

Generalized Recovery From Node Failure in Virtual Network Embedding

Nashid Shahriar, Reaz Ahmed, Shihabur Rahman Chowdhury, *Student Member, IEEE*, Aimal Khan, Raouf Boutaba, *Fellow, IEEE*, and Jeebak Mitra

Abstract—Network virtualization has evolved as a key enabling technology for offering the next generation network services. Recently, it is being rolled out in data center networks as a means to provide bandwidth guarantees to cloud applications. With increasing deployments of virtual networks (VNs) in commercial-grade networks with commodity hardware, VNs need to tackle failures in the underlying substrate network. In this paper, we study the problem of recovering a batch of VNs affected by a substrate node failure. The combinatorial possibilities of alternate embeddings of the failed virtual nodes and links of the VNs make the task of finding the most efficient recovery both non-trivial and intractable. Furthermore, any recovery approach ideally should not cause any service disruption for the unaffected parts of the VNs. We take into account these issues to design a generalized recovery approach that can achieve customized objectives such as fair treatment on the failed VNs, partial treatment based on priority, and so on. We provide integer linear programming (ILP) formulations for two variants of our recovery scheme, namely, fair recovery model and priority-based recovery model. We also propose a fast and scalable heuristic algorithm to tackle the computational complexity of the ILP solution. Evaluation results demonstrate that our heuristic performs close to the optimal solution and outperforms the state-of-the-art algorithm.

Index Terms—Network survivability and resilience, network virtualization, node failure, optimization techniques, proactive and reactive management, recovery approach, virtual network embedding.

I. INTRODUCTION

RAPID proliferation of the Internet is continuously increasing our dependence on networked services. Consequently, diverse Quality of Service (QoS) guarantees are required from the underlying network infrastructure. Network Virtualization (NV) [1] is evolving as a key technology for allowing a wide variety of online services with diverse

reliability and performance requirements to co-exist and seamlessly operate on top of the same network infrastructure. An Infrastructure Provider (InP) manages a network infrastructure also known as the Substrate Network (SN), and leases network slices in the form of Virtual Networks (VN) to multiple Service Providers (SPs). An SP offers customized services on top of its VN, and is free to deploy any technology and/or communication protocol in the VN. By allowing heterogeneous VNs to coexist on a shared SN, the goal of NV is to provide flexibility, diversity, security, and manageability. Several new challenges need to be addressed to achieve these goals.

An important challenge in NV is to efficiently allocate substrate resources to VNs. This is known as the VN embedding (VNE) problem [2] that maps virtual nodes and links of a VN request on substrate nodes and paths (a sequence of substrate links), respectively, while satisfying physical resource constraints. The VNE problem is \mathcal{NP} -hard and has been studied from various perspectives [2]. One particular aspect of VNE that has received much attention recently is Survivable Virtual Network Embedding (SVNE). SVNE approaches deal with substrate resource (i.e., nodes or links) failures that are not a rare event in large networks [3], [4]. Surviving failures is even more challenging in NV, since the shared nature of VNs exposes them to a more vulnerable state than that of a non-virtualized network. For instance, a link failure in the SN may cause multiple virtual links to fail, which may significantly degrade service performance and reliability of VNs.

A number of mechanisms have been proposed to increase VN reliability against substrate resource failures. These mechanisms can be broadly classified into two categories [5]: a) proactively provision disjoint redundant resources as backup [6], [7] and b) reactively re-embed the failed nodes and links of a VN on the available resources after a failure has occurred [8], [9]. Proactive approaches offer immediate recovery from failures at the expense of backup resource reservation [10]–[12]. However, preallocating backup resources for multiple failures resulting from a substrate node failure can be extremely expensive [11], [13]. Instead, an SP may prefer to reactively re-embed the failed part of its VN to avoid the huge cost of preallocated backup resources in a failure-prone SN. The SP can adopt a load balanced VN embedding strategy to leave higher amount of available resources during re-embedding. Moreover, in case of permanent substrate node failures (e.g., hardware malfunction), the nodes and links of all the affected VNs have to be re-embedded. These scenarios

Manuscript received January 3, 2017; revised April 2, 2017; accepted April 3, 2017. Date of publication April 13, 2017; date of current version June 9, 2017. This work was supported in part by Huawei Technologies and in part by an NSERC Collaborative Research and Development Grant. Additionally, this work benefited from the use of the CrySP RIPPLE Facility at the University of Waterloo. The associate editor coordinating the review of this paper and approving it for publication was H. Lutfiyya. (*Corresponding author: Nashid Shahriar.*)

N. Shahriar, R. Ahmed, S. R. Chowdhury, A. Khan, and R. Boutaba are with the School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: nshahria@uwaterloo.ca; r5ahmed@uwaterloo.ca; srchowdhury@uwaterloo.ca; a273khan@uwaterloo.ca; rboutaba@uwaterloo.ca).

J. Mitra is with Huawei Technologies Canada Research Center, Ottawa, ON K2K 3J1, Canada (e-mail: jeebak.mitra@huawei.com).

Digital Object Identifier 10.1109/TNSM.2017.2693404

motivate us to study the problem of re-embedding a batch of VNs affected by a substrate node failure.

While VN embedding is already intractable, the combinatorial number of sequences of VNs in a batch re-embedding further increases the complexity [14]. In addition, any solution must be significantly fast (*e.g.*, in the order of milliseconds) to meet the stringent timing requirement usually imposed by Service Level Agreements (SLAs). To meet such requirements, we consider the re-embedding of only the failed nodes and links of a batch of VNs without disrupting their unaffected parts. While re-embedding the failed part as well as the unaffected part of an active VN may incur lower costs of re-embedding [8], [15], it will require additional time for virtual node migration and virtual link re-configuration, which may increase VN downtime. Further, the reactive approaches in [9] and [16] opt for the recovery of all the failed nodes and links of a VN. If the SN does not have adequate resources to recover all the failed nodes and links, those approaches re-embed the complete VN, turning it inactive for a while and causing service disruptions.

Unlike existing approaches, we adopt a generalized recovery approach that allows partial recovery of an affected VN, while adhering to any SLA requirement. To demonstrate the versatility of our approach, we investigate 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. Then, we explore 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. It is worth discussing here that a partial recovery scheme may lead to skewness in the utilization of SN resources. To eliminate such skewness, VN reconfiguration mechanisms such as [17]–[20] should be performed periodically during an off-peak period.

In this paper, we focus on the problem of **R**ecovering from a **N**ode failure in **V**irtual **N**etwork **E**mbedding (*ReNoVatE*). It accepts a batch of VN failures resulting from a single substrate node failure, and produces alternate embeddings for the failed virtual nodes and links. The objective of FRM is to maximize the number of recovered virtual links across all the affected VNs, while minimizing total bandwidth required for recovery. On the other hand, PRM seeks to prioritize the recovery of affected VNs based on some predefined requirements and minimize total bandwidth for recovery. We formulate *ReNoVatE* as an Integer Linear Programming (ILP) based optimization model, namely *Opt-ReNoVatE*. Since *Opt-ReNoVatE* cannot scale to large instances of the problem, we devise an efficient heuristic algorithm, called *Fast-ReNoVatE*, to find satisfactory solutions within prescribed time limits. We evaluate *Fast-ReNoVatE* by extensive simulations and compare it with *Opt-ReNoVatE*, as well as with the most related state-of-the-art proposal in [9]. Evaluation results demonstrate that *Fast-ReNoVatE* performs close to *Opt-ReNoVatE* and outperforms the state-of-the-art solution in terms of i) number of

recovered virtual links, ii) cost of recovery, and iii) execution time.

This paper extends our initial work presented in [21] in several aspects. First, we modify our previous recovery approach to accommodate a variety of design alternatives such as prioritizing the affected VNs based on SLA requirements or impacts of failure or profits. Based on this modification, we propose two different recovery models, namely, fair recovery model and priority-based recovery model. Since these models are orthogonal to our proposed solutions, we incorporate the models in each of our solutions and evaluate them through rigorous simulations. Second, we perform more simulations and present additional results on the scalability aspects of our solutions to *ReNoVatE* and the state-of-the-art solution. Finally, we update related work section to include an in-depth discussion on the subtle differences between *ReNoVatE* and the state-of-the-art literature on SVNE.

The rest of the paper is organized as follows. Section II presents the related literature. In Section III, we present the system model and formally introduce the problem. Section IV presents two variants of *Opt-ReNoVatE* for solving the problem optimally. The heuristic solution, *Fast-ReNoVatE*, is presented in Section V. Evaluation results are presented in Section VI. Finally, we summarize our findings and conclude in Section VII.

II. RELATED WORK

Survivability of VNs during substrate failures was first addressed by Rahman *et al.* [22]. They formulated the problem of ensuring survivability in VNs under single SLink failure as a Mixed Integer Linear Program and proposed heuristics to obtain solutions in a reasonable time. A large body of research literature has been developed since to address different aspects of SVNE such as considering single SLink failure [23], [24], multiple SLink failures [25], single SNode failure [13], [26], regional failures [27]–[29], and ensuring certain levels of availability [30]–[32] among others. The approaches for ensuring survivability of VNs during substrate failures can be broadly classified into two categories, namely, proactive and reactive approaches. Proactive approaches provision redundant backup resources when a VN is embedded in the first place, whereas, reactive approaches take mitigation actions after a failure has occurred. In the following we discuss the proactive (Section II-A) and reactive (Section II-B) approaches from SVNE literature and contrast them with our solution for *ReNoVatE*.

A. Pro-Active Approaches

Pro-active approaches for SVNE preallocates backup resources to ensure Quality of Service for the VNs during one or more substrate failures. Most of the approaches focusing on substrate node (SNode) failure are pro-active [10], [11], [13]. Yu *et al.* [13] proposed a two-step method to recover a VN. The first step enhances the VN with backup virtual nodes (VNodes) and virtual links (VLinks), and the second step maps this enhanced VN on the SN. This approach, in the worst case, has to reserve a backup VNode for each VNode.

In contrast, [10] designed the enhanced VN with a failure-dependent strategy to reduce backup resources. Despite the resource efficiency of this approach, it is not practical due to the large number of migrations of working VNodes. Unlike these methods, [11] presented a joint optimization strategy for allocating primary and backup resources altogether. The location constrained SVNE, to address geographically-correlated SNode failures, has been studied in [12], [27], and [33]. While [33] adopts sequential embedding of working and backup VNodes, [12], [27] embeds them jointly to minimize total bandwidth. Recently, [34] and [35] proposed embedding primary and dedicated backup resources for each VNode and VLink in a VN simultaneously. There is a growing trend towards designing survivable resource allocation schemes for embedding virtual data centers (VDCs) in cloud [36]–[38]. For instance, [39] proposed a scheme for provisioning VDCs with backup virtual machines and links. Bodík *et al.* [40] proposed an optimization framework for improving survivability, while reducing total bandwidth consumption. Yeow *et al.* [38] defines the reliability level as a function of backup resources that are shared between VNs through opportunistic pooling. Finally, [31] provided a VDC embedding framework for achieving high VDC availability by considering heterogeneous failure rates.

