

# Protection Strategies for Virtual Network Functions Placement and Service Chains Provisioning\*

**Ali Hmaity**

*Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy*

**Marco Savi**

*FBK CREATE-NET, 38123 Povo, Trento, Italy*

**Francesco Musumeci, Massimo Tornatore, and Achille Pattavina**

*Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy*

**Telecom operators worldwide are witnessing squeezed profit margins mainly due to hyper-competition. Hence, new business models/strategies are needed to help operators reduce Operational and Capital Expenditures. In this context, the Network Function Virtualization (NFV) paradigm, which consists of running Virtual Instances of Network Functions (NFs) in Commercial-Off-The-Shelf (COTS) hardware, represents a solid alternative. Virtual Network Functions (VNFs) are then concatenated together in a sequential order to form service chains (SCs) that provide specific Internet services. In this article, we study different approaches to provision SCs with resiliency against single-link and single-node failures. We propose three Integer Linear Programming (ILP) models to jointly solve the problem of VNF placement and traffic routing, while guaranteeing resiliency against single-link and/or single-node failures. Specifically, we focus on the trade-off between the conflicting objectives of meeting SCs latency requirements and consolidating as many as possible VNFs in NFV-capable nodes. We show that providing resiliency against both single-link and single-node failures comes at twice the amount of resources in terms of NFV-capable nodes, and that for latency-critical services providing resiliency against single-node failures comes at the same cost with respect to resiliency against single-link and single-node failures. Finally, we discuss important insights about the deployment of bandwidth-intensive SCs. © 2017 Wiley Periodicals, Inc. NETWORKS, Vol. 000(00), 000–000 2017**

Received February 2017; accepted August 2017

Correspondence to: A. Hmaity; e-mail: ali.hmaity@polimi.it

\*A preliminary version of this article appeared in A. Hmaity, M. Savi, F. Musumeci, M. Tornatore, A. Pattavina, Virtual Network Function Placement For Resilient Service Chain Provisioning, in Proceedings of International Workshop on Resilient Networks Design and Modeling 2016 (RNDM 2016) [6].

DOI 10.1002/net.21782

Published online 00 00 2017 in Wiley Online Library

(wileyonlinelibrary.com).

© 2017 Wiley Periodicals, Inc.

**Keywords:** Network Function Virtualization; service chaining; protection strategies; Virtual Network Embedding; 5th Generation; edge computing

## 1. INTRODUCTION

Network operators rely on hardware appliances to provide Internet services. Such services are usually provided thanks to the adoption of a purpose-built hardware that implements specific network functions (i.e., Firewalls, Network Address Translator (NAT), Intrusion Detection Prevention System (IDPS), etc.)<sup>1</sup> within the network [21]. From the cost point of view, telecom operators are witnessing a decrease of the revenue-per-bit, which is envisioned to be even lower than the cost-per-bit, due to the competition from Over-The-Tops (OTTs). The applications introduced by OTTs (e.g., Voice-over-IP (VoIP)) leave the Internet Service Provider (ISP) responsible for only transporting the information, hence contributing heavily in their revenue decrease.

Network Function Virtualization (NFV) is a new architectural paradigm that was proposed to improve the flexibility of network service provisioning and reduce the time to market of new services [14]. NFV can revolutionize how network operators design their infrastructure, by leveraging virtualization, to separate software instances from hardware appliances, and decoupling functionalities from locations for faster service provisioning. NFV supports the instantiation of Virtual Network Functions (VNFs) through software virtualization techniques and runs them on Commercial-Off-The-Shelf (COTS) hardware. Hence, the virtualization of network functions opens the way to the provisioning of new services without the installation of new equipment. It is clear

<sup>1</sup> A list of acronyms to ease the reading is presented in the last page of this article.

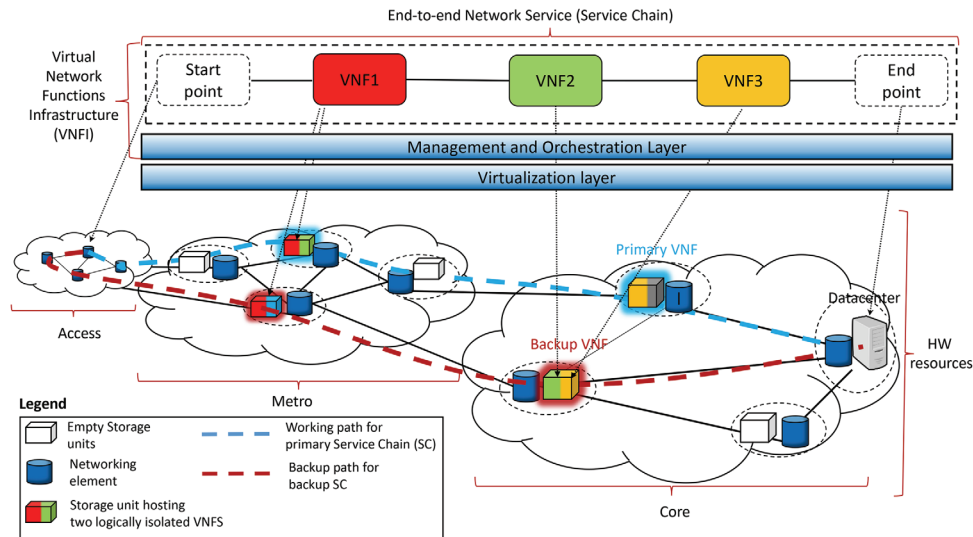


FIG. 1. Resilient deployment of an end-to-end network service exploiting redundant Virtual Network Functions, consolidated in multiple COTS hardware, distributed across metro/core networks. [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

that NFV brings a whole new dimension to the landscape of the telecommunication industry market due to the possibility of reducing capital investments, energy consumption by consolidating network functions, and by introducing tailored services based on customer needs.

Moreover, NFV simplifies service deployment by exploiting the concept of *service chaining* [5]: a service chain (SC) is a sequential concatenation of VNFs and/or hardware appliances to provide a specific Internet service (e.g., VoIP, Web Service, etc.) to the users. Deploying NFV solutions in operational networks requires solving multiple issues related to performance, availability, security, and survivability. One important key design in an NFV framework is the ability of the NFV-Management and Orchestration (NFV-MANO) component to ensure service continuity. Such an objective translates into many requirements that the Virtual Network Function Infrastructure (NFVI) must satisfy, among which are resiliency and geo-redundancy requirements. Hence, the deployment of SCs must meet a given resiliency level and aim to consolidate as much as possible the VNFs within *NFV-nodes* (i.e., those nodes in the physical network that can be used to instantiate VNFs), as an indiscriminate distribution of VNFs instances would lead to a cost increase.

Figure 1 clarifies some of the introduced concepts in more detail. In the top part an end-to-end network service is shown, represented through a SC composed of three VNFs, denoted in red, green, and yellow. Such a SC is deployed on top of the orchestration layer, which leverages the virtualization layer to instantiate redundant instances of the VNFs in the metro/core physical networks, and steers traffic according to the specific order required by the network service. Note that the primary and backup paths traverse, in an ordered manner, network nodes holding storage units and hosting an instance of the required VNF. For instance, the primary path, denoted by a dashed blue line, crosses a network node in the metro

network where a storage unit hosts the red and green VNFs, and successively the traffic is steered to another network node, in the core network, where it gets processed within another storage unit hosting the yellow VNF. Note that the storage unit hosting the yellow VNF hosts also an instance of another VNF, which is not required by the SC (denoted in gray). This is meant to explain that a storage unit can host multiple VNFs that are logically isolated, and that traffic gets processed only by the required VNFs. Finally, the placement and routing of traffic through VNFs, as early described, provides resilient deployment of the end-to-end network service (or SC) and hence satisfies redundancy requirements.

In this article, we address the issue of resiliency against single-link/node failures. We tackle this problem with the aim of investigating the trade-off between latency requirements and the amount of resources required in terms of *NFV-nodes*. To this objective, we propose three different Integer Linear Programming (ILP) models to jointly solve the VNF placement and SC traffic routing problems while guaranteeing resiliency against single-link failures, single-node failures, and single-node/link failures.

The main contributions of this work are the following:

- We show the amount of resources needed, for each resilient design, and compare them with the unprotected scenario, for SCs with different latency and bandwidth requirements.
- We observe that traffic processing by VNFs causes a variation of data-rate, represented by compression factors, and include this aspect into the optimization framework.
- We investigate the trade-off between node consolidation and the average hop count for different resilient design scenarios.
- We solve the ILP models considering the conflicting objectives of VNFs consolidation within *NFV-nodes* and load balancing and derive important insights on how different SCs are deployed in such scenario.

Numerical results indicate that providing resiliency against single-node/link failures comes at the same cost as the resiliency against single-node failures. Moreover, for latency-stringent SCs, we find that, to provide resiliency against single-link failures, the operator must place backup VNFs in physically disjoint locations. In addition, for SCs with a loose latency requirement, we observe that a trade-off exists between the average length of primary and backup paths. Finally, we analyze the effect of bandwidth requirement of two SCs with the same latency requirement and find that balancing the load on a physical link is beneficial for small values of node capacity, expressed in terms of number of Central Processing Unit (CPU) cores it is equipped with.

The rest of this article is organized as follows. Section 2 discusses the NFV and the service chaining concept and overviews related works. Section 3 discusses general requirements for resiliency and failure models in NFV, as per standards guidelines. In Section 4, we present the network model used, while in Section 5 we present the resilient design scenarios and discuss their failure prevention potential. In Section 6, the resilient SC provisioning problem is formally stated and the ILP models proposed to solve it are shown. In Section 7, we present the case studies and show the obtained numerical results. Finally, conclusions and future work are discussed in Section 8.

