Virtual Network Function Placement For Resilient Service Chain Provisioning

Ali Hmaity, Marco Savi, Francesco Musumeci, Massimo Tornatore, Achille Pattavina
Politecnico di Milano, Department of Electronics, Information and Bioengineering, Milan, Italy
E-mail: *firstname.lastname*@polimi.it

***Abstract*—Virtualization technologies are changing the way network operators deploy and manage Internet services. In particular in this study we focus on the new Network Function Virtualization (NFV) paradigm, which consists in instantiating Virtual Network Function (VNFs) in Commercial-Off-The-Shelf (COSTS) hardware. Adopting NFV network operators can dynamically instantiate Network Functions (NFs) based on current demands and network conditions, allowing to save capital and operational costs. Typically, VNFs are concatenated together in a sequential order to form Service Chains (SCs) that provide specific Internet Services to the users. In this paper we study different approaches to provide the resiliency of SCs against single-link and single-node failures. We propose three Integer Linear Programming (ILP) models to solve the VNF placement problem with the VNF service chaining while guaranteeing resiliency against single-node/link, single-link and single-node failures. Moreover we evaluate the impact of latency of SCs on the VNFs distribution. We show that providing resiliency against both single-link and single-node failures necessitates the activation of twice the amount of resources in terms of 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-nodes failures.**

I. Introduction

For network operators, offering new bandwidth-intensive and latency-constrained (e.g cloud gaming or video streaming)Internet services is a challenging task due to the adoption of proprietary hardware appliances, the high cost of offering, maintaining and integrating these services. Network Functions 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 [1]. 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 Function (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 that NFV brings a whole new dimension to the landscape of telecommunication industry market due to the possibility of reducing capital investments, energy consumption by consolidating network functions, and by introducing tailored services based on customers needs. Moreover, NFV simplifies service deployment by exploiting the concept of *service chaining* [2]: 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.

Before deploying NFV solutions in operational networks several challenges regarding performance, availability, security and survivability need to be tackled. In this work we focus on SCs resiliency against single-link/node failures. To the best of our knowledge, this is the first study to investigate the resiliency of SCs provisioning. Our main objective is to model a survivable placement of VNFs while minimizing resources, in terms of number of VNFs instances placed in the network. We develop three different Integer Linear Programming (ILP) models to solve the VNF placement problem with service chaining while guaranteeing resiliency against single-node/single-link, single-link and single-node failures. We show the amount of nodes needed to supply and compare these results with an Unprotected scenario. Furthermore, we investigate the effect of latency on the proposed protection schemes.

The rest of this paper is organized as follows. Section II discuss the NFV and the service-chaining concept and overviews existing works appeared in literature. In Section III we present the network model used, while in Section IV we present the resilient design scenarios and discuss their failure prevention potential. In Section V the resilient SCs provisioning problem is formally stated and the ILP models are shown. In Section VI we present the case-studies and show the obtained numerical results. Finally, conclusion and future work are discussed in Section VII.

II. Related works

NFV is still a concept under standardization. Currently, a number of standardization activities in the NFV area are carried by ETSI and IETF, [3] [4] [5]. ETSI has defined an architectural framework that enables VNFs to be deployed and executed on a Network Functions Virtualization Infrastructure (NFVI), which comprises commodity servers logically separated/partitioned by a software layer. Above the hypervisor layer, the component in charge of mapping the VNFs to physical resources, a VNF is mapped to a Virtual Machine (VM) in the NFVI and its deployment and management is handled by the Management and Orchestration (MANO) system [6].

The problem of embedding SCs into a physical infrastructure can be considered as an extended version of two NP-

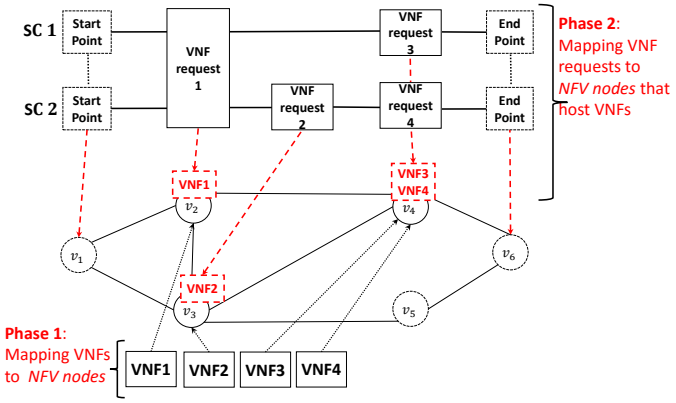

Fig. 1. Two service chains, each having different VNFs, embedded in the physical network.

hard problems: The Virtual Network Embedding (VNE) [7],[8] and Location-Routing Problems (LRP) [9]. The similarity with 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.The similarity with 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 this two planning tasks and solves them with the objective to reduce costs of nodes, edges or paths.

