Hybrid Backup Resource Optimization for VNF Placement Over Optical Transport Networks

João Pedro[1,2(✉)] and António Eira[1]

[1] Infinera Unipessoal Lda, Rua da Garagem 1, 2790-078 Carnaxide, Portugal
{JPedro,AEira}@Infinera.com
[2] Instituto de Telecomunicações, Instituto Superior Técnico, Avenida Rovisco Pais 1, 1049-001 Lisbon, Portugal

Abstract. The concept of edge computing is vital in the 5G ecosystem, as a means of introducing application awareness in the network and enabling constructs such as slicing to be effectively implemented. In this scope, an efficient infrastructure dimensioning requires visibility of both network and data-center resources. While this joint optimization is becoming increasingly common even at the optical layer, some aspects of the dimensioning remain siloed between the network/IT worlds. Survivability mechanisms are one such example, where protection for lightpaths and/or virtual network functions (VNFs) is typically provisioned independently, potentially incurring in resource overprovisioning. This paper investigates the merits of exploiting a hybrid strategy where backup resources are selectively distributed between the IT and optical layers in metro ring scenarios, according to specific service requirements such as latency and bandwidth. Critically, this analysis incorporates, through an integer linear programming (ILP) model, the effect of optical path performance on the cost efficiency of protection mechanisms, which is shown to greatly influence the optimal resource distribution in each deployment scenario.

Keywords: Routing and spectrum assignment · VNF placement · Network survivability

1 Introduction

The requirements of 5G services are transforming the way transport networks are architected. While capacity remains a key driver in their development, many emerging services relying e.g. on massive machine-to-machine type communications or crowdsourced video applications are introducing additional constraints with respect to service dynamicity, latency and availability. As a result, and in an effort to leverage existing central office assets and reduce bandwidth requirements towards the backbone, the computing resources required for these applications are more suitably co-located with metro aggregation nodes closer to end-users. These converged nodes mix data-center (DC) and virtualization capabilities with packet/optical transport network interfaces within the same physical location and switching infrastructure [1, 2].

Within this scope, the traditional capacity planning of transport networks becomes entangled with the dimensioning of distributed mini data-center infrastructures. The placement of specific virtual network functions (VNFs) such as firewalling, video processing, etc. defines the logical topology required on the transport network, and thus its overall bandwidth requirements [3]. Therefore, consolidation of DC nodes and optimized placement of VNFs must take into account the specific requirements introduced on the optical transport side, in order to balance the cost effectiveness of converged central office/data-center architectures.

This joint IT/optical dimensioning is being increasingly explored in the context of edge computing, by extending traditional VNF placement and virtual network embedding (VNE) problems from packet networks to the optical circuit switching domain [4, 5]. The typical modeling of application-specific service requests involves defining a service chain, consisting of an ordered set of VNFs that must be traversed by a flow, with each hop between VNFs being characterized by a required bandwidth and allowable latency, and each VNF requiring a set amount of IT resources (e.g. instantiated virtual machines and/or storage space) [6].

A less explored aspect of this problem is how reliability is ensured at the different layers. Typically, backup resources are provisioned independently at the optical/IT layers, by provisioning protection lightpaths in the transport network and/or replicating VNFs at alternative DC locations for redundancy. Adding reliability on one layer independently of the other reduces complexity, but at the expense of resource overprovisioning. In [7], the authors present the comparative benefits of adding reliability at each layer, based on latency and bandwidth requirements, as well as the prevalent type of failures (optical link or in the DC). In [8], an approach considering protection at both layers simultaneously is introduced and shown to reduce the network and computing requirements.

In this paper, we focus on the specific constraints inherent to optical line-side protection, and how they influence the optimal resiliency strategy. We present an ILP model to select, given a set of service chains, the lowest cost solution leveraging a hybrid of backup lightpaths and VNFs. The trade-offs between transponder costs, IT requirements and latency performance are also exposed through this analysis, particularly for the common scenario of metro aggregation rings. In the remainder of the paper, Sect. 2 presents the network scenario and possible resiliency options, Sect. 3 details the optimization model used, and Sect. 4 discusses the results obtained through network simulations. Section 5 concludes the paper and points towards new research directions.

