

Survivable Virtual Network Embedding

Nashid Shahriar*, Raouf Boutaba†

*Department of Computer Science, University of Regina; Email: Nashid.Shahriar@uregina.ca

†David R. Cheriton School of Computer Science, University of Waterloo; Email: rboutaba@uwaterloo.ca

Abstract—Network Virtualization (NV) is an enabling technology for the future Internet and next-generation communication networks. A fundamental problem in NV is to map the virtual nodes and virtual links of a VN to physical nodes and paths, respectively, known as the Virtual Network Embedding (VNE) problem. A VNE that can survive physical resource failures is known as the survivable VNE (SVNE) problem, and has received significant attention recently. In this thesis, we address variants of the SVNE problem with different bandwidth and reliability requirements for transport networks. Specifically, the thesis includes four main contributions. First, a connectivity-aware VNE approach that ensures VN connectivity without bandwidth guarantee in the face of multiple link failures. Second, a joint spare capacity allocation and VNE scheme that provides bandwidth guarantee against link failures by augmenting VNs with necessary spare capacity. Third, a generalized recovery mechanism to re-embed the VNs that are impacted by a physical node failure. Fourth, a reliable VNE scheme with dedicated protection that allows tuning of available bandwidth of a VN during a physical link failure. We show the effectiveness of the proposed SVNE schemes through extensive simulations.

I. INTRODUCTION

Network Virtualization (NV) is an enabler for the future Internet and next-generation communication networks to support emerging applications' Quality of Service (QoS) requirements and programmability of networks [1]. NV transforms a physical network from a static “one size fits all” architecture to a new paradigm where virtual networks are created using virtualization and software defined networking, with appropriate isolation, resources, and topology to serve a particular application or service. It also facilitates a new business model, namely, Network-as-a-Service (NaaS), which decouples applications and services from the networks supporting them. The 5th Generation (5G) mobile network operators or Infrastructure Providers (InPs) have adopted the NaaS model to partition their physical networks into multiple network slices and lease these slices to Service Providers (SPs) [2]. SPs use the leased network slices to offer tailored services satisfying application or service specific QoS requirements without any investment in deploying and managing a physical network [3].

The benefits of NV comes with additional resource management challenges for infrastructure providers as they have to keep virtual networks up and running throughout their lifespans. A fundamental problem in NV is to allocate physical resources to virtual networks, which is known as the Virtual Network Embedding (VNE) problem [4]. The VNE problem respectively maps virtual nodes and links of a virtual network to physical nodes and paths (a sequence of links) in a physical network, while satisfying resource budgets and achieving specific objectives such as minimizing resource usage [3].

Virtual networks cannot function properly in the face of network device or link failures that are inevitable in everyday operations [5], [6]. Such failures, if not treated properly, may

disrupt the services hosted by virtual networks, incurring high penalties in terms of revenue losses and SLA violations [7]. While accepting the possibility of failures, infrastructure providers have to ensure different levels of availability (*e.g.*, 99.999%, 99.9999%, etc.) for their virtual networks as part of SLAs [8]. Survivability mechanisms can greatly reduce the impact of failures, while minimizing downtime and upholding the reputation of a network provider. In this context, a VN embedding that can survive substrate failures is known as the survivable VNE (SVNE) [9], and has received significant attention from the research community.

This dissertation proposes original formulations of the SVNE problem with bandwidth and reliability requirements for transport networks that provide data transmission services over large geographical (metro, regional, or national) areas [10]. The considered transport technologies are: Transport Software Defined Networks (T-SDNs) [11] and Elastic Optical Networks (EONs) [12], expected to meet the demanding requirements of 5G networks. In these networks, failures (*e.g.*, device outage or fiber-cut) can result in significant amount of data loss [5], [6]. Existing SVNE strategies advocate for keeping failure management tasks transparent to the service provider and support a limited set of QoS requirements. In contrast, this dissertation argues for delegating failure management responsibilities to a virtual network operator. Offloading failure management to a virtual network has two benefits: i) enabling a variety of QoS requirements through different survivability models; ii) offering more opportunities to minimize resource consumption. This dissertation presents four novel survivability models to enable failure management at the virtual network layer as opposed to the physical network layer. Specifically, this dissertation makes the following contributions that we summarize in this paper:

- Guaranteeing virtual network connectivity against multiple link failures in T-SDNs (Section II).
- Jointly optimizing spare capacity allocation and SVNE to guarantee bandwidth in the presence of multiple link failures in T-SDNs (Section III).
- Re-embedding a batch of virtual networks to recover from a node failure in T-SDNs (Section IV).
- Reliable slicing of EONs to ensure fast fail-over of virtual networks against a link failure in EONs (Section V).

II. CONNECTIVITY-AWARE VIRTUAL NETWORK EMBEDDING (*CoViNE*)

The goal of *CoViNE* is to find a VN embedding that remains connected (without any bandwidth guarantee) in the presence of multiple substrate link failures in T-SDNs [13], [14]. Guaranteeing connectivity in the VN embedding will incur less resource overhead and reduced cost of leasing resources for a

VN, however, providing a weaker form of survivability. This survivability model is well-suited for VNs that carry best-effort traffic (e.g., file transfer and email communication) and can tolerate disruption during failure restoration. Upon failures, the affected VN traffic can be rerouted to alternate paths following any predefined policy, e.g., customer priority. *CoViNE*'s survivability model also allows to delegate failure handling responsibility to an SP, which can then use a Software Defined Network (SDN) controller to employ their own restoration techniques instead of simply relying on the InP [15]–[17].

