



HAL
open science

Resource Requirements for Reliable Service Function Chaining

Andrea Tomassilli, Nicolas Huin, Frédéric Giroire, Brigitte Jaumard

► **To cite this version:**

Andrea Tomassilli, Nicolas Huin, Frédéric Giroire, Brigitte Jaumard. Resource Requirements for Reliable Service Function Chaining. 2018 IEEE International Conference on Communications (ICC 2018), May 2018, Kansas City, United States. 10.1109/ICC.2018.8422774 . hal-01921096

HAL Id: hal-01921096

<https://inria.hal.science/hal-01921096>

Submitted on 13 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Resource Requirements for Reliable Service Function Chaining

Andrea Tomassilli, Nicolas Huin and Frédéric Giroire
Université Côte d’Azur, CNRS, Inria
Sophia Antipolis, France

Brigitte Jaumard
Concordia University
Montreal (Qc) Canada

Abstract—In the context of Software-Defined Networks (SDN), Network Function Virtualization (NFV) is a new network paradigm in which network functions are implemented in software as Virtual Network Functions (VNFs). To meet the demand, VNFs are next interconnected to form different complete end-to-end services, also known as a Service Function Chains (SFCs). We study the problem of deploying reliable Service Function Chains over a virtualized network function architecture. While there is a need for reliable service function chaining, there is a high cost to pay for it in terms of bandwidth and VNF processing requirements. We investigate two different protection mechanisms and discuss their resource requirements, as well as the latency of their paths. For each mechanism, we develop a scalable exact mathematical model using column generation.

I. INTRODUCTION

Software-defined networking (SDN), network functions virtualization (NFV), and network virtualization (NV) have provided new ways to design and manage networks. Over the past decade, we have seen significant differences in the applications and services we depend on to run the economies and the individual lives, and in the computing and storage solutions we rely on to manage all the required big data that is generated.

With Network Function Virtualization, network functions, such as firewalls, content filtering, or intrusion detection prevention system, can be executed on generic-purpose servers, in a specific order. For instance, a firewall may need to inspect incoming packets before encryption or compression takes place. The interconnection of network functions to provide end-to-end services is known as Service Function Chaining (SFC).

NFV gives network operators a great freedom to customize their networks and offers a chance to reduce both the capital expenditure and operational costs. Indeed, design choices such as the placement of the functions may have a significant impact on the overall expenditure. On the other hand, network operators are responsible for ensuring that the network provides all the services that users are expecting, with the agreed quality of service (QoS). Hence, different factors need to be taken into account during network design and management in order to optimize both the cost and the performance.

However, the underlying network that connects all of these things has remained virtually unchanged. Demands of the exploding number of devices using the network are stretching its limits, and network failures such as (multiple) link or node failures may have a significant impact on the QoS experienced by the customers and lead to SLA (Service Level Agreement) violations. Consequently, resiliency needs to be

strongly addressed while designing a network.

Network failures have been widely investigated (see, e.g., [1],[2],[3]). One of the key findings of [1] is that links experience about an order of magnitude more failures than devices. Moreover, according to their analysis, low-cost commodity switches are highly reliable. This finding is also confirmed by [3]. On the other hand, links are failure-prone. Indeed, in a monitored network, each link experiences in average 16 failures per year, considering a five years period [2]. Despite this, most of the failures have very short duration. The majority of link failures are solved within 5 minutes (the median time to repair is 13s). Another finding is that link failures tend to be isolated. The short time to repair and the absence of a relation between link failures motivate us to focus our attention on the *single link failure* scenario.

Fault management techniques can be grouped into two categories: *restoration* and *protection*. Restoration is a reactive approach in which a backup path is computed and established after a failure. Protection is a proactive technique in which capacity on links is reserved during connection setup. Restoration schemes are more efficient in utilizing capacity, but, on the other side, protection schemes have a faster restoration time and offer a guaranteed recovery [4]. We study the latter in this paper.

There are different protection schemes. In *dedicated protection*, some spare capacity is reserved for each backup path. This implies that the backup resources are used for at most one path. In *shared protection*, backup paths can share some link capacity if failures in their primary paths do not occur simultaneously. Thus, in shared protection, capacity is used more efficiently [5]. However, dedicated protection is often used by network operators because of its simplicity. We thus study both protection schemes in this paper.

Each protection scheme may have two different recovery mechanisms: a local repair (i.e., link protection) or an end-to-end repair (i.e., path-protection). Link protection schemes reroute the traffic around the failing link. In path protection schemes, the traffic is rerouted on a link-disjoint backup path. In this case, the backup path would be used in all the failure situations that involve links of the primary path. Path-protection mechanisms have been shown to lead to better resource utilization compared to link protection [6], [7].

We thus consider the problem of *providing for each demand, a primary and a link-disjoint backup path, under both dedicated and shared protection schemes*. Moreover, the problem also consists in *provisioning VNFs* in order to ensure that the *traversal order* of the network functions by each path is respected. This adds a challenge to the classical version of

the problem.

Our goal is to minimize the bandwidth requirements while ensuring that the delays on primary and backup paths stay below SLAs.