2   Network Scenario

The analysis in this work is focused on wavelength division multiplexing (WDM) metro rings, where each optical node may optionally be co-located with a DC possessing storage/compute resources. As Fig. 1 illustrates, service chains are deployed between a source and destination node by instantiating the set of required VNFs at one or more DCs. The logical topology established by the optical lightpaths over the WDM network must ensure that each set of VNFs in a chain can be traversed in the desired order. The placement of the VNFs across the network thus determines both the IT requirements at each DC, as well as the logical topology that supports it.
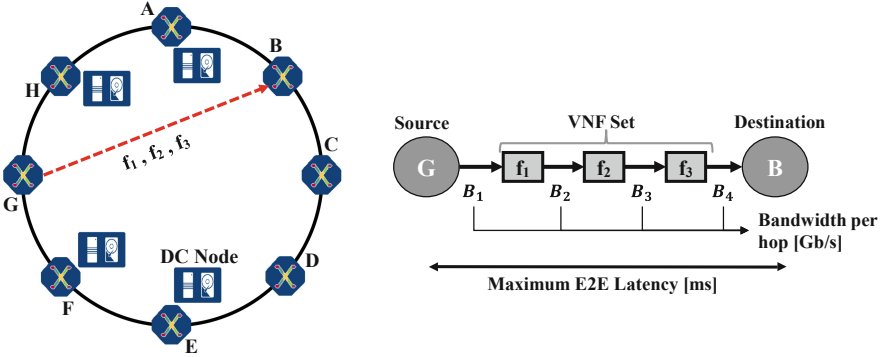
Fig. 1. Metro ring aggregation topology with co-located DCs and service chain characterization.

2.1 Survivability Mechanisms

Given the mission-critical nature of some 5G services, it is expected that reliability will be a key requirement in their deployment. Here, the focus is on guaranteeing that a chain can withstand any single optical link failure. In order to ensure this, it is possible to implement optical line-side protection, wherein a Y-cable splits the input signal to different directions, enabling the same transponder to be shared at the source/destination nodes of a lightpath. Alternatively, resilience against fiber link failures can be embedded in the application layer, by instantiating redundant VNFs at separate nodes and effectively creating a backup chain that is link-disjoint from the working one. Both options entail trade-offs with respect to the resource consumption (spectrum, transponders, storage/compute), and specific planning scenarios may favor one over the other. As such, combining both approaches in network dimensioning can potentially reduce the overall cost of providing reliable services. The following subsections detail the advantages and drawbacks of each single-layer technique, and showcases the motivation to consider a joint approach.

2.2 Hop Protection

Path protection is implemented by duplicating the signal at the source towards link-disjoint working and backup paths, while selecting the strongest signal at the receiver. This strategy overprovisions only the spectral resources needed on the backup path, since the transponders at the end-nodes are the same. In the context of a service chain, as Fig. 2(a) illustrates, every lightpath is protected by a backup. Hence, every VNF hop in a chain requiring optical connectivity is protected against a link failure. Note that sequential VNFs in a chain instantiated at a same node do not require WDM connectivity, and thus are not susceptible to fiber link failures.

