Modeling and Provisioning Highly Available NFV Services

Besmir Tola (PhD Candidate) and Yuming Jiang (Supervisor)
Department of Information Security and Communication Technology, NTNU, Norway
Email:{besmir.tola, yuming.jiang}@ntnu.no

***Abstract*—The ubiquitous deployment of middleboxes can hamper the network capability to be flexible, scalable and innovative to an extent that new and specialized services cannot be easily introduced in the network. Network Function Virtualization (NFV) promises to overcome these limitations by providing the ability to execute virtual instances of networking functions on top of a common physical network substrate. It moves data processing tasks from proprietary hardware middleboxes to virtualized entities that can run on commodity hardware. Together with Service Function Chaining, it enables the replacement of traditional network hardware appliances with softwarized Virtualized Network Function (VNF) chains. However, this major transformation brings new challenges and a crucial one is the ability to ensure the high-availability demands of carrier-grade services provided by NFV-enabled networks. This challenge is further exacerbated by the extreme availability levels that 5G use cases demand, e.g., ultra-reliable services. This work tackles the challenge by addressing the problem of how to assess and quantify the availability of NFV-supported services, and how to provision highly available services by means of fault-tolerant mechanisms that are both effective and resource efficient.**

***Index Terms*—NFV, Availability, Redundancy, Modeling, Resource allocation.**

I. Introduction & Motivation

Modern communication networks include a multitude of network functions, alias middleboxes, and they have become an integral part of network infrastructures [1]. They are typically expensive both in terms of investment and operation and usually are closed systems with little or no possibility to enable innovation [2]. Specialized network services may require traffic flows to go through a chain of network functions like a firewall, an IDS, and finally through a proxy. This mechanism is referred to as service function chaining (SFC), and traffic flow routes have to be manually set up for the desired sequence of middleboxes, which is not efficient for large infrastructures. Additionally, middleboxes are deployed in fixed positions which limit traffic routing paths from an efficient utilization of network resources, hence making the middleboxes potential bottlenecks.

A fast-emerging and prominent solution that promises to alleviate these limitations is Network Function Virtualization (NFV) [3]. It enables the decoupling of the network appliance software from purpose-built hardware and runs it in virtualized environments, which can be deployed on a range of industry standard server hardware. This way, virtualized network functions (VNFs) can offer many benefits such as reduced equipment cost, through consolidation and exploitation of server hardware, and introduce greater flexibility in deploying and operating network functions [3]. VNFs can be deployed anywhere on the network and an operator can optimize their location so that network resources are efficiently utilized. However, the "softwarization" of hardware-specific middleboxes poses several challenges and service availability represents a major concern that can undermine the success of NFV [3]–[5]. It is crucial that service providers adopting NFV can guarantee at least the same level of availability compared to traditional specialized hardware-based appliances, which have, through years of development, grown to mature and robust technologies satisfying 5'9s availability.

Availability is a critical design factor in NFV due to various concerns including: *(i)* the replacement of traditional purpose-built hardware enriched with built-in fault management mechanisms by off-the-shelf hardware whose failure intensities are potentially highers [4], *(ii)* software code developed for implementing VNFs may be less robust and more error-prone [4], *(iii)* the use of virtualization layers comes at the cost of increased system dynamics caused by the lack of direct control over the underlying physical hardware [6], *(iv)* any eventual abnormal execution of applications, e.g., resource overload, may lead to availability issues for third party services due to the sharing of a common physical infrastructure [5]. In addition, also the level of availability expectation of the imminent 5G cellular system, for which NFV represents an essential enabling technology [7], envisions highly demanding usage scenarios such as Ultra Reliable and Low Latency Communications (URLLC) that require beyond 5-nines availability.

Given the importance of NFV robustness, ETSI has provided guidelines with regard to availability requirements, models, and capabilities for end-to-end NFV-enabled services [6]. However, the included models, and their estimations, are derived from simple and basic models, which fail to capture the failure and recovery process dynamics, and the inter-dependencies between the different components involved in the end-to-end service delivery such as VNFs, virtualization layers, compute, storage, and internetworking elements (e.g., routers, links, switches). Consequently, it becomes important to evaluate and quantify the availability of NFV-enabled services through more realistic models that are able to capture the system behavior and include all the involved service elements. Assessing availability attributes will help identify critical elements within the NFV architecture and provide useful feedback on how to deploy, operate, and manage the network infrastructure for providing highly available services.

The basic principle that helps systems achieve high availability in the presence of faults, also known as fault-tolerance, relies on employing redundant resources to cope with failures of those providing the actual service. To achieve this, an operator needs also to plan for availability by orchestrating NFV resources such that the allocation of redundant resources provides effective protection against failures, service availability demands are fulfilled, network resources are efficiently utilized, and business profit is maximized. The NFV redundancy allocation is a challenging problem that involves a set of decisions on the *placement* of redundant instances running network functions and the *instance assignment* that determines the traffic routing of flows through specific functions requested by the service [8]. Both problems can be considered as extended versions of two $\mathcal{NP}$-hard problems [9]: the virtual network embedding and the location-based routing problems. In addition, redundancy can be costly, especially when high availability levels are demanded [9], and unless planned carefully it may significantly limit the network resource efficiency. Therefore, smart resource allocation decisions are necessary for optimizing the benefits of NFV.