## 2. RELATED WORK

NFV is a concept still under standardization. Currently, a number of standardization activities in the NFV area are carried by the European Telecommunications Standards Institute (ETSI) and Internet Engineering Task Force (IETF) [3].

### 2.1. The Service Chaining Problem

The problem of embedding SCs into a physical infrastructure is similar both to Virtual Network Embedding (VNE) [4, 18] and the Location-Routing Problem (LRP) [16]. Its similarity to VNE resides in the fact that SCs can be considered as *virtual networks* characterized by a chain topology where VNFs represent virtual nodes, chained together through virtual links that must be mapped to a physical path. Its similarity to LRP consists in jointly considering the problem of finding the optimal placement of VNFs, among a set of potential locations, along with the routing between VNFs. The LRP combines these two planning tasks and solves them with the objective of reducing costs of nodes, edges, or paths. Regarding the differences, the service chaining problem requires that the routing of traffic between the VNFs occurs according to a specific ordered sequence. Moreover, the sharing of VNFs between multiple SCs increases the number of combinatorial possibilities for the embedding of the SCs.

Several works dealing with the VNF placement and routing problems have appeared in the literature. Reference [13] formalizes the VNF and SC concepts and develops an ILP model for the optimal placement of VNFs and SCs. In [19],

an extended version of the model considers that the upscaling of an existing VNF introduces additional cost, whereas hosting multiple VNFs within the same physical nodes introduces context switching costs. Our model leverages and extends both of the above-mentioned works by including resiliency aspects. In [7], an online algorithm that considers jointly the Virtual Machine (VM) placement and routing is proposed. Finally, the authors in [1] formulate an ILP and a greedy algorithm for the VNF placement and routing problem, including traffic compression/decompression constraints, and adopting two different forwarding latency regimes. The obtained results draw interesting considerations on NFV deployment strategies. However, this work assumed a completely-reliable NFV infrastructure, which is not realistic. The authors in [15] focus on the deployment of VNFs in a hybrid environment where some NFs are virtualized and others use specialized hardware appliances. Finally, the authors in [10] propose an ILP and a game theory model to capture the competition for physical resources between network function instance allocation and routing. However, these last works do not consider any resiliency aspects.

### 2.2. Reliable NFV Deployment

The authors in [2] describe some NFV-related reliability issues and discuss the types of failures that may arise from both hardware (i.e., shutdown of physical machine, hardware issues, etc.) and software (i.e., cyber attacks, bugs, etc.). A more detailed discussion on reliability challenges in NFV network scenarios can be found in [3]. Network reliability in NFV-enabled networks is a new problem whose resolution has not yet attained maturity, even though few preliminary work has already been done. Reference [22] addresses the problem of Joint Topology Design Mapping (JTDM) in a Telco Cloud (TC) environment. The authors propose an efficient heuristic algorithm that leverages the feedback obtained from mapping the critical sub-topologies of a Service Function Chain (SFC), to better coordinate and jointly optimize the VNF combination and SFC mapping. They extend such an algorithm with dedicated and shared protection schemes and compare the results with a baseline scenario (i.e., unprotected). However, they do not consider latency requirements on the SFCs and the processing delay introduced from the sharing of VNFs. Reference [11] presents a framework for reliability evaluation of NFV deployment, and three heuristic algorithms to identify the minimum number of physical and logical nodes whose removal lead to the failure of an NFV deployment. Reference [8] proposes Software-Defined Networking (SDN) and NFV benchmarking test metrics for performance and reliability from an operator perspective. Reference [17] presents ILP and heuristic solutions that exploit multiple backup nodes for the purpose of provisioning each of the supported network services with reliability guarantees. However, the authors focus only on failures that might happen within the hardware hosting the VNF and, unlike our work, discard the possibility of link failures and the failure of other network elements within the nodes. Moreover, they

assume that the backup VNFs are placed in different physical machines than those hosting the corresponding primary VNFs, but in the same physical location, whereas our models consider also scenarios with disjoint physical locations between primary and backup VNFs. Reference [12] presents a VM placement method to achieve redundancy against host-server failures with a minimum set of servers. The idea is to minimize the resources while ensuring a certain protection level. With respect to our work, no consideration is made on the resource sharing and the performance requirements of the VNFs that run on the VMs. Moreover, the authors focus only on failures that occur within the physical nodes, while we include also failures of physical links. Finally, Ref. [20] proposes a model to describe the components of services along with a management system to deploy such an information model, with the objective of providing an automated and resilient deployment. Apart from the differences in the general approach, the authors in [20] focus on resiliency of a single VNF, whereas we consider the resiliency of the whole SC.

### 3. NFV ARCHITECTURE AND RESILIENCY GUIDELINES

In the following, we introduce the architecture used in this study and we highlight the role of some of its primary components. Successively, we discuss a few of the relevant resiliency guidelines and illustrate the possible VNF failure models, as introduced in [3].

#### 3.1. NFVI

The NFV architecture, shown in Figure 2, is a combination of both hardware and software resources making up the environment in which VNFs are deployed. The physical resources include COTS hardware, on top of which virtual resources are abstracted. The abstraction is achieved through the Virtualization layer (based on a hypervisor) which decouples the virtual resources from the underlying physical resources. NFV-MANO provides the necessary functionalities to provision VNFs, and all the related operations such as configuration, orchestration, and life-cycle management, and so forth. Moreover, MANO plays an important role in achieving a resilient NFV deployment. In the following, we discuss standard guidelines for a resilient deployment of VNFs and show the different failure models that arise.

#### 3.2. General Requirements for Resiliency

Different Internet services have different requirements in terms of service continuity and maximum tolerated latency. For instance, in case of a Web service, outages lasting seconds are tolerable and the user typically initiates retries, whereas in the telecom domain (i.e., phone calls) outages must last less than a certain expected level (i.e., few milliseconds). In the NFV framework, not every network function has the same requirements for resiliency. Consequently, the Virtualization of NFs needs to fulfill multiple design criteria, such as

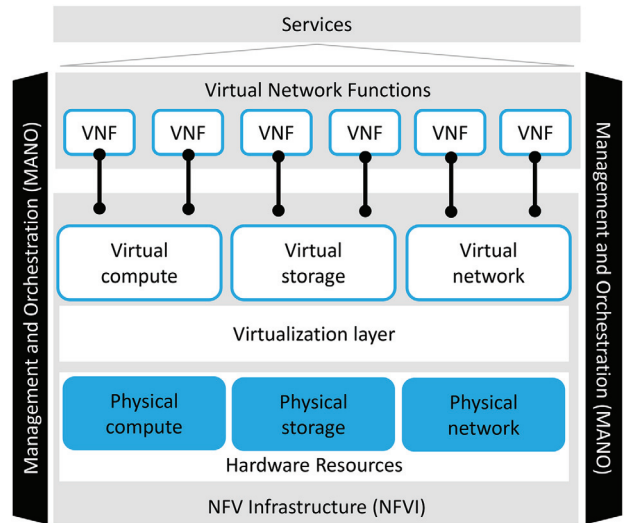


FIG. 2. Simple illustration of NFV architecture. [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

service continuity, automated recovery from failures, prevent single point of failures in the NFV infrastructure as well in the underlying infrastructure. Below we present some important resiliency requirements, according to ETSI guidelines [3]:

- The Virtualized Network Function (VNF) needs to ensure the availability of its part of the end-to-end service, just as in the case of a non-virtualized NF.
- The whole NFV framework must not contain a single point of failure with the potential to endanger service continuity. Thus, mechanisms to recreate any failed component to its state prior to failure, and to support recovery from total component failure, must exist.
- The Network Function Virtualization Infrastructure (NFVI) shall provide the necessary functionality to ensure high availability at the VNF level, such as failure notification and remediation.

Besides the relative availability of a service, the impact of failures is also an important aspect for network providers. To limit the potential of failure impacts, the VNF limitations in terms of number of parallel users allowed, parallel transactions to be handled, and so forth must be accurately defined. Our models follow such guidelines and allow one to evaluate the impact of different network parameters on resiliency.

#### 3.3. VNF Failure Modes

Depending on the type of VNF deployment, the impact of failure varies, hence the survivability method differs. In Figure 3, we show the non-virtualized deployment of NF (option 1). The straightforward approach to virtualize such environment is to take the network function software, run it into a VM image and execute on virtual resources provided by the hypervisor (option 2). This scenario adds a new failure mode to the existing ones since the failure on the hypervisor does not exist in the “*box-model*.” In addition, to achieve




FIG. 3. Deployment options of VNFs along with classical approach. [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

high hardware utilization, the physical resources are sliced into multiple virtual machines, so that different VNFs can be hosted by the same hardware (option 3). This design might cause performance degradation if the resource isolation for the different VNFs is not properly addressed. Finally, a VNF that is composed of multiple atomic VNFs (i.e., VNF components) can be hosted on different VMs running on the same or on different physical hardware (option 4). Again, new failure modes are introduced due to simultaneous failures of multiple VNF components caused by the failure of the underlying hardware or connectivity failures between VNF components (i.e., physical node/link failures). In this work, we assume that the VNFs are running on physical machines according to option 2 or option 3. Hence, the failure of physical nodes (i.e., hardware or hypervisor failures) would cause the failure of all the VMs running in that specific node. We also assume that the NFVI components do not constitute a single point of failure. In Section 5, we will discuss the possible redundancy modes to protect stateful and stateless VNFs and illustrate the possible protection designs for each category of VNFs.

## 4. NETWORK AND SC MODEL

### 4.1. Network Model

We model the physical network as a directed graph composed of a set of physical nodes (which can host VNFs or only act as forwarding nodes) and a set of physical links representing the set of fiber links. Each physical link is associated with a bandwidth capacity. The physical nodes equipped with COTS hardware are referred to as *NFV-nodes* and can have different amount of processing capacity in terms of number of CPU cores they are equipped with.