One attribute of this survivability scheme is that the achievable rate of a lightpath may be different between the working/backup paths, according to the physical characteristics of each path (distance, number of spans). In a protection mechanism, the lowest common denominator between both paths must be used (i.e. the lowest bit-rate) since both paths are active at all times [9]. This is particularly critical in the case of ring topologies,
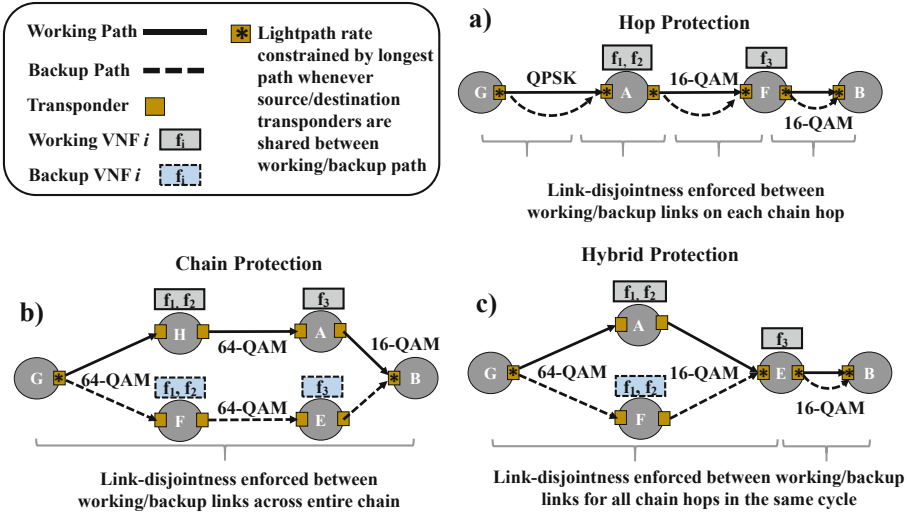
Fig. 2. Protection mechanisms for service chains: (a) Hop protection; (b) Chain protection; (c) Hybrid protection.

where the differences between working/backup paths are most extreme. The other issue affecting the performance of this scheme is end-to-end latency. For the whole chain, if a single link is used by multiple lightpaths on different VNF hops, then the backup path is triggered for each of them, further constraining the end-to-end latency budget.

2.3 Chain Protection

The main alternative to deploying protected lightpaths is to create an end-to-end alternative chain, replicating the required VNFs at different DCs, exemplified in Fig. 2(b). Ensuring survivability to link failures in this instance implies guaranteeing that the end-to-end (i.e. across the entire chain) working and backup links are disjoint, such that one chain is always available end-to-end. In a ring topology, such an approach has the benefit of ensuring the maximum latency is bounded by either the working or backup chain, regardless of which link may fail. Furthermore, the lightpaths are unprotected and hence can use the best transmission format unconstrained by servicing a protection path in simultaneous. As Fig. 2(b) shows, this may enable a higher average throughput in the deployed lightpaths. However, this strategy tends to be less efficient regarding resource provisioning. On the IT side, it requires additional storage/compute resources to duplicate all VNFs, although it provides further resilience against failures within the DCs. On the optical network side, creating an alternative chain through separate nodes may require additional transponders (if they are not already necessary for other chains).

2.4 Hybrid Protection

From the above descriptions, it is intuitive that there is a potential benefit in combining, for a single chain or a set of chains, both protection mechanisms selectively, in a way that

minimizes resource overprovisioning. Figure 2(c) provides an example of this approach, where chain protection is used in the first two VNF hops, and hop protection (backup path protection) is used in the final chain hop. As the example in Fig. 2(c) shows, applying chain protection to only a subset of all VNF hops creates a cycle between two nodes, formed by two sets of lightpaths that must be entirely link-disjoint (i.e. a single link cannot break both sub-chains simultaneously).

3 Optimization Model

The optimal protection strategy, even for a single isolated chain, depends on the combination of the service bandwidth, possible DC placements, latency constraints, etc. As the examples in Fig. 2 illustrate, particular line-side protection setups can be detrimental to bandwidth efficiency, such that the best solution is either avoiding those configurations (working/backup lightpaths with accentuated performance differences) either through optimized placement of the VNFs, or by introducing VNF redundancy at specific portions of the chain. This highly multifactorial problem structure results in a complex optimization challenge, which must address a survivable VNE problem over an optical infrastructure (i.e. solving routing and spectrum assignment on top of the VNF placement), further considering the optical performance constraints of backup path protection.

Table 1. Parameters and sets.