Ref [10] formalizes the VNF and SC concepts and develops an ILP model fro the optimal placement of VNF and SCs. An extended version of this model [11], 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 the above mentioned works. Ref [12] develops an ILP model for the efficient placement of VNFs considering processing-resource sharing. In [13] an Online algorithm that considers jointly the VM placement and routing is proposed. Finally authors in [14] focus on the deployment of VNFs in a hybrid environment where some NFs are virtualized and others are use specific hardware appliances. In this work we focus on resiliency of deployed SCs and the impact of the QoS (Quality of Service) requirements on the protection scheme adopted. Some research efforts have focused on resiliency of VMs. Ref [15] 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 in order to provide 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 failure that occur within physical nodes, while we include also failures of physical links. Finally, Ref [16] proposes a model to describe the components of services along with a management system to deploy such information model, with the objective to provide an automated and resilient deployment. A part from the differences in the general approach in Ref [16] focuses on resiliency of single VNF, Whereas we consider the resiliency of the whole SC.

III. Service chains and network model

A. 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 link 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 VMs that they can host.

B. Service chains model

Service chains are composed by sequential concatenation of multiple VNFs. To deploy a SC, an operator need to find the right placement of VNFs into the *NFV nodes* (VNF placement process) 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 Fig. 1, two SCs composed of different VNFs have both 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 VNF.

C. 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 that host that VNF. Moreover, we assume that each service corresponds to one SC modeled through a simple line graph composed by 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[1]. In order to simplify the modeling, the concept of requests are decoupled from the VNFs that compose the Service chains. In other words, as shown in Fig. 1 (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 those *NFV nodes* that hosts the requested VNFs (phase 2). The same apply for the mapping of end-points, which we assume have fixed location,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

IV. Resilient design protection schemes

One important aspect for network operators is to guarantee service continuity in case of failures. To achieve such objec-

[1]We use the term *virtual node* to indicate the start/end point and the VNFs composing the SC and refer to to the segment used to chain two consecutive VNFs within the same SC as *virtual link*.

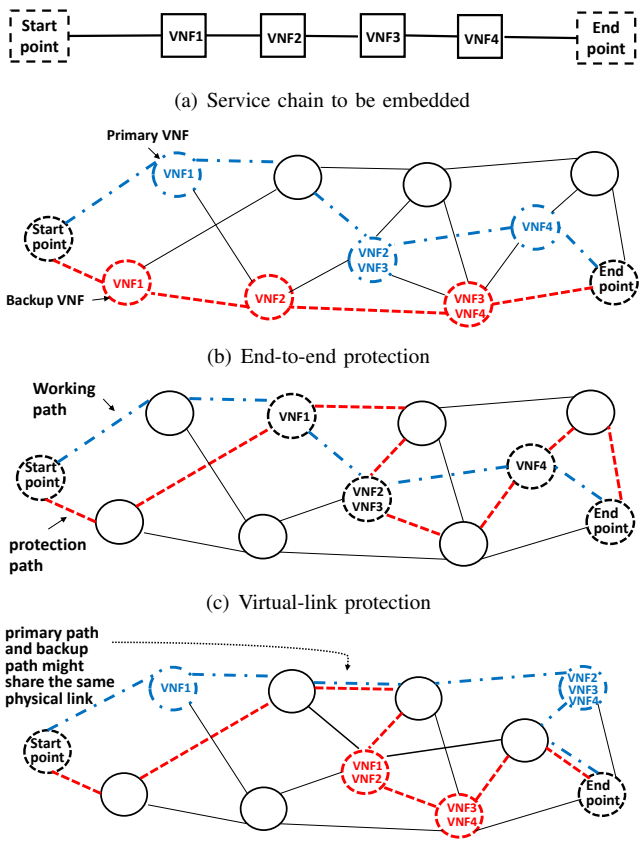

(a) Service chain to be embedded

(b) End-to-end protection

(c) Virtual-link protection

(d) Virtual-node protection

Fig. 2. Proposed protection schemes.

tive, resiliency must be taken into account in the design phase. This means deploying redundancy VNFs instances, which are kept in standby mode and activated upon the occurrence of a node or link failure that comprises the service of the primary VNFs. The redundancy schemes depend on the type of failures. In this section we present the three protection schemes proposed in this work. The redundancy approaches are divided in the following two categories:

a) On-Site Redundancy: Critical VNFs supporting critical services and customers require fast switchover to backup VNFs in order to ensure availability. In order 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 function can tolerate a 1:1 redundancy. The main benefits from a centralized redundancy is to reduce switchover time, which allow to speed up the recovery process, and reduce the amount of VNF internal state information that need to be transfered 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.

b) Off-Site Redundancy: A off-site redundancy architecture involves having redundant VNFs placed in (hot or cold) standby mode in selected 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 NFVI-Points-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 resiliency guarantees targeted the operator can choose between an on-site or an off-site redundancy approach [17].