Our contributions are as follows:

- To the best of our knowledge, we are the first to propose a *scalable exact method* to solve the problem of *reliable service function chaining*. The method is based on a decomposition model using column generation.
- The model allowed us to solve the problem with dedicated and shared protection schemes for networks with up to 1000 traffic requests.
- We also studied the *costs in terms of bandwidth and computation requirements* of both protection schemes and for networks of different sizes. When service function chaining is considered, dedicated protection requires three times more bandwidth and two times more processing than without protection. The ratios drop to 1.5 and 1.25 for shared protection.
- We additionally study the impact of the number of nodes of the network being able to host VNFs on the bandwidth requirements and on the delays of both primary and backup paths.

The paper is organized as follows. We study two different protection schemes: Dedicated Protection, in Section IV-B and Shared Backup Path Protection, in Section IV-C. For each of them, we propose both a compact Integer Linear Program (ILP) formulation and a decomposition model. In Section V, we compare the models and show the superiority of the decomposition models over the compact ILP formulations in terms of scalability. We then study the impact of design choices, such as the number of VNF nodes and the kind of protection, on the experienced delay by both the primary and backup paths of the demands, as well as the impact on the bandwidth requirements.

II. RELATED WORK

The design of survivable networks has been widely studied in the network literature (see, e.g., [8],[6]). However, when dealing with NFV and SFCs, an additional challenge is to map network functions to nodes and to guarantee that the execution order of the network functions is respected in both primary and backup paths.

The problem of guaranteeing service continuity in Service Function Chain scenario has started to be investigated recently. Both *restoration* [9], [10] and *protection* [11], [12], [13] techniques have been investigated. In [9], the authors address VNF placement and chaining in the presence of physical link failures. The proposed algorithm makes use of a Monte-Carlo Tree Search algorithm to place VNFs and simultaneously steer traffic flows across them. When a link fails, the algorithm reactively re-maps the failed virtual links in other substrate paths. In Hmaity *et al.* [10], the authors consider the problem of recovering the traffic path after the failure of a network function. In their proposed solution, an alternative VNF is selected, in a greedy manner, to replace the failed one and then the communication is ensured by allocating a path between the new VNF and its neighbors.

In [12], the authors consider node and link failures and they propose a heuristic algorithm with the goal of meeting the client's reliability requirements. They propose two algorithms. The first one is based on dedicated protection and the second

one on shared protection. In [11], the authors propose a compact ILP model in order to provide resiliency against single node, virtual link, and single node/single virtual link failure scenarios. They aim to reduce the number of VNF nodes used. The difference with our work is that the authors consider *link protection*, while we look into *path protection*, and their ILP models are not scalable.

[13] discusses measures on how to backup resources in order to protect network services from failures. They consider both node and link failures and propose a resource allocation algorithm heuristic-based that aims at keeping the number of physical resources allocated to VNF chains low.

The *main difference* with our work is that we propose a scalable exact decomposition model to provide reliable service function chaining. (Other differences is that using path protection in order to minimize network bandwidth was also not considered in this setting.) Column generation techniques have been shown to be effective in dealing with both Service Function Chaining [14], [15] and failure protection [16]. In [14], the authors propose a decomposition modeling for the SFC Problem with the goal of optimizing the bandwidth. Through extensive numerical experiments, they show that their model can solve exactly and in an efficient way the problem. We here extend their results to the case of unreliable networks in the case of single link failure scenario.

III. PROBLEM AND NOTATIONS

We model the network as a graph $G = (V, L)$, where V represents the set of nodes and L the set of links. A request is modeled as a quadruple (v_s, v_d, c, D_{sd}^c) with v_s the source, v_d the destination, $c = f_1^c, f_2^c, \dots, f_{n_c}^c$ the sequence of VNFs that need to be performed with n_c the chain length, and D_{sd}^c the required units of bandwidth.

Each network function f has associated processing requirements per unit of bandwidth, denoted with Δ_f . Namely, given a request (v_s, v_d, c, D_{sd}^c) , the number of cores needed to process the i -th function of the chain c is equal to $D_{sd}^c \cdot \Delta_{f_i}^c$. Different chains may have different maximum tolerated latency. For instance, the latency requirement of Video Streaming is less stringent than Online Gaming. Following a similar idea as in [11], we associate to each chain c a maximum tolerated delay, denoted as $\phi(c)$. Each network function f is associated with a processing latency per unit of bandwidth ρ_f^u , which also depends on the node in which the function is performed, and each link with a transmission and propagation latency λ_l .

Not all nodes may be enabled to run virtual functions. We denote by $V^{\text{VNF}} \subseteq V$ the set of VNF-enabled nodes equipped with Commercial Off The Shelf (COTS) hardware. We are given for each node $v \in V^{\text{VNF}}$ a capacity CAP_v , representing the amount of available resources, such as CPU, memory, and disk. Similarly, for each link $\ell \in L$, we are given the transport capacity CAP_ℓ .

The optimization task is to minimize the amount of bandwidth used in the network. At the same time, the problem consists in providing to each demand an edge disjoint backup path and to guarantee that the traversing order of the functions is respected in both primary and backup paths. Both node and link capacities must be respected, as well as the maximum tolerated latency for each request.

As in [14], to model the function ordering problem, we use a

layered G^L graph with n_c+1 layers. We denote by $u(i)$ the copy of node u in layer i . The paths for demand D_{sd}^c starts from node $v_s(0)$ in layer 0 and ends at node $v_d(n_c)$ in layer n_c . Layer i corresponds to nodes of the paths encountered after the i^{th} function of the service chain.

Using link $(u(i), v(i))$ on G^L , implies using link (u, v) on G . On the other hand, using link $(u(i), u(i+1))$ implies using the $(i+1) - th$ function of the chain at node u .