A. Problem Statement

We represent an SN as an undirected graph, $G = (V, E)$, where V and E denote the set of Substrate Nodes (SNodes) and Substrate Links (SLinks), respectively. Bandwidth capacity of an SLink $(u, v) \in E$ is b_{uv} , while the cost of allocating one unit of bandwidth in (u, v) is C_{uv} . A VN is represented as an undirected graph $\bar{G} = (\bar{V}, \bar{E})$, where \bar{V} and \bar{E} denote the set of Virtual Nodes (VNodes) and Virtual Links (VLinks), respectively. Each VLink $(\bar{u}, \bar{v}) \in \bar{E}$ has bandwidth requirement $b_{\bar{u}\bar{v}}$. Each VNode $\bar{u} \in \bar{V}$ has a location constraint, $L(\bar{u}) \subseteq V$, that denotes the set of SNodes where \bar{u} can be embedded. We represent the location constraint with the binary variable $\ell_{\bar{u}u}$ that is set to 1 if $\bar{u} \in \bar{V}$ can be mapped to $u \in V$, 0 otherwise. Let Q^{uv} represent a path in the SN between a pair of SNodes $u \in V$ and $v \in V$ such that $u \neq v$. Given an SN $G = (V, E)$, a VN $\bar{G} = (\bar{V}, \bar{E})$, and location constraints $L(\bar{u}), \forall \bar{u} \in \bar{V}$, *CoViNE* finds an embedding that

- provides a function $f : \bar{V} \rightarrow V$ to map every VNode $\bar{u} \in \bar{V}$ to exactly one SNode $u \in V$ while satisfying the location constraint and incurring no overlap, i.e., $\forall \bar{u}, \bar{v} \in \bar{V} \wedge \bar{u} \neq \bar{v} \implies f(\bar{u}) \neq f(\bar{v})$ and $\forall \bar{u} \in \bar{V} f(\bar{u}) \in L(\bar{u})$,
- provides a function $g : \bar{E} \rightarrow 2^E$ to map each VLink $(\bar{u}, \bar{v}) \in \bar{E}$ to a substrate path $Q^{f(\bar{u})f(\bar{v})}$ with sufficient bandwidth to satisfy the VLink demand $b_{\bar{u}\bar{v}}$,
- ensures the connectivity in \bar{G} in the presence of up to k SLink failures in G ,
- minimizes the total cost of embedding in terms of substrate bandwidth consumption.

$$\sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}} \sum_{\forall (u, v) \in Q^{f(\bar{u})f(\bar{v})}} C_{uv} \times b_{\bar{u}\bar{v}} \quad (1)$$

B. Contribution

A VN embedding remains connected during k SLink failures if the following two necessary conditions are met: i) the VN is $k + 1$ edge-connected following the definition of $k + 1$ edge-connected graphs, implying that the size of each edge-cut is at least $k + 1$, $\forall C_i \in \bar{C}^{\bar{G}}, |\bar{C}_i| \geq k + 1$, ii) the VLinks in each edge-cut C_i are embedded on at least $k + 1$ edge-disjoint paths in the SN. If the given VN \bar{G} lacks $k + 1$ connectivity, we propose to augment \bar{G} with parallel VLinks to satisfy the first condition. To characterize the edge-disjointness relationship among the VLinks, we develop *conflicting set* abstraction based on a theoretical analysis of *CoViNE*. *Conflicting set* abstraction allows us to satisfy second condition without enumerating exponential number of edge-cuts of a VN. Based on this foundation, we propose three novel solutions to *CoViNE*:

CoViNE-opt. An Integer Linear Program (ILP) formulation that jointly optimizes VN augmentation, disjointness constraints computation, and embedding of the augmented VN to optimally solve *CoViNE*. *CoViNE-opt* has an exponential number of variables and constraints that severely limits its scalability. To scale to larger problem instances, we decompose *CoViNE* into three sub-problems: (i) augmenting the VN with zero or more virtual links to make it $k + 1$ -edge connected; (ii) computing the set of virtual links to be embedded disjointly for ensuring connectivity against k substrate link failures; and (iii) embedding the VN while satisfying the aforementioned disjointness constraints. The following two approaches (i.e., *CoViNE-ILP* and *CoViNE-fast*) sequentially solve all the sub-problems of *CoViNE* in a more scalable manner, however, without guaranteeing an optimal solution.

CoViNE-ILP. Employs a *conflicting set*-based heuristic that solves sub-problems (i) and (ii) in polynomial time. The *conflicting set* abstraction allows to generate a polynomial number of variables and constraints to be used by an ILP for solving sub-problem (iii). The complexity of this ILP limits its applicability to substrate networks of few hundred nodes.

CoViNE-fast. Uses heuristics for all three sub-problems of *CoViNE* to scale to larger problem instances.

C. Results

We performed extensive simulations to evaluate the optimality and scalability of the proposed solutions under single and double Sink failures. We also compare our solutions with a VNE approach that does not guarantee VN connectivity upon failures. Through our simulation study, we show that virtual network connectivity is ensured against single link failures with about 25% additional resources, compared to a VNE that does not guarantee connectivity. We also show that *CoViNE* restores more (close to 100%) bandwidth for higher priority traffic than VNE without any connectivity guarantee in the presence of multiple SLink failures as presented in Fig. 1.