Preallocated backup resources remain unused during normal operations and hence reduce resource utilization. In contrast, ReNoVatE takes a reactive approach and allocates resources only when a failure has occurred, thus improving on the resource utilization.

B. Reactive Approaches

Reactive approaches, on the other hand, do not preallocate any backup resources. Chang *et al.* [8] proposed a migration aware VN re-embedding algorithm to recover from an SNode failure. The algorithm allows the migration of some active VNodes and VLinks to free up some substrate resources, thus facilitating the re-mapping of the failed VNodes and VLinks. Cai *et al.* [15] addressed the problem of optimally upgrading the existing VN in a highly evolving SN. Their goal is to minimize the upgrading cost, in addition to migration and remapping costs, while satisfying QoS constraints. Both of these approaches may need a chain of migrations to converge, thus disrupting ongoing communication in the VN. Reactive approaches to recover from geographically correlated failures has been proposed in [28] and [29]. However, unlike ReNoVatE, [28], [29] allow the VNs to operate at a degraded Quality of Service. Recently, reactive recovery approaches have been proposed for multi-cast VNs. Unlike traditional VNs, multi-cast VNs take the Quality of Service requirements of multi-cast services (*e.g.*, bounded delay for multi-cast completion) into account. Ayoubi *et al.* [41] has shown that restoring multi-cast trees with delay constraints during a single substrate node failure is NP-hard for general graphs. Polynomial time heuristic algorithms for restoring multi-cast trees in VNs during single substrate node, single substrate link and multiple substrate node and link failures have been proposed in [42]–[44], respectively. In contrast to

the reactive approaches for restoring multi-cast VNs, our proposal does not make any assumption about the structure of SN and communication pattern in the VN. Another line of works [22], [45], originated in optical networks, proposed a hybrid mechanism to select from a set of precomputed detours for recovery during failure. However, in a highly saturated SN, this mechanism may not find adequate resources left for the recovery.

Distributed reactive approaches including, [16] and [46], proposed multi-agent based algorithms to dynamically adapt the VNs in response to SN failures. When an agent detects a failure of another agent in the same cluster, the agents within the same cluster collaborate with each other to re-provision the failed VNodes and VLinks. These approaches may generate sub-optimal solutions due to the lack of global knowledge of the SN. Finally, [9] proposed a greedy algorithm to find alternate substrate resources for the affected part of a VN. In case of resource inadequacy, this approach requires remapping the entire VN resulting in prolonged service unavailability. Nonetheless, existing reactive algorithms focus on re-embedding only a single VN, whereas *ReNoVatE*, for the first time, considers partial re-embedding of a set of affected VNs to improve recovery performance.

III. SYSTEM MODEL AND PROBLEM STATEMENT

We first present a mathematical representation of a substrate network, virtual network, and types of failure (Section III-A). Then, we formally define the problem (Section III-B). A glossary of key notations used in the paper is provided in Table I.

A. System Model

1) *Substrate Network (SN)*: We represent the SN as an undirected graph, $G = (V, E)$, where V and E denote the set of SNodes and Substrate Links (SLinks), respectively. The set of neighbors of an SNode $u \in V$ is denoted by $\mathcal{N}(u)$. Bandwidth capacity and residual bandwidth of an SLink $(u, v) \in E$ are represented by b_{uv} and r_{uv} , respectively, while the cost of allocating one unit of bandwidth in (u, v) is C_{uv} . V^f and E^f represent the set of failed SNodes and SLinks, respectively. P_{uv} represents a path between SNodes u and v .

2) *Virtual Network (VN)*: 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. The set of neighbors of a VNode $\bar{u} \in \bar{V}_i$ is denoted by $\mathcal{N}(\bar{u})$. Each VLink $(\bar{u}, \bar{v}) \in \bar{E}_i$ has a bandwidth demand $b_{\bar{u}\bar{v}}$. We associate a penalty $\pi_{\bar{u}\bar{v}}$ to each VLink $(\bar{u}, \bar{v}) \in \bar{E}_i$, where $\pi_{\bar{u}\bar{v}}$ represents the penalty due to resource 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 $u \in L_i(\bar{u})$. We represent this location constraint with a binary variable $\ell_{\bar{u}u}$, defined as:

$$\ell_{\bar{u}u} = \begin{cases} 1 & \text{iff } \bar{u} \in \bar{V}_i \text{ can be provisioned on } u \in V, \\ 0 & \text{otherwise.} \end{cases}$$

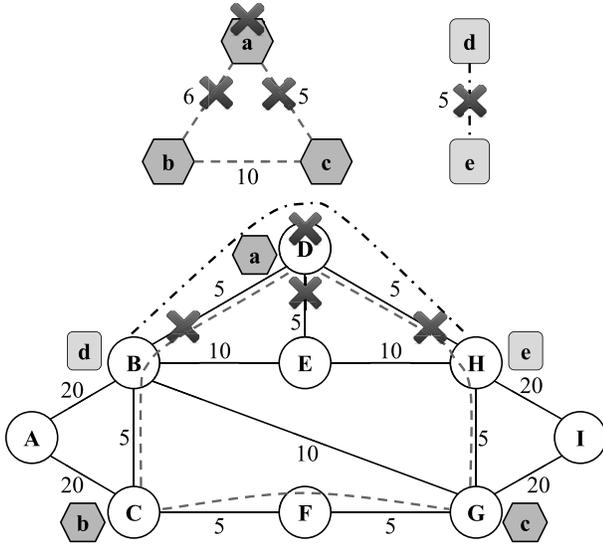


Fig. 1. VN embedding and impact of failure.

3) *Types of Failure*: 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 .

Adjacent VLinks: The set of VLinks adjacent to the failed VNode $\bar{u} \in \bar{V}_i^f$ is represented by $\bar{E}_i^f = \{(\bar{u}, \bar{v}) | \bar{u} \in \bar{V}_i^f \wedge \bar{v} \in \mathcal{N}(\bar{u})\}$.

Independent VLinks: The set of VLinks that have failed due to the failure of some SLinks on their mapped substrate paths is denoted by $\bar{E}_i^f = \{(\bar{u}, \bar{v}) | (u, v) \in g(\bar{u}\bar{v}) \wedge (u, v) \in E^f \wedge \bar{u} \notin \bar{V}_i^f \wedge \bar{v} \notin \bar{V}_i^f\}$.

Finally, $\bar{V}^f = \{\cup \bar{V}_i^f\}$, $\bar{E}^f = \{\cup \bar{E}_i^f\}$, and $\bar{E}^f = \{\cup \bar{E}_i^f\}$ represent the set of failed VNodes, VLinks, and Independent VLinks of all the VNs in \bar{G} , respectively. Fig. 1 illustrates the embedding of two VNs, \bar{G}_1 with $\bar{V}_1 = \{a, b, c\}$ and \bar{G}_2 with $\bar{V}_2 = \{d, e\}$ on the SN G shown in the bottom. The numbers next to a VLink and an SLink represent the VLink demand and SLink residual bandwidth, respectively. The VNode mapping, i.e., $f(\cdot)$ is shown by placing a VNode beside its mapped SNode and the VLink mapping, i.e., $g(\cdot)$ is depicted by dashed paths between SNodes. For instance, $f(a) = \{D\}$, $f(b) = \{C\}$, $f(c) = \{G\}$, $f(d) = \{B\}$, $f(e) = \{H\}$ and $g(ab) = \{DB, BC\}$, $g(ac) = \{DH, HG\}$, $g(bc) = \{CF, FG\}$, and $g(de) = \{BD, DH\}$. We now show the impact of an SNode failure with $V^f = \{D\}$ and $E^f = \{DB, DE, DH\}$. VN \bar{G}_1 experiences a VNode failure with $\bar{V}_1^f = \{a\}$. Consequently, the VLinks adjacent to a fail leading to $\bar{E}_1^f = \{ab, ac\}$. Note there is no Independent VLink failures in \bar{G}_1 , and so $\bar{E}_1^f = \bar{E}_1^f$. On the other hand, $\bar{V}_2^f = \emptyset$ since no VNode of \bar{G}_2 is mapped on D . However, the failure of D results in an independent VLink failure yielding $\bar{E}_2^f = \bar{E}_2^f = \{de\}$. Hence, any recovery algorithm should re-embed all the affected VNodes and VLinks from \bar{G}_1 and \bar{G}_2 leaving unaffected part such as VLink bc undisrupted.

TABLE I
SUMMARY OF KEY NOTATIONS

$G = (V, E)$	Substrate Network
b_{uv}	Bandwidth capacity of SLink $(u, v) \in E$
r_{uv}	Residual bandwidth of SLink $(u, v) \in E$
V^f	Set of failed SNodes
E^f	Set of failed SLinks
C_{uv}	Cost of unit bandwidth on SLink $(u, v) \in E$
$\bar{G}_i = (\bar{V}_i, \bar{E}_i)$	i -th Virtual Network Request
$b_{i\bar{u}\bar{v}}$	Bandwidth requirement of virtual link $(\bar{u}, \bar{v}) \in \bar{E}_i$
$L_i(\bar{u})$	Location constraint set for virtual node $\bar{u} \in \bar{V}_i$
$\ell_{i\bar{u}u} \in \{0, 1\}$	$\ell_{i\bar{u}u} = 1$ if $u \in L_i(\bar{u})$, $u \in V$, $\bar{u} \in \bar{V}_i$
$f(\bar{u}) \in V$	Embedding of VNode \bar{u}
$g(\bar{u}\bar{v}) \in 2^E$	Embedding of VLink (\bar{u}, \bar{v})
\bar{V}_i^f	Set of failed VNodes from VN \bar{V}_i
\bar{E}_i^f	Set of failed VLinks from VN \bar{V}_i
\bar{E}_i^f	Set of failed VLinks adjacent to a failed VNode
\bar{E}_i^f	Set of VLinks not adjacent to a failed VNode
$x_{i\bar{u}\bar{v}} \in \{0, 1\}$	$x_{i\bar{u}\bar{v}} = 1$ if $(u, v) \in E$ is on the embedded substrate path for $(\bar{u}, \bar{v}) \in \bar{E}_i$
$y_{i\bar{u}u} \in \{0, 1\}$	$y_{i\bar{u}u} = 1$ if $\bar{u} \in \bar{V}_i$ is mapped to $u \in V$
$z_{i\bar{u}\bar{v}} \in \{0, 1\}$	$z_{i\bar{u}\bar{v}} = 1$ iff $(\bar{u}, \bar{v}) \in \bar{E}_i^f$ is mapped to any substrate path
$\pi_{i\bar{u}\bar{v}}$	Penalty due to resource unavailability of $(\bar{u}, \bar{v}) \in \bar{E}_i^f$

B. Problem Statement

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 such that the re-embedding achieves the following objectives:

- Primary objective differs based 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.
 - Secondary objective is to minimize the total cost of re-embedding in terms of SLink bandwidth consumption.
- 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.