### 4.2. SC Model

SCs are composed by sequential concatenation of multiple VNFs. To deploy a SC, an operator needs to find a feasible placement of VNFs into the *NFV-nodes* in the physical network and chain them through a physical path. Different SCs can share multiple VNFs and different VNFs can be placed into the same physical *NFV-node*. As shown in Figure 4, two SCs composed of different VNFs both have as start point the physical node  $v_1$  and as end point the physical node  $v_6$ . In addition, VNF1 is shared among the two SCs and mapped to physical node  $v_2$  which shall be equipped with enough processing capacity to host such a VNF. Finally, we assume that

each VNF is assigned a fraction of CPU cores of the VM. To simplify the model, we assume that all the VNFs consolidated into one node run each on a single VM and use a fraction of CPU cores of such VM, as per NFV deployment modes shown in Figure 3.

### 4.3. VNF Model

Generally, a VNF is an abstracted object that performs operations on input traffic. Each VNF has a processing capability which corresponds to the number of CPU cores that are assigned to the VM hosting the VNF. Moreover, we assume that each service corresponds to one SC modeled through a simple chain graph composed of a pair of start/end points, a set of virtual nodes representing the VNFs and a set of virtual links chaining consecutive VNFs requests within the SC.<sup>2</sup> In order to simplify the modeling, the concept of *VNF request* is decoupled from the concept of *VNF instance*. In other words, as shown in Figure 4 (phase 1 and 2), a SC is considered as a chain of VNF requests. In order to deploy SCs in the network, VNF instances are mapped to *NFV-nodes* (phase 1) and successively, VNF requests are mapped to the *NFV-nodes* hosting the requested VNFs (phase 2). The same applies for the mapping of end points, which we assume to have a fixed location that is known a priori, and that they cannot host VNFs. Furthermore, we assume the each SC serves aggregated traffic of a set of users requesting a specific service from a specific physical location.

## 5. PROTECTION SCHEMES FOR SERVICE CHAINING

In this section, we discuss the possible redundancy strategies for a resilient SC provisioning against single-node, single-link, and single-node/link failures.

### 5.1. On-Site Redundancy

Critical VNFs supporting critical services and customers require fast switchover to backup VNFs to ensure high availability. Whereas to ensure latency expectation, backup VNFs need to be instantiated on-site (i.e., centralized redundancy). Critical VNFs may necessitate a 1 + 1 level of redundancy while less critical functions can tolerate a 1:1 redundancy. The main benefits from a centralized redundancy is to reduce switchover time, which allows one to speed up the recovery process, and to reduce the amount of VNF internal state information that needs to be transferred from primary to backup VNFs. Note that this approach does not provide resiliency against node failures, since primary and backup VNFs share the same physical location.

<sup>2</sup> We use the term *virtual node* to indicate the start/end point and the VNFs composing the SC and the term *virtual link* to refer to the segment used to chain two consecutive VNFs within the same SC.



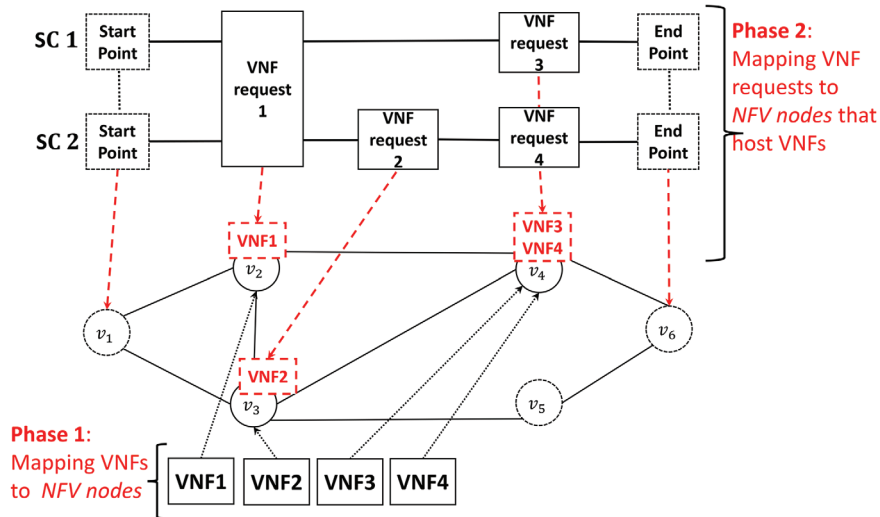


FIG. 4. Two SCs, each having different VNFs, embedded in the physical network. [Color figure can be viewed at wileyonlinelibrary.com]

## 5.2. Off-Site Redundancy

An off-site redundancy architecture involves having redundant VNFs placed in (hot or cold) standby mode in remote locations or NFVI nodes in the network operator's serving region. The intent is to instantiate them when there are failed VNFs in many Network Function Virtualization Infrastructure Point-of-Presence (NFVI-PoP). Moreover, this approach can guarantee resiliency against link and node failures, since backup VNFs do not share the same physical locations as primary VNFs. Hence, based on the service criticalness and the targeted resiliency guarantees, the operator can choose between an on-site or an off-site redundancy approach [3]. In this work we propose three resiliency protection schemes:

**5.2.1. End-to-End Protection (E2E-P).** This protection scheme consists of an end-to-end protection of the entire SC. The idea behind such design is to have a SC that is resilient against single-link and single-node failures. To achieve such goal, a primary SC is embedded in the physical network to support the related service in normal conditions. Such SC is protected through a backup SC which has its VNFs embedded in different physical locations. The physical paths used to chain primary and backup VNFs must also be node disjoint. Figure 5b shows an example of such a protection scheme, where the SC illustrated in Figure 5a, composed of four VNFs, is embedded into the physical network. This protection scheme can be considered as an *off-site redundancy* strategy since all backup VNFs are instantiated in different locations from where the primary ones are hosted. In this case, both redundancy strategies 1+1 and 1:1 are possible, depending on the service latency requirement and operators' design objective in terms of resource utilization. Note that both primary and backup physical paths resulting from the embedding must meet the latency requirement of the service.

**5.2.2. Virtual-Link Protection (VI-P).** The second protection scheme that can be considered as an on-site redundancy protection scheme, with the objective of protecting the virtual links used to concatenate the VNFs of a certain SC and of providing resiliency against physical link failures. Each virtual link of the SCs is embedded through two physical paths, one primary path and one backup path, which must not share any physical link, while different primary/backup virtual links of the same SC can share physical links. An example of such scenario is shown in Figure 5c.

**5.2.3. Virtual-Node Protection (Vn-P).** The last protection scheme provides resiliency against single-node failure. Each VNF composing the SC is instantiated in two disjoint physical locations, whereas the physical paths used to concatenate the primary and backup VNFs might share physical links. This protection scheme suits operators' needs when failures occur in nodes with higher probability with respect to links. An example of this scenario is shown in Figure 5d.

## 6. PROBLEM STATEMENT

In the following, we formally state the problem of resilient SC provisioning and show the ILP models used to design each protection scheme.

### 6.1. Modeling the Physical Topology

We model the physical network as a directed graph  $G = (V, E)$  where  $V$  represents the set of physical nodes  $v \in V$ , which can host VNFs or act as forwarding nodes, while  $E$  represents the set of physical links  $(v, v') \in E$ , which model high-capacity fiber links. Each physical link is associated with a latency contribution due to signal transmission and propagation, denoted by  $\lambda(v, v')$ , and a bandwidth capacity  $\beta(v, v')$ . The physical nodes equipped with COTS hardware

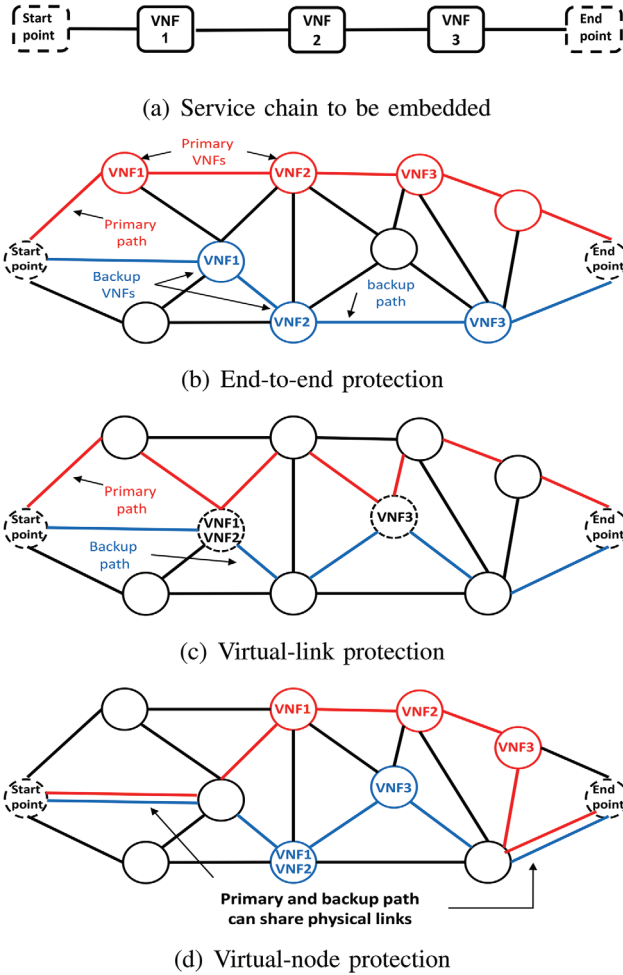


FIG. 5. Proposed protection schemes. [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

are referred to as *NFV-nodes* and can have different amount of processing capacity in terms of number of CPU cores they can host. Finally, we consider a processing-related latency  $\omega(v) : v \in V$ , introduced by the *NFV-nodes*. This latency contribution is called context switching latency and it is proportional to the number of SCs sharing the same VNF [19]. Hence, if a VNF is shared among a high number of SCs, the context switching latency would impact more on the total latency.