$Src/Dst(p)$	Source/Destination node of path p
$p \in D$	Set of dummy paths between the same nodes
$B_{s,i}$	Required bandwidth [Gb/s] of the i^{th} hop of chain s
M	Big-M value for constraint activation/de-activation
$ITcap_n$	Maximum IT resource capacity of node n
$IT_{s,i}$	Required IT resource units for the i^{th} hop of chain s
$Linkcap$	Maximum number of channels per link
$maxLat_s$	Maximum allowable end-to-end latency for chain s [ms]
Lat_p	Propagation latency in path p [ms]
$maxIT$	Maximum number of DC nodes

As an exploratory approach to evaluate the potential benefits of combining chain and hop protection, we model the problem through an ILP formulation, which enables all of the interdependencies to be considered jointly, even if limited in computational complexity to small/medium sized networks. Particularly, as outlined in Sect. 2, metro aggregation rings are an interesting case study due to their relevance in the 5G/edge-computing landscape, as well as the optical performance differential that naturally arises between working/protection lightpaths in a ring topology. The model minimizes the transponder count for a set of service chains, imposing that both a working and a backup chain must be provisioned. The paths of the backup chain determine whether redundant

Table 2. Variables.

$x_{s,i,p}/y_{s,i,p} \in [0,1]$	Binary variable equal to 1 if chain s uses working/backup path p on the i^{th} hop
$w_{s,i,p} \in \mathbb{N}^0$	Number of transponders required to carry the traffic of chain s on path p for the i^{th} hop
$b_{s,i,p}^{src}/b_{s,i,p}^{dst} \in \mathbb{N}^0$	Number of transponders required for protection at the source/destination node of path p for the i^{th} hop of chain s
$z_{s,i,n} \in [0,1]$	Binary variable equal to 1 if the working/backup paths on the i^{th} hop of chain s converge at node n
$v_{n,f} \in \mathbb{N}^0$	Number of IT resources required for VNF f at node n
$c_n \in [0,1]$	Binary variable equal to one if node n has a DC

VNFs are placed at alternative nodes, or if working lightpaths are simply path protected. As a simplifying assumption, only maximum number of channels per fiber restrictions are considered (no spectrum assignment). The model's parameters and variables are defined in Tables 1 and 2, respectively.

The ILP model can thus be formally defined as:

$$\min \sum\nolimits_{s,i,p} 2w_{s,i,p} + b_{s,i,p}^{src} + b_{s,i,p}^{dst} \tag{1}$$

subject to:

$$\sum\nolimits_{p} x_{s,i,p} = 1, \quad \forall s, i = 1 \tag{2}$$

$$\sum\nolimits_{p} y_{s,i,p} = 1, \forall s, i = 1 \tag{3}$$

$$\sum\nolimits_{p:n=Dst(p)} x_{s,i,p} = \sum\nolimits_{p:n=Src(p)} x_{s,i+1,p}, \quad \forall s, i > 1, n \tag{4}$$

$$\sum\nolimits_{p:n=Dst(p)} y_{s,i,p} = \sum\nolimits_{p:n=Src(p)} y_{s,i+1,p}, \quad \forall s, i > 1, n \tag{5}$$

$$0 \leq \sum\nolimits_{p:n=Dst(p)} x_{s,i,p} + \sum\nolimits_{p:n=Dst(p)} y_{s,i,p} - 2z_{s,i,n} \leq 1, \forall s, i, n \tag{6}$$

$$\sum\nolimits_{\substack{p \ni l}} x_{s,q,p} + \sum\nolimits_{\substack{m \in [k:i] \\ p \ni l}} y_{s,m,p} - M \sum\nolimits_{n,m \in [k:i-1]} z_{s,m,n} \leq 1,$$
$$\forall s, l, i, k \in [1:i], q \in [k:i] \tag{7}$$

$$\sum\nolimits_{\substack{p \ni l}} y_{s,q,p} + \sum\nolimits_{\substack{m \in [k:i] \\ p \ni l}} x_{s,m,p} - M \sum\nolimits_{n,m \in [k:i-1]} z_{s,m,n} \leq 1,$$
$$\forall s, l, i, k \in [1:i], q \in [k:i] \tag{8}$$

