SURVIVABLE VNE

By Sedef Savas

Networks Lab, May 25, 2017

SURVEY ON SURVIVABLE VIRTUAL NETWORK EMBEDDING PROBLEM AND SOLUTIONS

Sandra Herker, Ashiq Khan, Xueli An DOCOMO Communications Laboratories Europe GmbH Munich, Germany

ICNS 2013 : The Ninth International Conference on Networking and Services

Introduction

VNE problem deals with finding a mapping of a virtual network request onto the substrate network/physical network (NP-hard)

Common algos: backtracking, simulated annealing, and approximation

Objective: mapping a VN to a substrate while guaranteeing the VNs survivability in the event of failures.

A. The Virtual Network Embedding Problem

a) Node mapping: One virtual node needs to be mapped to exactly one substrate node (NP-hard). Greedy methods are used.
b) Link mapping: k-shortest path or multi-commodity flow algorithms.

B. The Survivable Virtual Network Embedding

- 70 % of the unplanned link failures are single link failures.
- Link failures happen ten times more than node failures per day.

The task is to embed a virtual network that can deal with virtual and substrate network failures, after the failure, the virtual network is still operating.

One option: extend the virtual network graph with backup nodes/links.

In the survivable mapping, virtual nodes of one virtual network should not be mapped on the same substrate node. Due to the fact, that a possible failure of this substrate node could affect several virtual nodes. For links, different virtual links should use distinct paths in the substrate network.

ALGORITHMS FOR THE SURVIVABLE VIRTUAL NETWORK EMBEDDING

A. Survivable VN Embedding against Link Failures

1) Link restoration and protection methods:

backup paths for each substrate link are calculated (SVNE-Boutaba), restoration OR

Shared Pre-Allocation approach, backup bandwidth for each substrate link is pre-allocated during the configuration phase before any VN request arrives. Since the bandwidth pre-allocation only needs to be done once and not for every VN request, there is less computing done during the VN embedding phase.

Advantage: backup bandwidth is already allocated before the failure happens and not after the failure. Disadvantage of the Shared Pre-Allocation approach is that backup bandwidth is reserved independent of the VN requests and may not be used at any time if few VN requests arrive.

2) Path protection methods with node migration: Instead of backing up the each primary link, each end-to-end primary substrate path is protected by a backup path. Migratory shared protection, migrates and maps a VN node to another substrate node to increase the resource efficiency when a failure occurs.

B. Survivable VN Embedding against Node Failures

TABLE I
SUMMARY OF THE SURVIVABLE EMBEDDING ALGORITHMS

References	Survivability	Type of failure	Optimization Objective	Survivable failure mechanism
Survivable virtual network embedding [13]	Link	Single substrate link failure	Maximize revenue for Infrastructure Provide	reactive, after failure (Restoration)
Shared backup network provision for virtual network embedding [14]	Link	Single substrate link failure	Maximize revenue/accepting VN requests	proactive, before failure (Protection)
Migration based protection for virtual infrastructure survivability for link failure [15]	Link	Single substrate link failure	Minimize sum of costs	before failure
QoSMap: Achieving Quality and Resilience through Overlay Construction [16]	Link	Single substrate link failure	Minimize delay and additional resources for backup	before failure
An overlay mapping model for achieving enhanced QoS and resilience performance [17] / An overlay mapping model for achieving enhanced QoS and resilience performance [18]	Link	Single substrate link failure	Minimize delay and additional resources for backup	before failure
Survivable virtual infrastructure mapping in a federated computing and networking system under single regional failures [21]	Node	Single regional failure	Minimize sum of cost	before failure
Cost efficient design of survivable virtual infrastructure to recover from facility node failures [19]	Node	Single facility node failure	Minimize sum of cost	before failure
A novel two-step approach to surviving facility failures [20]	Node	Single facility node failure	Minimize resources/total cost	before failure
Location-constrained survivable network virtualization [22]	Node	Single facility node failure	Minimize resources	before failure
Designing and embedding reliable virtual infrastructures [23]	Node	Single substrate node failure	Minimize amount of resources used	before failure
Survivable virtual infrastructure mapping in virtualized data centers [24]	Node	single server failure	Minimize operational cost	before failure
Adaptive virtual network provisioning [25]	Node or Link	single node failure or single link failure	-	after failure

Connectivity-aware Virtual Network Embedding

IFIP Networking Conference (IFIP Networking) and Workshops, May 2016

Nashid Shahriar∗ , Reaz Ahmed∗ , Shihabur Rahman Chowdhury∗ , Md. Mashrur Alam Khan∗ , Raouf Boutaba∗ , Jeebak Mitra† , and Feng Zeng† ∗David R. Cheriton

∗ University of Waterloo

†Huawei Technologies

Patented by Huawei | Feb, 2017

Connectivity-aware Virtual Network Embedding (CoViNE)

The problem of ensuring virtual network (VN) connectivity in presence of multiple link failures in the substrate network (SN) is not well investigated.