## 6.2. VNF and SCs Modeling

Each  $VNF f \in F$  has a processing requirement which corresponds to the fraction of CPU cores that are assigned to the VM that hosts the VNF  $f$ . We assume that a VNF shared among different SCs runs on a VM with enough capacity in terms of CPUs. Moreover, we consider that each service corresponds to one SC modeled through a simple chain graph  $S^c = (E^c \cup U^c \cup G^c)$  where  $E^c$  is the set of end points of the SC,  $U^c$  is the set of VNF requests  $u$ , while  $G^c$  is the set of virtual links  $(u, u')$  chaining requests  $u$  and  $u' \in U^c$ . In order to simplify the modeling, VNFs are mapped to VNF

TABLE 1. Parameters description for the ILP models.

Parameter	Domain	Description
$\eta_u^c$	$c \in C, u \in U^c$	Physical start/end point Where $u$ is mapped for SC $c$
$\gamma_{u,u'}$	$(u, u') \in G^c$	Bandwidth requirement of the virtual link Concatenating VNF request $u$ and VNF request $u'$
$\beta_{v,v'}$	$(v, v') \in E$	Bandwidth capacity of physical link $(v, v')$
$\lambda_{v,v'}$	$(v, v') \in E$	Latency of physical link $(v, v')$
$\omega_v$	$v \in V$	Context switching latency of node $v$
$\tau_u^c \in F$	$c \in C, u \in U^c$	VNF $f$ requested by request $u$ by SC $c$
$\phi^c$	$c \in C$	Maximum tolerated latency for SC $c$
$N_{req}(f)$	$f \in F$	Maximum number of requests of different SCs that VNF $f$ can handle
$N_{VM}(v)$	$v \in V$	Maximum number of virtual machines that node $v$ can host
$M$		Big-M parameter

requests through a mapping parameter  $\gamma_u^c$  specifying the network function  $f \in F$  requested by VNF request  $u \in U^c$ , while VNF requests are mapped to physical nodes through a decision variable. The same applies for the mapping of end points, which we assume are fixed locations and known a priori. Furthermore, we consider that each virtual link composing the SC is characterized by a bandwidth requirement  $\gamma(u, u') : u, u' \in U^c, c \in C$ , and that each SC is associated with a maximum tolerated latency (end-to-end), referred to as  $\phi(c) : c \in C$ .

## 6.3. ILP Models

We now formulate the ILP models for resilient placement of VNFs. In Tables 1 and 2, we summarize the considered parameters and variables. Given a physical topology and a set of SCs to be deployed in the network, we want to find an optimal placement of VNFs such that:

- The number of *VNF nodes* is minimized;
- Latency and bandwidth requirements of SCs are met;
- Resiliency is achieved according to one of the above-mentioned scenarios (see Fig. 5 of Section 5).

### 6.3.1. Objective Function.

$$\text{Minimize } \sum_{v \in V} a_v \quad (1)$$

The objective function aims at minimizing the number of active *NFV-nodes*, which are an indicator of both the Capital Expenditures (CapEx) and OpEx in a telecom operator network. In this context, CapEx are due to the physical hardware necessary to virtualize network functions, while OpEx include all the costs due to network management operations such as energy consumption, monitoring, and so forth.

We consider three classes of constraints to solve this problem, namely: placement constraints, routing constraints, and performance constraints. Due to space limitation we show only the constraints for the E2E-P protection scenario and

TABLE 2. Variables description for the ILP models.

Variable	Domain	Description
$m_{u,v}^c \in \{0, 1\}$	$u \in U^c$ $c \in C$ $v \in V$	Binary variable equal to 1 iff the primary VNF request $u$ of SC $c$ is mapped to physical node $v$
$n_{u,v}^c \in \{0, 1\}$	$u \in U^c$ $c \in C$ $v \in V$	Binary variable equal to 1 iff the backup VNF request $u$ of SC $c$ is mapped to physical node $v$
$x_{v,v',x,y,u,u'}^c \in \{0, 1\}$	$c \in C$ $(v, v') \in E$ $(u, u') \in G^c$ $x, y \in V$	Binary variable equal to 1 iff the physical link $(v, v')$ belongs to the path between $x$ and $y$ where primary VNFs requests $u$ and $u'$ for SC $c$ are mapped, otherwise 0
$y_{v,v',x,y,u,u'}^c \in \{0, 1\}$	$c \in C$ $(v, v') \in E$ $(u, u') \in G^c$ $x, y \in V$	Binary variable equal to 1 iff the physical link $(v, v')$ belongs to the path between $x$ and $y$ where backup VNFs requests $u$ and $u'$ for SC $c$ are mapped, otherwise 0
$i_{f,v} \in \{0, 1\}$	$f \in F$ $v \in V$	Binary variable equal to 1 iff VNF $f$ is hosted by physical node $v$ otherwise 0
$a_v \in \{0, 1\}$	$v \in V$	Binary variable equal to 1 iff node $v$ hosts at least one VNF.

give a brief description of what differs in the other two scenarios, that is, V1-P and Vn-P.

**6.3.2. Placement Constraints.** Constraints (2a) and (2b) force each primary/backup VNF to be mapped to one single node. Constraints (2c) and (2d) ensure that a primary/backup VNF request of VNF  $f$  can be mapped to physical node  $v$  only if a VNF  $f$  is already mapped to such a physical node. Constraint (2e) enforces that primary and backup VNF request  $u$  cannot be mapped to the same node (node disjointness).

$$\sum_{v \in V} m_{u,v}^c = 1 \quad \forall c \in C, u \in U^c \quad (2a)$$

$$\sum_{v \in V} n_{u,v}^c = 1 \quad \forall c \in C, u \in U^c \quad (2b)$$

$$i_{f,v} \leq \sum_{u \in U^c: \gamma_u^c = f} m_{u,v}^c + n_{u,v}^c \quad \forall f \in F, v \in V \quad (2c)$$

$$\sum_{u \in U^c: \gamma_u^c = f} m_{u,v}^c + n_{u,v}^c \leq M \cdot i_{f,v} \quad \forall f \in F, v \in V \quad (2d)$$

$$m_{u,v}^c + n_{u,v}^c \leq 1 \quad \forall u \in U^c, c \in C, v \in V : v \neq \eta_u^c \quad (2e)$$

**6.3.3. Routing Constraints.** Constraints (3c)–(3f) enforce flow conservation, for primary/backup VNF requests, on the intermediate physical nodes that do not host any VNF. In particular, constraints (3c) and (3d) ensure that for any intermediate node  $\omega$  along the physical path between  $x$  and  $y$ ,

if one of the incoming links belongs to the primary/backup physical path, then also one of its outgoing links belongs to the physical path. Constraints (3e) and (3f) avoid the use of multiple incoming (outgoing) links of the intermediate node. Finally, constraint (3g) ensures that a physical link  $(v, v')$  is either part of the primary physical path or in the backup physical path used for the embedding of all VNF request of SC  $c$ .

$$\sum_{(v,x) \in E: v \in V} w_{v,x,x,y,u,u'}^c = \sum_{(v,x) \in E: v \in V} p_{v,x,x,y,u,u'}^c = 0$$

$$\forall c \in C, x \in V, y \in V : x \neq y, (u, u') \in G^c \quad (3a)$$

$$\sum_{(y,v) \in E: v \in V} w_{y,v,x,y,u,u'}^c = \sum_{(y,v) \in E: v \in V} p_{y,v,x,y,u,u'}^c = 0$$

$$\forall c \in C, x \in V, y \in V : x \neq y, (u, u') \in G^c \quad (3b)$$

$$\sum_{(v,w) \in E: v \in V} w_{v,w,x,y,u,u'}^c = \sum_{(w,v') \in E: v \in V} w_{w,v',x,y,u,u'}^c$$

$$\forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (3c)$$

$$\sum_{(v,w) \in E: v \in V} p_{v,w,x,y,u,u'}^c = \sum_{(w,v') \in E: v \in V} p_{w,v',x,y,u,u'}^c$$

$$\forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (3d)$$

$$\sum_{(v,w) \in E: v \in V} w_{v,w,x,y,u,u'}^c \leq 1$$

$$\forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (3e)$$

$$\sum_{(v,w) \in E: v \in V} p_{v,w,x,y,u,u'}^c \leq 1$$

$$\forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (3f)$$

$$\sum_{(u,u') \in G^c} w_{v,v',x,y,u,u'}^c + p_{v,v',x,y,u,u'}^c \leq 1$$

$$\forall c \in C, x, y, v, v' \in V : (v, v') \wedge (v', v) \in E \quad (3g)$$

Note that constraints (4a)–(4d) contain products of binary variables that we linearize to solve the ILP models. In addition, when mapping primary/backup VNF requests on a physical path with source  $x$  and destination  $y$ , incoming links of node  $x$  and outgoing links of node  $y$  are not considered. This is represented by constraints (3a) and (3b), respectively.

$$\sum_{(x,v) \in E: x, y \in V} w_{x,v,x,y,u,u'}^c \cdot m_{u,x}^c \cdot m_{u',y}^c = 1$$

$$\forall c \in C, (u, u') \in G^c \quad (4a)$$



TABLE 3. ILP formulations of the proposed protection scenarios.