IV. OPTIMIZATION MODELS

We now present the optimization models for the dedicated and shared protection schemes, in Section IV-B and IV-C respectively. For each scheme, we present both a compact ILP formulation and a decomposition model.

A. Generalities on Column Generation

The column generation technique is an efficient algorithm used to solve problems with a large number of variables. The main idea is to decompose the original problem into a Restricted Master Problem (RMP), with only a small subset of configurations, and one or several subproblems, called Pricing Problems. They are solved in turn iteratively using the dual multipliers of the RMP. The goal of the pricing problems is to find new configurations with negative reduced cost that would improve the current optimal solution z_{LP} . The new configurations found are added iteratively to the RMP.

If no configuration with negative reduced cost exists, then the solution of the linear relaxation of the RMP z_{LP}^* is optimal. In order to derive an ILP solution \tilde{z}_{ILP} , we use an ILP solver on the last RMP.

B. Dedicated Protection

We consider here a dedicated protection scheme, also known as 1+1 protection. The capacity for the backup path is fully reserved. This is a method often used by operators in the case each demand is load balanced over both paths used at less than 50%. When a failure involving the primary path happens, the traffic on the failing path is switched to the second path.

Model NFV_ILP_DP

In the case of dedicated protection, the problem consists in finding two sd -paths in G^L for each demand. Note that the paths have to be edge disjoint (i.e., if a path uses a link l for a layer of G^L then the other path cannot use l in any layer of G^L) but not node disjoint.

Variables:

- $\varphi_{\ell,p}^{sd,c,i}, \varphi_{\ell,b}^{sd,c,i} \in \{0,1\}$, where $\varphi_{\ell,p}^{sd,c,i} = 1$ ($\varphi_{\ell,b}^{sd,c,i} = 1$) if (v_s, v_d, c, D_{sd}^c) is provisioned on link ℓ for the primary (backup) path.
- $a_{v,p}^{sd,c,i}, a_{v,b}^{sd,c,i} \in \{0,1\}$ where $a_{v,p}^{sd,c,i} = 1$ ($a_{v,b}^{sd,c,i} = 1$) if f_{i+1}^c is installed on node v for the primary (backup) path. If $v \notin V^{VNF}$, $a_{v,1}^{sd,c,i}$ and $a_{v,2}^{sd,c,i}$ are set to 0.

Objective: minimization of the bandwidth used in the network

$$\min \sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} D_{sd}^c \sum_{\ell \in L} \sum_{i=0}^{n_c} (\varphi_{\ell,p}^{sd,c,i} + \varphi_{\ell,b}^{sd,c,i}) \quad (1)$$

Constraints: both primary and backup paths must satisfy the following constraints. They are written for the general case.

Flow Conservation: for all $(v_s, v_d) \in \mathcal{SD}$, $c \in C_{sd}$,

$$\sum_{\ell \in \omega^+(v)} \varphi_{\ell}^{sd,c,0} - \sum_{\ell \in \omega^-(v)} \varphi_{\ell}^{sd,c,0} + a_v^{sd,c,0} = \begin{cases} 1 & \text{if } v = v_s \\ 0 & \text{else} \end{cases} \quad (2)$$

$$\sum_{\ell \in \omega^+(v)} \varphi_{\ell}^{sd,c,n^c} - \sum_{\ell \in \omega^-(v)} \varphi_{\ell}^{sd,c,n^c} - a_v^{sd,c,n^c} = \begin{cases} -1 & \text{if } v = v_d \\ 0 & \text{else} \end{cases} \quad (3)$$

$$\sum_{\ell \in \omega^+(u)} \varphi_{\ell}^{sd,c,i} - \sum_{\ell \in \omega^-(u)} \varphi_{\ell}^{sd,c,i} + a_v^{sd,c,i} - a_v^{sd,c,i-1} = 0. \quad 0 < i < n_c \quad (4)$$

Link capacity: for all $\ell \in L$,

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} D_{sd}^c \sum_{i=0}^{n_c} (\varphi_{\ell,p}^{sd,c,i} + \varphi_{\ell,b}^{sd,c,i}) \leq \text{CAP}_{\ell}. \quad (5)$$

Node capacity: for all $v \in V^{VNF}$,

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} \sum_{i=0}^{n_c-1} D_{sd}^c \Delta f_i^c (a_{v,p}^{sd,c,i} + a_{v,b}^{sd,c,i}) \leq \text{CAP}_v. \quad (6)$$

Latency: for all $(v_s, v_d) \in \mathcal{SD}$, $c \in C_{sd}$,

$$\sum_{\ell \in L} \sum_{i=0}^{n_c} \varphi_{\ell}^{sd,c,i} \lambda_{\ell} + \sum_{i=0}^{n_c-1} D_{sd}^c \rho_{f_i^c}^v a_v^{sd,c,i} \leq \phi(c). \quad (7)$$

In order to guarantee that the paths are edge disjoint, we add the following constraint. For all $(v_s, v_d) \in \mathcal{SD}$, $c \in C_{sd}$, $\ell \in L$,

$$\sum_{i=0}^{n_c} \varphi_{\ell,p}^{sd,c,i} + \sum_{i=0}^{n_c} \varphi_{\ell,b}^{sd,c,i} \leq 1. \quad (8)$$

Model NFV_CG_DP

We now propose a decomposition model for the dedicated protection scenario. Each configuration consists of a Service Path. A Service Path for a request (v_s, v_d, c, D_{sd}^c) is composed of: (i) a path, i.e., an ordered set of nodes from the source to the destination, and (ii) a set of locations for the VNFs in the SFC request. The goal of the Master Problem is thus to select a pair of configurations, i.e., Service Paths for each request.