Accordingly, the overarching theme of this work is to propose methods and tools to abstract, estimate, and analyze availability of end-to-end NFV-driven services aiming at identifying availability flaws, effective redundant mechanisms, and critical system elements that pose threats to service resilience. Furthermore, the research focus is further extended on the orchestration of redundant NFV resources in an efficient and scalable way such that the provisioning of highly available services can be achieved.

II. Related Work and Open Challenges

In this section, we briefly present the closest related work on availability modeling and provisioning of NFV services.

A. Availability Modeling of NFV-based Services

Gonzalez *et. al* [10] propose a Stochastic Activity Networks (SAN) model for assessing the steady-state availability of a virtual Evolved Packet Core as a composition of elements. The work analyzes the system availability but they assume the same SAN model for the different VNF submodels, despite this not being a realistic case. In addition, their findings are related to the delivery of network functions and as the authors state, the availability of end users' services needs to integrate also the data center network topology.

A two-level hierarchical availability model of a network service in NFV architectures proposed in [11] aggregates reliability block diagrams (RBDs) on the higher level and stochastic reward nets on the lower level. Leveraging this model, the authors evaluate the steady-state availability and perform a sensitivity analysis to determine the most critical parameters influencing the service availability. Using the same approach, the authors model and assess the availability of an NFV-oriented IP multimedia subsystem in [12]. In [13], the same authors exploit universal generating functions to evaluate system availability of a virtualized SFC. The work considers a multi-tenant SFC where the VNFs that compose the SFC are shared among multiple tenants. In addition, a sensitivity analysis investigates the range of failure and repair nominal variations that the best configuration can still handle.

A hierarchical model based on stochastic petri nets and RDBs is proposed for the availability analysis of a generic SFC in [14]. The model incorporates software rejuvenation mechanisms and VM live migration services. However, the analysis is limited to the availability evaluation of the system and lacks insights into the failure and repair dynamics of single system components and their impact on the system availability.

B. Availability (Reliability)-aware Resource Allocation

Fan *et al.* present a heuristic algorithm [15] and an optimized Integer Linear Programming (ILP) model [16] aiming at minimizing the employed physical resources for hosting network functions while satisfying reliability requirements. In a later work [17], they propose a framework for minimizing resource usage while providing SFC availability demands. Nevertheless, only VM failures are considered while hosting node and link availabilities are ignored. Moreover, they adopt only on-site redundancy and disregard the impact that correlated failures may have in a network.

Reference [9] constructs three different ILP models for the VNF placement and service chaining with protection against single node/link, single link, and single node failures with the objective of balancing link bandwidths. The evaluation shows that providing protection against the considered scenarios comes with at least twice the amount of resources being deployed into the network. However, the models only place the VNFs without verifying that availability requirements are met with the assignment of the backup chains.

In [18], the reliability-aware service chaining problem is formulated through an ILP, dubbed REACH. The idea of the algorithm is similar to the one proposed in [16], where repeatedly, the least available VNF of a service chain is provided with backup until the chain reliability demand is satisfied. The same authors propose an ILP model and customized heuristics, with the objective of guaranteeing carrier-grade reliability while taking into consideration the sharing of adjacent VNFs in [19]. The sharing approach is similar to the one adopted in [15] where for the purpose of protecting two adjacent instances the maximum amount of resources among them is allocated in a single node. Nonetheless, the achieved service reliability is calculated considering only physical node reliability and does not regard VNF instances and network elements such as forwarding nodes and links.

The work in [20] employs a cost-aware importance measure to select the set of VNFs that require backups. On similar lines, the work in [21] proposes an algorithm for reliability-guaranteed VNF redundancy allocation that is based on a criticality importance measure (CIM). The scheme exploits the measure for finding the best suited VNFs to protect and factors in the computational costs so that the output results in a cost-efficient and reliability-guaranteed placement. In [22], VNF redundancy is allocated with the objective of reducing resource

consumption while assuming heterogeneous VNF resource requirements. However, similar to [18], also the investigations in [20] and [22] are limited to three-nines SFC availability requests, which is far below the high expectations that carrier-grade services have.

A Mixed Integer Programming (MIP) for ensuring high availability when placing VNFs is formulated in [23]. The developed MIP model works only for small problem instances; hence the authors propose two heuristic approaches; a solution based on bin-packing and a meta-heuristic consisting in a variable neighborhood search. An extension of the work is carried out in [8] by proposing a flexible VNF placement, which enables VNF protection only if needed, and the redundant VNF can also be shared among multiple active instances. Although the authors highlight a possible extension of the shared protection to cope with correlated failures, they do not investigate their relative impact. In addition, both works perform only VNF placement and do not address the assignment problem, i.e., service chain composition.