$$w_{s,i,p} \geq \frac{x_{s,i,p} B_{s,i}}{Cap_p}, \quad \forall s, i, p \tag{9}$$

$$w_{s,i,p} \geq \frac{\sum_{p'} y_{s,i,p'} B_{s,i}}{Cap_{p'}} - M\left(1 - \sum_n z_{s,i-1,n} + \sum_{p' \neq p} x_{s,i,p'}\right), \quad \forall s, i, p \quad (10)$$

$$w_{s,i,p} \geq \sum_{p'} \frac{y_{s,i,p'} B_{s,i}}{Cap_{p'}} - M\left(1 - \sum_n z_{s,i,n} + \sum_{p' \neq p} x_{s,i,p'}\right), \quad \forall s, i, p \quad (11)$$

$$\sum_{p \in D} x_{s,i,p} + \sum_n z_{s,i,n} - \sum_{p \in D} y_{s,i,p} \leq 1, \quad \forall s, i \quad (12)$$

$$b_{s,i,p}^{src} \leq M\left(1 - \sum_n z_{s,i-1,n}\right), \quad \forall s, i, p \quad (13)$$

$$b_{s,i,p}^{src} \geq \frac{y_{s,i,p} B_{s,i}}{Cap_p} - M\left(1 - \left(\sum_n z_{s,i,n} - z_{s,i-1,n}\right)\right), \quad \forall s, i, p \quad (14)$$

$$b_{s,i,p}^{src} \geq \sum_{p'} \frac{x_{s,i,p'} B_{s,i}}{Cap_{p'}} - M\left(1 - \left(\sum_n z_{s,i,n} - z_{s,i-1,n}\right) + \sum_{p' \neq p} y_{s,i,p'}\right), \quad \forall s, i, p$$
$$(15)$$

$$b_{s,i,p}^{src} \geq \frac{y_{s,i,p} B_{s,i}}{Cap_p} - M\left(\sum_n z_{s,i,n} + z_{s,i-1,n}\right), \quad \forall s, i, p \quad (16)$$

$$b_{s,i,p}^{dst} \leq M\left(1 - \sum_n z_{s,i,n} + \sum_{p' \in D} x_{s,i,p'}\right), \quad \forall s, i, p \quad (17)$$

$$b_{s,i,p}^{dst} \geq \frac{y_{s,i,p} B_{s,i}}{Cap_p} - M\left(1 - \left(\sum_n z_{s,i-1,n} - z_{s,i,n}\right)\right), \quad \forall s, i, p \quad (18)$$

$$b_{s,i,p}^{dst} \geq \sum_{p'} \frac{x_{s,i,p'} B_{s,i}}{Cap_{p'}} - M\left(1 - \left(\sum_n z_{s,i-1,n} - z_{s,i,n}\right) + \sum_{p' \neq p} y_{s,i,p'}\right), \quad \forall s, i, p$$
$$(19)$$

$$b_{s,i,p}^{dst} \geq \frac{y_{s,i,p} B_{s,i}}{Cap_p} - M\left(\sum_n z_{s,i,n} + z_{s,i-1,n}\right), \quad \forall s, i, p \quad (20)$$

$$b_{s,i,p}^{dst} \geq \frac{y_{s,i,p} B_{s,i}}{Cap_p} - M\left(2 - \sum_n z_{s,i,n} - \sum_{p' \in D} x_{s,i,p'}\right), \quad \forall s, i, p \quad (21)$$

$$\sum_{\substack{(s,i) \in f \\ p:n=Dst(p)}} \left(x_{s,i,p} + y_{s,i,p}\right) IT_{s,i} - \sum_{(s,i) \in f} z_{s,i,n} IT_{s,i} \leq v_{n,f}, \quad \forall n, f \quad (22)$$