IV. ILP FORMULATION: *Opt-ReNoVatE*

We provide an ILP formulation, *Opt-ReNoVatE*, based on the *Multi-commodity Flow Problem* formulation of *ReNoVatE*. We first present the decision variables (Section IV-A). Then, we introduce the constraints (Section IV-B) followed by the objective functions of the two variants of *Opt-ReNoVatE* (Section IV-C and Section IV-D). Finally, we describe the complexity of the problem in Section IV-E.

A. Decision Variables

The following decision variables indicate VNode and VLink embedding of a VN $\bar{G}_i \in \bar{G}$ on an SN G .

$$y_{\bar{u}u} = \begin{cases} 1 & \text{iff } \bar{u} \in \bar{V}_i \text{ is mapped to } u \in V, \\ 0 & \text{otherwise.} \end{cases}$$

$$x_{\bar{u}\bar{v}}^{uv} = \begin{cases} 1 & \text{iff } (\bar{u}, \bar{v}) \in \bar{E}_i \text{ is mapped to } (u, v) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

The objective of *Opt-ReNoVatE* is to recover as many failed VLinks in \bar{E}^f as possible to mitigate the impact of failure. It may be possible that not all VLinks in \bar{E}^f can be re-embedded due to substrate resource limitation. The following decision variable defines which VLinks are re-embedded:

$$z_{\bar{u}\bar{v}} = \begin{cases} 1 & \text{iff } (\bar{u}, \bar{v}) \in \bar{E}_i^f \text{ is mapped to any substrate path} \\ 0 & \text{otherwise.} \end{cases}$$

B. Constraints

1) *Intactness of Unaffected VNodes and VLinks*: The mapping of VNodes and VLinks that are not affected by the substrate failure remains unchanged. Constraints (1) and (2) ensure that unaffected VNodes and VLinks are not re-embedded.

$$\forall \bar{G}_i \in \bar{G}, \forall \bar{u} \in \bar{V}_i \setminus \bar{V}_i^f : y_{\bar{u}f(\bar{u})} = 1 \quad (1)$$

$$\forall \bar{G}_i \in \bar{G}, \forall (\bar{u}, \bar{v}) \in \bar{E}_i \setminus \bar{E}_i^f, \forall (u, v) \in g(\bar{u}\bar{v}) : x_{uv}^{\bar{u}\bar{v}} = 1. \quad (2)$$

2) *Exclusion of Failed SNodes and SLinks From Re-Embedding*: The failed VNodes or VLinks cannot use any of the failed SNodes or SLinks during re-embedding. Constraint (3) ensures that the failed VNodes are not re-embedded on the failed SNodes, and (4) ensures that the failed VLinks are not re-embedded on substrate paths containing a failed SLink.

$$\forall \bar{G}_i \in \bar{G}, \forall \bar{u} \in \bar{V}_i^f, \forall u \in V_f : y_{\bar{u}u} = 0 \quad (3)$$

$$\forall \bar{G}_i \in \bar{G}, \forall (\bar{u}, \bar{v}) \in \bar{E}_i^f, \forall (u, v) \in E_f : x_{uv}^{\bar{u}\bar{v}} = 0. \quad (4)$$

3) *Link Mapping Constraints*: Constraint (5) prevents overcommitment of SLink bandwidth. Constraint (6) ensures that the in-flow and out-flow of each SNode is equal except at the SNodes where the endpoints of a failed VLink are embedded. Finally, constraint (7) ensures that if a VLink (\bar{u}, \bar{v}) is selected to be re-embedded due to the failure of \bar{u} , there is some flow from the SNode u where \bar{v} is embedded already.

$$\forall (u, v) \in E : \sum_{\bar{G}_i \in \bar{G}} \sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i} x_{uv}^{\bar{u}\bar{v}} \times b_{\bar{u}\bar{v}} \leq b_{uv} \quad (5)$$

$$\forall \bar{G}_i \in \bar{G}, \forall (\bar{u}, \bar{v}) \in \bar{E}_i^f, \forall u \in V \setminus f(\bar{v}) : \sum_{\forall v \in \mathcal{N}(u)} (x_{uv}^{\bar{u}\bar{v}} - x_{vu}^{\bar{u}\bar{v}}) \leq y_{\bar{u}u} - y_{\bar{v}v} \quad (6)$$

$$\forall \bar{G}_i \in \bar{G}, \forall (\bar{u}, \bar{v}) \in \bar{E}_i^f, \forall u \in f(\bar{v}) : \sum_{\forall v \in \mathcal{N}(u)} (x_{uv}^{\bar{u}\bar{v}} - x_{vu}^{\bar{u}\bar{v}}) = z_{\bar{u}\bar{v}}. \quad (7)$$

4) *Node Mapping Constraints*: First, constraint (8) ensures that re-embedding of a failed VNode should be done according to the provided location constraint set. Second, constraint (9) makes sure that a VNode should be mapped to at most an SNode in the SN. Third, constraint (10) enforces that an SNode will not host more than one VNodes from the same VN. Finally, constraint (11) ensures that if a VLink $(\bar{u}, \bar{v}) \in \bar{E}_i^f$ is selected to be re-embedded due to the failure of \bar{u} , the VNode \bar{u} must be re-embedded on an SNode according to the location constraint. Here, λ is a very large integer that turns the left side of (11) into a fraction between 0 and 1 when any of the $z_{\bar{u}\bar{v}}$ is 1. This enforces the right side of (11) to become 1, thus ensuring the failed VNode to be re-embedded.

$$\forall \bar{G}_i \in \bar{G}, \forall \bar{u} \in \bar{V}_i^f, \forall u \in V : y_{\bar{u}u} \leq \ell_{\bar{u}u} \quad (8)$$

$$\forall \bar{G}_i \in \bar{G}, \forall \bar{u} \in \bar{V}_i^f : \sum_{u \in V} y_{\bar{u}u} \leq 1 \quad (9)$$

$$\forall \bar{G}_i \in \bar{G}, \forall u \in V : \sum_{\bar{u} \in \bar{V}_i} y_{\bar{u}u} \leq 1 \quad (10)$$

$$\forall \bar{G}_i \in \bar{G}, \forall \bar{u} \in \bar{V}_i^f : \frac{1}{\lambda} \sum_{\bar{v} \in \mathcal{N}(\bar{u})} z_{\bar{u}\bar{v}} \leq \sum_{\forall u \in V} y_{\bar{u}u}. \quad (11)$$

C. Fair Recovery Model

The fair recovery model (FRM) treats all the failed VLinks equally while recovering them. Hence, the objective of this model is to recover as many failed VLinks as possible. Following the problem statement, the objective function (12) of FRM has two components. The first and primary component maximizes the number of re-embedded failed VLinks. In other words, it minimizes the total number of un-recovered VLinks as shown in (12). It is obvious to observe that the minimum value of the primary component is zero. However, there can be multiple solutions that achieve the minimum value. Hence, we need a secondary component in the objective function to break ties among multiple solutions having the same value for the primary objective. Therefore, the secondary component minimizes the total cost of provisioning bandwidth for re-embedding the failed VLinks on substrate paths. A weight factor w is multiplied to the second component to impose the relative weight to the components of (12). The value of w is chosen to be a very small fraction so that it comes into effect only to break ties among multiple solutions that have the same value for the primary objective. In this way, w prefers the number of recovered VLinks over the cost of re-embedding.

$$\text{minimize} \left(|\bar{E}^f| - \sum_{\bar{G}_i \in \bar{G}} \sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} z_{\bar{u}\bar{v}} \right) + w \left(\sum_{\bar{G}_i \in \bar{G}} \sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} \sum_{\forall (u, v) \in E} x_{uv}^{\bar{u}\bar{v}} \times C_{uv} \times b_{\bar{u}\bar{v}} \right). \quad (12)$$

D. Priority-Based Recovery Model

The priority-based recovery model (PRM) prioritizes some failed VLinks over others to satisfy the SLA requirements or to minimize the impact of failure or to maximize profits.

To impose such priorities, we utilize the penalty parameter $\pi_{i\bar{u}\bar{v}}$ associated to each VLink $(\bar{u}, \bar{v}) \in \bar{E}_i^f$. Each $\pi_{i\bar{u}\bar{v}}$ can be given as an input based on SLA requirement violation penalty or can be computed based on the impact of failure. For instance, a VLink with strict SLA requirement may have a higher value of the penalty than that of a VLink with less-stringent SLA requirement. The objective of PRM is to minimize the total penalty across all the failed VLinks. Similar to the FRM case, the objective function (13) has two components. The primary component minimizes the total penalty incurred by the VLinks that are not recovered. The second one minimizes the total cost of provisioning bandwidth for re-embedding the failed VLinks on substrate paths. A weight factor w is multiplied to the second component to impose the relative weight to the components of (13). The value of w is chosen to be a very small fraction so that it comes into effect only to break ties among multiple solutions that have the same value for the primary objective. In this way, w prioritizes minimizing total penalty over the cost of re-embedding.

$$\begin{aligned} & \text{minimize} \left(\sum_{\forall \bar{G}_i \in \bar{G}} \sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} (1 - z_{i\bar{u}\bar{v}}) \times \pi_{i\bar{u}\bar{v}} \right) \\ & + w \left(\sum_{\forall \bar{G}_i \in \bar{G}} \sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} \sum_{\forall (u, v) \in E} x_{uv}^{i\bar{u}\bar{v}} \times C_{uv} \times b_{i\bar{u}\bar{v}} \right) \quad (13) \end{aligned}$$

Note that, if we assume $\pi_{i\bar{u}\bar{v}} = 1$, $\forall \bar{G}_i \in \bar{G}, \forall (\bar{u}, \bar{v}) \in \bar{E}_i^f$, (13) reduces to (12). Therefore, objective function of PRM, *i.e.*, (13) encompasses objective functions of both FRM and PRM, and can be considered as the generalized recovery model.