	Unprotected	End-to-end protection	VI-protection	Vn-protection
Objective function	Minimize $\sum_{v \in V} a_v$			
Constraints	(2a) (5a) (4b) (4c) (3a)–(3c) (3e) (7a) (7b) (6a) (9a)–(9c)	(2a)–(2e) (5a)–(5b) (4a)–(4d) (3a)–(3g) (7a)–(6d)	(2a) (5a)–(5b) (4a)–(4d) (3a)–(3g) (6a)–(6d) (8a)–(8c)	(2a)–(2e) (5a)–(5b) (4a)–(4d) (3a)–(3f) (7a)–(6d) (9a)–(9c)

$$\sum_{(v,y) \in E: x,y \in V} w_{v,y,x,y,u,u'}^c \cdot m_{u,x}^c \cdot m_{u',y}^c = 1$$

$$\forall c \in C, (u, u') \in G^c \quad (4b)$$

$$\sum_{(x,v) \in E: x,y \in V} p_{x,v,x,y,u,u'}^c \cdot n_{u,x}^c \cdot n_{u',y}^c = 1$$

$$\forall c \in C, (u, u') \in G^c \quad (4c)$$

$$\sum_{(v,y) \in E: x,y \in V} p_{v,y,x,y,u,u'}^c \cdot n_{u,x}^c \cdot n_{u',y}^c = 1$$

$$\forall c \in C, (u, u') \in G^c \quad (4d)$$

$$a_v \leq \sum_{f \in F} i_{f,v} \quad \forall v \in V \quad (7b)$$

$$\sum_{c \in C} \sum_{\substack{(u,u') \in G^c \\ x,v \in V}} (w_{v,v',x,y,u,u'}^c + p_{v,v',x,y,u,u'}^c) \cdot \beta_{u,u'} \leq C_{v,v'}$$

$$\forall (v, v') \in E \quad (7c)$$

$$\sigma_w^c = \sum_{v \in V, u \in U^c} m_{u,v}^c \cdot \omega_v \quad \forall c \in C \quad (7d)$$

$$\sigma_p^c = \sum_{v \in V, u \in U^c} n_{u,v}^c \cdot \omega_v \quad \forall c \in C \quad (7e)$$

Constraints (5a) and (5b) ensure that a physical link  $(v, v')$  can belong to a path between two nodes  $x$  and  $y$  for a virtual link  $(u, u')$  of the SC  $c$  only if two consecutive primary (backup) VNF requests  $u$  and  $u'$  are mapped to these nodes, respectively.

$$w_{v,v',x,y,u,u'}^c \leq m_{u,x}^c \cdot m_{u',y}^c$$

$$\forall c \in C, (v, v') \in E, x, y \in V, (u, u') \in G^c \quad (5a)$$

$$p_{v,v',x,y,u,u'}^c \leq n_{u,x}^c \cdot n_{u',y}^c$$

$$\forall c \in C, (v, v') \in E, x, y \in V, (u, u') \in G^c \quad (5b)$$

$$\sum_{\substack{x,v \in V \\ (u,u') \in G^c \\ (v,v') \in E}} (w_{v,v',x,y,u,u'}^c \cdot \lambda_{v,v'}) + \sigma_w^c \leq \phi_c \quad \forall c \in C \quad (6a)$$

$$\sum_{\substack{x,v \in V \\ (u,u') \in G^c \\ (v,v') \in E}} (p_{v,v',x,y,u,u'}^c \cdot \lambda_{v,v'}) + \sigma_p^c \leq \phi_c \quad \forall c \in C \quad (6b)$$

$$\sum_{f \in F} i_{f,v} \leq N_{VM}(v) \quad \forall v \in V \quad (6c)$$

$$\sum_{\substack{c \in C \\ u \in U^c: \gamma_u^c = f}} m_{u,v}^c + n_{u,v}^c \leq N_{req}(f) \quad \forall v \in V, f \in F \quad (6d)$$

**6.3.4. Latency and Capacity Constraints.** The maximum latency of primary/backup embedding of SC  $c$  is enforced by constraints (6a)–(6b). Finally, the maximum number of CPU cores that the *NFV-node*  $v$  can host is bounded using constraint (6c), and the number of parallel requests that a given VNF can serve is bounded using constraint (6d).

$$\sum_{f \in F} i_{f,v} \leq M \cdot a_v \quad \forall v \in V \quad (7a)$$

Constraints (7a)–(7b) select the active *NFV-nodes*. A node is considered active if it hosts at least one single VNF. Constraint (7c) ensures that link capacity is not exceeded, whereas constraints (7d) and (7e) compute the context switching latency contribution  $\sigma_w^c$  and  $\sigma_p^c$  for primary and backup embedding of SC  $c$ , respectively.

#### 6.4. Additional Modeling Constraints

In the following, we illustrate the constraints used to model the *VI-P* and *Vn-P*.

**6.4.1. Virtual-Link Protection.** With respect to the *E2E-P*, the *VI-P* scenario ensures that the primary and backup physical path used to map a certain virtual link of a SC do not share any physical link and avoid closed loops. This is ensured using the constraints (3g) and (8a)–(8b). See Table 3.

Regarding the placement of primary/backup VNFs, since they share the same physical location, we can reduce the problem complexity by using only one placement variable ( $m_{u,v}^c$ ) to indicate the placement of both primary and backup VNFs. However, we assume that each of these VNFs is placed within a different physical machine. Regarding the physical paths, the latency constraint should be met from source to destination, independent of which path is used. An illustrative example is provided in Figure 6 considering the embedding of the SC shown in 6(a).

$$w_{v,v',x,y,u,u'}^c + w_{v',v,x,y,u,u'}^c \leq 1$$

$$\forall c \in C, (u, u') \in G^c, x, y \in V : x \neq y, (v, v') \in E \quad (8a)$$

$$p_{v,v',x,y,u,u'}^c + p_{v',v,x,y,u,u'}^c \leq 1$$

$$\forall c \in C, (u, u') \in G^c, x, y \in V : x \neq y, (v, v') \in E \quad (8b)$$

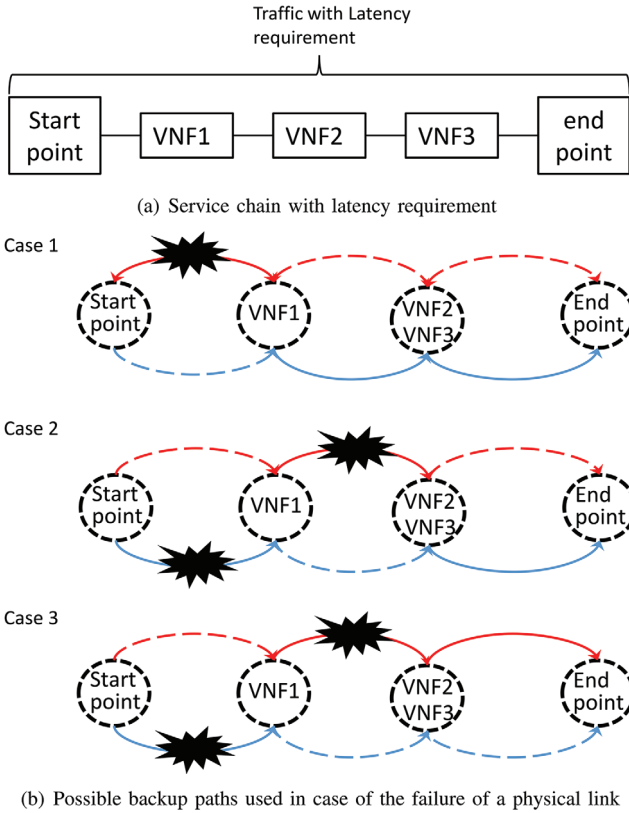


FIG. 6. Possible backup paths in the  $VI-P$  design scenario. [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

$$\sum_{\substack{x,y \\ (u,u')}} \sum_{(v,v')} (w_{v,v',x,y,u,u'}^c + p_{v,v',x,y,u,u'}^c) \cdot \lambda_{v,v'} + \sigma_p^c \leq \phi_c \quad \forall c \in C \quad (8c)$$

We assume that the embedding process resulted in  $VNF1$  placed in one node, while  $VNF2$  and  $VNF3$  are consolidated in a second node. According to  $VI-P$ , every pair of nodes from start to end points are connected using a pair of disjoint paths (i.e., red and blue paths). The embedding of virtual links can result in one single physical link carrying the primary and backup embedding of different virtual links. Hence, different physical paths can be used to transport the traffic from start to end point. In Figure 6b (case 1), the failure of a physical link causes the failure of the primary virtual link between the start point and  $VNF1$ . The backup path (dashed lines) must meet the latency requirement.

Similarly, in case 2 and case 3 (Fig. 6b), we assume that the failure of one physical link causes the failure of the backup path of the first virtual link and the primary path of the second virtual link. In this case, two possible end-to-end paths are possible (dashed lines) and both of these options must satisfy the latency requirement. Equation (8c) ensures that the latency requirements are met in all three cases. Please note that the paths between the starting point and an  $NFV$ -node or between two consecutive  $NFV$ -nodes are multi-hop paths.

Intermediate nodes were omitted in the figure for the sake of simplicity.