In this work we propose three resiliency protection schemes. The first 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 and it 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 be node disjoint. Fig. 2(b) shows an example of such protection scheme, where a the SC illustrated in Fig. 2(a), composed fo four VNFs, is to 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. We refer to this protection strategy as *End-to-end protection (E2E-P)*.

The second protection scheme can be considered as an *On-site redundancy* protection scheme, with the objective to protect the virtual links used to concatenate the VNFs of a certain SC. Hence, 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 common physical links. An example of such scenario is shown in Fig. 2(c). We refer to this protection scheme as *Virtual-link protection (Vl-P)*.

Finally, 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' need when failures occur in nodes with higher probability with respect to links. An example of this scenario is shown in Fig. 2(d). We refer to this scenario as *Virtual-node protection (Vn-P)*.

V. Problem Statement

A. 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 with $\lambda(v, v')$ and a bandwidth capacity

$\beta(v, v')$. The physical nodes equipped with COTS hardware are reffered as *NFV nodes* and can have different amount of processing capacity in terms of number of Virtual machine that they can host. Finally, we consider a processing-related latency $\omega(v) : v \in V$, introduced by *NFV nodes*. This latency contribution is proportional to the number of SCs sharing the same VNF, hence, if a VNF is shared among a high number of SCs, the context switching latency would impact more the total latency.

B. VNF and service chains Modeling

Generally, a VNF is an abstracted object that performs operations on input traffic. Each VNF $f \in F$ has a processing capability which corresponds to the number of CPU Cores that are assigned to the VM that host the VNF f. We assume that a VNF shared among different SCs must run on a VM with enough capacity in terms of CPUs and that each VNF require one CPU core of the VM.

Moreover, we assume that each service corresponds to one SC modeled through a simple line graph $S^c = (E^c \cup U^c G^c)$ 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 the concept of requests are decoupled from the actual network functions that compose the Service chains. In other words VNFs are mapped to requests through a mapping parameter γ_u^c that specify the network function $f \in F$ requested by request $u \in U^c$, while requests are mapped to physical nodes through a decision variable. The same apply for the mapping of end-points, which we assume are fixed location and known a priori. Furthermore, we assume the each SC serve a set of users requesting a specific service from a specific physical location, and that each virtual link composing the SC is characterized by a bandwidth requirement $\gamma(u, u') : u, u' \in U^c, c \in C$. In addition, each SC is associated with a maximum tolerated latency, referred to as $\phi(c) : c \in C$.

TABLE I
PARAMETERS DESCRIPTION FOR THE ILP MODEL

Parameter	Domain	Description
η_u^c	$c \in C, u \in U^c$	Physical start/end point where u is mapped for SC c
γ_u^c	$c \in C$ $u \in G^c$	Network function requests u for SC c, $\gamma_u^c \in F$
$\beta_{v,v'}$	$(v, v') \in E$	Bandwidth capacity of physical link (v, v')
$\lambda_{v,v'}$	$(v, v') \in E$	Latency of physical link (v, v')
$\omega_v \in E$	$v \in V$	contest switching latency of node v.
$\tau_u^c \in F$	$c \in C, u \in U^c$	VNF f requested by request u in the SC c
ϕ^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

C. ILP models

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

- The number of *VNF nodes* is minimized.
- Latency requirements of SCs are met.
- Resiliency is achieved according to the goals of the above mentioned scenarios (see Fig. 2 of section IV).

TABLE II
VARIABLES DESCRIPTION FOR THE ILP MODELS

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

Objective function

$$Minimize \sum_{v \in V} a_v \tag{1}$$

We consider three types of constraints to solve this problems, namely: Placement constraint, routing constraints and performance constraints. Due to space limitation we show only the constraints for the E2E-P protection scenario and give a brief description of what differs in the other two scenarios, Vl-P and Vn-P.

Placement constraints

Constraints (2a) and (2b) force each primary/backup VNF to be mapped to one single node. Equations 2c) and (2d state that a corresponding VNF f is mapped to physical node v only if there is a primary/backup VNF request. Constraint (2e) enforces that primary and backup VNF request u cannot be mapped to the same node (node disjointness).