$$\sum_f v_{n,f} \leq c_n IT cap_n, \quad \forall n \quad (23)$$

$$\sum_{s,i,p \ni l} x_{s,i,p} + y_{s,i,p} \leq Linkcap_l, \quad \forall l \quad (24)$$

$$\sum_{i,p} x_{s,i,p} Lat_p \leq maxLat_s, \forall s \quad (25)$$

$$\sum_n c_n \leq maxIT \quad (26)$$

The objective function (1) minimizes the total amount of transponders required for working and backup chains. Constraints (2–5) implement flow conservation for the first and subsequent hops in working and backup chains. Constraint (6) identifies if a cycle is closed when the working and backup paths of the same chain hop converge on the same node. Constraint (7) imposes that, for any sequence of hops forming a cycle, no

link of the backup sub-chain may overlap with a link on the working one. Constraint
(8) imposes the same condition, iterating instead over all working links. Constraints
(9–12) enforce that the lightpath bit-rate on a given chain's hop is bound by the smallest
rate achievable between the working/backup paths, whenever the source or destination
nodes are shared between the paths (i.e. $z_{s,i-1,n}$ or $z_{s,i,n} = 1$). Constraints (13–16)
impose that backup transponders are required on the source node of each hop, whenever
the source node is not shared between the working/backup chains, or the lightpath bit-
rates are different between them. Constraints (17–20) enforce the same restriction on
the destination node of each VNF hop, requiring $b_{s,i,p}^{dst}$ backup transponders whenever a
backup chain does not share the same destination node as the working one for the i^{th} hop,
or the working/backup lightpaths for that hop have different bit-rates. Constraint (21)
covers the special case where the working lightpath is a dummy (i.e. sequential VNFs
are instantiated at the same node), but a backup chain converges with the working one
coming from a different node. Constraint (22) instantiates the required VNF capacity
for each hop at the destination node of each hop's active lightpath. Constraint (23) limits
the instantiated IT resources at node n to $ITcap_n$, provided n is an active DC node.
Constraint (24) limits the number of lightpaths per link. Constraint (25) imposes that the
sum of propagation latencies for each lightpath cannot exceed the end-to-end allowable
latency of the chain. Finally, constraint (26) sets a hard limit on the number of nodes
that may have co-located DCs.

The described model can decide, for each VNF hop, if protection should be imple-
mented at the optical or application layer. The single-layer protection cases can be
obtained by simple manipulation of the $z_{s,i,n}$ variables. For chain protection, all $z_{s,i,n}$
are forced to zero, except for the chain's last hop (the working/backup chains only con-
verge at the destination node). In order to emulate the hop protection case, it must be
imposed that every cycle must close at every VNF hop on a single node:

$$\sum_n z_{s,i,n} = 1, \quad \forall s, i \tag{27}$$

Latency restrictions are enforced end-to-end across an entire chain. This applies only to
the working chain, since the actual end-to-end path of a complete chain in the event of
a link failure depends on which specific link has failed. Thus, the latency performance
of backup chains in each scenario is the object of study in the following Section.

The total number of variables in the ILP model is $5 * \left(|S| * |H| * \frac{N(N-1)}{2} * k \right) +$
$|S| * |H| * |N| + N * |F| + |N|$, where $|S|$ is the number of service chain instances,
$|H|$ is the (average) number of VNF hops per chain, $|N|$ is the number of nodes in the
network, k is the number of candidate paths per node-pair, and $|F|$ is the number of
VNFs in the scenario. Overall, the biggest complexity driver is the number of nodes,
since the variable count evolves with $O(N^2)$ due to having to model all candidate paths
between arbitrary node-pairs (for every chain hop). Note that, in the specific case of ring
network topologies analyzed here, k always equals 2.