The set of configurations must be chosen in such a way that: (i) each request is associated to a pair of edge-disjoint configurations; (ii) node and link capacities are respected and (iii) the overall required bandwidth is minimized.

• $\pi \in \Pi_{sd}^c$ is a service path from s to d . A service path is composed of a path and a set of node/function pairs (v, f) expressing that the function f is installed on node v .

• $a_{v,\pi}^f \in \{0,1\}$, where $a_{v,\pi}^f = 1$ if f is installed on node v for service path $\pi \in \Pi_{sd}^c$ w.r.t sd, c .

• $\delta_{\ell}^{\pi} \in \{0,1\}$, where $\delta_{\ell}^{\pi} = 1$ if link ℓ belongs to path π .

Variables:

- $y_{\pi,p}^{sd,c}, y_{\pi,b}^{sd,c} \geq 0$, where $y_{\pi,p}^{sd,c} = 1$ ($y_{\pi,b}^{sd,c} = 1$) if the request from v_s to v_d for service chain c is forwarded through service path π for the primary (backup) path.

Objective

$$\min \sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} \sum_{\pi \in \Pi_{sd}^c} D_{sd}^c \text{LEN}(\pi) (y_{\pi,p}^{sd,c} + y_{\pi,b}^{sd,c}) \quad (9)$$

One primary and one backup path per demand and per chain:

$$\sum_{\pi \in \Pi_{sd}^c} y_{\pi,p}^{sd,c} \geq 1. \quad (10a) \quad \sum_{\pi \in \Pi_{sd}^c} y_{\pi,b}^{sd,c} \geq 1. \quad (10b)$$

Edge disjoint primary and backup path per demand, per chain and per link:

$$\sum_{\pi \in \Pi_{sd}^c} \delta_\ell^\pi (y_{\pi,p}^{sd,c} + y_{\pi,b}^{sd,c}) \leq 1. \quad (11)$$

Link capacity: for all $\ell \in L$,

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} \sum_{\pi \in \Pi_{sd}^c} D_{sd}^c \delta_\ell^\pi (y_{\pi,p}^{sd,c} + y_{\pi,b}^{sd,c}) \leq \text{CAP}_\ell. \quad (12)$$

Node capacity: for all $v \in V^{\text{VNF}}$,

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} \sum_{f \in F_c} \sum_{\pi \in \Pi_{sd}^c} \Delta_f D_{sd}^c a_{v,\pi}^f (y_{\pi,p}^{sd,c} + y_{\pi,b}^{sd,c}) \leq \text{CAP}_v. \quad (13)$$

The role of the pricing problem is to generate a valid *Service Path* for a given request. Once again, the formulation uses the layer graph (G^L) introduced in Section ???. We denote by $u^{(j)}$ the vector of dual variables of constraints (j) in the RMP. Note that these values are given as input to the pricing problem in the column generation solution process.

Variables:

- $a_v^i \in \{0, 1\}$, where $a_{v,f_i} = 1$ if f_i^{st} is installed on node v .
- φ_ℓ^i , where $\varphi_\ell^i = 1$ if the flow forwarded on link ℓ on layer i .

For each request (v_s, v_d, c, D_{sd}^c) , we use two pricing problems to generate a primary and a backup service path. A Service Path generated by the pricing problems must respect constraints (2)-(7) of the model NFV_ILP_DP presented before. The two paths are then added to the collection of service paths Π_{sd} . The only difference between the two sub-problems for each request relies in the objective function of the Pricing Problem. The objective function of the pricing problem for a primary path can be written as follows.

$$\begin{aligned} \min \quad & D_{sd}^c \sum_{\ell \in L} \sum_{i=0}^{n_c} \varphi_\ell^i - u_{sd,p}^{(10a)} - \sum_{\ell \in L} \sum_{i=0}^{n_c} \varphi_\ell^i u_{sd}^{(11)} \\ & + \sum_{\ell \in L} u_\ell^{(12)} D_{sd}^c \sum_{i=0}^{n_c} \varphi_\ell^i + D_{sd}^c \sum_{v \in V} u_v^{(13)} \sum_{i=0}^{n_c} \Delta_f a_{v,f_i} \quad (14) \end{aligned}$$

C. Shared Protection

We now consider a shared protection scheme, also known as 1:1 protection. The capacities for the backup paths are reserved in case of a single link failure. In this case, the network resources may be shared among different failure scenarios. For each failure scenario, we guarantee that link and node resources are not exceeded.

We denote by Ω the set of all the possible failure situations. Since we are considering only single link failures, $\Omega = L \cup \emptyset$.

Model NFV_ILP_SP

In the shared protection case, the objective changes. Indeed, while in the dedicated protection case the required bandwidth depends on the length of the paths, this is no longer true here. Let $x_\ell \geq 0$ be the bandwidth requirements of link $\ell \in L$.

The objective is thus:

$$\min \sum_{\ell \in L} x_\ell \quad (15)$$

In addition to the variables introduced in the dedicated protection scheme, we now define two new kinds of variables. Their goal is to tell us, given a failure situation ω which link the backup paths use and on which node a function will be performed.