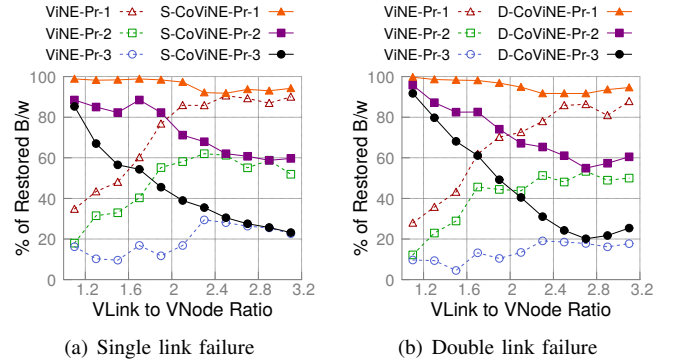


Fig. 1: Restored bandwidth of our approaches

III. JOINT SPARE CAPACITY ALLOCATION AND VIRTUAL NETWORK EMBEDDING

Virtual networks for bandwidth-savvy applications may need bandwidth guarantee even in the presence of failures in T-SDNs. One way of providing such guarantee is to allocate spare bandwidth on VLinks (as opposed to doing so on the

SN) so that all the traffic impacted by a failure can be rerouted within the VN [18], [19]. With this type of survivability, InPs can offload failure management tasks to SPs by augmenting VNs with sufficient spare capacity for backup and embedding the VNs in a way that primary and backup VN resources are not affected by the same substrate link failure. When a substrate link fails, it is the SP's responsibility to reroute the affected traffic to the pre-allocated backup resources within the VN. Independently addressing spare bandwidth allocation and VNE may lead to sub-optimal solutions. In this work, we study the joint optimization problem of computing spare bandwidth allocation and VNE with the objective of guaranteeing VN survivability under multiple substrate link failures and minimizing resource usage in the SN.

A. Problem Statement

Given an SN $G = (V, E)$, a VN $\bar{G} = (\bar{V}, \bar{E})$ and location constraints $L(\bar{u}), \forall \bar{u} \in \bar{V}, :$

- For each VLink $(\bar{u}, \bar{v}) \in \bar{E}$, allocate spare capacity along a set of K backup virtual paths (VPaths) $\bar{P}_{\bar{u}\bar{v}}^K = \{\bar{P}_{\bar{u}\bar{v}}^k | 1 \leq k \leq K\}$ in the VN, where $\bar{P}_{\bar{u}\bar{v}}^k$ is the k th VPath between \bar{u} and \bar{v} such that $\bar{P}_{\bar{u}\bar{v}}^k$ is edge disjoint from (\bar{u}, \bar{v}) and from each $\bar{P}_{\bar{u}\bar{v}}^j \in \bar{P}_{\bar{u}\bar{v}}^K$ with $j \neq k$, and $b_{\bar{u}\bar{v}}$ spare bandwidth is available on the VLinks in $\bar{P}_{\bar{u}\bar{v}}^k$ after (\bar{u}, \bar{v}) is affected by an SLink failure.
- Map each VNode $\bar{v} \in \bar{V}$ to exactly one SNode, $u \in V$ as per location constraints. Multiple VNodes from the same VN should not be mapped to the same SNode. However, multiple VNodes from different VNs can share an SNode.
- Map each VLink (\bar{u}, \bar{v}) to a non-empty substrate path (SPath) $P_{\bar{u}\bar{v}}$ having sufficient bandwidth to accommodate the primary demand of (\bar{u}, \bar{v}) and the spare backup bandwidth $S_{\bar{u}\bar{v}}$ allocated on (\bar{u}, \bar{v}) . A VLink (\bar{u}, \bar{v}) and the VLinks on its VPath $\bar{P}_{\bar{u}\bar{v}}^k$ are edge disjointly mapped on the SN to ensure that SLink failures do not affect them at the same time. Similarly, two VLinks present in the two VPaths, such as $\bar{P}_{\bar{u}\bar{v}}^k$ and $\bar{P}_{\bar{u}\bar{v}}^j$ where $j \neq k$, of the same VLink (\bar{u}, \bar{v}) are mapped on edge disjoint SPaths to eliminate the risk of both the VPaths failing together.
- Minimize the total cost of allocating bandwidth on the SN to embed the VN equipped with spare bandwidth.

$$\sum_{\forall(\bar{u}, \bar{v}) \in \bar{E}} \sum_{\forall(u, v) \in P_{\bar{u}\bar{v}}} C_{uv} \times (b_{\bar{u}\bar{v}} + S_{\bar{u}\bar{v}}) \quad (2)$$

B. Contribution

VLinks that share at least one SLink on their mapped SPaths share the same risk since all of them can be impacted if the shared SLink fails. In a context where only single SLink failure is considered, a set of VLinks belong to the same shared risk group (SRG) if and only if they share at least one SLink on their mapped SPaths. In contrast, VLinks that do not share any SLink on their mapped SPaths belong to different SRGs. To represent the SRGs, we partition the VLinks into a number of SRGs represented by the set $D = \{d_1, d_2, d_3, \dots, d_{|D|}\}$, where $|D| \leq |\bar{E}|$. A VLink belongs to exactly one SRG $d_i \in D$ and shares at least one SLink on its mapped SPath with other VLinks in d_i . Based on how the VLinks form different

SRGs during VN embedding, the requirement for spare backup capacity on the VLinks can be different.