A multi-tenancy approach is proposed in [24]. It allows the sharing of backup resources by multiple service requests and the analysis shows that it outperforms the single-tenancy approaches. Nonetheless, this approach is constrained by the placement of backup chains onto the same computing node, thus limiting the resource efficiency as backup chains are prevented from utilizing different hosting nodes and, in addition, the same node represents a single point of failure for chain(s).

C. Open Challenges

From the review of the state of the art we can identify several open challenges regarding NFV service availability modeling and redundancy allocation as follows.

1) Network-aware Modeling including Element Inter-dependencies: In [6], ETSI emphasizes that a correct availability evaluation should incorporate all the service elements and components involved in the end-to-end delivery. None of the related works have performed a comprehensive assessment of end-to-end NFV service availability since they lack key service elements such as physical network links and forwarding/routing devices, which are essential networking elements inter-connecting the VNFs that compose a service chain [6], [10]. Henceforth, integrating all the network components in the availability model remains a fundamental endeavor for a detailed end-to-end service availability assessment. Moreover, the inter-dependencies of these elements need to be fine-grained/reflected. To illustrate, a crash of the operating system (OS) brings down the VNF software yet, the reboot of the OS is not enough for the system to be considered operational unless the VNF software is restarted too. In addition, the failure of the OS should not influence the status of the underlying hardware since the latter may fail independently on whether the OS is running or not. Such inter-dependencies are often omitted in the models that the state-of-art presents.

2) NFV-MANO System Availability Assessment: The NFV Management and Orchestration, briefly MANO, is a logically-centralized entity that maintains a global view of the network and is responsible for the correct orchestration and management of end-to-end services. Any misoperation or logical/physical faults in the MANO components may jeopardize and severely impact the provisioning of network services [25]. Despite the importance and criticality that the MANO system has on the service continuity [26], most of the related works focus only on the data plane availability. The only related work that models the availability of the MANO [10] relies on a very simple model and lacks a thorough analysis of MANO availability and the factors that mostly impact it. Consequently, it is pivotal to investigate and identify these factors through a proper evaluation of MANO availability using more realistic models derived from standardized implementations.

3) Network Topology Dependencies Impact: In [6] ETSI recommends anti-affinity placement policies, i.e., deployment of primary and backup VNFs into separate computing nodes, so that they will not experience a simultaneous failure. However, this may not be enough because the operator needs to ensure that a failure of a primary VNF will not impact its respective backup VNFs (and vice versa), or both should not be subject to a common failure mode. For example, in a data center network, a top-of-the-rack switch failure will impact the connectivity of all the servers placed on the rack. In particular, for cases where high-availability levels are demanded, which in turn may require more than one backup instance, the failure of a primary resource should not impact any of the relative backup instances. Several of the related works, see for example [15], [18], [20], [22], consider the deployment of primary and backup VNFs into different nodes, yet they do not check any topological dependencies that may affect both primary and backup chains. Therefore, unless carefully designed, such correlation may undermine the benefits of redundancy.

4) Efficient Resource Utilization: If not accurately planned, the number of required backup instances may grow up to an unsustainable level, i.e., more than twice the required resources for primary allocation. This can be further exacerbated when high-availability levels are required [16]. The works in [15], [19], [24] employ a mechanism that protects two adjacent primary VNFs of the same service chain by allocating a backup VNF that is shared among the two. Yet, all approaches perform backup sharing of one VNF instance with dedicated capacity being the backup capacity equal to the sum of the capacities of the respective primary resources. The *resource overbuild*, as one important figure of merit for resource efficiency [27], which expresses the amount of extra resources needed for providing protection as a percentage of the required amount without protection, would equal 100% in case only one backup instance is required. Therefore, designing and adopting approaches that achieve lower resource overbuild would provide significant benefits to network operators both in terms of employed resources and increased capacity in accommodating new requests.

III. Approaches & Contributions

Our work and contributions consist of a set of approaches aiming at addressing the above-mentioned gaps. We propose

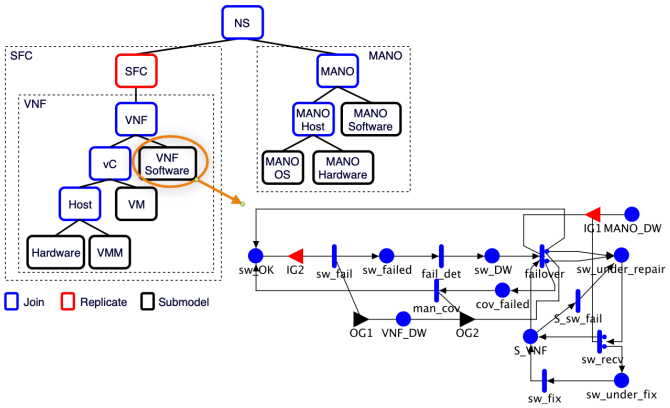

Fig. 1. Replicate and Join composition model of a virtualized network service (NS) and the submodel of the VNF software.

comprehensive availability models and redundancy allocation strategies that integrate network topological structure correlation and resource-efficient features. Details of each contribution are discussed in the following subsections.