Variables:

- $z_{\ell,\omega}^{sd,c} \in \{0, 1\}$, where $z_{\ell,\omega}^{sd,c} = 1$ if the request uses link ℓ in the backup path in the failure situation ω .
- $z_{i,v,\omega}^{sd,c} \in \{0, 1\}$, where $z_{i,v,\omega}^{sd,c} = 1$ if the request uses function the i^{th} function of the chain C_{sd} on node v in the backup path in the failure situation ω .

We now describe the modified constraints, with respect to NFV_ILP_DP. Link Capacity: for all $\ell \in L$, $\omega \in \Omega$

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} D_{sd}^c \left(\sum_{i=0}^{n_c} \varphi_\ell^{sd,c,i} + z_{\ell,\omega}^{sd,c} \right) \leq x_\ell \leq \text{CAP}_\ell. \quad (16)$$

Node Capacity: for all $v \in V$, $v \in V^{\text{VNF}}$, $\omega \in \Omega$

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in C_{sd}} D_{sd}^c \sum_{i=0}^{n_c-1} \Delta_{f_i^c} (a_v^{sd,c,i} + z_{i,v,\omega}^{sd,c}) \leq \text{CAP}_v. \quad (17)$$

Model NFV_CG_SP

Following a similar idea as in [17], with each $\pi \in \Pi_{sd}$ we now represent a configuration as a service paths pairs (π_p, π_b) from s to d . The reason relies on the fact that, by using the same model as NFV_CG_DP, in order to ensure that node and link capacities are not exceeded in any failure situation $\omega \in \Omega$, we would need additional variables and constraints. This would lead to an increase in the size and consequently, to the complexity of the Reduced Master Problem. The price to pay for this choice is an increase in the complexity of the Pricing Problems, that would lead to a higher resolution time with respect to the dedicated protection case, as we will see in Section V. More specifically, while the Pricing Problem in NFV_CG_DP reduces to be a Shortest Path Problem on the layered graph G^L , in this case solving the Pricing Problem is NP-Hard [17].

Each $\pi \in \Pi_{sd}$ is associated with a binary value SWITCH_π^ω telling if in failure situation ω the primary path cannot be used (i.e., if the failure involves a link that belongs to the primary path).

Variables:

- $y_\pi^{sd,c} \geq 0$, where $y_\pi^{sd,c} = 1$ if demand from v_s to v_d for service chain c uses $\pi = \{\pi_p, \pi_b\}$ as pair of service paths.
- $x_\ell \geq 0$, is the bandwidth required on link $\ell \in L$.

Objective

$$\min \sum_{\ell \in L} x_\ell \quad (18)$$

Exactly one path pair per demand and per chain:

$$\sum_{\pi \in \Pi_{sd}^c} y_\pi^{sd,c} \geq 1. \quad (19)$$

Link capacity: for all $\ell \in L$ and failure situations ω :

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in \mathcal{C}_{sd}} \sum_{\pi \in \Pi_{sd}^c} y_\pi^{sd,c} D_{sd}^c (\delta_\ell^{\pi_p} + \text{SWITCH}_\pi^\omega \delta_\ell^{\pi_b}) \leq x_\ell \leq \text{CAP}_\ell. \quad (20)$$

Node capacity: for all $v \in V^{VNF}$ and failure situations ω ,

$$\sum_{(v_s, v_d) \in \mathcal{SD}} \sum_{c \in \mathcal{C}_{sd}} \sum_{\pi \in \Pi_{sd}^c} \sum_{f \in \mathcal{F}_c} y_\pi^{sd,c} \Delta_f D_{sd}^c (a_{v, \pi_p}^f + \text{SWITCH}_\pi^\omega a_{v, \pi_b}^f) \leq \text{CAP}_v. \quad (21)$$

In this case, the role of the Pricing Problem is to generate a pair of valid *Service Paths* for a given request. The path pair $\pi = (\pi_p, \pi_b)$ has to be link-disjoint but not node-disjoint. Given the layered graph (G^L), if one of the two paths uses link ℓ in some of the layers, the other path cannot use link ℓ in any of the layer of G^L .

We look at each iteration at the minimum cost path pair according to the dual values provided by the Restricted Master Problem. As in the dedicated protection scheme, the pricing problem is expressed as an ILP and solved independently for each demand and chain.

Variables:

- $a_{v,p}^i, a_{v,b}^i \in \{0, 1\}$, where $a_{v,p}^i = 1$ ($a_{v,b}^i = 1$) if f_{i+1}^c is installed on node v in the primary (backup) path.
- $\varphi_{\ell,p}^i, \varphi_{\ell,b}^i \in \{0, 1\}$, where $\varphi_{\ell,p}^i = 1$ ($\varphi_{\ell,b}^i = 1$) if the flow is forwarded on link ℓ on layer i in the primary (backup) path.
- $\gamma_{\ell,\omega} \in \{0, 1\}$, where $\gamma_{\ell,\omega} = 1$ if the primary path needs to switch to the backup path in the failure situation ω and the backup path uses link ℓ .

$$\begin{aligned} \min \quad & -u^{(19)} + \sum_{\ell \in L} \sum_{\omega \in \Omega} u_{\ell,\omega}^{(20)} \left(\sum_{i=0}^{n_c-1} \varphi_{\ell,p}^i + \gamma_{\ell,\omega} \right) \\ & + D_{sd}^c \sum_{v \in V} u_v^{(21)} \sum_{i=0}^{n_c-1} \Delta_f (a_{v,p}^i + a_{v,b}^i) \end{aligned} \quad (22)$$

V. EXPERIMENTAL STUDY

In this section, we evaluate the performance of the four proposed models. We compare the time performance of the ILP models with their respective decomposition models. Moreover, we evaluate the trade-off between an efficient allocation of primary paths bandwidth and the total amount of required bandwidth needed to guarantee the protection.