Assume that a VLink $(\bar{u}, \bar{v}) \in \bar{E}$ is present on the backup VPaths of a set of VLinks $\bar{H}_{\bar{u}\bar{v}} \subseteq \bar{E}$. If the VLinks in $\bar{H}_{\bar{u}\bar{v}}$ are in the same SRG, they share at least one SLink on their mapped SPaths whose failure can affect all the VLinks in $\bar{H}_{\bar{u}\bar{v}}$. Therefore, spare capacity $S_{\bar{u}\bar{v}}$ should be sufficient to support the bandwidth requirement of all the affected VLinks. If the VLinks in $\bar{H}_{\bar{u}\bar{v}}$ belong to different SRGs, they do not share any SLink on their mapped SPaths. At most one of the VLinks will be affected by a single SLink failure. Therefore, $S_{\bar{u}\bar{v}}$ should be sufficient to support the maximum bandwidth requirement of these VLinks. In general, if the VLinks in $\bar{H}_{\bar{u}\bar{v}}$ form a set of $D = \{d_1, d_2, d_3, \dots, d_{|D|}\}$ SRGs, we can generalize the spare backup bandwidth allocated to (\bar{u}, \bar{v}) as:

$$S_{\bar{u}\bar{v}} = \max_{\forall d_i \in D} \left(\sum_{\forall(\bar{x}, \bar{y}) \in \bar{H}_{\bar{u}\bar{v}}} d_i^{\bar{x}\bar{y}} b_{\bar{x}\bar{y}} \right) \quad (3)$$

We formulate a joint optimization model using a Quadratic Integer Program (QIP) to optimally solve spare capacity allocation and survivable VN embedding simultaneously. We transform the QIP into an ILP without sacrificing its optimality. We provide two more ILP formulations for solving two extreme cases of spare capacity sharing. We present a mathematical analysis that dictates how the topological properties of the SN affect the level of spare capacity sharing. The ILP formulations for the joint optimization problem are not scalable to large problem instances. Hence, we devise an efficient heuristic algorithm to tackle the computational complexity of the ILP-based solutions. The algorithm leverages a novel spare bandwidth sharing model to estimate the spare capacity and computes embedding based on the estimated spare capacity. In the final step, the algorithm re-optimizes spare capacity allocation based on the final embedding information.

C. Results

We perform simulations to evaluate our solutions for single and double link failures. Simulation results show that ILP formulations for the two special cases can provide upper and approximately lower bounds of the solution spectrum. Moreover, the heuristic allocates $\sim 21\%$ additional resources compared to the approximated lower bound, while executing several orders of magnitude faster. We also perform a quantitative comparison between the SVNE with VN level protection and the traditional SVNE with SN level protection [20]. Evaluation results shows that our solution decreases resource usage by 30% compared to SVNE solutions that rely on the SN to perform failure management as presented in Fig. 2.

IV. RECOVERY FROM NODE FAILURE IN VIRTUAL NETWORK EMBEDDING (*ReNoVatE*)

ReNoVatE takes a batch of VN failures resulting from a single SNode failure, and produces alternate embeddings for the failed VNodes and VLinks [21], [22]. The impact of an SNode failure is more drastic than that of an SLink failure since a node failure affects embedding of all the VNodes and VLinks of all VNs passing through the failed

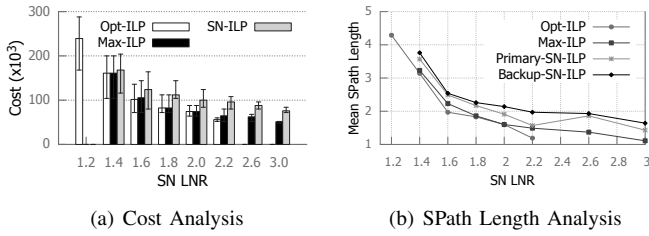


Fig. 2: Comparison with SN level Survivability of [20]

node. Preallocating backup resources for multiple failures resulting from an SNode failure can be expensive [23], [24]. Instead, an SP may prefer to reactively re-embed the failed part of its VN to avoid huge cost of preallocated backup resources in a failure-prone SN. Such reactive approaches can be adopted by a VN whose traffic can tolerate non-negligible service disruptions (*e.g.*, non-realtime services). We adopt a generalized recovery approach that can result in two different recovery models. The first one is a fair recovery model (FRM) that maximizes the number of recoveries across all the affected VNs. This can be the sought-after choice of an InP who wants to treat all the affected VNs fairly in a resource constrained SN. Second is a priority-based recovery model (PRM) that takes into account an InP's preference during recovery. PRM allows an InP to prioritize the recovery of affected VNs based on SLA strictness, impacts of failure, profits, and so on to achieve its goal.