**6.4.2. Virtual-Node Protection.** For the  $Vn-P$  scenario, only the node-disjointness constraint applies and no disjointness constraints between primary/backup physical paths are needed since they can share physical links. In addition, Equations (2c), (2d), and (7c) are substituted by the following constraints:

$$i_{f,v} \leq \sum_{u \in U^c: \gamma_u^c = f} m_{u,v}^c \quad \forall f \in F, v \in V \quad (9a)$$

$$\sum_{u \in U^c: \gamma_u^c = f} m_{u,v}^c \leq M \cdot i_{f,v} \quad \forall f \in F, v \in V \quad (9b)$$

$$\sum_{\substack{c \in C \\ (u,u') \in G^c \\ x,v \in V}} w_{v,v',x,y,u,u'}^c \cdot \beta_{u,u'} \leq C_{v,v'} \quad \forall (v,v') \in E \quad (9c)$$

Please refer to Table 3 for a detailed description of the constraints used in each design scenario.

$$V_u = |V| \cdot (|C| \cdot |U^c| + |C| \|E'\| |V| |G^c| + |F| + 1) \quad (10a)$$

$$C_u = |C| \cdot (|U^c| + |E'\| |V|^2 |G^c| + 2|G^c| + 3 \cdot |V|^2 |G^c| + 1) + 2|V| \cdot (2|F| + 1) + |E'| \quad (10b)$$

$$V_{e2e} = V_u + \alpha_{e2e} \quad (10c)$$

$$C_{e2e} = C_u + \beta_{e2e} \quad (10d)$$

$$V_{vl-p} = V_u + \alpha_{vl-p} \quad (10e)$$

$$C_{vl-p} = C_u + \beta_{vl-p} \quad (10f)$$

$$V_{vn-p} = V_u + \alpha_{vn-p} \quad (10g)$$

$$C_{vn-p} = C_u + \beta_{vn-p} \quad (10h)$$

## 6.5. Computational Complexity

In this section, we compute the total number of variables and constraints of each design scenario. The number of variables of the  $E2E$  and the  $Vn-P$  scenarios is the same, while it differs slightly in the case of  $VI-P$  and *unprotected* scenarios, as no backup of nodes or nodes/links is required. The number of constraints is slightly different in each scenario. However, this difference does not affect the overall complexity, which is the same for all designs, and it is of order  $O(|G^c| \cdot |E'| \cdot |C| \cdot |V|^2)$ .<sup>3</sup> Equations (10a)–(10h) compute the number of variables and constraints of each design scenario based on the values of the unprotected scenario, denoted by  $N_u$  and  $C_u$ , and using the equations denoted by  $\alpha$  and  $\beta$ , computed for each scenario, in Equations (11a)–(11f).

<sup>3</sup> Please note that, to simplify the modeling and in order to allow multiple VNFs of the same SCs to be hosted in the same  $NFV$  node, we assumed that each of these nodes has a self-loop link with infinite bandwidth. Such self-loops links were included in the complexity computation by considering that  $|E'| = |E| \cup |N|$ , as the number of self-loop links is equal to the number of physical nodes.

$$\alpha_{e2e} = |V| \cdot (|C| \cdot |U^c| + |C||E' ||V||G^c|) \quad (11a)$$

$$\beta_{e2e} = |C| \cdot (|U^c| + |E' ||V||G^c| + 2|G^c| + 3) \quad (11b)$$

$$+ |V| \cdot (|U^c| + |V|^2|G|^2 + |V||E'| + |F| + 1)$$

$$\alpha_{vl-p} = \alpha_{e2e} - |C||E' ||V||G^c| \quad (11c)$$

$$\beta_{vl-p} = \beta_{e2e} + (|V|^2 \cdot (|C||G^c| + 1)) \quad (11d)$$

$$\alpha_{vn-p} = \alpha_{e2e} \quad (11e)$$

$$\beta_{vn-p} = \beta_{e2e} + 2|V|^2|G|^2 \quad (11f)$$

## 7. CASE STUDY AND RESULTS

In this section, we present and discuss the results of the ILP models shown in Section 6. To solve the ILP problems we used CPLEX 12.6.1.0 installed on hardware platforms equipped with an  $8 \times 2$  GHz processor and 8 GB RAM. To evaluate the impact of latency requirements on the protection scenarios we investigated the embedding of four types of services chains, with different processing requirements and latency constraints, namely: web-service (WS), video streaming (VS), VoIP and online-gaming (OG). The maximum end-to-end tolerated latency for these services has been set to 500 ms for WS, 100 ms for both VS and VoIP, and 60 ms for OG similarly to [19]. Table 4 shows the VNFs composing the SCs, their bandwidth requirements and maximum allowed latency. We consider heterogeneous and homogeneous traffic scenarios. In the heterogeneous scenario, five different SC requests randomly selected from the SCs in Table 4 are considered. The type of SCs in this case is randomly selected at each ILP run. In the homogeneous scenario, five SC requests of the same type are considered. The start/end points, for both traffic scenarios, are randomly selected for each SC request, at each ILP run. Moreover, we assume that all physical nodes can act as *NFV-nodes* and that the start/end points of SCs requests cannot host VNFs. As for the physical topology, we considered the National Scientific Foundation Network (NSFNET) network with 14 nodes and 22 bidirectional links. Each *NFV-node* is assumed to have the same capacity in terms of CPU cores. We set the context switching delay to 4 ms per VNF [9, 19]<sup>4</sup> and set the link capacity to be equal to 2.5 Gbps. We assume that each SC aggregates the traffic of 2,000 users. We also assume that the bandwidth requirement of virtual links chaining VNFs varies according to a compression factor [1].<sup>5</sup> Such a value is randomly selected at each ILP run. The results, shown in Figure 7, were obtained averaging the results of 10 instances, solved within 5% of the optimal solution, for each value of *NFV-node* capacity and each protection scenario, while considering different start/end points pairs at each ILP run. Figure 7b–e show the average number

<sup>4</sup> Note that the provisioning of SCs introduces other latency contributions due to the upscaling of the capacity of VNFs and hypervisor processing of the VNF requests [19].

<sup>5</sup> Normally, NFs expand or compress the input traffic, based on the performed task. In this work, we used random compression factors ranging between 0 and 1, given that no reference is available for such values.

TABLE 4. Performance requirements for the SCs.

Service chain	Chained VNFs	$\beta$	$\phi_c$
Web-service	NAT-FW-TM-WOC-idps	100 kbit/s	500 ms
Video streaming	NAT-FW-TM-VOC-IDPS	4 Mbit/s	100 ms
VoIP	NAT-FW-TM-FW-NAT	64 kbit/s	100 ms
Online-gaming	NAT-FW-VOC-WOC-IDPS	50 kbit/s	60 ms

Network Address Translator (NAT), Firewall (FW), Traffic Monitor (TM), WAN Optimization Controller (WOC), Intrusion Detection Prevention System (IDPS), Video Optimization Controller (VOC).

of active *NFV-nodes* needed to support the proposed protection scenarios for different values of node capacity (number of CPU cores it can host), for the WS, VoIP, VS, OG and heterogeneous traffic scenarios, respectively. In the following, we analyze the effect of latency and node capacity for the different traffic scenarios.

### 7.1. Impact of Latency

Figure 7b presents the number of active nodes for the less stringent SC in terms of latency. We observe that all protection scenarios are possible and that the VI-P scenario requires the same number of *NFV-nodes* activated for the *Unprotected* scenario (baseline). We note that a SC with low requirements on latency can be protected against single-link failures (VI-P) with no additional *NFV-nodes* with respect to the *Unprotected* case. On the other hand, providing protection against both single-link and single failure (E2E-P) requires the activation of twice the number of *NFV-nodes* when node capacity is greater than 10 CPU cores per *NFV-node*, and more than twice when the node capacity is less than 10 CPU cores. This is due to the fact that, when decreasing node capacity, more *NFV-nodes* must be activated to provide the off-site resiliency scenarios. Moreover, increasing the capacity by a factor of five reduces the number of active *NFV-nodes* by 33% in case of off-site redundancy protection (E2E-P, Vn-P) and 80% in case of on-site redundancy protection (VI-P). We also observe that the resources required to supply end-to-end protection (E2E-P) is almost the same with respect to protection against single-node failures (Vn-P), independent of the capacity values, meaning that in case the operator chooses to place backup VNFs off-site, the protection against both link and node failures comes at the same cost, in terms of *NFV-nodes*, with respect to protection against node failures.

Figure 7a and c show the results obtained by solving the VNF placement of VoIP and VS SCs, which have an average latency requirement. For both SCs, we observe that all scenarios are possible except for the VI-P scenario, which leads to infeasible solution for node capacity values less than 6 CPU cores per *NFV-node*. This is mainly due to the fact that VI-P has a stringent link disjointness constraint that, in case of VNFs distribution among a high number of nodes, increases the latency of physical paths needed to chain the VNFs and consequently leads to violation of the latency constraint.

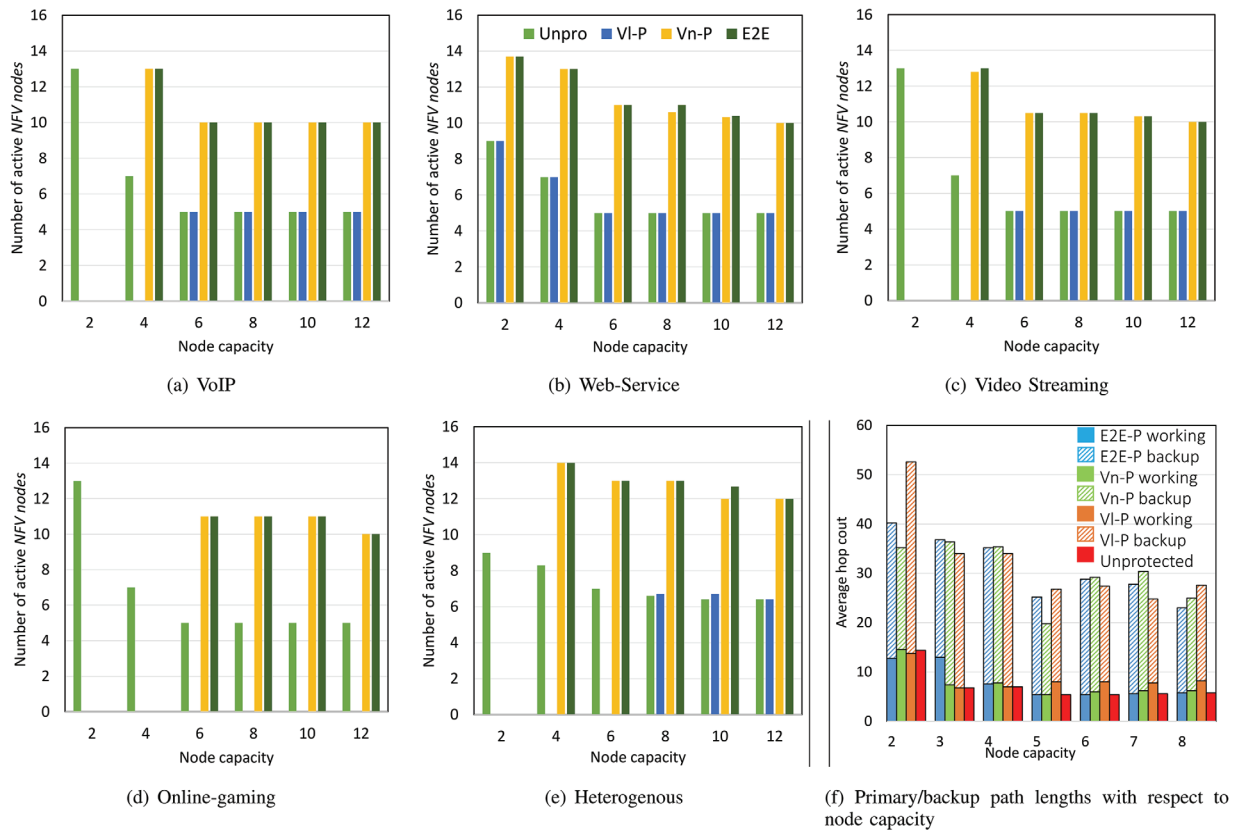


FIG. 7. Comparison of the proposed protection scenarios for different latency requirements. [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