Service Chain	Chained VNFs	% traffic
Web Service	NAT-FW-TM-WOC-IDPS	18.2%
VoIP	NAT-FW-TM-FW-NAT	11.8%
Video Streaming	NAT-FW-TM-VOC-IDPS	69.9%
Online Gaming	NAT-FW-VOC-WOC-IDPS	0.1%

TABLE I: Service Chain Requirements [18]

Data Sets. We conduct experiments on three network topologies from SNDlib [19]: `pdh` (11 nodes, 34 links), `geant` (22 nodes, 36 links) and `germany50` (50 nodes, 88 links). The number of requests varies according to the network size. All experiments are run on an Intel Xeon E5520 with 24GB of RAM. We consider 200 requests for `pdh`, 400 for `geant`, and 1000 for `germany50`. Network load is the same for all the networks and is set to 1TB of data. The network traffic is divided into four common categories of traffic: Web Services, VoIP, Video Streaming and Online Gaming. Each traffic category is associated with a service function chain of five network functions. The traffic loads and the associated chains are given in Table I.

For each network, we limit the nodes able to host VNFs. We consider different numbers and study the impact of the design choice on the delay and the required bandwidth. Nodes able to host VNFs are chosen according to their *betweenness centrality*, defined as the number of paths going through the node when considering the shortest paths between all pairs of nodes. It measures the relative importance of a node in a graph.

Compact ILPs vs. CG Models. In Figure 1, we compare the compact ILPs vs. the CG model for both dedicated and shared protection, on the `pdh` network. All nodes are assumed VNF enabled, and we consider an increasing number of demands from 4 to 200. The figure demonstrates the limits regarding the computing time of a compact ILP model. For 60 demands, the time needed to find an exact solution with the compact ILP model exceeds 30 min for shared protection and 25 min for dedicated protection. Hence, the compact ILP models are not suitable for large instances due to their limited scalability. On the other hand, these results indicate that the decomposition models are fairly efficient. For 50 demands, 3 min are enough for both protection schemes. With larger values, the difference between the two decomposition models can be clearly seen. Indeed, for 200 demands CG_DP requires 11 min, while CG_SP takes 20 min, almost twice the time. This is due to the fact that the models for shared protection are more complex. Indeed, all failure scenarios have to be considered in the model, in order to share the backup bandwidth when possible.

Performance of CG Models. Table II summarizes the results of the decomposition models for dedicated and shared protection. We present the results for 3 different values of the number of VNF enabled nodes. Each traffic instance corresponds to 1 TB of traffic.

We provide the number of generated columns, the value of the ILP solution (\tilde{z}_{ILP}) and the accuracy ε , defined as the ratio $(\tilde{z}_{\text{ILP}} - z_{\text{LP}})/z_{\text{LP}}$. For most of the instances, the value of the ILP solution coincides with the value of the linear relaxation (z_{LP}). In any case, the solution accuracy never exceeds 4%. The number of generated columns is similar in the two models. However, there is a fundamental difference in terms of complexity. We recall that, in the CG_DP model, a column cor-

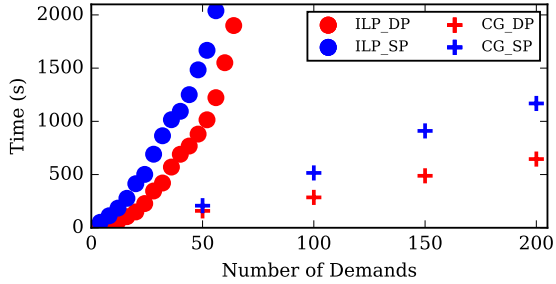


Fig. 1: Execution time of the 4 models on the pdh network.

Network	# traffic requests	# VNF nodes	# generated columns		$\bar{\epsilon}_{ILP}$		ϵ	
			CG_DP	CG_SP	CG_DP	CG_SP	CG_DP	CG_SP
pdh	200	2	501	790	4,400	3,030.17	0	1.6×10^{-2}
		3	438	778	4,080	2,694.11	0	2.1×10^{-2}
		4	413	751	3,680	2,328.67	0	3.2×10^{-2}
geant	400	3	1,185	1,016	7,190	5,141.88	0	3.8×10^{-5}
		5	1,094	1,040	6,650	4,844.69	0	8.2×10^{-4}
		7	1,076	1,034	6,390	4,651.37	0	8.6×10^{-4}
germany50	1000	5	3,654	2,701	9,270	7,253.11	0	2.4×10^{-4}
		10	3,338	2,771	9,188	6,563.46	0	4.6×10^{-3}
		15	3,165	2,674	8,800	6,198.77	0	1.7×10^{-6}

TABLE II: Numerical results for CG_DP and CG_SP

responds to a service path, and that, in the master problem, we look for a pair of service paths for each demand. Conversely, in the CG_SP model, a column corresponds to a pair of service paths and then, for each demand, only one column is selected. Hence, the problem is very hard in the shared protection case with respect to the dedicated one.