E. Complexity of the Problem

The binary nature of the decision variables and the flow constraints of *Opt-ReNoVatE* prevent any VLink from being mapped to multiple substrate paths. This restricts the re-embedding of Independent VLinks to the \mathcal{NP} -hard *Multi-commodity Unsplittable Flow Problem* [47]. The VNode re-embedding follows from VLink re-embedding since there are no costs associated with VNodes. Therefore, the re-embedding of a VNode and its adjacent VLinks of a VN becomes the \mathcal{NP} -hard *Single-source Unsplittable Flow Problem* with unknown source [48]. When there are a batch of affected VNs, computing the best sequence of VNs from a combinatorial number of sequences to maximize the number of recovered VLinks makes the problem computationally intractable.

V. HEURISTIC SOLUTION: *Fast-ReNoVatE*

Due to the intractability of *Opt-ReNoVatE*, we resort to a heuristic algorithm, *Fast-ReNoVatE*, to find feasible solutions in reasonable time. *Fast-ReNoVatE* re-embeds the failed VNodes and their Adjacent VLinks (Section V-A) and Independent VLinks (Section V-B) of the affected VNs efficiently.

Algorithm 1 VNodes Recovery

```

1: function VNODES RECOVERY( $G, \bar{G}, V^f, E^f, recovery - model$ )
2:   if  $recovery - model = FRM$  then
3:      $\bar{G} \leftarrow \text{Sort } \bar{G}_i \in \bar{G}$  in increasing order of
        $\sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} b_{i\bar{u}\bar{v}}$ 
4:     else if  $recovery - model = PRM$  then
5:        $\bar{G} \leftarrow \text{Sort } \bar{G}_i \in \bar{G}$  in decreasing order of
        $\sum_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} \pi_{i\bar{u}\bar{v}}$ 
6:     end if
7:      $max_{\mathcal{R}} \leftarrow 0$ 
8:     for all  $\bar{G}_i \in \bar{G}$  do
9:        $\bar{u} \leftarrow$  failed virtual node in  $\bar{G}_i$ ,  $best_{\bar{u}} \leftarrow NIL$ 
10:      for all  $l \in L_i(\bar{u}) \setminus V^f$  do
11:         $\mathcal{P}_l \leftarrow \text{Max-Paths}(G, \bar{G}_i, \bar{u}, l, \bar{E}_i^f, V^f, E^f)$ 
12:        if  $\mathcal{R}(\mathcal{P}_l) > max_{\mathcal{R}}$  or  $(\mathcal{R}(\mathcal{P}_l) = max_{\mathcal{R}} \text{ and } C(\mathcal{P}_l) \leq C(M))$  then
13:           $M \leftarrow \mathcal{P}_l$ ,  $best_{\bar{u}} \leftarrow l$ 
14:        end if
15:      end for
16:      if  $best_{\bar{u}} \neq NIL$  then
17:        map  $\bar{u}$  to  $best_{\bar{u}}$ 
18:         $\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f$  : map  $(\bar{u}, \bar{v})$  to  $M[(\bar{u}, \bar{v})]$ 
19:      end if
20:    end for
21: end function

```

A. Recovery of VNodes and Adjacent VLinks

The inputs to the VNode Recovery algorithm (Algorithm 1) comprise an SN G , a set of affected VNs \bar{G} that are embedded on G , the recovery model to be applied, and a set of failed SNodes V^f and failed SLinks E^f . The purpose of the recovery model is to select the proper model between the two possibilities, *i.e.*, FRM and PRM. Algorithm 1 initially sorts the affected VNs in \bar{G} and \bar{G} represents this sorted order. The sorting function depends on the recovery model being adopted. In FRM, the function sorts the affected VNs in \bar{G} in increasing order of the total lost bandwidth in their Adjacent VLinks. The total lost bandwidth of a VN, $\bar{G}_i \in \bar{G}$ is computed as the summation of the bandwidth demands for all the failed VLinks in \bar{E}_i^f of \bar{G}_i . On the other hand, PRM sorts the affected VNs in \bar{G} in decreasing order of the total penalty in their Adjacent VLinks. The total penalty of a VN, $\bar{G}_i \in \bar{G}$ is computed as the summation of the penalty $\pi_{i\bar{u}\bar{v}}$ of all the failed VLinks in \bar{E}_i^f of \bar{G}_i . Intuitively, the algorithm proceeds to recover the VNs in \bar{G} in the sorted order of \bar{G} to increase the number of recovered VLinks in FRM or to decrease the cumulative penalty in PRM. For each VN \bar{G}_i , the algorithm tries to re-embed the failed VNode $\bar{u} \in \bar{V}_i^f$ and the VLinks adjacent to \bar{u} , *i.e.*, \bar{E}_i^f to an SNode present in the location constraint set of \bar{u} , *i.e.*, $L_i(\bar{u})$ and to substrate paths, respectively. To accomplish this, it iterates over all the SNodes $l \in L_i(\bar{u}) \setminus V^f$, and selects the SNode, $best_{\bar{u}}$ that maximizes the cardinality of the set of

Algorithm 2 Max-Paths

```

1: function MAX-PATHS( $G, \bar{G}_i, \bar{u}, l, \bar{E}_i^f, V^f, E^f$ )
2:    $V \leftarrow V \setminus V^f \cup \{S\}, E \leftarrow E \setminus E^f$ 
3:    $\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f : E \leftarrow E \setminus E^f \cup \{(f(\bar{v}), S)\}$ 
4:    $max\_bw \leftarrow \max_{\forall (\bar{u}, \bar{v}) \in \bar{E}_i^f} \{b_{i\bar{u}\bar{v}}\}$ 
5:    $\forall (u, v) \in E$  s.t.  $u = S$  or  $v = S$  :
6:      $flow_{uv} \leftarrow 0, c_{uv} \leftarrow 1$ 
7:    $\forall (u, v) \in E$  s.t.  $r_{uv} > 0$  and  $u \neq S$  and  $v \neq S$  :
8:      $flow_{uv} \leftarrow flow_{vu} \leftarrow 0$ 
9:      $c_{uv} \leftarrow c_{vu} \leftarrow \lfloor \frac{r_{uv}}{max\_bw} \rfloor$ 
10:   $G_r \leftarrow G, \mathcal{P} \leftarrow \phi, \mathcal{P}_l \leftarrow \phi$ 
11:  while  $\exists$  augmenting path  $P$  from  $l$  to  $S$  in  $G_r$  do
12:     $C_f(P) \leftarrow \min_{\forall (u,v) \in P} \{c_{uv}\}$ 
13:    Augment  $C_f(P)$  units flow in  $G_r$  along  $P$ 
14:    Update residual capacity in  $G_r$  along  $P$ 
15:    Remove links with residual capacity  $\leq 0$ 
16:    for all  $P_i \in \mathcal{P}$  do
17:      if  $P$  cancels the flow along  $(u, v) \in P_i$  then
18:         $\tau \leftarrow$  Sub paths in  $P_i \setminus (u, v)$ 
19:         $\nu \leftarrow$  Sub paths in  $P \setminus (v, u)$ 
20:         $P \leftarrow$  Path formed by segments in  $\tau$  and  $\nu$ 
21:         $P_i \leftarrow$  Path formed by segments in  $\tau$  and  $\nu$ 
22:      end if
23:    end for
24:     $\mathcal{P} \leftarrow \mathcal{P} \cup P$ 
25:  end while
26:   $\forall P_i \in \mathcal{P}, \forall (\bar{u}, \bar{v}) \in \bar{E}_i^f$  :
27:     $P_i \leftarrow$  Path containing  $(f(\bar{v}), S)$ 
28:     $\mathcal{P}_l[(\bar{u}, \bar{v})] \leftarrow P_i \setminus (f(\bar{v}), S)$ 
29:  return  $\mathcal{P}_l$ 
30: end function

```

Algorithm 3 VLinks-Recovery

```

1: function VLINKS-RECOVERY( $G, \bar{E}^f, V^f, E^f, recovery - model$ )
2:    $\mathcal{P} \leftarrow \phi$ 
3:    $V \leftarrow V \setminus V^f, E \leftarrow E \setminus E^f$ 
4:   if  $recovery - model = FRM$  then
5:      $\bar{E}^f \leftarrow$  Sort  $(\bar{x}, \bar{y}) \in \bar{E}^f$  in increasing order of  $b_{i\bar{x}\bar{y}}$ 
6:   else if  $recovery - model = PRM$  then
7:      $\bar{E}^f \leftarrow$  Sort  $(\bar{x}, \bar{y}) \in \bar{E}^f$  in decreasing order of  $\pi_{i\bar{x}\bar{y}}$ 
8:   end if
9:   for all  $(\bar{x}, \bar{y}) \in \bar{E}^f$  do
10:     $\mathcal{P}[(\bar{x}, \bar{y})] \leftarrow MCP(G, f(\bar{x}), f(\bar{y}), b_{i\bar{x}\bar{y}})$ 
11:    map  $(\bar{x}, \bar{y})$  to  $\mathcal{P}[(\bar{x}, \bar{y})]$ 
12:   end for
13: end function

```

Algorithm 3 first sorts the VLinks in \bar{E}^f based on the recovery model, and \bar{E}^f represents this order. FRM sorts the VLinks in increasing order of their bandwidth demands, whereas PRM sorts them in decreasing order of their penalties. Such sorting orders help FRM (or, PRM) maximize (or, minimize) the number of recoveries (or, the total incurred penalties). According to this order, the algorithm computes alternate substrate path

for each VLink in $(\bar{x}, \bar{y}) \in \bar{E}^f$ in the subgraph induced by excluding the failed SNode and SLinks from G . For a particular VLink (\bar{x}, \bar{y}) , the algorithm finds the minimum cost path from $f(\bar{x})$ to $f(\bar{y})$ using the procedure *MCP*, and adds it to the set \mathcal{P} . The procedure *MCP* uses a modified version of *Dijkstra's shortest path* algorithm [50] to take into account SLink residual capacity and VLink demand while computing the minimum cost path.

To illustrate, Fig. 2 depicts the recovery of Independent VLink de of VN \bar{G}_2 , embedded on SN G according to Fig. 1. Algorithm 3 should find alternate substrate paths between B and H for VLink de . The max flow based heuristic may return substrate path between B and G (or, between H and C), and is thus inappropriate. Hence, Algorithm 3 recovers de through the minimum cost path $\{BG, GH\}$ that has sufficient bandwidth to satisfy the demand of de .