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

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

$$i_{f,v} \leq \sum_{u\in U^c:\gamma^c_u=f} m^c_{u,v} + n^c_{u,v} \qquad \forall f \in F, v \in V \quad (2c)$$

$$\sum_{u\in U^c:\gamma^c_u=f} m^c_{u,v} + n^c_{u,v} \leq M \cdot i_{f,v} \qquad \forall f \in F, v \in V \quad (2d)$$

$$m^c_{u,v} + n^c_{u,v} \leq 1 \qquad \forall u \in U^c, c \in C, v \in V : v \neq \eta^c_u \quad (2e)$$

Routing constraints

Constraints (3a) [(3b)] 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. Note that equations (3a)-(4d) contain products of binary variables that we linearize in order to solve the ILP models.

$$w^c_{v,v',x,y,u,u'} \leq m^c_{u,x} \cdot m^c_{u',y} \qquad \forall c \in C, (v, v') \in E, x, y \in V, (u, u') \in G^c \quad (3a)$$

$$p^c_{v,v',x,y,u,u'} \leq n^c_{u,x} \cdot n^c_{u',y} \qquad \forall c \in C, (v, v') \in E, x, y \in V, (u, u') \in G^c \quad (3b)$$

Equations (4a)-(4b) [(4c)-(4d)] are source and destinations constraints for primary and backup VNF requests, respectively. They ensure that a virtual link starts in node x where primary [backup] start-point request u of SC c is mapped, and that the virtual link end in node y where primary [backup] end-point requests u' of SC c is mapped.

$$\sum_{(x,v)\in E: x,y\in V} w^c_{x,v,x,y,u,u'} \cdot m^c_{u,x} \cdot m^c_{u',y} = 1 \qquad \forall c \in C, (u, u') \in G^c \quad (4a)$$

$$\sum_{(v,y)\in E: x,y\in V} w^c_{v,y,x,y,u,u'} \cdot m^c_{u,x} \cdot m^c_{u',y} = 1 \qquad \forall c \in C, (u, u') \in G^c \quad (4b)$$

$$\sum_{(x,v)\in E: x,y\in V} p^c_{x,v,x,y,u,u'} \cdot n^c_{u,x} \cdot n^c_{u',y} = 1 \qquad \forall c \in C, (u, u') \in G^c \quad (4c)$$

$$\sum_{(v,y)\in E: x,y\in V} p^c_{v,y,x,y,u,u'} \cdot n^c_{u,x} \cdot n^c_{u',y} = 1 \qquad \forall c \in C, (u, u') \in G^c \quad (4d)$$

During the mapping of primary/backup VNF requests on a physical path between x and y incoming links for the node x are not considered, constraint (5a), and no outgoing link for node y is considered (constraint (5b)

$$\sum_{(v,x)\in E: v\in V} w^c_{v,x,x,y,u,u'} = \sum_{(v,x)\in E: v\in V} p^c_{v,x,x,y,u,u'} = 0 \qquad \forall c \in C, x \in V, y \in V : x \neq y, (u, u') \in G^c \quad (5a)$$

$$\sum_{(y,v)\in E: v\in V} w^c_{y,v,x,y,u,u'} = \sum_{(y,v)\in E: v\in V} p^c_{y,v,x,y,u,u'} = 0 \qquad \forall c \in C, x \in V, y \in V : x \neq y, (u, u') \in G^c \quad (5b)$$

Constraints (6a)-(6d) are transit constraints for primary/backup VNF requests. In particular, constraints (6a) and (6b) ensure that for any intermediate node ω within the physical path between x and y, if one of the incoming links belong to the primary/backup physical path, then also one of its outgoing links belong to the physical path. While constraints (6c) [(6d)] avoid the use of multiple incoming [outgoing] links of the intermediate node.

$$\sum_{(v,w)\in E: v\in V} w^c_{v,w,x,y,u,u'} = \sum_{(w,v')\in E: v\in V} w^c_{w,v',x,y,u,u'} \qquad \forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (6a)$$

$$\sum_{(v,w)\in E: v\in V} p^c_{v,w,x,y,u,u'} = \sum_{(w,v')\in E: v\in V} p^c_{w,v',x,y,u,u'} \qquad \forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (6b)$$

$$\sum_{(v,w)\in E: v\in V} w^c_{v,w,x,y,u,u'} \leq 1 \qquad \forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (6c)$$

$$\sum_{(v,w)\in E: v\in V} p^c_{v,w,x,y,u,u'} \leq 1 \qquad \forall c \in C, w \in V, x, y \in V : x \neq w, y \neq w, (u, u') \in G^c \quad (6d)$$