Network	DP	SP
pdh	2	1.26
geant	2	1.28
Germany	2	1.38

TABLE III: Average ratio between the processing requirements of dedicated and shared protection over the processing requirements without protection.

Bandwidth and Processing Requirements. As expected, there is a relationship between the number of VNF nodes and the bandwidth needed for both the primary paths and the global protection. Indeed, a larger number of VNF nodes allows more flexibility to find shorter paths.

VNFs nodes are expensive for both purchase and maintenance (e.g., hardware, software licenses, energy consumption, and maintenance). Thus, it is necessary to find the right trade-off between bandwidth and number of VNF nodes. For example, in pdh, using two nodes instead of four leads to an increase in the total required bandwidth of about 20% in the dedicated protection case and 30% in the shared protection case. Similar results are observed for the other networks. Even if the solution computed by the two protection schemes in terms of bandwidth used by primary paths is almost the same, there is a noticeable difference in terms of total bandwidth requirements. CG_DP requires on average about 40% more bandwidth than CG_SP. Similar results are found for the processing requirements. About 60% more processing units are required by the dedicated protection scheme than the shared one.

In Figure 2, we show the bandwidth requirements for the 3

networks without any protection strategy compared with the bandwidth requirements of the dedicated and shared protection schemes. In the case of Dedicated Protection, we may need up to 3 times more bandwidth than the one needed if we do not consider protection. In the case of Shared Protection, the price to pay is less than twice. Note that we put a limit to the latency of the paths in order not to violate the SLA requirements. Hence, we expect the savings opportunities of the shared protection to be even larger in the general case.

Delay. In Figure 3 we show the delay of the primary and backup paths for the three networks in the case of dedicated and shared protection. In order to compute the link delays, we used the distances given by the geographical coordinates provided in SNDlib.

The delays of the primary paths tend to be close between the two different protection schemes with a maximum delay of 7.2, 9 and 18 ms for pdh, geant, and germany50 respectively. The delay distributions of the primary paths slightly change when varying the number of allowed VNF nodes and tend to be homogeneous among them.

However, this is not true for the backup paths. In the dedicated protection case, paths are interested in using shorter paths in order to minimize the bandwidth requirements. In the shared protection case, this is not true. In fact, backup paths may find convenient to increase their lengths in order to share as much as possible and, thus, reduce the bandwidth requirements. This can be observed in the results. For example, the delay for a backup path in the dedicated protection case for germany50 never exceeds 20 ms while it may go up to 40 ms in the shared protection case. Hence, particular attention should be paid to paths' latencies when considering shared path protection.

VI. CONCLUSION

We provided exact methods to obtain *reliable Service Function Chains* against single-link failure. We considered two different protection schemes, dedicated and shared path protection, providing for each of them a scalable decomposition ILP model. The models are very general and can be easily extended to the case of node-disjoint paths or to deal with multiple failures. We showed the limits of the ILP based approaches and the time efficiency of the decomposition models.

We implemented and evaluated the models on 3 network topologies with different sizes, studying the bandwidth requirements for the protection, as well as their latency robustness. We also studied the trade-off between the network bandwidth requirement to guarantee the protection and the number of VNF capable nodes.

REFERENCES

- [1] P. Gill, N. Jain, and N. Nagappan, "Understanding network failures in data centers: measurement, analysis, and implications," in *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, 2011.
- [2] D. Turner, K. Levchenko, A. Snoeren, and S. Savage, "California fault lines: understanding the causes and impact of network failures," in *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4, 2010.
- [3] R. Potharaju and N. Jain, "Demystifying the dark side of the middle: a field study of middlebox failures in datacenters," in *Internet Measurement Conference*, 2013, pp. 9–22.
- [4] L. Sahasrabudde, S. Ramamurthy, and B. Mukherjee, "Fault management in IP-over-WDM networks: WDM protection versus IP restoration," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 1, pp. 21–33, 2002.

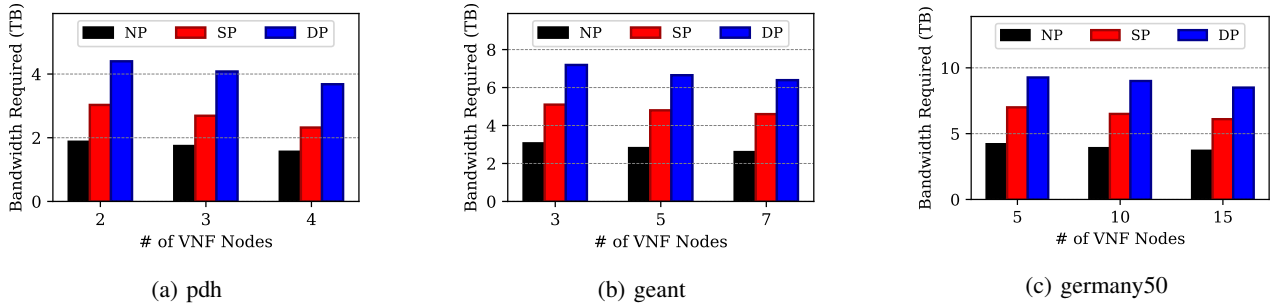


Fig. 2: Bandwidth requirements: no protection (NP) scheme vs. dedicated (DP) and shared (SP) protection schemes