A. Problem Statement

In this problem, V^f and E^f represent the set of failed SNodes and SLinks, respectively. P_{uv} represents a path between SNodes u and v . We denote the set of VNs embedded on the SN G as $\bar{G} = \{\bar{G}_1, \bar{G}_2, \dots, \bar{G}_{|\bar{G}|}\}$. Each VN $\bar{G}_i \in \bar{G}$ is represented as an undirected graph $\bar{G}_i = (\bar{V}_i, \bar{E}_i)$, where \bar{V}_i and \bar{E}_i are the sets of VNodes and VLinks of \bar{G}_i , respectively. Each VLink $(\bar{u}, \bar{v}) \in \bar{E}_i$ has a bandwidth demand $b_{i\bar{u}\bar{v}}$ and a penalty $\pi_{i\bar{u}\bar{v}}$ that represents the revenue loss due to unavailability of (\bar{u}, \bar{v}) . Each VN \bar{G}_i has a set of location constraints, $L_i = \{L_i(\bar{u}) | L_i(\bar{u}) \subseteq V, \forall \bar{u} \in \bar{V}_i\}$, such that a VNode $\bar{u} \in \bar{V}_i$ can only be mapped to an SNode $u \in L_i(\bar{u})$. Let, $f(\bar{u})$ and $g(\bar{u}\bar{v})$ denote the SNode and substrate path where \bar{u} and (\bar{u}, \bar{v}) have been embedded, respectively. An SNode failure results in a set of VNode and VLink failures of a VN \bar{G}_i defined as $\bar{V}_i^f = \{\bar{u} \in \bar{V}_i | f(\bar{u}) \subseteq V^f\}$ and $\bar{E}_i^f = \{(\bar{u}, \bar{v}) \in \bar{E}_i | (u, v) \in g(\bar{u}\bar{v}) \wedge (u, v) \in E^f\}$, respectively. There are two types of VLinks in \bar{E}_i^f : i) adjacent VLinks: VLinks adjacent to the failed VNode $\bar{u} \in \bar{V}_i^f$ is represented by $\bar{E}_i^{f,adj}$; ii) Independent VLinks: VLinks that have failed due to the failure of some SLinks on their mapped substrate paths is denoted by $\bar{E}_i^{f,ind}$. Finally, $\bar{V}^f = \{\bigcup \bar{V}_i^f\}$, $\bar{E}^f = \{\bigcup \bar{E}_i^f\}$, and $\bar{E}^f = \{\bigcup \bar{E}_i^{f,ind}\}$ represent the set of failed VNodes, VLinks, and independent VLinks of all the VNs in \bar{G} , respectively. Given an SN $G = (V, E)$, a failed SNode implying $|V_f| = 1$, and a set of affected VNs \bar{G} embedded on G , re-embed the failed VNodes in \bar{V}^f and the failed VLinks

in \bar{E}^f on G as per PRM or FRM. The binary decision variable $z_{i\bar{u}\bar{v}}$ indicates if a failed VLink (\bar{u}, \bar{v}) is re-embedded or not.

B. Contribution

We formulate *ReNoVatE* as an ILP based optimization model with a primary and a secondary objectives. The primary objective depends on the recovery model being chosen. For FRM, primary objective is to maximize the total number of recovered VLinks across all the affected VNs. In PRM, primary objective is to minimize the total penalty for all the failed VLinks that remain unrecovered. The secondary objective is to minimize the total cost of re-embedding in terms of SLink bandwidth consumption and used to break ties for the primary objective.

$$\begin{aligned} \text{minimize} & \left(\sum_{\forall \bar{G}_i \in \bar{G}} \sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} (1 - z_{i\bar{u}\bar{v}}) \times \pi_{i\bar{u}\bar{v}} \right) \\ & + w \left(\sum_{\forall \bar{G}_i \in \bar{G}} \sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} \sum_{\forall (u, v) \in E} C_{uv} \times b_{i\bar{u}\bar{v}} \right) \end{aligned} \quad (4)$$

These objectives are subject to the following constraints:

- a failed VNode $\bar{u} \in \bar{V}_i^f$ is re-embedded on exactly one SNode, $v \in L_i(\bar{u})$. In addition, multiple VNodes of the same VN cannot be mapped to an SNode. However, multiple VNodes from different VNs can share an SNode.
- a failed VLink $(\bar{u}, \bar{v}) \in \bar{E}_i^f$ is re-embedded on a substrate path $P_{f(\bar{u})f(\bar{v})}$ having sufficient bandwidth to accommodate the demand of the VLink. The re-embedding cannot use a substrate path containing the failed SNode.
- VNodes and VLinks not affected by the SNode failure are not re-embedded.

Since the optimization model cannot scale to large instances of the problem, we devise an efficient heuristic algorithm to find satisfactory solutions within prescribed time limits. The heuristic augments the SN with a pseudo sink SNode and applies a Max-flow (*e.g.*, *Edmonds-Karp*) algorithm to achieve the objective of PRM or FRM.

C. Results

We evaluate our solutions through extensive simulations and compare them with the most related state-of-the-art proposal in the literature [25]. Our evaluation results suggest that FRM-based solutions fail to take into account variety of recovery requirements. In contrast, PRM-based solutions can prioritize the affected VNs based on SLA requirements, impacts of failure or profits, and adhere to that priority during recovery. Our evaluation results demonstrate that our heuristic algorithm performs close to the ILP based optimal solution and outperforms the state-of-the-art solution in [25], in terms of i) number of recovered virtual links, ii) cost of recovery, and iii) execution time as shown in Fig. 3.