C. Running Time Analysis

The most expensive step in Algorithm 1 is the computation of the maximum paths using Algorithm 2. The core part of Algorithm 2 follows the steps of *Edmonds-Karp* Algorithm. If *Edmonds-Karp* Algorithm computes augmenting paths using Breadth-first Search, it runs in $O(|V||E|^2)$ time. Since Algorithm 2 is invoked $|L(\bar{u})|$ times for recovering a VN, and there are $|\bar{V}^f|$ number of affected VNs, total running time yields $O(|L(\bar{u})||\bar{V}^f||V||E|^2)$. In contrast, Algorithm 3 invokes *Dijkstra's shortest path* algorithm for each VLink in \bar{E}^f . Since *Dijkstra's shortest path* algorithm runs in $O(|E| + |V| \log |V|)$ time, Algorithm 3 requires $O(|\bar{E}^f|(|E| + |V| \log |V|))$ time.

VI. EVALUATION

In this section, we first present the compared approaches in Section VI-A followed by a description of the evaluation metrics in Section VI-B. Then, we describe the simulation setup in Section VI-C and VN embedding data generation method in Section VI-D. Finally, we present our evaluation results found through extensive simulation in Section VI-E.

A. Compared Algorithms

We first demonstrate the performance of FRM-based recovery by comparing *Opt-ReNoVatE* and *Fast-ReNoVatE* with an implementation of dynamic recovery [9], called *Dyn-Recovery*. For a fair comparison, we selected a related work that takes a similar approach as we did, i.e., a reactive recovery scheme that is executed centrally with a global knowledge about what are embedded on the substrate network. In addition, we exclude the last step of re-embedding the entire VN from the implementation of *Dyn-Recovery* in the event of resource inadequacy. Although we evaluate all three algorithms in small size networks, we cannot evaluate *Opt-ReNoVatE* for large networks because of the inherent complexity of ILP-solvers. Alternatively, we present *Fast-ReNoVatE's* baseline performance assuming the SLinks of an SN have infinite bandwidth, and refer to it as *Fast-ReNoVatE-INF*. We compare *Fast-ReNoVatE* with *Fast-ReNoVatE-INF* to demonstrate the impact of residual bandwidth and the possible partitioning in a substrate network on the recovery from an SNode

TABLE II
SUMMARY OF SIMULATION PARAMETERS

Scenario	Figure	SNodes	SLinks	VNodes/VN	VLinks/VN	VNs	SN Utilization	VLink BW	Total Failed VLinks
Small Scale	Fig. 3, 6	50	90	5	8	10-32	20% - 75%	10% of SLink	250 - 866
	Fig. 4	20-65	37-118	5	8	8-19	53%	15% of SLink	178 - 519
Large Scale	Fig. 5, 7	1000	1798	3-15	2-30	93-563	20% - 80%	10% of SLink	4712 - 13546

failure. Finally, we analyze the impact of priority-based recovery through a rigorous comparison between the two proposed recovery models, *i.e.*, FRM and PRM. Since these models are orthogonal to our solutions, *i.e.*, *Opt-ReNoVatE* and *Fast-ReNoVatE*, we compare the two models on each of our solutions. We differentiate between the two variants of *Opt-ReNoVatE* as *Opt-ReNoVatE-FRM* and *Opt-ReNoVatE-PRM*. Similarly, *Fast-ReNoVatE-FRM* and *Fast-ReNoVatE-PRM* represent the variants of *Fast-ReNoVatE* for FRM and PRM, respectively.

B. Performance Metrics

1) *Recovery Efficiency*: The fraction of successfully recovered VLinks over all failed VLinks expressed in percentage.

2) *Recovery Cost*: The average cost of provisioning bandwidth along a substrate path times the cost of allocating one unit bandwidth for re-embedding each failed VLink.

3) *Execution Time*: The time required for an algorithm to find the solution for all the VNs affected by an SNode failure.

4) *Normalized Penalty*: Total incurred penalty normalized with the total number of VLinks that remain unrecovered.

C. Simulation Setup

We implement *Opt-ReNoVatE* using IBM ILOG CPLEX C++ library; and *Fast-ReNoVatE* and *Dyn-Recovery* [9] using C++. We evaluate the algorithms on both small and large scale networks as summarized in Table II. For each problem instance in this table, we generate 5 random SNs by taking the number of SNodes and link-to-node ratio as inputs, and randomly creating SLinks between SNodes. Then, we generate a number of VNs for each SN, and embed the VNs on the SN following the procedure described in Section VI-D. We select a random SNode and its one hop neighbor SNodes as the location constraint set of a VNode. Afterwards, we simulate an SNode failure by removing each SNode in the SN one-by-one and execute the recovery algorithms being evaluated on the affected VNs. Finally, we measure the performance metrics of the compared algorithms by taking averages across all SNode failures for all 5 similar problem instances. The simulations are performed on a server with 2 Intel Xeon E5-2650 (8 cores @ 2.0GHz, each) processors and 256GB of RAM.

D. VN Embedding Method

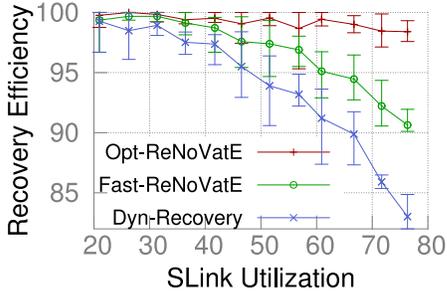
VN generation and embedding are done simultaneously to overcome the issue of creating VNs that do not have feasible embeddings. The embedding of VNs on an SN is done in a random but load-balanced manner to achieve uniform distribution of VNs across the SN. Load balancing has been considered as one of the key criteria for VN embedding in previous work [2]. Achieving higher utilization (beyond 70%)

of SN requires a denser embedding which is done by relaxing the load-balancing criteria. The embedding of a VN starts by randomly selecting a source SNode that has free capacity for a VNode. This forms the first VNode (source VNode) of the VN currently being embedded. For each randomly assigned neighbor of this VNode, a random destination SNode within several hop distances from the source SNode is selected as a candidate for embedding the subsequent VNode (destination VNode). The source and destination VNodes are joined using a VLink embedded on the shortest path between source and destination SNodes. If the path is not found, the destination VNode is moved to a different destination SNode. If the embedding is successful, the bandwidth demand of the VLink is generated and subtracted from the SLink residual capacity along the embedding path. To evaluate PRM, the penalty associated with the VLink is generated using a uniform random distribution between 1 and the total number of failed VLinks. In the case of FRM, the penalty of each VLink is set to 1 to enforce impartial recovery. These steps are repeated until all VNodes are embedded, generating both the VN and its embedding on the SN, simultaneously.

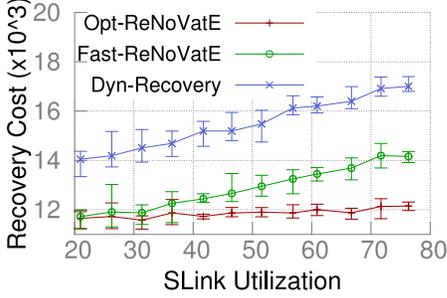
E. Results

1) *Small Scale Evaluations*: In small scale settings, we evaluate *Opt-ReNoVatE*, *Fast-ReNoVatE*, and *Dyn-Recovery* focusing on the following aspects: (i) by varying number of embedded VNs to achieve different SLink utilizations in the same SN (Fig. 3) (ii) by varying SN sizes while keeping the VN size and SLink utilization fixed (Fig. 4). We now present the impact of these aspects on our performance metrics.

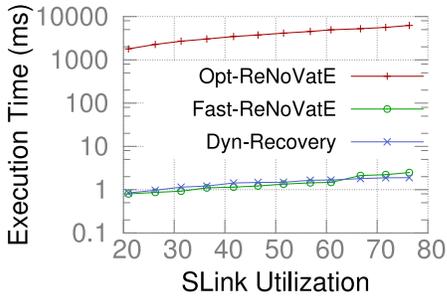
a) *Recovery efficiency*: Fig. 3(a) shows that the recovery efficiencies of *Opt-ReNoVatE*, *Fast-ReNoVatE*, and *Dyn-Recovery* decrease with the increase in SLink utilization. As the utilization increases, more VNs are affected by the SNode failure, and less bandwidth is left for recovery resulting in the gradual decrease in the number of recovered VLinks. Further, the impact of utilization is more profound in the higher utilization cases due to the lack of load balanced embedding. In the higher utilized cases, the recovery efficiencies of *Fast-ReNoVatE* is $\sim 6\%$ better than those of *Dyn-Recovery* and $\sim 3\%$ worse than those of *Opt-ReNoVatE*. The reason behind *Fast-ReNoVatE*'s worse performance is that *Fast-ReNoVatE* recovers adjacent VLinks and independent VLinks in one particular order. In contrast, *Opt-ReNoVatE* recovers all the failed VLinks at once by exploring all possible sequences resulting in the optimal solution. Fig. 4(a) compares the recovery efficiencies of the three approaches for different SN sizes. As observed in the figure, recovery efficiencies of all three approaches increase slightly with the increase in SN size. This is due to the higher path diversity augmented by the higher



(a) Recovery Performance



(b) Cost Analysis

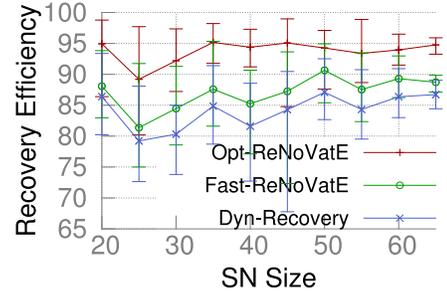


(c) Scalability Analysis

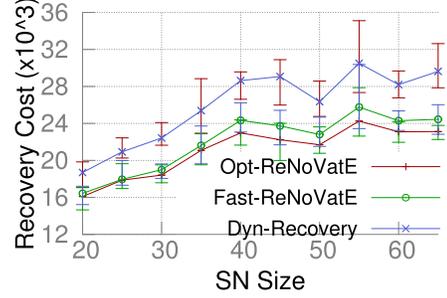
Fig. 3. Small scale performance by varying SLink utilization.

number of SLinks in the larger SN. All the three approaches utilize the path diversity in the SN to recover more failed VLinks resulting in the increased recovery efficiencies.