A. Network-aware Availability Modeling including Element Inter-dependencies

Availability modeling techniques fall into the two main categories, static and dynamic models, also referred to as non-state space and state-space models, respectively [28]. Among them, state-space models offer a more powerful modeling capability for capturing system dynamics, component dependencies, and common mode failures [29]. Due to this, we exploit stochastic activity networks, Petri-net based dynamic models, to abstract system elements involved in end-to-end network services. Specifically, in [30] we propose a hierarchically-composed model for NFV-enabled service function chains. The SFC model features various failure modes including hardware, operating system, hypervisor, VNF and MANO software, and each component's failure and repair dynamics are abstracted through individual SAN models. The relationship and dependencies of the components are modeled through a state-sharing Replicate and Join formalism implemented in the Möbius software tool [31]. Figure 1 illustrates the hierarchical model and depicts an example of a SAN model abstracting the VNF software dynamics. More details on the model design, the used formalism, and numerical evaluation can be found in [30].

Second, despite the importance and criticality of the MANO in orchestrating the service and in particular on fault management [26], most of related works focus only on the data plane availability. Henceforth, it becomes essential to evaluate and assess the factors that can mostly impact the MANO robustness. To this end, [32] represents another contribution of this work which proposes a SAN-based availability model of a cloud-native MANO implementation, i.e., containerized, subject to different deployment options and fault-tolerant mechanisms that are inspired by standard Docker guidelines. The model is further extended in [33] with the integration of software rejuvenation mechanisms, as proactive maintenance to cope with the software aging phenomenon. Experimental tri-

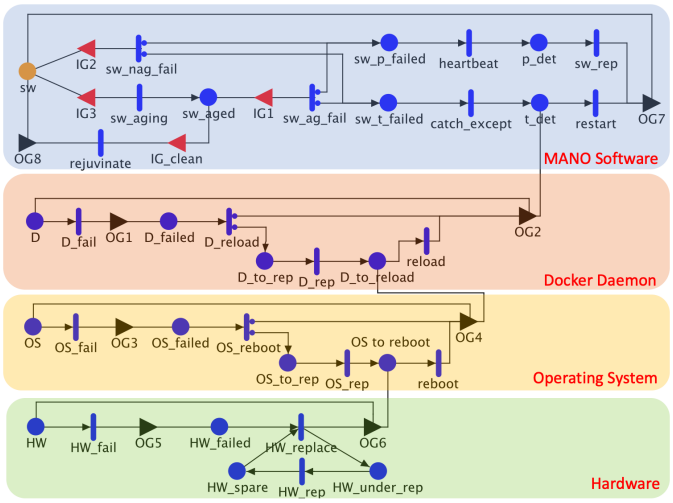

Fig. 2. SAN availability model of the MANO with software rejuvenation.

als on a real MANO testbed helped retrieve empirical recovery parameters regarding a number of system elements. Moreover, extending [32], the work proposes two additional models for portraying distributed MANO deployments that feature redundancy configurations on both hardware and software levels. Figure 2 depicts the MANO model of the basic deployment option. Further details on the model characterization and the additional deployments can be found in [33].

Third, an end-to-end NFV-enabled service, where both ends are customers, is a composition of several functional blocks, which can be connected in series or parallel to construct a network service chain [6]. These functional blocks include not only the VNFs and the supporting infrastructure (e.g., virtualized hardware) but also networking and inter-working equipment. This is particularly important since a key characteristic of NFV is the ability to flexibly and dynamically deploy VNFs anywhere in the network, and an operator can interconnect them for realizing specialized services. However, from an availability perspective, this flexibility imposes connectivity requirements among elements since the service will be available if and only if all the elements of the requested service are available. Therefore, availability has to be estimated based on the availability of all these blocks.

To address the lack of integrating network links and forwarding devices in the availability models, a subsequent contribution of our work is given in [34]. The work proposes a two-level approach. On the low level, SAN-based dynamic models abstract each of the NFV components and interconnecting devices, such as forwarding devices and links. The high level exploits a minimal-cut set analysis to derive end-to-end service availability by including both data and control/management plane. The minimal-cut set analysis allows the definition of connectivity requirements by defining the sets of elements that are strictly required to be operational for an end-to-end service to be considered available [28]. The dynamic models characterize and quantify the steady-state availability of each of the involved elements and the end-to-end service availability is evaluated by merging the two levels.