V. RELIABLE SLICING OF ELASTIC OPTICAL NETWORKS

The SVNE problem in EONs is compounded by a number of tunable transmission parameters (*e.g.*, modulation format, baud rate, and error correction codes) and a larger solution space for spectrum allocation. Compared to T-SDNs, EONs have additional constraints, such as spectrum continuity and

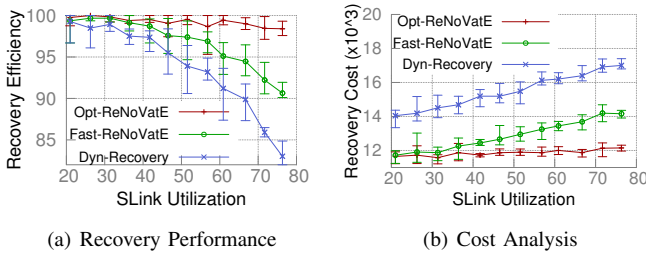


Fig. 3: Performance of *ReNoVatE*

contiguity constraints, imposed by the physical characteristics of optical devices and the properties of light. Furthermore, lightpaths in EONs carry huge volumes of data, and hence, even a short-lived outage can cause a significant traffic loss for virtual networks, necessitating a fast fail-over (*e.g.*, within 50 milliseconds [26]) capability. To meet this requirement, dedicated protection is the appropriate survivability option, although it incurs a 100% resource overhead that remains unused most of the time. To minimize resource footprint of dedicated protection, we leverage two techniques: bandwidth squeezing rate (BSR) that allows a virtual network operator to tune the amount of available bandwidth in case of failures and multi-path provisioning with demand splitting [27], [28]. These two techniques allow to significantly reduce spectrum usage for providing dedicated protection, especially in the case of fully-flexible EONs.

A. Problem Statement

The substrate EON is an undirected graph $G = (V, E)$, where V and E are the set of optical SNodes and optical SLinks, respectively. SLinks are bi-directional, *i.e.*, adjacent optical nodes are connected by one optical fiber in each direction. The optical frequency spectrum on each SLink $e = (u, v) \in E$ is divided into equal-width frequency slots represented by the set S and enumerated as $1, 2, \dots, |S|$. \mathcal{P} and $\mathcal{P}_{uv}^k \subset \mathcal{P}$ represent the set of all paths in G and the set of k -shortest paths between nodes $u, v \in V$, respectively. The number of SLinks and the physical length of a path p in kilometers are represented by $|p|$ and $len(p)$, respectively.

The following transmission parameters can be configured on a path p with length $len(p)$ to enable data transmission with different data-rates $d \in \mathcal{D}$: *baud-rate* or *symbol-rate*, b , *modulation format*, m , and *FEC overhead*, f , selected from the set of possible values \mathcal{B} , \mathcal{M} , and \mathcal{F} , respectively. We use a tuple $t = (d, b, m, f) \in \mathcal{T} = (\mathcal{D} \times \mathcal{B} \times \mathcal{M} \times \mathcal{F})$ to represent a transmission configuration that dictates the combination of $b \in \mathcal{B}$, $m \in \mathcal{M}$, and $f \in \mathcal{F}$ yielding a data-rate $d \in \mathcal{D}$. A *reach table* \mathcal{R} , computed based on physical layer characteristics, specifies the maximum length of a path (*i.e.*, the *reach* r_t) capable of retaining a satisfactory optical signal to noise ratio when configured according to a transmission configuration $t \in \mathcal{T}$. Finally, n_t denotes the number of slots required for a transmission configuration $t \in \mathcal{T}$, which is dependent on the parameters of t .

Similar to previous sections, the VN is represented by an undirected graph $\tilde{G} = (\tilde{V}, \tilde{E})$. The function $\tau: \tilde{V} \rightarrow V$ represents VNode to SNode mapping and is an input to our problem

(a common assumption for optical network virtualization [29]). Each virtual link $\bar{e} \in \tilde{E}$ has a bandwidth requirement $\bar{\beta}_{\bar{e}}$ and a bandwidth squeezing requirement $0 < BSR_{\bar{e}} \leq 100$, which indicates the percentage of original bandwidth that should be available after an SLink fails. We allow VLinks to be mapped on multiple SPaths (similar to [30], [31]), each with a lower data-rate than $\bar{\beta}_{\bar{e}}$. Splitting $\bar{\beta}_{\bar{e}}$ over multiple SPaths is a feasible way to support higher data-rates (*e.g.*, ≥ 400 Gbps) that limit the number of usable paths due to their shorter reaches. However, we restrict the number of VLink splits to maximum q (≥ 1). Given an SN G , a reach table \mathcal{R} , and a VN \tilde{G} with given VNode mapping function τ , Compute the link embedding function $\gamma: \tilde{E} \rightarrow \chi: \chi \subset \mathcal{P} \times \mathcal{T} \times S^2$ and $1 \leq |\chi| \leq q$, *i.e.*, compute up to a maximum of q splits for each VLink $\bar{e} \in \tilde{E}$ such that $0.01 \times BSR_{\bar{e}} \times \bar{\beta}_{\bar{e}}$ bandwidth is available during an SLink failure and at least $\bar{\beta}_{\bar{e}}$ bandwidth is available during the rest of the time.