4   Results and Discussion

The three protection mechanisms outlined in Sect. 2 were comparatively evaluated with
the ILP model. The network scenarios consisted of ring topologies with total lengths

of 200 and 400 km. For each case, 5- and 10-node rings were considered, with evenly spaced spans. In order to enforce different levels of DC consolidation, $maxIT$ was set to 40% or 80% of the total node count. The transponders are assumed to be modulation format adaptive, operating on a 75 GHz grid with BPSK, QPSK, and 8/16/32/64-QAM (between 100 and 600 Gb/s bit-rates). For each format, the reach is obtained with a performance estimation approach detailed in [10]. The service chain profiles are taken from [11]. In each simulation run, 10 Tb/s of requested traffic (summing over all VNF hops of every chain) are generated uniformly between all nodes, and each network scenario is evaluated by averaging 10 independent runs.

Figure 3 illustrates the number of transponders required in each scenario for the three protection methods. Chain protection is considerably less efficient, requiring between 23% and 94% more transponders than hop protection. This strategy is particularly inefficient when rings are shorter and there are less available DC sites. The main reason behind this is that optical performance differences between two paths around a ring are less pronounced with both smaller rings and less nodes. Therefore, the fact that chain protection requires additional transponders for backup vastly outweighs having improved average lightpath bit-rates.
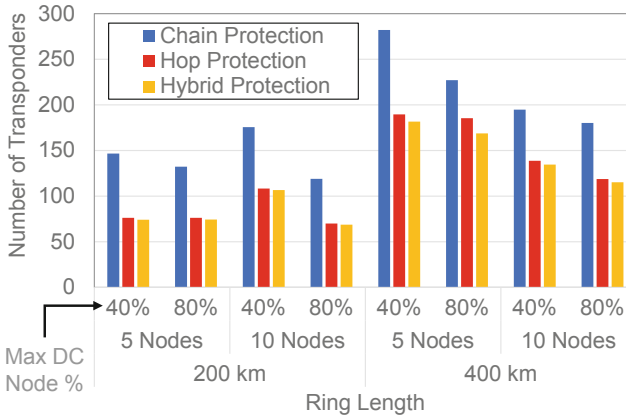

Fig. 3. Number of transponders per ring topology and protection mechanism.

Looking at the comparison between hop and hybrid protection, we find that the latter is able to improve the average transponder utilization in most scenarios. For smaller rings, this benefit is modest, standing at 3% on average for 200 km topologies. As outlined above, this is due to the unsuitability of chain protection mechanisms in these scenarios, which make hop protection the best strategy in the vast majority of chain hops. However, when considering 400 km rings, hybrid protection can save up to 9% transponders relative to hop protection. In these scenarios, protecting every lightpath has a toll on achievable throughput that makes mixing both baseline protection schemes more attractive in terms of cost efficiency.

Figure 4 shows the IT capacity requirements of each method for the same network scenarios. Naturally, chain protection requires the most resources, since it forcibly duplicates all VNFs at every node. Although it is clearly less resource efficient, it should be mentioned that it does provide an additional degree of resilience against failures within the DC. What is interesting to analyze is the comparative difference between hop and hybrid protection schemes. We find that, compared with Fig. 3, VNF capacity in the hybrid scheme is higher precisely in the cases where there was a higher benefit in saved transponders. This occurs because in such cases the model provisions a higher share of chain hops with VNF redundancy to reduce optical interfaces, at the expense of replicated VNFs. On the most extreme case (5-node ring with 400 km), the 9% saved transponders are obtained through an additional 24% IT resources provisioned.


Fig. 4. IT resource unit requirements per ring topology and protection mechanism.

The final aspect to analyze is how each resiliency mechanism affects latency. The working chain's latency is straightforwardly given by the routing paths selected by the ILP model. For the backup chain calculation, we simulate a failure on every network link, and compute the worst-case end-to-end latency for each output of the ILP model. The results are shown in Fig. 5, which displays the average working/backup chain latency in each case.

The analysis reveals that using chain protection increases the working latency by an average of 20%. However, when fiber link failures occur, the backup latency is on average 5% smaller. For working chains, VNF replication requires a higher spread of functions across the available DCs, which implies a higher average number of physical hops per chain. However, a link failure automatically forces the hop protection case to route around the ring in the opposite direction, significantly degrading latency performance, particularly in longer rings with many nodes and few DCs (where VNFs are further apart).