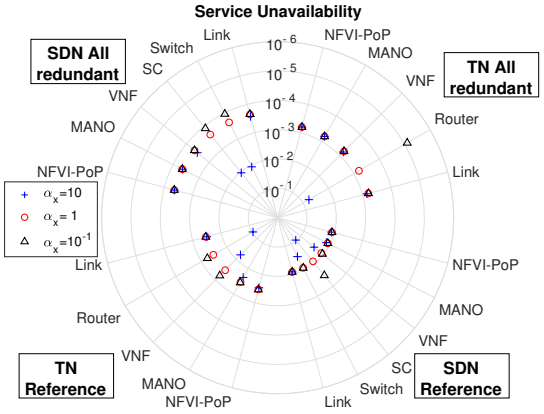

Fig. 3. Service unavailability of the traditional (TN) vs. SDN-integrated network for varying element unavailability factor α_x.

The *inclusion-exclusion principle*, which is a probabilistic technique to obtain the elements in a union of finite sets, is used to merge the two levels and define the end-to-end service unavailability as follows:

$$U_{NS} = P\left(\bigcup_{i=1}^{n} C_i\right) = \sum_{k=1}^{n}(-1)^{k-1} \sum_{0\neq I\subseteq[n],|I|=k} P\left(\bigcap_{i\in I} C_i\right) \quad (1)$$

where $C_1, C_2, ..., C_n$ are the minimal-cut sets and $P(C_i)$ is the probability of set C_i.

Knowing the unavailability of each element composing the minimal-cut sets, which are derived from the individual SAN models, allows to compute the service unavailability defined in Equation 1. Moreover, the proposed models also include the integration of SDN (software-defined networking)-enabled networks, which support network dynamic programmability through the separation of the control and data plane of network devices [7]. The presence of the SDN controller increases connectivity requirements for an available service since SDN switches need to have a working path with their controller. The work evaluates various case studies of both traditional and SDN-integrated NFV services with different level of element redundancies such as single and double protection of involved elements, i.e., VNFs, MANO, and SDN controllers.

Besides the models proposed, the works presented above perform extended investigation on the failure and repair dynamics, and carry out sensitivity analysis aiming at identifying critical failure parameters, service elements, and the best redundancy configurations for ensuring highly available services. As mere example, Fig. 3 shows service unavailability of the traditional and SDN-integrated network for the *Reference* and redundant deployments, as a function of varying by one order of magnitude the element unavailabilities, i.e., α_x. As expected, the SDN-integrated service unavailability is higher compared to the traditional deployment due to the increase of connectivity requirements, $1.2{\cdot}10^{-2}$ vs. $2.9{\cdot}10^{-3}$, respectively. Moreover, obtaining carrier-grade availability with traditional networks (TN) requires not only redundant NFV elements such as VNFs, NFVI points of presence (PoP) and MANO, but also more robust IP routers as shown in the top-right corner of the figure. Additional results can be found in [30], [32]–[34].

B. Resource-efficient Redundancy Allocation in NFV

Availability models can be used to construct, assess, and compare various settings that simulate practical fault-tolerant configurations but an operator needs to know also how to provision these redundant resources in the network such that the targeted availability values can be achieved.

The NFV redundancy allocation problem, also known as the availability (reliability)-aware VNF deployment, is an NFV resource allocation problem that defines how to provision redundant resources while satisfying availability demands of SFC requests [8], [24]. It consists of the optimal *placement* of backup VNFs into the network substrate and the optimal *assignment* of backup functions to network traffic flows that demand specific service chains, i.e., backup chain composition. In addition, depending on an operator's goals, there could be other system optimization constraints such as node capacities, link bandwidth, service delay or operating costs. This is a well-known $\mathcal{NP}$-hard problem [8], [16], [24] and quite often ad-hoc heuristics are proposed to address the limitation of problem intractability. Yet, none of the previous works have considered the impact that due to network topological dependencies the failure of one node can have on other nodes of the network , i.e., correlation of nodes as a result of the topological structure. An effective redundancy allocation must regard correlated failures so that primary and backup resources cannot be affected by the same failure event [4], [35].

Our work in [36] develops two mathematical models for allocating redundant resources with the objective of minimizing resource consumption, i.e., required number of backup instances, while satisfying, among others, flow carrier-grade availability demands. In particular, two ILP models, coined *AllOne* and *AllAny*, formulate the redundancy problem and exploit a novel algorithm that enables the quantification of the structural dependency among network nodes as a result of the network topology. The algorithm employs a metric proposed in [37] called *node dependency index* $DI(i|n)$ which measures the average level of dependency that node i has on node n in connecting with the other nodes of the network. To this end, given a network represented as an undirected connected graph $G(\mathcal{N}, \mathcal{L})$, where $\mathcal{N}$ denotes the set of nodes and $\mathcal{L}$ the set of links, we introduce $\mathcal{C}(i)$ as *the set of critical nodes of node* i, i.e., node i highly depends on. Node i is said to highly depend on a node $n \in \mathcal{C}(i)$, or in other words, n is critical to i, if $DI(i|n)$ is above a given threshold t_{DI}:

$$\mathcal{C}(i) = \{n \mid DI(i|n) > t_{DI}, n \in \mathcal{N}^{-i}\} \quad (2)$$

If $\mathcal{C}(i)$ is empty, it means that i is not highly dependent on the other nodes. For example, in a full mesh network, all nodes are structurally independent of each other as the failure of one node does not affect the connectivity among the others.