Finally, constraint (7a) ensures that a physical link (v, v') is whether part of the primary physical path or in the backup physical path used for the embedding of all VNF request of SC c.

$$\sum_{(u,u')\in G^c} w^c_{v,v',x,y,u,u'} + p^c_{v,v',x,y,u,u'} \leq 1 \qquad \forall c \in C, x, y, v, v' \in V : (v, v') \wedge (v', v) \in E \quad (7a)$$

Latency and capacity constraints

$$\sum_{f\in F} i_{f,v} \leq M.a_v \qquad \forall v \in V \quad (8a)$$

$$a_v \leq \sum_{f\in F} i_{f,v} \qquad \forall v \in V \quad (8b)$$

$$\sum_{\substack{c\in C \\ (u,u')\in G^c \\ x,v\in V}} (w^c_{v,v',x,y,u,u'} + p^c_{v,v',x,y,u,u'}) \cdot \beta_{u,u'} \leq C_{v,v'} \qquad \forall (v, v') \in E \quad (8c)$$

$$\sigma^c_w = \sum_{v\in V, u\in U^c} m^c_{u,v} \cdot \omega_v \qquad \forall c \in C \quad (8d)$$

$$\sigma^c_p = \sum_{v\in V, u\in U^c} n^c_{u,v} \cdot \omega_v \qquad \forall c \in C \quad (8e)$$

$$\sum_{\substack{x,v\in V \\ (u,u')\in G^c \\ (v,v')\in E}} (w^c_{v,v',x,y,u,u'} \cdot \lambda_{v,v'}) + \sigma^c_w \leq \phi_c \qquad \forall c \in C \quad (8f)$$

$$\sum_{\substack{x,v\in V \\ (u,u')\in G^c \\ (v,v')\in E}} (p^c_{v,v',x,y,u,u'} \cdot \lambda_{v,v'}) + \sigma^c_p \leq \phi_c \qquad \forall c \in C \quad (8g)$$

$$\sum_{f\in F} i_{f,v} \leq N_{VM}(v) \qquad \forall v \in V \quad (8h)$$

$$\sum_{\substack{c\in C \\ u\in U^c:\gamma^c_u=f}} m^c_{u,v} + n^c_{u,v} \leq N_{req}(f) \qquad \forall v \in V, f \in F \quad (8i)$$

Constraints (8a)-(8b) select the active *NFV nodes*. A node is considered active if it hosts at least one single VNF. Constraint (8c) ensures that link capacity is not exceeded, whereas constraints (8d) and (8e) compute the context switching latency contribution σ_w^c and σ_p^c for primary and backup embedding of SC c, respectively. The maximum latency of primary/backup embedding of SC c are constrained in (8f)-(8g). Finally, the maximum number of VMs that node v can host is bounded by (8h), and the number of parallel requests that a given VNF can serve is constrained in (8i).

D. Modeling other scenarios

With respect to the E2E-P, in the Vl-P we must ensure that the primary and backup physical path used to map a certain virtual link of a SC do not share any physical link. and no node disjointness constraint is required. Here the link disjointness is applied considering only one single virtual link at the time. Finally, for the Vn-P scenario, only the node disjointness constraint apply and no disjointness constraints between primary/backup physical paths are needed since they can use the same physical links.

E. Problem complexity

The total number of variables and constraints of the E2E-P optimization problems can be calculated using the following formulas:

$$N_{vars} = |V| \cdot (2 \cdot |C| \dot{|} U^c| + 2 \cdot |C||E||V||G^c| + |F| + 1) \quad (9)$$

$$\begin{aligned} N_{const} = |C| \cdot (2 \cdot |U^c| + |U^c||V| + 2 \cdot |E||G^c||V|^2 + \\ 4 \cdot |G^c| + |V|^2 + 4 \cdot |V|^3 |G^c| + |V|^2 |E| + 4) + \\ 3 \cdot |V| \cdot (|F| + 1) + |E| \end{aligned} \quad (10)$$

In both equations we observe that the dominant term for variables and constraints is $2 \cdot |E||G^c||V|^2$. Thus, the problem complexity, for all proposed protection scenarios, given by the sum of the number of variables and number of constraints is in the order of $O(|G^c| \cdot |E| \cdot |C| \cdot |V|^2)$.