B. Contribution

We develop an ILP formulation to optimally solve SVNE on EON by jointly considering all the flexible transmission parameters, while minimizing total spectrum usage. For each split, the formulation select a SPath and a suitable transmission configuration $t \in \mathcal{T}$ from the reach table \mathcal{R} , and allocate a contiguous segment of slots represented by the starting and ending slot index on each SLink along the SPath. Note that the same SPath can be used multiple times as the splits of a VLink following the reasoning in [31]. $\chi_{\bar{e}i} = (p, t, s_b, s_t) | 1 \leq i \leq q$ represents the i -th split, where $\chi_{\bar{e}i}^{(p)}$ and $|\chi_{\bar{e}i}^{(p)}|$ denote the selected SPath and the number of SLinks on the SPath $\chi_{\bar{e}i}^{(p)}$, respectively. In addition, allocation of spectrum slots for the i -th split begins at index $\chi_{\bar{e}i}^{(s_b)}$ and ends at index $\chi_{\bar{e}i}^{(s_t)}$ along each SLink in the SPath $\chi_{\bar{e}i}^{(p)}$. The ILP formulation has the following objective where ω is a tie-breaker:

$$\sum_{\forall \bar{e} \in \tilde{E}} \sum_{i=1}^q ((\chi_{\bar{e}i}^{(s_t)} - \chi_{\bar{e}i}^{(s_b)} + 1) \times |\chi_{\bar{e}i}^{(p)}| + \omega \times \{1 \text{ if } \chi_{\bar{e}i} \neq \phi\}) \quad (5)$$

The first part of the objective minimizes total number of slots required to provision the VN and second part minimizes the total number of splits. The objective function is subject to substrate resource constraints, slot exclusiveness constraints (*i.e.*, a spectrum slot is allocated to only one split) and spectral contiguity (*i.e.*, the allocated slots of each split are always adjacent to each other) and continuity (*i.e.*, the same sequence of slots are allocated on each SLink along an SPath) constraints on the lightpaths. Given the intractability of the ILP formulation, we propose a heuristic algorithm to solve larger instances of the problem. The heuristic uses dynamic programming and pruning to reduce computational complexity.

C. Results

We perform simulations using realistic network topologies, which provide valuable insight into how different levels of BSR and path diversity in the EON can impact the extent of backup resource savings for dedicated protection. We also analyze the steady state behavior of our heuristic solution

using a discrete event simulator. Our evaluation shows that by using multi-path provisioning, it is possible to guarantee up to 40% of the requested bandwidth of a VN during failure (*i.e.*, $BSR \leq 40\%$) while using as low as 10% additional spectrum resources. Consequently, VN blocking ratio for $BSR \leq 40\%$ remains very similar to that of the case with no backup as shown in Fig. 4.

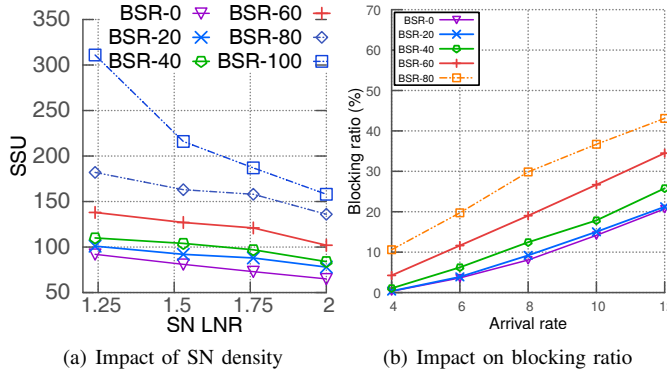


Fig. 4: Analysis of our reliability model

VI. CONCLUSION

As infrastructure providers are rolling out virtual networks as services, they are facing a number of challenges. One of the major challenges is to ensure virtual network survivability against failures in large physical infrastructures. Existing virtual network survivability literature fails to address the complexity and scale of the problem, resulting in over-simplified and impractical solutions. In this paper, we described four survivability models introduced to realize the vision of enabling failure management at the virtual network level (as opposed to at the physical network level). We addressed four key research challenges in survivable network virtualization and proposed novel solutions to solve them. We have also demonstrated the superiority of our proposed solutions through extensive simulations using realistic networks. We believe that these contributions can set the stage for further research specially in the area of automated failure management for future networks.

VII. FINAL REMARKS

The thesis is available at <http://hdl.handle.net/10012/16085> [10]. The work conducted during the course of the dissertation research has been published in [13], [14], [18], [19], [21], [22], [27], [28], [31]. Part of the work presented in this thesis has received the IEEE/ACM CNSM 2019 Best Paper Award, and the thesis has been recognized with the 2020 PhD Alumni Gold Medal from the University of Waterloo.

REFERENCES

- [1] N. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Computer Networks*, vol. 54, no. 5, pp. 862–876, 2010.
- [2] I. Afolabi, T. Taleb, K. Samdanis, A. Ksentini, and H. Flinck, "Network slicing and softwareization: A survey on principles, enabling technologies, and solutions," *IEEE CST*, vol. 20, no. 3, pp. 2429–2453, 2018.
- [3] A. Fischer, J. F. Botero, M. T. Beck, H. De Meer, and X. Hesselbach, "Virtual Network Embedding: A Survey," *IEEE CST*, Feb 2013.