It is worth noting that $DI(j|i)$ has two directions, i.e., $DI(i|j)$ – dependence of i on j and $DI(j|i)$ – that of j on i, which obviously can quantify different impacts. Consequently, to minimize the correlation among nodes i and j we should avoid $j \in \mathcal{C}(i)$ or $i \in \mathcal{C}(j)$. We call this the *first-level*

dependency among nodes. In addition, a critical node n of i may also be critical to another node k i.e., $n \in \mathcal{C}(i)$ and $n \in \mathcal{C}(k)$. In this case, both nodes i and k depend on the same node n. As a result, the failure of such a critical node, e.g. n, may result in the unavailability of those nodes that depend on it, e.g., i and k, hence presenting a structural correlation that we refer to as the *second-level dependency* among nodes. An implication of this is that: we should avoid using k as a backup for node i, even though k is not in $\mathcal{C}(i)$. Under these considerations, Algorithm 1 presents the generic flow of the algorithm for finding the set of network-structurally correlated nodes with node i called $\hat{\mathcal{B}}_i$. The set is initially empty and the algorithm starts by finding the set of nodes based on the first-level dependency in both directions (Lines 1-2 and Lines 3-5 respectively). Then, nodes that have the second-level of dependency described above are added (Lines 6-9). The set $\hat{\mathcal{B}}_i$ represents the set that shall not be used to host backup instances of the functions that are running in node i, and such information is used both on the placement and assignment of VNFs to service chains. Our investigation shows that for the considered network topologies, up to 30% more flows are able to reach the targeted 5'9s availabilities compared to cases where no structural correlation is considered in the model.

Algorithm 1 Finding network-structurally correlated nodes

Input: $G(\mathcal{N}, \mathcal{L})$, t_{DI}
Output: $\hat{\mathcal{B}}_i$

```
1: Find C(i) using (2)
2: Insert C(i) to B̂_i
3: for j ∈ N^{-i} do
4:     if i ∈ C(j) then
5:         Insert j to B̂_i
6:     if j ∈ C(i) then
7:         for k ∈ N^{-j} do
8:             if j ∈ C(k) then
9:                 Insert k to B̂_i
10: return B̂_i
```

The models *AllOne* and *AllAny* differ in that the former considers the placement of backup instances of a service chain into one physical node, given sufficient node capacity, whereas the latter considers multiple nodes for hosting backup instances of a given chain. The work further investigates the resource efficiency of the models under different resource constraints and flows' targeted availabilities. Additional details on the models and numerical results can be found in [36].

The optimization problems formulated in [36] find optimal placements and assignments of VNFs to service chains but problem complexity is $\mathcal{NP}$-hard, hence being intractable for large-scale problems. This limits the algorithmic applicability and its scalability. To address this issue the final contribution of this work is presented in [38]. The work proposes customized heuristics that assemble the fundamental aspects of the original problem including structural dependency, heterogeneity of nodes and VNF instances, and system resource constraints, and scales well even for large scale setups. A distinctive feature of the algorithm, coined CoShare, is the introduction of a novel idea, referred to as *NF shared reservation*, for achieving higher

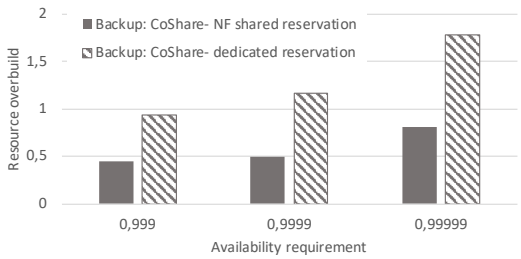

(a) Resource overbuild.

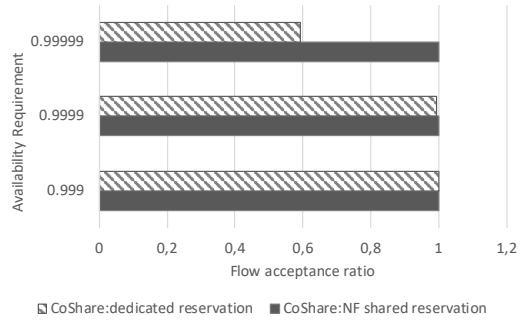

(b) Flow acceptance ratio.