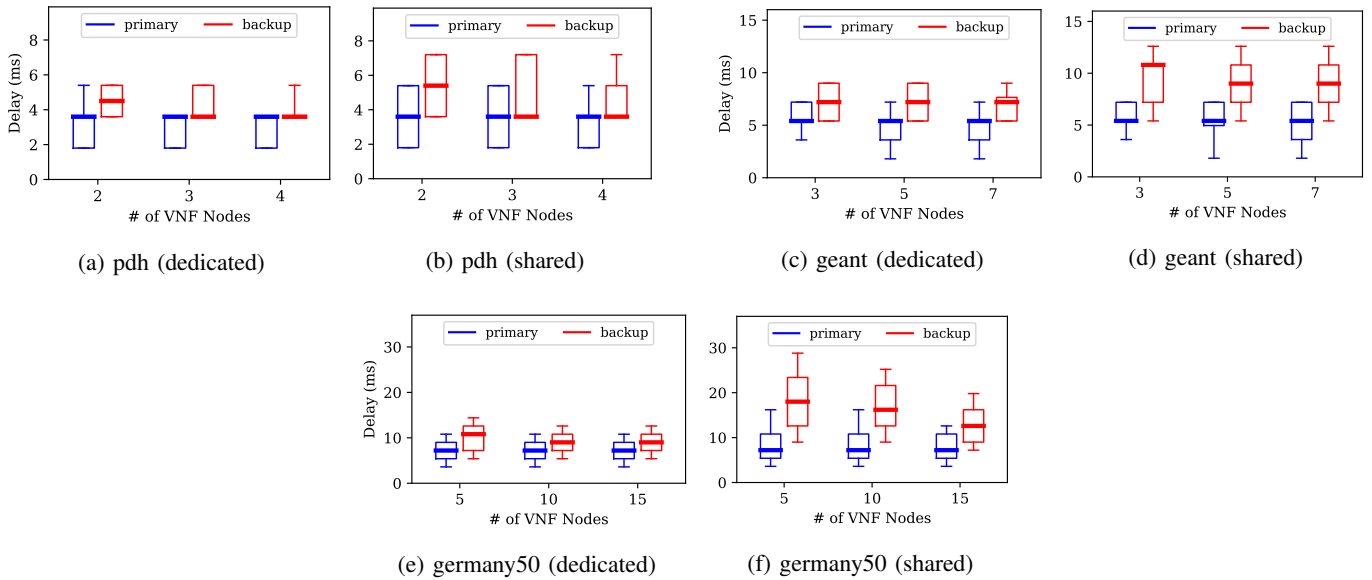


Fig. 3: Primary and backup path delay distributions under the two protection schemes vs. the number of VNF nodes with 1TB offered load. Boxes are defined by the first and third quartiles. Ends of the whiskers correspond to the first and ninth deciles.

- [5] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE network*, vol. 14, no. 6, pp. 16–23, 2000.
- [6] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks. Part I - protection," in *Annual Joint Conference of the IEEE Computer and Communications Societies - INFOCOM*, vol. 2, 1999, pp. 744–751.
- [7] R. Iraschko, M. MacGregor, and W. Grover, "Optimal capacity placement for path restoration in STM or ATM mesh-survivable networks," *IEEE/ACM Transactions on Networking*, vol. 6, no. 3, 1998.
- [8] M. Alanyali and E. Ayanoglu, "Provisioning algorithms for WDM optical networks," *IEEE/ACM Transactions On Networking*, vol. 7, no. 5, 1999.
- [9] O. Soualah, M. Mechtri, C. Ghribi, and D. Zeglache, "A link failure recovery algorithm for virtual network function chaining," in *IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 2017.
- [10] S.-I. Lee and M. M.-K. Shin, "A self-recovery scheme for service function chaining," in *International Conference on Information and Communication Technology Convergence (ICTC)*, 2015, pp. 108–112.
- [11] A. Hmaity, M. Savi, F. Musumeci, M. Tornatore, and A. Pattavina, "Protection strategies for virtual network functions placement and service chains provisioning," *Networks*, pp. 1–15, 2017.
- [12] 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 Network*, vol. 30, no. 3, 2016.
- [13] M. T. Beck, J. F. Botero, and K. Samelin, "Resilient allocation of service function chains," in *Network Function Virtualization and Software Defined Networks (NFV-SDN), IEEE Conference on*. IEEE, 2016.
- [14] N. Huin, B. Jaumard, and F. Giroire, "Optimization of Network Service Chain Provisioning," in *IEEE International Conference on Communications 2017*, Paris, France, May 2017.
- [15] N. Huin, A. Tomassilli, F. Giroire, and B. Jaumard, "Energy-efficient service function chain provisioning," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 10, no. 2, 2018.
- [16] Y. Agarwal and P. Venkateshan, "Survivable network design with shared-protection routing," *European Journal of Operational Research*, vol. 238, no. 3, pp. 836–845, 2014.
- [17] T. Stidsen, B. Petersen, K. B. Rasmussen, S. Spoorendonk, M. Zachariassen, F. Rambach, and M. Kiese, "Optimal routing with single backup path protection," in *International Network Optimization Conference (INOC)*, 2007.
- [18] M. Savi, M. Tornatore, and G. Verticale, "Impact of processing costs on service chain placement in network functions virtualization," in *Network Function Virtualization and Software Defined Network Conference (NFV-SDN)*, 2015, pp. 191–197.
- [19] S. Orłowski, R. Wessäly, M. Pióro, and A. Tomaszewski, "Sndlib 1.0survivable network design library," *Networks*, vol. 55, no. 3, 2010.