VI. Case study and results

In this section we present and discuss the results of the ILP models shown in section V. To solve the ILP problems CPLEX 12.6.1.0 installed on hardware platform equipped with 8×2 GHz processor and 8 Gb of RAM. In order to evaluate the impact of latency requirements on the protection scenarios we investigated the embedding of two types of services chains: The first one with stringent latency requirement (On-line Gaming) and second one with non stringent latency requirements (Web Service). The maximum end-to-end tolerated latency for these services has be set to 500 ms for Web-service and 60 ms for Online-Gaming [11]. Table III shows the VNFs composing both SCs, their bandwidth requirements and maximum allowed latency. Due to the hardness of the optimization problem, we considered only two SCs in each optimization run and solved the ILP models for two homogeneous cases, when 2 SCs of the same type are embedded in the network, and for one heterogeneous case, when the two SC types are embedded in the network. As physical topology we considered the NSFNET network (14 nodes and 22 bidirectional links). In addition, we assume that all the physical nodes are *NFV nodes* and can act as start/end points of SCs. Each *NFV node* is assumed to have the same capacity in terms of VMs they can accomodate. We set the context switching delay to 4 ms per VNF and assume that link capacity is 1 Gbps (i.e., link capacity is not a strict constraint). Moreover, we assume that the bandwidth requirements of virtual links chaining VNFs is the same for the whole SC (i.e, data rate do not change at the output of the VNFs). These results were obtained averaging the results of 10 instances, for each value of nodes capacity and each protection scenario, considering different pairs of start/end points at each instance. Fig. 3(a), Fig. 3(b) and Fig. 3(c) show the average number of active nodes needed to support of the proposed protection scenarios for different values of node capacity (number of VMs that a node can host).

TABLE III
Performance Requirements for the Service Chains

Service Chain	Chained VNFs	β	ϕ_c
Web-Service	NAT-FW-TM-WOC-IDPS	100 kbit/s	500 ms
Online-Gaming	NAT-FW-VOC-WOC-IDPS	50 kbit/s	60 ms

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

A. Impact of latency

Fig. 3(a) presents the number of active nodes for the less stringent SC in terms of latency (Web-Service). We observe that all protection scenarios are possible and that the Vl-P scenario activates the same amount of Unprotected Scenario. We note that a service with low requirements on latency can be protected against single-link failures (Vl-P) with no additional *NFV nodes* with respect to the *Unprotected* case (baseline). On the other hand, providing protection against both single-link and single failure (E2E-P) requires the activation of twice the amount of *NFV nodes*. In case of SCs with high latency requirements, in Fig. 3(b), we observe that all scenarios lead to infeasible solutions when only two VMs are allowed per node, mainly due to the fact that distributing VNFs among high number of nodes increases the latency of physical paths needed to chain the VNFs and consequently violates the latency constraint. We also observe that the unprotected scenario, considered as baseline case, requires at least three VM per *VNF node* to meet latency requirement. Different results were obtained for the Vl-P case which is infeasible independently from node capacity. This means that the operator is constrained to place backup VNFs "Off-site" to provide resiliency against only single-link failures, when only latency critical SCs are deployed. In this case, it is preferable to provide resiliency against both node and link failures (E2E-P) rather than provide protection against only node failures (Vn-P) since both scenarios activates the same number of *NFV nodes* independently from node capacity. For the heterogeneous scenario shown in Fig. 3(c), all protection scenarios are possible with at