Fig. 4. CoShare with shared reservation performance.

resource efficiency. The concept is based on the consideration that flows that do not share a node or link on their primary path thus, do not fail simultaneously upon a single failure, can share capacity of one or more backup instances. This is because for non-structurally correlated nodes, the failure of one node will not impact the other(s) and thus, the common shared capacity can be used among flows that are hosted onto non-structurally correlated nodes. These flows are called independent flows and such consideration is explicitly taken into account in both redundancy placement and assignment. In addition, for comparison, CoShare considers also the case where *dedicated* reserved capacity is allocated for backup instances.

CoShare performs a bin packing of the instances onto nodes that have been precedently categorized and prioritized based on their availability and structural correlation. Moreover, the algorithm also checks and prioritizes the placement of the NF types that are mostly requested by flows such that the highest number of flows can be accommodated into the network. Finally, the heuristics performs the actual bin packing by allocating backup instances into nodes that have sufficient capacity based on the set of priorities on both nodes and instance types. Subsequently, the algorithm performs the assignment of backup instances to service chains, also referred to as flow routing, for satisfying flow's requested availability. It employs a weight-based approach for applying the shared reservation mechanism where out of the feasible backup chains, i.e., the set of candidate backup chains, it assigns to a given flow f the chain that maximizes the utilization of resources so as to minimize the total number of backup instances required. Due to space limitations we omit showing the algorithmic pseudo-code but more details can be found in [38].

Remarkably, CoShare allows halving the required number

of backup instances compared to both *dedicated reservation* and the optimized solutions in [36], while still satisfying high availability demands. For 700 flows requiring 5'9s availability, the *shared reservation* resource overbuild results in 93% compared to 178% of the *dedicated reservation*, refer to Fig. 4(a). In addition, Fig. 4(b) shows that *shared reservation* achieves a 100% flow acceptance ratio, i.e., number of accommodated over the total number of service requests, compared to 59% of the *dedicated* variant. Furthermore, for the experimented networks, the optimized solution requires more than 12 minutes while CoShare generates the solution in less than one second.

IV. Final Remarks

With more and more end users taking for granted network and service continuity, ensuring highly available NFV-based services becomes of paramount importance for embracing the promising benefits of NFV. To this end, this work proposes solutions for modeling, assessing, and providing highly available NFV-based services. The proposed solutions consist of specific modeling approaches for a more realistic abstraction of end-to-end NFV services by featuring essential aspects such as element inter-dependencies, network connectivity requirements, and important elements such as the MANO. Moreover, the work proposes resource-efficient and scalable solutions that perform the allocation of redundant resources for satisfying high-availability demands.

References

[1] J. Sherry *et al.*, "Making middleboxes someone else's problem: Network processing as a cloud service," in *Proceedings of the ACM SIGCOMM 2012*, p. 13–24.

[2] V. Sekar *et al.*, "The middlebox manifesto: Enabling innovation in middlebox deployment," in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, 2011.

[3] R. Mijumbi *et al.*, "Network function virtualization: State-of-the-art and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 236–262, 2016.

[4] B. Han *et al.*, "On the resiliency of virtual network functions," *IEEE Communications Magazine*, vol. 55, no. 7, pp. 152–157, 2017.

[5] D. Cotroneo, L. De Simone, and R. Natella, "NFV-Bench: A Dependability Benchmark for Network Function Virtualization Systems," *IEEE TNSM*, vol. 14, no. 4, pp. 934–948, 2017.

[6] ETSI, "Network Functions Virtualisation (NFV); Reliability; Report on Models and Features for End-to-End Reliability," 2016.

[7] B. Blanco *et al.*, "Technology pillars in the architecture of future 5G mobile networks: NFV, MEC and SDN," *Computer Standards & Interfaces*, vol. 54, pp. 216–228, 2017.

[8] M. Casazza, M. Bouet, and S. Secci, "Availability-driven NFV orchestration," *Computer Networks*, vol. 155, pp. 47–61, 2019.

[9] A. Hmaity *et al.*, "Protection strategies for virtual network functions placement and service chains provisioning," *Networks*, vol. 70, pp. 373–387, 2017.

[10] A. Gonzalez *et al.*, "Service Availability in the NFV Virtualized Evolved Packet Core," in *GLOBECOM*. IEEE, 2015, pp. 1–6.

[11] M. Di Mauro *et al.*, "Service function chaining deployed in an NFV environment: An availability modeling," in *IEEE Conference on Standards for Communications and Networking*. IEEE, 2017, pp. 42–47.

[12] ——, "IP multimedia subsystem in an NFV environment: Availability evaluation and sensitivity analysis," in *2018 IEEE Conference on NFV-SDN*. IEEE, 2018, pp. 1–6.

[13] M. Di Mauro, M. Longo, and F. Postiglione, "Availability evaluation of multi-tenant service function chaining infrastructures by multidimensional universal generating function," *IEEE Transactions on Services Computing*, 2018.