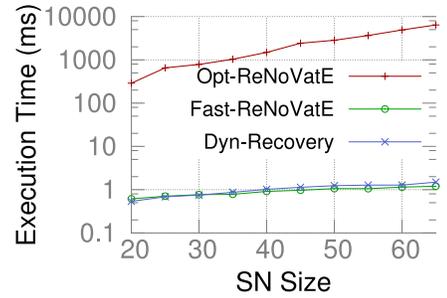
b) Recovery cost: We show the recovery cost in Fig. 3(b) against different SLink utilizations. In higher utilization cases, more SLinks become saturated in terms of bandwidth, and all the algorithms have to select longer substrate paths to recover resulting in increased costs. Further, the costs of *Opt-ReNoVatE* are the least among the three, since it selects the most suitable paths through exhaustive search. In contrast, *Fast-ReNoVatE* iterates over all the affected VNs and the independent VLinks in a greedy manner, sometimes preferring less suitable paths resulting in more costs than *Opt-ReNoVatE*. Finally, *Dyn-Recovery* does not consider the cost of a path while recovering a VLink. It may select a longer path than that selected by *Fast-ReNoVatE* leading to a much larger cost than *Fast-ReNoVatE*. On average, *Fast-ReNoVatE* incurs $\sim 7\%$ more cost than *Opt-ReNoVatE* and $\sim 20\%$ less cost than *Dyn-Recovery*. Fig. 4(b) presents the recovery cost against different SN sizes. This figure shows that the recovery costs of all three



(a) Recovery Performance



(b) Cost Analysis



(c) Scalability Analysis

Fig. 4. Small scale performance by varying SN size.

approaches gradually increase with the increase in SN size. This is due to the fact that the two end SNodes of a failed VLink are embedded far apart from one another in a larger SN. Therefore, all the approaches recover the failed VLinks using longer substrate paths leading to higher recovery costs. However, the differences of recovery costs among the three approaches follow the same pattern as in the SLink utilization cases.

c) Execution time: To demonstrate the scalability of the compared algorithms, we report their execution times on different problem instances. Fig. 3(c) presents the execution times by varying utilization on a fixed SN, whereas Fig. 4(c) presents the execution times in SNs with varying sizes while keeping the utilization fixed at $\sim 53\%$. According to Table II, both problem size and the total number of failed VLinks increase with the increase in utilization and SN size. Consequently, the execution times grow for all the approaches, however, the increase is exponential for *Opt-ReNoVatE*. As it turns out, *Fast-ReNoVatE* and *Dyn-Recovery* take less than 3ms and 2ms, respectively, in the largest problem instances, whereas

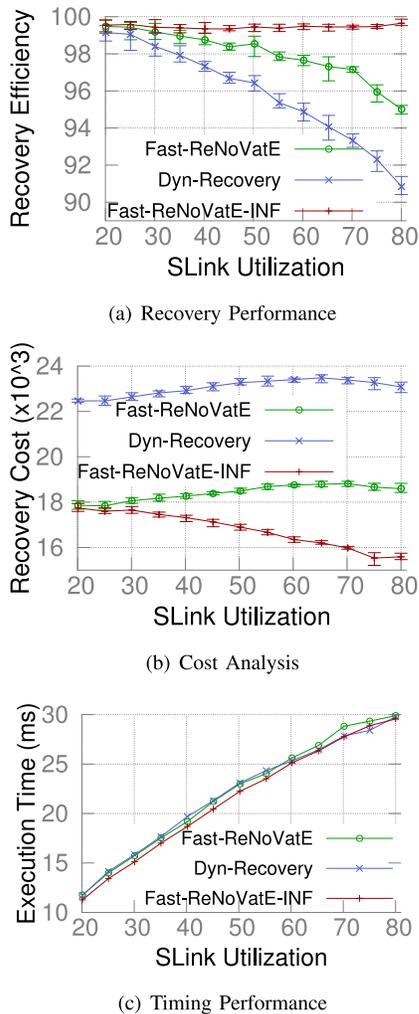
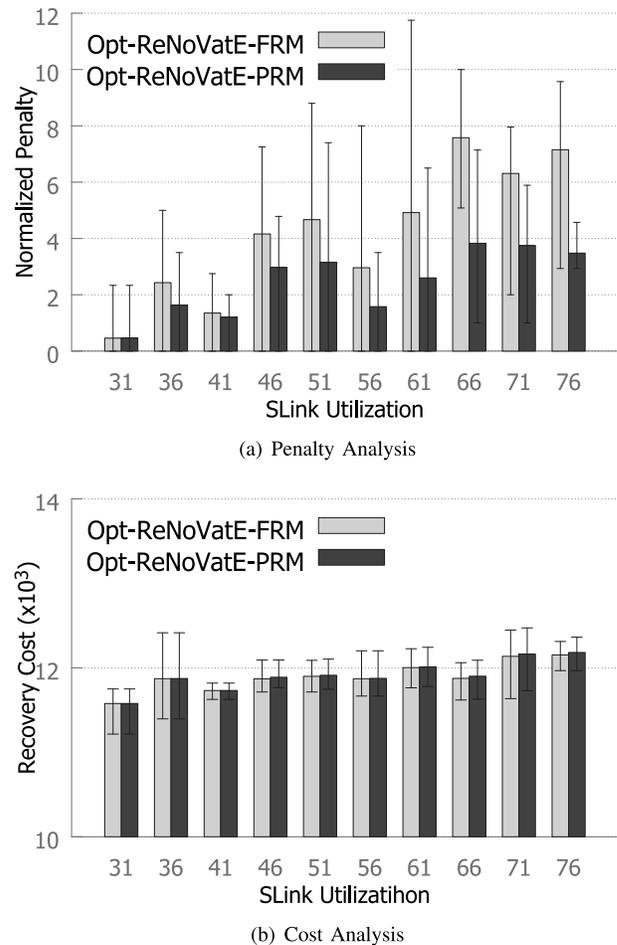


Fig. 5. Performance of large scale testcases.

Opt-ReNoVatE takes $\sim 6s$ on the same problem instances. The slightly higher execution times of *Fast-ReNoVatE* compared to *Dyn-Recovery* are due to the extra iterations of *Fast-ReNoVatE* to avoid bottleneck SLinks to achieve higher recovery efficiencies. Despite that, *Fast-ReNoVatE* is 400x–2000x faster than *Opt-ReNoVatE* depending on problem instance. Finally, *Opt-ReNoVatE* could not scale to more than 65 SNode SNs as shown in Fig. 4(c).

2) *Large Scale Evaluations*: In large scale settings, we evaluate *Fast-ReNoVatE*, *Dyn-Recovery*, and *Fast-ReNoVatE-INF* by varying SLink utilizations in the same SN (Fig. 5).

a) *Recovery efficiency*: Fig. 5(a) shows that the recovery efficiencies of *Fast-ReNoVatE* are $\sim 6\%$ better than those of *Dyn-Recovery* and $\sim 2.5\%$ worse than those of *Fast-ReNoVatE-INF*. Similar to the small scale results, recovery efficiencies of *Fast-ReNoVatE* and *Dyn-Recovery* decrease with the increase in SLink utilization whereas recovery efficiencies remain almost the same for *Fast-ReNoVatE-INF*. The near constant recovery efficiencies of *Fast-ReNoVatE-INF* confirms that the reason of failing to recover is the insufficiency of bandwidth in SLinks. In other words, if there were adequate bandwidth in the SLinks, *Fast-ReNoVatE* could recover


 Fig. 6. Impact of adding priority on *Opt-ReNoVatE*.

$\sim 99\%$ of the failed VLinks. The very small percentage of un-recovered VLinks of *Fast-ReNoVatE-INF* is due to the partitioning in the SN caused by the SNode failure. In these cases, it is not possible to recover a failed VLink even if there is sufficient bandwidth.

b) *Recovery cost*: Fig. 5(b) shows the recovery cost in large networks. For the same reasons discussed in the case of small scale networks, *Dyn-Recovery* incurs the largest amount of cost, and the costs of *Fast-ReNoVatE* and *Dyn-Recovery* rise with the increase in SLink utilization. In contrast, there is no effect of residual bandwidth in finding an alternate path in *Fast-ReNoVatE-INF*, and it can select the minimum cost path resulting in the least costs. This is true for *Fast-ReNoVatE* in very low utilized SNs. Further, the two end nodes of the failed VLink are embedded closely to each other in a highly utilized SN. Hence, *Fast-ReNoVatE-INF* recovers the VLink with shorter path leading to the decrease of costs with the increase in SLink utilization. The counterintuitive behavior of Fig. 5(b) from SLink utilization of 70 to 80 is due to relaxing the load-balancing criteria as explained in Section VI-D. The denser embedding to achieve higher utilization maps the VNodes of a VN closer to one another, thus requiring lower recovery cost than a sparser one.

c) *Execution time*: Fig. 5(c) shows that *Fast-ReNoVatE* has similar timing performance to *Dyn-Recovery*, and both are

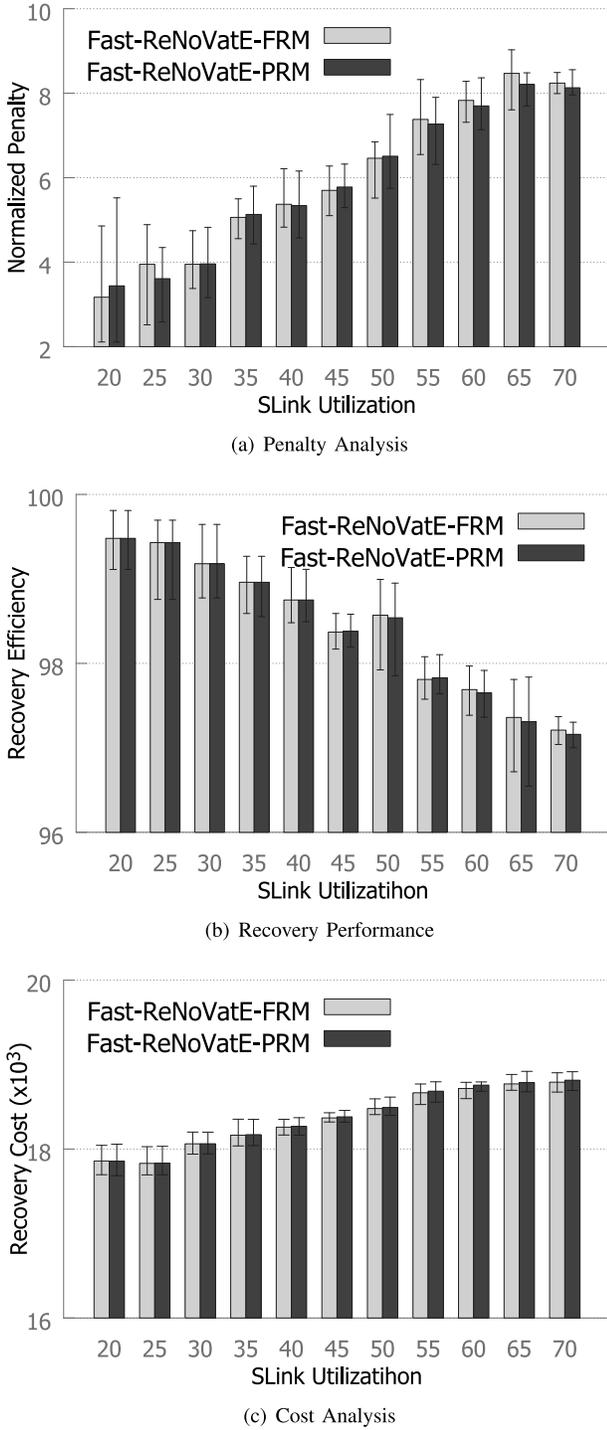


Fig. 7. Impact of adding priority on Fast-ReNoVatE.

able to find a solution in less than 30ms even in the highest utilized SN.