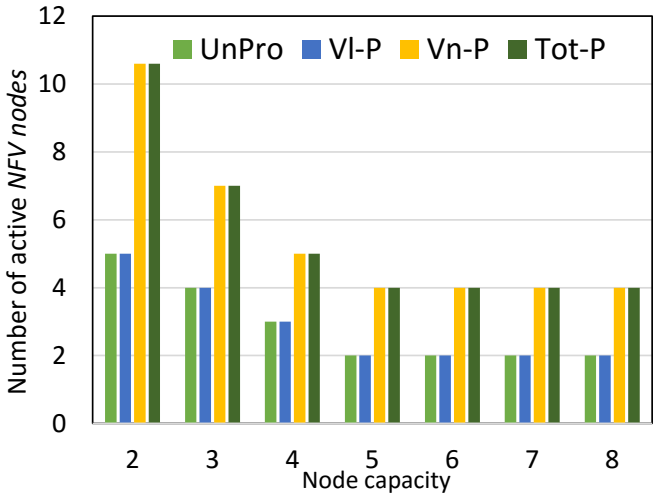

(a) Web-service

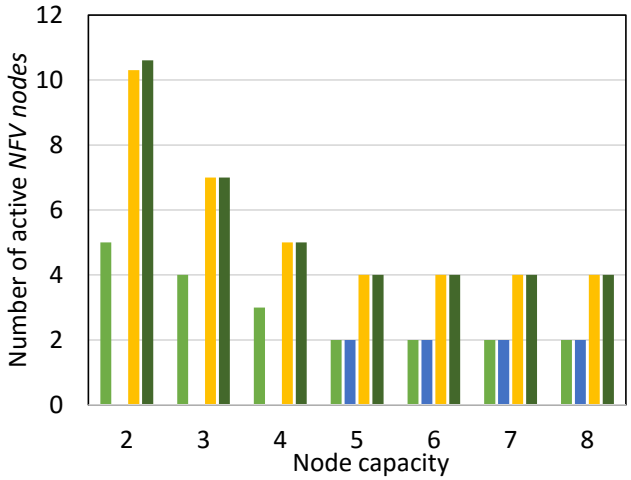

(b) Heterogeneous

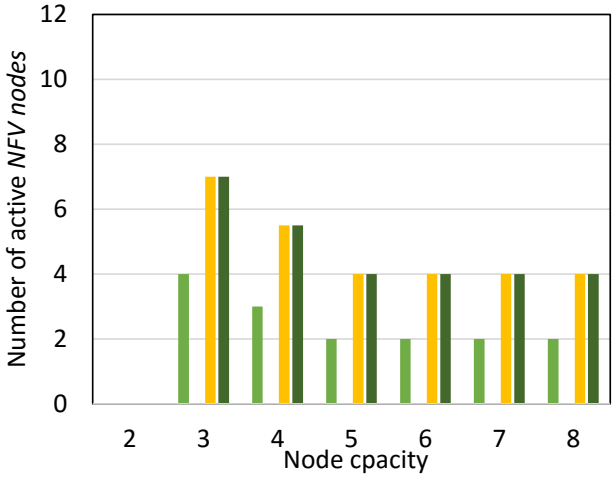

(c) Online-gaming

Fig. 3. Comparison of the proposed protection scenarios for different latency requirements

least 2 VMs except from the Vl-P scenario which is only possible starting from 5 VMs. In terms of latency, it means that deploying SCs with different latency requirements and sharing VNFs between SCs can guarantee resiliency with a small number of VMs, and consequently less failure impact within *NFV nodes*. On the other hand, for Vn-P and E2E-P protection scenarios, deploying the same SCs or different SCs in terms of latency requirements does not impact resources in terms of *NFV nodes* as similar results were obtained in both homogeneous and heterogeneous cases starting from 3 VMs per node, except from the case of On-line gaming SCs when 2 VMs are allowed per node.

B. *Effect of node capacity*

As can be seen from Fig. 3(a), 3(b) and 3(c), increasing node capacity allows to decrease the number of *NFV nodes* irrespectively from the type of SCs deployed. In general, we observe that the number of active nodes is halved for all protection scenarios, when increasing the number of VMs per node from 2 to 5. Further increase of node capacity does not impact the number of active nodes, which means that VNF consolidation is limited by latency, as consolidating more VNFs into less nodes would increase the impact of context switching latency.

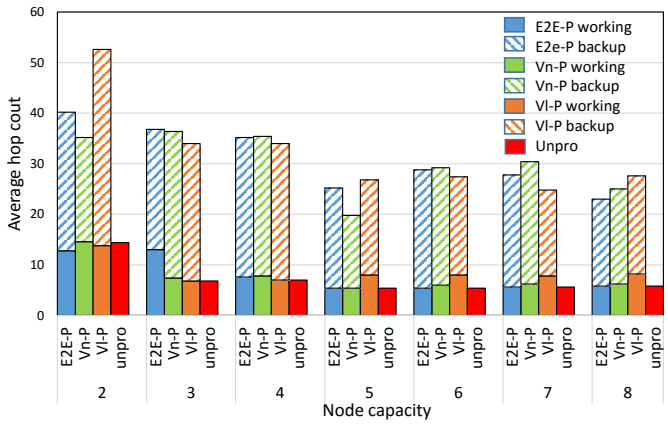

Fig. 4. primary/backup path lengths with respect to node capacity.

C. *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 Fig.4 we show the primary/backup paths lengths when 2 SCs with low requirements on latency are deployed. These results were obtained by averaging the paths lengths of 5 start/end point pairs randomly selected and tested for all protection scenarios. We observe that at the increasing of node capacity the length of the primary path does not change significantly, for all protection strategies. Different results are observable in case of backup primary paths, where it is clear that increasing node capacity does not mean reducing backup paths lengths. This is shown by the fact that allowing more than 5 VMs per nodes does not reduce the

average backup path length, meaning that a trade-off between consolidation of VNFs and link capacity exist.