- [4] M. Chowdhury, M. R. Rahman, and R. Boutaba, "Vineyard: Virtual network embedding algorithms with coordinated node and link mapping," *IEEE/ACM Trans. on Netw.*, vol. 20, no. 1, pp. 206–219, 2012.
- [5] P. Gill, N. Jain, and N. Nagappan, "Understanding network failures in data centers: measurement, analysis, and implications," in *ACM SIGCOMM CCR*, vol. 41, no. 4, 2011, pp. 350–361.
- [6] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot, "Characterization of failures in an ip backbone," in *INFOCOM*, vol. 4. IEEE, 2004, pp. 2307–2317.
- [7] R. Govindan, I. Minei, M. Kallahalla, B. Koley, and A. Vahdat, "Evolve or die: High-availability design principles drawn from googles network infrastructure," in *ACM SIGCOMM*, 2016, pp. 58–72.
- [8] A. Gupta and J. Shute, "High-availability at massive scale: Building google's data infrastructure for ads," in *Real-Time Business Intelligence and Analytics*. Springer, 2015, pp. 63–81.
- [9] S. Herker, A. Khan, and X. An, "Survey on survivable virtual network embedding problem and solutions," in *ICNS*, 2013.
- [10] N. Shahriar, "Survivable virtual network embedding in transport networks," 2020. [Online]. Available: <http://hdl.handle.net/10012/16085>
- [11] R. Alvizu *et al.*, "Comprehensive survey on t-sdn: Software-defined networking for transport networks," *IEEE CST*, vol. 19, no. 4, pp. 2232–2283, 2017.
- [12] R. Boutaba, N. Shahriar, and S. Fathi, "Elastic optical networking for 5G transport," *Springer JNSM*, vol. 25, no. 4, pp. 819–847, 2017.
- [13] N. Shahriar, R. Ahmed, S. R. Chowdhury, M. M. A. Khan, R. Boutaba, J. Mitra, and F. Zeng, "Connectivity-aware virtual network embedding," in *IFIP Networking Conference*. IEEE, 2016, pp. 46–54.
- [14] N. Shahriar *et al.*, "Virtual network embedding with guaranteed connectivity under multiple substrate link failures," *IEEE Trans. on Commun.*, vol. 68, no. 2, pp. 1025–1043, 2020.
- [15] D. Stamatelakis and W. D. Grover, "IP layer restoration and network planning based on virtual protection cycles," *IEEE JSAC*, vol. 18, no. 10, pp. 1938–1949, 2006.
- [16] A. E. Kamal *et al.*, "Overlay protection against link failures using network coding," *IEEE/ACM TNet*, vol. 19, no. 4, pp. 1071–1084, 2011.
- [17] A. Haider and R. J. Harris, "Recovery techniques in next generation networks," *IEEE CST*, vol. 9, no. 1-4, pp. 2–17, 2007.
- [18] N. Shahriar, R. Ahmed, S. R. Chowdhury, A. Khan, R. Boutaba, J. Mitra, and L. Liu, "Joint backup capacity allocation and embedding for survivable virtual networks," in *IFIP Networking Conference*, 2017.
- [19] N. Shahriar *et al.*, "Virtual network survivability through joint spare capacity allocation and embedding," *IEEE JSAC*, vol. 36, no. 3, pp. 502–518, 2018.
- [20] Y. Chen, J. Li, T. Wo, C. Hu, and W. Liu, "Resilient virtual network service provision in network virtualization environments," in *IEEE 16th ICPADS*, 2010, pp. 51–58.
- [21] N. Shahriar, R. Ahmed, A. Khan, S. R. Chowdhury, R. Boutaba, and J. Mitra, "ReNoVatE: Recovery from node failure in virtual network embedding," in *IEEE/ACM CNSM*, 2016, pp. 19–27.
- [22] N. Shahriar *et al.*, "Generalized recovery from node failure in virtual network embedding," *IEEE TNSM*, vol. 14, no. 2, pp. 261–274, 2017.
- [23] Q. Hu, Y. Wang, and X. Cao, "Survivable network virtualization for single facility node failure: A network flow perspective," *Optical Switching and Networking*, vol. 10, no. 4, pp. 406–415, 2013.
- [24] H. Yu *et al.*, "Cost efficient design of survivable virtual infrastructure to recover from facility node failures," in *IEEE ICC*, 2011, pp. 1–6.
- [25] L. Bo, T. Huang, X.-c. SUN, J.-y. CHEN, and Y.-j. LIU, "Dynamic Recovery for Survivable Virtual Network Embedding," *Journal of China Universities of Posts and Telecommunications*, vol. 21, pp. 77–84, 2014.
- [26] S. Ramamurthy and B. Mukherjee, "Survivable wdm mesh networks. part i-protection," in *IEEE INFOCOM*, vol. 2, 1999, pp. 744–751.
- [27] N. Shahriar, S. Taeb, S. R. Chowdhury, M. Zulfiqar, M. Tornatore, R. Boutaba, J. Mitra, and M. Hemmati, "Reliable slicing of 5G transport networks with dedicated protection," in *IEEE/ACM CNSM*, 2019.
- [28] N. Shahriar *et al.*, "Reliable slicing of 5G transport networks with bandwidth squeezing and multi-path provisioning," *IEEE TNSM*, 2020.
- [29] G. Zhang *et al.*, "A survey on OFDM-based elastic core optical networking," *IEEE CST*, vol. 15, no. 1, pp. 65–87, 2012.
- [30] A. Pagès *et al.*, "Optimal route, spectrum, and modulation level assignment in split-spectrum-enabled dynamic elastic optical networks," *IEEE/OSA JOCN*, vol. 6, no. 2, pp. 114–126, 2014.
- [31] N. Shahriar *et al.*, "Achieving a fully-flexible virtual network embedding in elastic optical networks," in *IEEE INFOCOM*, 2019, pp. 1756–1764.