3) *Comparison Between FRM and PRM*: In this section, we compare FRM and PRM to demonstrate the impact of priority based recovery on *Opt-ReNoVatE* and *Fast-ReNoVatE*, respectively. Additionally, evaluations on *Opt-ReNoVatE* are done for small scale problem instances, whereas, *Fast-ReNoVatE* is evaluated on large scale topologies.

a) *Impact on Opt-ReNoVatE*: Fig. 6 presents normalized penalties for the two variants of *Opt-ReNoVatE*,

i.e., *Opt-ReNoVatE-FRM* and *Opt-ReNoVatE-PRM*. Although *Opt-ReNoVatE-FRM* and *Opt-ReNoVatE-PRM* employ two separate objective functions as presented in Section IV, their performances in terms of recovery efficiency are found similar. Hence, we do not report those results. However, they recover different VLinks in the face of resource inadequacy. This difference can be clearly observed in their corresponding normalized penalties and recovery costs shown in Fig. 6(a) and Fig. 6(b), respectively. As seen in Fig. 6(a), the normalized penalties incurred by *Opt-ReNoVatE-FRM* remain always higher than those incurred by *Opt-ReNoVatE-PRM*. Even though both of them recover the same number of VLinks, *Opt-ReNoVatE-PRM* prioritizes the failed VLinks having higher penalties in order to minimize the total penalty. On the other hand, *Opt-ReNoVatE-FRM* remains indifferent to the failed VLinks while recovery leading to the higher normalized penalty. Furthermore, the normalized penalties increase with the increase in SN utilizations for both models. This is due to the higher number of VLinks being affected by an SNode failure and the lower residual capacities left for recovery in a higher utilized SN. However, *Opt-ReNoVatE-FRM* suffers more than *Opt-ReNoVatE-PRM* due to *Opt-ReNoVatE-FRM*'s impartial treatment on the failed VLinks. In contrast, *Opt-ReNoVatE-PRM* prioritizes the failed VLinks with higher penalties even though they incur higher recovery cost. This accounts for the slightly higher recovery costs incurred by *Opt-ReNoVatE-PRM* as seen in Fig. 6(b).

b) *Impact on Fast-ReNoVatE*: Fig. 7 demonstrates the performances of the two variants of *Fast-ReNoVatE*, *i.e.*, *Fast-ReNoVatE-FRM* and *Fast-ReNoVatE-PRM* in terms of our evaluation metrics. Fig. 7(a) shows the normalized penalties incurred by the two models by varying SN utilizations in large scale topologies. The graph of Fig. 7(a) follows similar trend as we observe in the case of *Opt-ReNoVatE* in Fig. 6 and reinforces the argument presented in the previous paragraph. However, the deviations between *Fast-ReNoVatE-FRM* and *Fast-ReNoVatE-PRM* are less profound than those of the models of *Opt-ReNoVatE*. The reason behind such difference is that *Fast-ReNoVatE-PRM* employs priority on the VN level rather than on the VLink level as employed by *Opt-ReNoVatE-PRM*. Unlike the models of *Opt-ReNoVatE*, *Fast-ReNoVatE-FRM* and *Fast-ReNoVatE-PRM* show slight differences in terms of recovery efficiency and cost. Therefore, we present those results in Fig. 7(b) and Fig. 7(c). According to these figures, the recovery efficiencies (or, costs) of *Fast-ReNoVatE-PRM* are slightly lower (or, higher) than those of *Fast-ReNoVatE-FRM*. This stems from the fact that *Fast-ReNoVatE-PRM* prioritizes the VLinks with higher penalties despite being recovered in longer substrate paths. Such recovery increases the recovery cost and reduces the residual capacity of the SLinks leaving less room for re-embedding some VLinks. This accounts for the higher recovery costs and lower recovery efficiencies of *Fast-ReNoVatE-PRM* than those of *Fast-ReNoVatE-FRM*.

c) *Impact on execution time*: Since FRM and PRM are orthogonal to our solutions *Opt-ReNoVatE* and *Fast-ReNoVatE*, incorporating the models to our solutions does not bring any significant change. Hence, the complexity of *Opt-ReNoVatE* presented in Section IV-E and the running

time analysis of *Fast-ReNoVatE* presented in Section V-C remain the same irrespective of the model being chosen. In our simulations, the execution times of *Opt-ReNoVatE-PRM* and *Fast-ReNoVatE-PRM* are found similar to their FRM counterparts. Hence, we refrain from reporting them in this paper.

VII. CONCLUSION

In this paper, we have addressed the problem of generalized recovery of a batch of affected VNs, resulting from a single substrate node failure. We have formulated the problem as an Integer Linear Programming (ILP) model, *Opt-ReNoVatE* and presented an efficient heuristic algorithm, *Fast-ReNoVatE*, to tackle the computational complexity. We have evaluated *Fast-ReNoVatE*, *Opt-ReNoVatE*, and a state-of-the-art solution, *Dyn-Recovery* in both small and large scale networks. Evaluation results demonstrate that *Fast-ReNoVatE* can recover $\sim 6\%$ more VLinks than *Dyn-Recovery* and $\sim 3\%$ less VLinks than *Opt-ReNoVatE* in high utilization scenarios. In terms of scalability, *Fast-ReNoVatE* is several orders of magnitude faster than *Opt-ReNoVatE*, and has comparable performance with *Dyn-Recovery*. In large scale networks, we have compared *Fast-ReNoVatE* with *Dyn-Recovery* and a baseline case of infinite bandwidth SN. These results demonstrate that *Fast-ReNoVatE* is able to recover $\sim 99\%$ of the failed VNs if the SN has adequate residual capacity, and has similar timing performance to *Dyn-Recovery*. Furthermore, we have investigated two variants of our recovery scheme, namely, fair recovery model (FRM) and priority-based recovery model (PRM). 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.

In the future, we plan to extend this work to accommodate recovery with over-subscribed bandwidth allocation for the failed VLinks. In addition, we intend to study the problem in a real testbed environment and evaluate our solutions through a prototype implementation.

ACKNOWLEDGMENT

The authors would like to thank our shepherd Luciano Paschoal Gaspar and the anonymous CNSM 2016 reviewers for their valuable feedback and Nicole Keshav for writing suggestions that helped improve the quality of this paper.

REFERENCES

- [1] N. M. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Comput. Netw.*, vol. 54, no. 5, pp. 862–876, Apr. 2010.
- [2] A. Fischer, J. F. Botero, M. T. Beck, H. D. Meer, and X. Hesselbach, "Virtual network embedding: A survey," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 4, pp. 1888–1906, 4th Quart., 2013.
- [3] P. Gill, N. Jain, and N. Nagappan, "Understanding network failures in data centers: Measurement, analysis, and implications," in *Proc. ACM SIGCOMM*, vol. 41. Toronto, ON, Canada, Aug. 2011, pp. 350–361.
- [4] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Chuah, and C. Diot, "Characterization of failures in an IP backbone," in *Proc. INFOCOM*, vol. 4. Mar. 2004, pp. 123–133.
- [5] S. Herker, A. Khan, and X. An, "Survey on survivable virtual network embedding problem and solutions," in *Proc. ICNS*, 2013, pp. 99–104.

- [6] T. Guo, N. Wang, K. Moessner, and R. Tafazolli, "Shared backup network provision for virtual network embedding," in *Proc. IEEE ICC*, Kyoto, Japan, Jun. 2011, pp. 1–5.
- [7] M. R. Rahman and R. Boutaba, "SVNE: Survivable virtual network embedding algorithms for network virtualization," *IEEE Trans. Netw. Service Manag.*, vol. 10, no. 2, pp. 105–118, Jun. 2013.
- [8] X. Chang, J. K. Muppala, B. Wang, J. Liu, and L. Sun, "Migration cost aware virtual network re-embedding in presence of resource failures," in *Proc. 18th IEEE Int. Conf. Netw. (ICON)*, Singapore, 2012, pp. 24–29.
- [9] L. Bo, H. Tao, S. Xiao-Chuan, C. Jian-Ya, and L. Yun-Jie, "Dynamic recovery for survivable virtual network embedding," *J. China Univ. Posts Telecommun.*, vol. 21, no. 3, pp. 77–84, Jun. 2014.
- [10] B. Guo *et al.*, "Survivable virtual network design and embedding to survive a facility node failure," *J. Lightw. Technol.*, vol. 32, no. 3, pp. 483–493, Feb. 1, 2014.
- [11] Q. Hu, Y. Wang, and X. Cao, "Survivable network virtualization for single facility node failure: A network flow perspective," *Opt. Switching Netw.*, vol. 10, no. 4, pp. 406–415, 2013.
- [12] H. Jiang, L. Gong, and Z. W. Zuqing, "Efficient joint approaches for location-constrained survivable virtual network embedding," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Austin, TX, USA, 2014, pp. 1810–1815.
- [13] H. Yu, V. Anand, C. Qiao, and G. Sun, "Cost efficient design of survivable virtual infrastructure to recover from facility node failures," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Kyoto, Japan, 2011, pp. 1–6.
- [14] O. Soualah, I. Fajjari, N. Aitsaadi, and A. Mellouk, "A batch approach for a survivable virtual network embedding based on Monte-Carlo tree search," in *Proc. IFIP/IEEE Int. Symp. Integr. Netw. Manag. (IM)*, Ottawa, ON, Canada, 2015, pp. 36–43.
- [15] Z. Cai, F. Liu, N. Xiao, Q. Liu, and Z. Wang, "Virtual network embedding for evolving networks," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Miami, FL, USA, 2010, pp. 1–5.
- [16] I. Houidi, W. Louati, D. Zeghlache, P. Papadimitriou, and L. Mathy, "Adaptive virtual network provisioning," in *Proc. 2nd ACM SIGCOMM Workshop Virtualized Infrastruct. Syst. Archit.*, New Delhi, India, 2010, pp. 41–48.
- [17] S. R. Chowdhury *et al.*, "Revine: Reallocation of virtual network embedding to eliminate substrate bottlenecks," in *Proc. IEEE/IFIP Integr. Netw. Manag. Symp. (IM)*, 2017.
- [18] S. B. Masti and S. V. Raghavan, "Simulated annealing algorithm for virtual network reconfiguration," in *Proc. 8th EURO-NGI Conf. Next Gener. Internet (NGI)*, Karlskrona, Sweden, 2012, pp. 95–102.
- [19] P. N. Tran and A. Timm-Giel, "Reconfiguration of virtual network mapping considering service disruption," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Budapest, Hungary, 2013, pp. 3487–3492.
- [20] P. N. Tran, L. Casucci, and A. Timm-Giel, "Optimal mapping of virtual networks considering reactive reconfiguration," in *Proc. IEEE 1st Int. Conf. Cloud Netw. (CLOUDNET)*, Paris, France, 2012, pp. 35–40.
- [21] N. Shahriar *et al.*, "ReNoVatE: Recovery from node failure in virtual network embedding," in *Proc. IEEE/ACM/IFIP Conf. Netw. Service Manag. (CNSM)*, Montreal, QC, Canada, 2016, pp. 19–27.
- [22] M. R. Rahman, I. Aib, and R. Boutaba, "Survivable virtual network embedding," in *Networking*. Heidelberg, Germany: Springer, 2010, pp. 40–52.
- [23] T. Guo, N. Wang, K. Moessner, and R. Tafazolli, "Shared backup network provision for virtual network embedding," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Kyoto, Japan, 2011, pp. 1–5.
- [24] M. M. A. Khan, N. Shahriar, R. Ahmed, and R. Boutaba, "Multi-path link embedding for survivability in virtual networks," *IEEE Trans. Netw. Service Manag.*, vol. 13, no. 2, pp. 253–266, Jun. 2016.
- [25] N. Shahriar *et al.*, "Connectivity-aware virtual network embedding," in *Proc. 15th IFIP Netw. Conf. (NETWORKING)*, Vienna, Austria, 2016, pp. 46–54.
- [26] Y. Chen, J. Li, T. Wo, C. Hu, and W. Liu, "Resilient virtual network service provision in network virtualization environments," in *Proc. IEEE 16th Int. Conf. Parallel Distrib. Syst. (ICPADS)*, 2010, pp. 51–58.
- [27] H. Yu *et al.*, "Survivable virtual infrastructure mapping in a federated computing and networking system under single regional failures," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Miami, FL, USA, 2010, pp. 1–6.
- [28] X. Liu, Y. Wang, A. Xiao, X. Qiu, and W. Li, "Disaster-prediction based virtual network mapping against multiple regional failures," in *Proc. IFIP/IEEE Int. Symp. Integr. Netw. Manag. (IM)*, Ottawa, ON, Canada, 2015, pp. 371–378.
- [29] M. Pourvali *et al.*, "Progressive recovery for network virtualization after large-scale disasters," in *Proc. IEEE Int. Conf. Comput. Netw. Commun. (ICNC)*, 2016, pp. 1–5.