[14] E. Guedes and P. Maciel, "Stochastic Model for Availability Analysis of Service Function Chains using Rejuvenation and Live Migration," in *2019 IEEE ISSREW*, 2019, pp. 211–217.

[15] J. Fan *et al.*, "GREP: Guaranteeing reliability with enhanced protection in NFV," in *Proceedings of the 2015 ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization*, pp. 13–18.

[16] J. Fan, M. Jiang, and C. Qiao, "Carrier-grade availability-aware mapping of service function chains with on-site backups," in *IEEE/ACM 25th IWQoS*. IEEE, 2017, pp. 1–10.

[17] J. Fan *et al.*, "A framework for provisioning availability of NFV in data center networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 10, pp. 2246–2259, 2018.

[18] L. Qu *et al.*, "A Reliability-Aware Network Service Chain Provisioning With Delay Guarantees in NFV-Enabled Enterprise Datacenter Networks," *IEEE Trans. on Network and Service Management*, vol. 14, no. 3, 2017.

[19] L. Qu, M. Khabbaz, and C. Assi, "Reliability-aware service chaining in carrier-grade softwarized networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 3, pp. 558–573, 2018.

[20] W. Ding, H. Yu, and S. Luo, "Enhancing the reliability of services in NFV with the cost-efficient redundancy scheme," in *IEEE International Conference on Communications (ICC)*. IEEE, 2017, pp. 1–6.

[21] N.-T. Dinh and Y. Kim, "An efficient reliability guaranteed deployment scheme for service function chains," *IEEE Access*, vol. 7, pp. 46 491–46 505, 2019.

[22] J. Zhang *et al.*, "RABA: Resource-Aware Backup Allocation For A Chain of Virtual Network Functions," in *INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 1918–1926.

[23] M. Casazza *et al.*, "Securing virtual network function placement with high availability guarantees," in *2017 IFIP Networking Conference and Workshops*. IEEE, 2017, pp. 1–9.

[24] D. Li *et al.*, "Availability Aware VNF Deployment in Datacenter Through Shared Redundancy and Multi-Tenancy," *IEEE Trans. on Network and Service Management*, vol. 16, no. 4, pp. 1651–1664, 2019.

[25] A. Gonzalez *et al.*, "Dependability of the NFV orchestrator: State of the art and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, pp. 3307–3329, 2018.

[26] ETSI, "Network Function Virtualisation (NFV); Reliability; Report on the resilience of NFV-MANO critical capabilities," 2017.

[27] G. Li, D. Wang, C. Kalmanek, and R. Doverspike, "Efficient distributed path selection for shared restoration connections," in *Proceedings. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 1. IEEE, 2002, pp. 140–149.

[28] K. S. Trivedi and A. Bobbio, *Reliability and availability engineering: modeling, analysis, and applications*. Cambridge Univ. Press, 2017.

[29] M. Malhotra and K. S. Trivedi, "Power-hierarchy of dependability-model types," *IEEE Transactions on Reliability*, vol. 43, pp. 493–502, 1994.

[30] B. Tola *et al.*, "Modeling and evaluating NFV-enabled network services under different availability modes," in *2019 International Conference on the Design of Reliable Communication Networks (DRCN)*, 2019.

[31] "Möbius: Model-based environment for validation of system reliability, availability, security and performance," "https://www.mobius.illinois.edu".

[32] B. Tola, Y. Jiang, and B. E. Helvik, "On the resilience of the NFV-MANO: An availability model of a cloud-native architecture," in *2020 16th International Conference on DRCN*, 2020, pp. 1–7.

[33] ——, "Model-driven availability assessment of the NFV-MANO with software rejuvenation," *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, 2021.

[34] B. Tola, G. Nencioni, and B. E. Helvik, "Network-aware availability modeling of an end-to-end NFV-enabled service," *IEEE Transactions on Network and Service Management*, vol. 16, no. 4, 2019.

[35] E. Zhai *et al.*, "Heading off correlated failures through independence-as-a-service," in *11th USENIX Symposium on OSDI*, 2014, pp. 317–334.

[36] Y. T. Woldeyohannes, B. Tola, and Y. Jiang, "Towards carrier-grade service provisioning in NFV," in *2019 15th International Conference on the Design of Reliable Communication Networks (DRCN)*, 2019.

[37] Y. T. Woldeyohannes and Y. Jiang, "Measures for network structural dependency analysis," *IEEE Communications Letters*, vol. 22, no. 10, pp. 2052–2055, 2018.

[38] Y. T. Woldeyohannes, B. Tola, Y. Jiang, and K. K. Ramakrishnan, "Coshare: An efficient approach for redundancy allocation in NFV," *IEEE/ACM Transactions on Networking*, vol. 30, no. 3, pp. 1014–1028, 2022.