Figure 7d shows the results obtained when SCs with the most stringent latency requirement (OG) are embedded into the network. We observe that, for node capacity greater than 6 CPU cores, all scenarios are possible except for the VI-P scenario which is infeasible independent of node capacity. For capacity values less than 6 CPU cores, we observe that all protection strategies lead to an infeasible solution. This means that, for latency-stringent SCs, in order to provide protection against node/link failures, each *NFV-node* must be equipped with a given amount of capacity, as the distribution of VNFs across multiple nodes leads to the violation of the latency constraint. Moreover, for E2E-P and Vn-P, doubling the capacity leads to a tiny decrease of active *NFV-nodes* (around 10%), mainly due to the fact that increasing consolidation causes an increase in the context switching latency and hence to violations of the latency constraint. Finally, since the only feasible protection scenarios are E2E-P and Vn-P, the operator is constrained to place backup VNFs off-site for providing resiliency against only single-link failures, when latency critical SCs are deployed.

Finally, for the heterogeneous traffic scenario, shown in Figure 7e, all protection scenarios are possible starting from 8 CPU cores per *NFV-node* and lead to infeasible solutions at 2 CPU cores per *NFV-node*. With respect to the OG case, we observe that, when SCs with different requirements are deployed, protection against single-link failures on-site can be provided starting from 8 CPU cores per *NFV-node*. In

general, we observe that the heterogeneous traffic scenario requires more resources with respect to the homogeneous scenarios, but meets latency requirements while allowing a better VNF consolidation. This means that deploying SCs with different latency requirements (as happens in real networks) guarantees resiliency with a smaller number of CPU cores per *NFV-node*, and consequently less failure impact within *NFV-nodes*.

## 7.2. Impact of Node Capacity

In terms of capacity, we observe that for WS (Fig. 7b) and heterogeneous deployment of SCs (Fig. 7e), while increasing node capacity, the off-site redundancy protection strategies increase the number of active *NFV-nodes* from 69% up to 96% for WS and from 52% up to 120% for heterogeneous deployment, with respect to the unprotected scenario. Whereas, for the on-site redundancy protection scenario (VI-P), we observe that increasing the capacity more than 6 and 8 CPU cores does not bring any benefit in terms of consolidation. This is mainly due to the fact that consolidation of VNFs is limited by the context switching latency. The same claim is valid for the VoIP and VS SCs (Fig. 7a and c), where we observe that increasing node capacity more than 6 CPU cores per node does not affect the number of active *NFV-nodes*, where it leads to a feasible solution. Finally, comparing the outcome obtained for VoIP and VS SCs, shown in

TABLE 5. ILP run-times (seconds) for homogeneous and heterogeneous SCs deployment, under fixed *NFV-node* capacity and different protection strategies.

	Unprotected	VI-P	Vn-P	E2E-P
Heterogeneous	1034	4812	3671	3342
Web-service	697	1247	971	1037
Video streaming	654	38067	36036	21348
VoIP	835	39225	30040	38630
Online-gaming	832	49463	31869	36481

Figure 7a and c, we observe that the impact of different bandwidth requirements of both SCs is slightly noticeable (around 2%). Later in this section, we will relax the constraint on the number of parallel requests that a VNF instance can serve and evaluate the impact of bandwidth requirement on both SCs, under different values of node capacity and different optimization targets. In general, for both traffic scenarios, we observe that VNF consolidation is limited by latency, as consolidating more VNFs into less nodes would increase the impact of context switching latency.

### 7.3. Impact of Node Capacity on the Average Hop Count

We analyzed the impact of node capacity on the average length of primary/backup physical paths of all proposed protection strategies. In Figure 7f, we show the primary/backup path lengths when 2 Web Service (WS) SCs are deployed. These results were obtained by averaging the path lengths of 5 start/end point pairs randomly selected and tested for all protection scenarios. We observe that for all protection strategies, when increasing the node capacity, the length of the primary path does not change significantly. For backup paths, we observe that increasing node capacity does not mean reducing backup path lengths, meaning that a trade-off between consolidation of VNFs and the average path length exists.

### 7.4. ILP Runtime

Table 5 shows the ILP runtime, in seconds, for homogeneous and heterogeneous SC deployment for the proposed protection design scenarios. These values were obtained from a set of experiments where node capacity is equal to 10 VMs per *NFV-node*. Generally, we observe that the more stringent the latency requirement of the SC, the longer solving the ILP model takes. Moreover, by comparing the different protection strategies, we observe that the *VI-P* scenario requires the largest amount of time, regardless of the type of SC that is provisioned. This is due to the fact that *VI-P* computes a pair of node-disjoint paths between every two consecutive VNFs, with respect to E2E-P, which imposes a less stringent disjointness constraint.

### 7.5. Impact of Different Optimization Targets

We consider the VNF placement of two different SCs, with the same latency requirements but different bandwidth

requirements (VoIP and VS), when the optimization target is to consolidate VNFs (CON) and when optimizing with the objective of balancing the load on physical links (LB). We run the ILP model for E2E-P for different number of users and different values of node capacity.

The objective here is to analyze the effect of increasing number of users (we consider that the two SCs aggregate the traffic all the users at each ILP run) on the average link occupation and on the number of active *NFV-nodes* when different objective functions are targeted. In addition, for this set of experiments, we assume that the data-rate between different VNFs of the same SC is fixed, and relax the constraint in Equation (6d), which limits the maximum number of parallel requests that a VNF instance can serve.

To solve the ILP model with a load balancing objective, we use the same formulation in Section 6 for the E2E-P, define an integer variable  $\mu \in [0, 1]$  to account for the maximum load of any edge and substitute the objective function Equation (1) and the link capacity constraint Equation (7c) with the objective function and link capacity constraints in Equations (12a) and (12b), respectively:

$$\text{Minimize } \mu \quad (12a)$$

$$\sum_{c \in C} \sum_{\substack{(u,u') \in C^c \\ x,v \in V}} (w_{v,v',x,y,u,u'}^c + p_{v,v',x,y,u,u'}^c) \cdot \beta_{u,u'} \leq C_{v,v'} \cdot \mu \\ \forall (v, v') \in E \quad (12b)$$

In Figure 8a and c, we show the average link occupation for the VoIP and VS SCs and compare the average number of active *NFV-nodes* for both optimization targets in Figure 8e, when each *NFV-node* is equipped with 2 CPU cores. In case of VoIP, LB shows a decrease in average link utilization from 28% up to 38% at the expense of tripling the number of *NFV-nodes* with respect to results obtained in CON. Whereas, in case of VS, LB decreases the load on physical link from 27% up to 37% while doubling the number *NFV-nodes* with respect to CON. We also observe that the number of *NFV-nodes* activated for both VoIP and VS under CON does not change, independent of the number of users. Generally, we observe that the VoIP activates 60% fewer *NFV-nodes* than VS. Instead, when targeting LB, VoIP activates up to 10% *NFV-nodes*, which is due to the different bandwidth requirements of VoIP and VS.