Fig. 5. Working chain latency and worst-case backup chain latency per ring topology and protection mechanism.

5   Conclusion and Future Work

This paper presented an ILP model to comparatively evaluate protection strategies for service chains based on provisioning backup resources exclusively at the optical layer, application layer, or both. The analysis showed that, although lightpath protection is the best option in terms of resource provisioning efficiency for the majority of VNF hops, in select cases combining this strategy with VNF replication can further reduce the overall solution cost. This is particularly true in scenario instances where there are significant optical performance differentials between working/backup paths, which can hinder overall throughput when working lightpath rates must be aligned by the backup ones. This is the case for reasonably large metro aggregation rings. The latency analysis concluded that VNF replication in ring topologies presents lower latencies in case of link failures, at the expense of additional IT resources.

Future expansions in order to further comprehend the potential of deploying such hybrid protection mechanisms should include the possibility of evaluating larger topologies, including meshed patterns. Additionally, the impact of client signal grooming on transponder utilization and end-to-end latency can also have a key effect on network efficiency. Realizing an optimization framework that can efficiently address all these joint factors is thus a challenging research prospect.

References

1. Peterson, L., et al.: Central office re-architected as a data center. IEEE Commun. Mag. **54**(10), 96–101 (2016)

2. Yan, Y., Shu, Y., Saridis, G., Rofoee, B., Zervas, G., Simeonidou, D.: FPGA-based optical programmable switch and interface card for disaggregated OPS/OCS data centre networks. In: 41st European Conference on Optical Communication (ECOC), pp. 361–363. IEEE, Valencia (2015)
3. Li, F., Sun, W., Yue, S., Hu, W.: Trading storage for bandwidth – a simulation study of optical circuit switching with massive storage at network edge. In: 19th International Conference on Transparent Optical Networks (ICTON), paper We.D3.5. IEEE, Girona (2017)
4. Xia, M., Shirazipour, M., Zhang, Y., Green, H., Takacs, A.: Network function placement for NFV chaining in packet/optical datacenters. IEEE/OSA J. Lightwave Technol. **33**(8), 1565–1570 (2015)
5. Fang, W., Zeng, M., Liu, X., Lu, W., Zhu, Z.: Joint spectrum and IT resource allocation for efficient VNF service chaining in inter-datacenter elastic optical networks. IEEE Commun. Lett. **20**(8), 1539–1542 (2016)
6. Herrera, J., Botero, J.: Resource allocation in NFV: a comprehensive survey. IEEE Trans. Netw. Serv. Manage. **13**(3), 518–532 (2016)
7. Hmaity, A., Savi, M., Musumeci, F., Tornatore, M., Pattavina, A.: Virtual network function placement for resilient service chain provisioning. In: 8th International Workshop on Resilient Networks Design and Modeling (RNDM), pp. 245–252. IEEE, Halmstad (2016)
8. Kong, J., et al.: Guaranteed-availability network function virtualization with network protection and VNF replication. In: IEEE Global Communications Conference (GLOBECOM), pp. 1–6. IEEE, Singapore (2017)
9. Hai, D., Morvan, M., Gravey, P.: An efficient network-side path protection scheme in OFDM-based elastic optical networks. Int. J. Commun. Syst. **31**(1), paper e3410 (2017)
10. Eira, A., Costa, N., Pedro, J.: On the capacity and scalability of metro transport architectures for ubiquitous service delivery. In: 20th International Conference on Transparent Optical Networks (ICTON), paper Mo.D3.5. IEEE, Bucharest (2018)
11. Savi, M., Hmaity, A., Verticale, G., Höst, S., Tornatore, M.: To distribute or not to distribute? Impact of latency on virtual network function distribution at the edge of FMC networks. In: 18th International Conference on Transparent Optical Networks (ICTON), paper We.C3.4. IEEE, Trento (2016)