- [30] M. G. Rabbani, M. F. Zhani, and R. Boutaba, "On achieving high survivability in virtualized data centers," *IEICE Trans. Commun.*, vol. 97, no. 1, pp. 10–18, 2014.
- [31] Q. Zhang, M. F. Zhani, M. Jabri, and R. Boutaba, "Venice: Reliable virtual data center embedding in clouds," in *Proc. IEEE INFOCOM*, Toronto, ON, Canada, Apr. 2014, pp. 289–297.
- [32] H. Jiang, Y. Wang, L. Gong, and Z. Zhu, "Availability-aware survivable virtual network embedding in optical datacenter networks," *IEEE/OSA J. Opt. Commun. Netw.*, vol. 7, no. 12, pp. 1160–1171, Dec. 2015.
- [33] Q. Hu, Y. Wang, and X. Cao, "Location-constrained survivable network virtualization," in *Proc. 35th IEEE Sarnoff Symp. (SARNOFF)*, Newark, NJ, USA, 2012, pp. 1–5.
- [34] S. R. Chowdhury *et al.*, "Protecting virtual networks with drone," in *Proc. IEEE/IFIP Netw. Oper. Manag. Symp.*, Istanbul, Turkey, Apr. 2016, pp. 78–86.
- [35] S. R. Chowdhury *et al.*, "Dedicated protection for survivable virtual network embedding," *IEEE Trans. Netw. Service Manag.*, vol. 13, no. 4, pp. 913–926, Dec. 2016.
- [36] S. Agarwal *et al.*, "Volley: Automated data placement for geo-distributed cloud services," in *Proc. NSDI*, San Jose, CA, USA, 2010, pp. 17–32.
- [37] N. Bansal *et al.*, "Towards optimal resource allocation in partial-fault tolerant applications," in *Proc. IEEE INFOCOM*, 2008, pp. 1319–1327.
- [38] W.-L. Yeow, C. Westphal, and U. C. Kozat, "Designing and embedding reliable virtual infrastructures," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 2, pp. 57–64, 2011.
- [39] J. Xu, J. Tang, K. Kwiat, W. Zhang, and G. Xue, "Survivable virtual infrastructure mapping in virtualized data centers," in *Proc. IEEE CLOUD*, Honolulu, HI, USA, 2012, pp. 196–203.
- [40] P. Bodík *et al.*, "Surviving failures in bandwidth-constrained datacenters," in *Proc. ACM SIGCOMM*, Helsinki, Finland, 2012, pp. 431–442.
- [41] S. Ayoubi, C. Assi, L. Narayanan, and K. Shaban, "Optimal polynomial time algorithm for restoring multicast cloud services," *IEEE Commun. Lett.*, vol. 20, no. 8, pp. 1543–1546, Aug. 2016.
- [42] S. Ayoubi, C. Assi, Y. Chen, T. Khalifa, and K. B. Shaban, "Restoration methods for cloud multicast virtual networks," *J. Netw. Comput. Appl.*, vol. 78, pp. 180–190, Jan. 2016.
- [43] A. M. Ghaleb, T. Khalifa, S. Ayoubi, and K. B. Shaban, "Surviving link failures in multicast VN embedded applications," in *Proc. IEEE/IFIP Netw. Oper. Manag. Symp. (NOMS)*, Istanbul, Turkey, 2016, pp. 645–651.
- [44] A. M. Ghaleb, T. Khalifa, S. Ayoubi, K. B. Shaban, and C. Assi, "Surviving multiple failures in multicast virtual networks with virtual machines migration," *IEEE Trans. Netw. Service Manag.*, vol. 13, no. 4, pp. 899–912, Dec. 2016.
- [45] R. R. Oliveira *et al.*, "No more backups: Toward efficient embedding of survivable virtual networks," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2013, pp. 2128–2132.
- [46] M. Soares and E. Madeira, "A multi-agent architecture for autonomic management of virtual networks," in *Proc. IEEE Netw. Oper. Manag. Symp. (NOMS)*, Budapest, Hungary, Apr. 2012, pp. 1183–1186.
- [47] S. Even, A. Itai, and A. Shamir, "On the complexity of time table and multi-commodity flow problems," in *Proc. 16th Annu. Symp. Found. Comput. Sci.*, 1975, pp. 184–193.
- [48] Y. Dinitz, N. Garg, and M. X. Goemans, "On the single-source unsplittable flow problem," *Combinatorica*, vol. 19, no. 1, pp. 17–41, 1999.
- [49] J. Edmonds and R. M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," *J. ACM*, vol. 19, no. 2, pp. 248–264, Apr. 1972. [Online]. Available: <http://doi.acm.org/10.1145/321694.321699>
- [50] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959.



Nashid Shahriar received the B.Sc. and M.Sc. degrees in computer science and engineering from the Bangladesh University of Engineering and Technology in 2009 and 2011, respectively. He is currently pursuing the Ph.D. degree with the School of Computer Science, University of Waterloo. His research interest includes network virtualization and network function virtualization. He is a recipient of David R. Cheriton Graduate Scholarship at the University of Waterloo.



Reaz Ahmed received the B.Sc. and M.Sc. degrees in computer science from the Bangladesh University of Engineering and Technology in 2000 and 2002, respectively, and the Ph.D. degree in computer science from the University of Waterloo, in 2007. His research interests include future Internet architectures, information-centric networks, network virtualization, and content sharing peer-to-peer networks with focus on search flexibility, efficiency, and robustness. He was a recipient of the IEEE Fred W. Ellersick Award in 2008.



Shihabur Rahman Chowdhury received the B.Sc. degree in computer science and engineering from the Bangladesh University of Engineering and Technology. He is currently pursuing the Ph.D. degree with the School of Computer Science, University of Waterloo. His research interests include virtualization and softwarization of computer networks. He was a recipient of the Graduate Excellence Scholarship, the Ontario Graduate Scholarship, the Presidents Graduate Scholarship, and the GoBell Scholarship at the University of Waterloo.



Aimal Khan received the B.Sc. degree in computer science and engineering from Pakistan. He is currently pursuing the master's degree in mathematics with the School of Computer Science, University of Waterloo. His research interest include network virtualization and Internet of Things.



Raouf Boutaba (F'12) received the M.Sc. and Ph.D. degrees in computer science from the University Pierre & Marie Curie, Paris, in 1990 and 1994, respectively. He is currently a Professor of Computer Science with the University of Waterloo, Canada. His research interests include resource and service management in networks and distributed systems. He was a recipient of several best paper awards and recognitions including the Technical Achievement Award of the IEEE Communications Society Technical Committee on Information Infrastructure and Networking, the Donald W. McLellan Meritorious Service Award of the IEEE Communications Society, the Premiers Research Excellence Award, the IEEE ComSoc Hal Sobol, Fred W. Ellersick, Joe LociCero, Dan Stokesbury, Salah Aidarous Awards, and the IEEE Canada McNaughton Gold Medal. He is the founding Editor-in-Chief of the IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT from 2007 to 2010 and on the editorial boards of other journals. He is a fellow of the Engineering Institute of Canada and the Canadian Academy of Engineering.



Jeebak Mitra received the M.A.Sc. and Ph.D. degrees in electrical engineering from the University of British Columbia, Canada, in 2005 and 2010, respectively. From 2010 to 2011, he was a Senior DSP Engineer with Wimatek Systems leading the system level design for an LTE baseband. From 2011 to 2012, he was a Team Leader for the Wireless DSP Team, BLINQ Networks, Ottawa, working on backhaul products for small cells. He has been a Staff Engineer with Huawei Technologies Canada Research Center, Ottawa, in the areas of algorithm design and implementation of high-speed ASICs for optical transceivers and flexible optical network design, since 2013. His research interests lie broadly in the areas of physical layer design aspects for fiber-optic and wireless networks with an emphasis on high-performance communication systems. He was a recipient of the Best Student Paper Award at the IEEE Canadian Conference in Electrical and Computer Engineering in 2009. He regularly serves as a Reviewer for several conferences and journals in related areas.