In the second set of experiments, we increase the node capacity by a factor of three and show the average link occupation for VoIP and VS in Figure 8b and d, respectively. We observe that, when deploying VoIP, both objective functions lead to the same average link utilization, at a high number of users, as shown in Figure 8f. The reason is due to the fact that once a VNF instance is activated, it can be used by all SCs. Hence, all SCs redirect the traffic through it. Hence, we obtain the same values of link occupation, independent of the optimization target. Whereas, in case of VS, the gain obtained from LB ranges between 12% and 20%, due to the fact that the VS is more bandwidth-intensive, hence balancing the load on physical links brings small benefit, as shown in



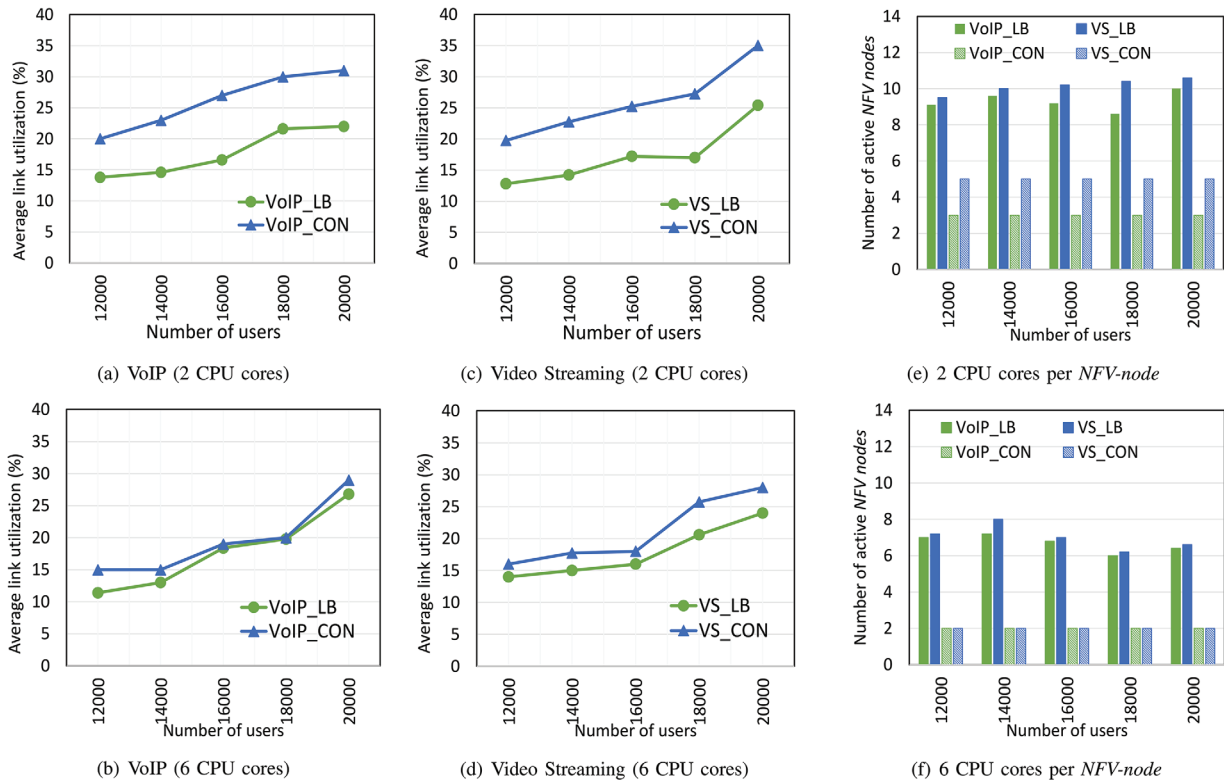


FIG. 8. Comparison of two objective functions when deploying two SCs with the same latency requirement and different bandwidth requirements. [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

Figure 8b. Intuitively, in terms of *NFV-nodes*, the increase of node capacity translates into better consolidation. Both SCs reduce the number of active *NFV-nodes* under CON by more than 3 times and activate the minimum number of *NFV-nodes* needed to support E2E-P. This is also due to the fact that an active VNF can be used by unlimited SCs. Consequently, the paths used to concatenate the VNFs are longer, which justifies the low gain achieved with LB, when node capacity is high.

We generally observe that LB is beneficial for low values of node capacity, while CON brings more benefit when *NFV-nodes* are equipped with a higher number of CPU cores.

## 8. CONCLUSION

In this work, we proposed three different protection strategies to provide resilient SCs deployment against single-node, single-link, and single-node/link failures. We reported the formulation for all the design scenarios, solved the ILP models considering a small number of SCs with different latency requirements, and found that a trade-off exists between node capacity and latency of the deployed SCs. Moreover, we analyzed the effect of *NFV-node* capacity on the average primary/backup path lengths. Finally, we solved one of the proposed ILP models considering two different SCs with equal latency constraints and different bandwidth requirements under two conflicting objectives, to analyze the effects of bandwidth requirements on the distribution of VNFs. In our small-scale scenario, we found that:

- To provide resiliency to SCs against single-link and single-node failures, twice the number of *NFV-nodes* are needed with respect to the unprotected scenarios and the case where only single-link failures are targeted. However, decreasing node capacity requires more than twice the number of *NFV-nodes* to be activated.
- Capacity allocation in *NFV-node* must be done taking into consideration the type of deployed SCs.
- By adopting heterogeneous traffic and allowing VNF sharing, resiliency to link failures can be provided at low values of node capacity.
- Increasing node capacity does not cause the reduction of the average path lengths.
- Bandwidth intensive SCs benefit more from consolidation when the node capacity is high, while load balancing is beneficial at small values of node capacity.

Owing to the complexity of this problem, future steps of this work aim at developing heuristic algorithms to solve the problem of SC provisioning for large instances in reasonable time. We also plan to solve this problem under dynamic conditions, while targeting the optimization of different cost functions.

## ACKNOWLEDGMENT

This article is based upon work from Cooperation in Science and Technology (COST) Action CA15127 (“Resilient Communication services prOTecting end-user applications from DISaster-based failures (RECODIS)”) supported by COST (European Cooperation in Science and Technology)

and on work from METRO-HAUL (METRO High bandwidth, 5th Generation Application-aware optical network, with edge storage, compute and low Latency) in the context of Horizon 2020 call, H2020-ICT-2016-2, supported by European Union (EU) for research, development and innovation.

## REFERENCES

- [1] B. Addis, D. Belabed, M. Bouet, and S. Secci, Virtual network functions placement and routing optimization, 2015 IEEE 4th Int Conf Cloud Netw (CloudNet), Oct 2015, pp. 171–177.
- [2] D. Cotroneo, L.D. Simone, A.K. Iannillo, A. Lanzaro, R. Natella, J. Fan, and W. Ping, Network function virtualization: Challenges and directions for reliability assurance, 2014 IEEE Int Symp Software Reliab Eng Workshops, Nov 2014, pp. 37–42.
- [3] ETSI, GS NFV-REL 001 v1. 1.1: Network functions virtualization NFV; resiliency requirements, Technical report, ETSI industry Specification Group (ISG) Network Functions Virtualization (NFV), 2015.
- [4] A. Fischer, J.F. Botero, M. Till Beck, H. De Meer, and X. Hesselbach, Virtual Network Embedding (VNE): A survey, IEEE Commun Surveys & Tutorials 15 (2013), 1888–1906.
- [5] J. Halpern and C. Pignataro, Service Function Chaining (SFC) Architecture, Technical report, European Telecommunication Standards Institute (ETSI), Service Functions Chaining (SFC) Framework, 2015.
- [6] A. Hmaity, M. Savi, F. Musumeci, M. Tornatore, and A. Pattavina, Virtual network function placement for resilient service chain provisioning, 2016 8th Int Workshop Resilient Netw Des Model (RNDM), Sept 2016, pp. 245–252.
- [7] J.W. Jiang, T. Lan, S. Ha, M. Chen, and M. Chiang, Joint VM placement and routing for data center traffic engineering, IEEE Conf Informat Commun (INFOCOM), IEEE 2012, pp. 2876–2880.
- [8] T. Kim, T. Koo, and E. Paik, SDN and NFV benchmarking for performance and reliability, 2015 17th Asia-Pacific Netw Oper Manage Symp (APNOMS), Aug 2015, pp. 600–603.
- [9] C. Li, C. Ding, and K. Shen, Quantifying the cost of context switch, Proc 2007 Workshop Exp Comput Sci (WECS), 2007, p. 2.
- [10] T. Lin, Z. Zhou, M. Tornatore, and B. Mukherjee, Demand-aware network function placement, J Lightwave Technol 34 (2016), 2590–2600.
- [11] J. Liu, Z. Jiang, N. Kato, O. Akashi, and A. Takahara, Reliability evaluation for NFV deployment of future mobile broadband networks, IEEE Wireless Commun 23 (June 2016), 90–96.
- [12] F. Machida, M. Kawato, and Y. Maeno, Redundant virtual machine placement for fault-tolerant consolidated server clusters, IEEE Network Oper Manage Symp (NOMS), April 2010, pp. 32–39.
- [13] S. Mehraghdam, M. Keller, and H. Karl, Specifying and placing chains of virtual network functions, IEEE 3rd Int Conf Cloud Netw (CloudNet), IEEE 2014, pp. 7–13.
- [14] R. Mijumbi, J. Serrat, J.L. Gorricho, N. Bouten, F.D. Turck, and R. Boutaba, Network function virtualization: State-of-the-art and research challenges, IEEE Commun Surv Tutor 18 (2016), 236–262.
- [15] H. Moens and F. De Turck, VNF-P: A model for efficient placement of virtualized network functions, 10th Int Conference Network Service Manage (CNSM), IEEE 2014, pp. 418–423.
- [16] C. Prodhon, and C. Prins, A survey of recent research on location-routing problems, Eur J Oper Res 238 (2014), 1–17.
- [17] L. Qu, C. Assi, K. Shaban, and M. Khabbaz, Reliability-aware service provisioning in NFV-enabled enterprise data-center networks, 2016 12th Int Conf Netw Service Manage (CNSM), Oct 2016, pp. 153–159.
- [18] M.R. Rahman, and R. Boutaba, SVNE: Survivable virtual network embedding algorithms for network virtualization, IEEE Trans Netw Service Manage 10 (2013), 105–118.
- [19] M. Savi, M. Tornatore, and G. Verticale, Impact of processing costs on service chain placement in network functions virtualization, IEEE Conf Netw Funct Virtualiz Softw Defined Netw (NFV-SDN), Nov 2015, pp. 191–197.
- [20] M. Scholler, M. Stiernerling, A. Ripke, and R. Bless, Resilient deployment of virtual network functions, 5th Int Congress Ultra Modern Telecommun Control Syst Workshops (ICUMT), IEEE 2013, pp. 208–214.
- [21] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar, Making middleboxes someone else’s problem: Network processing as a cloud service, ACM SIGCOMM Comput Commun Rev 42 (2012), 13–24.
- [22] Z. Ye, X. Cao, J. Wang, H. Yu, and C. Qiao, Joint topology design and mapping of service function chains for efficient, scalable, and reliable network functions virtualization, IEEE Netw 30 (2016), 81–87.