiemerling, A. Ripke, and R. Bless. Resilient Deployment of Virtual Network Functions. In *IEEE RNDM*, 2013.

[19] A. Jarray and A. Karmouch. VCG auction-based approach for efficient Virtual Network embedding. *IFIP/IEEE International Symposium on Integrated Network Management*, pages 609–615, 2013.

[20] A. Gember, R. Viswanathan, C. Prakash, R. Grandl, J. Khalid, S. Das, and Aditya Akella. OpenNF: Enabling Innovation in Network Function Control. In *ACM SIGCOMM*, Chicago, Illinois, August 2014.

[21] Ilenia Tinnirello, Giuseppe Bianchi, Pierluigi Gallo, Domenico Garlisi, Francesco Giuliano, and Francesco Gringoli. Wireless mac processors: Programming mac protocols on commodity hardware. In *Proceedings IEEE INFOCOM 2012*, pages 1269–1277. IEEE, 2012.

[22] Giuseppe Bianchi, Marco Bonola, Antonio Capone, and Carmelo Cascone. Openstate: programming platform-independent stateful openflow applications inside the switch. *ACM SIGCOMM Computer Communication Review*, 44(2):44–51, 2014.

[23] Jan Broeckhove. Network aware scheduling for virtual machine workloads with interference models. *IEEE Trans. on Services Computing*, in press, 2014.

[24] D.P. Bertsekas and R.G. Gallagher. *Data Networks*. Prentice Hall, second edition, 1992.

[25] Bernard Fortz and Mikkel Thorup. Internet traffic engineering by optimizing ospf weights. *IEEE INFOCOM*, pages 519–528, 2000.

[26] Mosharaf Chowdhury, Muntasir Raihan Rahman, and Raouf Boutaba. Vineyard: virtual network embedding algorithms with coordinated node and link mapping. *IEEE/ACM Trans. on Networking*, 20(1):206–219, 2012.

[27] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, seventh edition, 2009.

[28] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, July 1965.

[29] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Netw.*, 1(5):510–521, 1993.

[30] E. Altman, J. Elias, and F. Martignon. A game theoretic framework for joint routing and pricing in networks with elastic demands. In *ICST/ACM VALUETOOLS*, pages 1–10, 2009.

[31] Yong Zhu and Mostafa H Ammar. Algorithms for assigning substrate network resources to virtual network components. In *IEEE INFOCOM*, pages 1–12, 2006.

[32] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of STACS'99*, pages 404–413, 1999.

Jocelyne Elias is Associate Professor at Paris Descartes University since September 2010. She held a Post-doc position at the Department of Information and Mathematical Methods of University of Bergamo (2009-2010). She obtained her Ph.D. in Information and Communication Technology at the Department of Electronics and Information of Politecnico di Milano in 2009. Her main research interests include network optimization, and in particular modeling and performance evaluation of networks (Cognitive Radio, Wireless, Overlay and Wired Networks), as well as the application of Game Theory to resource allocation, spectrum access, and pricing problems.

Fabio Martignon received the M.S. and the Ph.D. degrees in telecommunication engineering from the Politecnico di Milano in October 2001 and May 2005, respectively. He has been associate professor at University of Bergamo, and he is now full professor in LRI (Laboratory for Computer Science) at Paris Sud University, and member of Institut Universitaire de France. His current research activities include cognitive radio networks, content-centric networks, network planning and game theory applications to networking problems.

Stefano Paris is an Associate Professor at the Department of Computer Science (LIPADE) at Paris Descartes University. He received his M.S. degree in Computer Engineering from University of Bergamo in 2007, and the Ph.D. in Information Engineering from Politecnico di Milano in 2011. His main research interests include topics related to optimization, performance evaluation and security of wireless networks. He is a member of the IEEE Communication Society.

Jianping Wang is an associate professor in the Department of Computer Science at City University of Hong Kong. She received the B.S. and the M.S. degrees in computer science from Nankai University, Tianjin, China in 1996 and 1999, respectively, and the Ph.D. degree in computer science from the University of Texas at Dallas in 2003. Jianping's research interests include dependable networking, optical networks, cloud computing, service oriented networking and data center networks.