VII. Conclusions

In this work we proposed three different protection strategies to provide resilient SCs deployment against single-node, single-link, single-node and single-link failures. We reported the formulation of one of them through ILP, solved the ILP models considering a small number of SCs with different latency requirements, and found that a trade-off between node capacity and latency of the deployed SCs. In our small-scale scenario, we conclude that in order to provide resiliency to SCs against single-link and single-node failures up to 107% more *NFV nodes* are needed with respect to the unprotected scenarios and the case where only single-link failures are targeted. Future steps of this work aim at developing an heuristic model to allow solving larger instances (large number of SCs) in reasonable time. We also aim at extending the proposed models with a shared protection scheme.

Acknowledgment

This research has received fundings from the European Community Seventh Framework Program FP7/2013-2015 under grant agreement no. 317762 COMBO project, and from COST ACTION 15127 RECODIS (Resilient Communication Services Protecting end-user Applications from Disaster-based Failures).

References

[1] Network Functions Virtualisation, "Draft ETSI GS NFV-SEC 001 v0.2.1 (2014-06)," 2014.

[2] J. Halpern and C. Pignataro, "Service Function Chaining (SFC) Architecture," Tech. Rep., 2015.

[3] R. Guerzoni *et al.*, "Network functions virtualisation: an introduction, benefits, enablers, challenges and call for action, introductory white paper," in *SDN and OpenFlow World Congress*, 2012.

[4] W. Liu, H. Li, O. Huang, M. Boucadair, N. Leymann, Z. Cao, Q. Sun, and C. Pham, "Service Function Chaining (SFC) general use cases," European Telecommunication Stadards Institute (ETSI), Service Functions Chaining (SFC) framework, Tech. Rep., 2014.

[5] M. Boucadair, C. Jacquenet, R. Parker, D. Lopez, J. Guichard, and C. Pignataro, "Service function chaining: Framework & architecture," ETSI-Service Functions Chaining (SFC) framework, Tech. Rep., 2013.

[6] R. Mijumbi, J. Serrat, J. l. Gorricho, S. Latre, M. Charalambides, and D. Lopez, "Management and orchestration challenges in network functions virtualization," *IEEE Communications Magazine*, vol. 54, no. 1, pp. 98–105, January 2016.

[7] A. Fischer, J. F. Botero, M. Till Beck, H. De Meer, and X. Hesselbach, "Virtual Network Embedding (VNE): A survey," *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 4, pp. 1888–1906, 2013.

[8] M. R. Rahman and R. Boutaba, "SVNE: Survivable virtual network embedding algorithms for network virtualization," *IEEE Transactions on Network and Service Management*, vol. 10, no. 2, pp. 105–118, June 2013.

[9] C. Prodhon and C. Prins, "A survey of recent research on location-routing problems," *European Journal of Operational Research*, vol. 238, no. 1, pp. 1–17, 2014.

[10] S. Mehraghdam, M. Keller, and H. Karl, "Specifying and placing chains of virtual network functions," in *IEEE 3rd International Conference on Cloud Networking (CloudNet)*. IEEE, 2014, pp. 7–13.

[11] M. Savi, M. Tornatore, and G. Verticale, "Impact of processing costs on service chain placement in network functions virtualization," in *IEEE Conference on Network Function Virtualization and Software Defined Network (NFV-SDN)*, Nov 2015, pp. 191–197.

[12] I. Cerrato, M. Annarumma, and F. Risso, "Supporting fine-grained network functions through intel DPDK," in *Third European Workshop on Software Defined Networks (EWSDN)*. IEEE, 2014, pp. 1–6.

[13] J. W. Jiang, T. Lan, S. Ha, M. Chen, and M. Chiang, "Joint vm placement and routing for data center traffic engineering," in *IEEE Conference on Information Communication (INFOCOM)*. IEEE, 2012, pp. 2876–2880.

[14] H. Moens and F. De Turck, "VNF-P: A model for efficient placement of virtualized network functions," in *10th International Conference on Network and Service Management (CNSM)*. IEEE, 2014, pp. 418–423.

[15] F. Machida, M. Kawato, and Y. Maeno, "Redundant virtual machine placement for fault-tolerant consolidated server clusters," in *IEEE Network Operations and Management Symposium (NOMS)*, April 2010, pp. 32–39.

[16] M. Scholler, M. Stiemerling, A. Ripke, and R. Bless, "Resilient deployment of virtual network functions," in *the 5th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*. IEEE, 2013, pp. 208–214.

[17] ETSI, "GS NFV-REL 001 v1. 1.1: Network functions virtualisation(nfv); resiliency requirements," ETSI industry Specfication Group (ISG) Network Functions Virtualisation (NFV), Tech. Rep., 2015.