Solving CoViNE will enable a VN operator to perform failure recovery without depending on the SN provider, similar to the IP restoration mechanisms in IP-over-WDM networks.

There are two steps in solving CoViNE: i) finding the virtual links that should be embedded disjointly, and ii) finding a substrate resource efficient embedding that ensures the virtual link disjointness constraint.

Connectivty- aware Virtual *Network* Embedding (*CoViNE*)

Majority of the works on SVNE focus on **link failures**, as they occur more frequently than node failures. SVNE approaches, allocate redundant bandwidth for each (or selected) virtual link(s), either proactively while computing the embedding or reactively after a failure occurs.

Goal: find a VN embedding that can ensure connectivity in a VN topology in presence of multiple substrate link failures. In contrast, SVNE approaches focus on guaranteeing virtual link demand in presence of failure(s).

Who Benefits?

SVNE approaches assume that SN-providers hide physical failures by over-provisioning a fraction of each virtual link's bandwidth, which in turn incurs additional cost to VN-operators.

In contrast, this work ensures VN connectivity, which is a weaker form of survivability incurring lesser resource overhead and reduced cost of leasing.

Intend is to empower a VN-operator to handle link failures according to its internal policy, e.g., customer priority. A VN- operator can plan and over-provision different amounts of bandwidth in each virtual link to handle failures according to its needs, instead of blindly relying on the SN-provider that over provisions fixed bandwidth for each virtual link as in SVNE approaches.

Connectivity-aware embedding will enable a VN-operator to reroute traffic on failed virtual links, which can be done using any IP link restoration protocol.

Similarity to IP over WDM

CoViNE problem is equally applicable in IP-over-WDM domain.

The problem of ensuring IP layer connectivity in presence of a single WDM link failure is known as *link survivable mapping*.

Two variations of the problem have been studied in IP-over-WDM literature:

i) *weakly link survivable mapping* (WLSM) ensures IP-layer connectivity;

ii) *strong link survivable* mapping guarantees both connectivity and bandwidth of the failed IP link(s) in presence of a single WDM link failure. WLSM, which considers single link failure, is merely a special case of *CoViNE*.

Multi-link Failure

Focus on multiple (up to double) link failures, since it is not a rare in large transport networks.

1) Repairing a failed link can take long time. Chances of a second link failure is not negligible given the high Mean- Time-to-Repair (MTTR).
2) Some inter-datacenter links destined to different places may be physically routed together for some distance, and a backhaul failure may cause multiple physical links to fail.
    - ~12% of the failures in inter-datacenter transport networks are double link failures in SN.

Two conditions for surviving multiple (k) link failures: i) VN topology must be k + 1 edge connected, and ii) the embedding algorithm must ensure at least k + 1 edge-disjoint paths in SN between every pair of virtual nodes.

- The first condition can be satisfied by augmenting the VN with new links [*].
- To satisfy the 2. condition: embed all virtual links of a k + 1 edge connected VN onto disjoint paths in the SN. However, this is an **NP-complete problem and requires unsatisfiable number of disjointness constraints.**

[*] K. Thulasiraman *et al.*, “Logical topology augmentation for guaranteed survivability under multiple failures in ip-over-wdm optical networks,” *Optical Switching and Networking*, vol. 7, no. 4, pp. 206–214, 2010.

Related Work

Existing cut set based approaches suffer from poor scalability.

IP-over-WDM literature focus on ensuring connectivity of IP links under WDM link failures. But, most: ILP (Their formulation explores exponential number of cut sets in the VN and routes all the VLinks of a cut set on disjoint WDM paths.), heuristics does not deal with multiple link failure or requires specific VN properties. (one identifies a set of spanning trees of the VN and computes a shortest-path based routing of the VLinks such that at least one of the spanning trees survives after an SLink failure).

Most heuristic schemes either focus on single link failure, or fail to deal with arbitrary VN topologies .

[*] Augment VLinks until a complete subgraph of k+1 VNodes is constructed and the remaining VNodes are k + 1 edge connected to the subgraph. Their solution maps any k of the VLinks incident to a VNode onto disjoint paths. Requires a large number of virtual links to be embedded disjointly, hence, is not resource efficient.

[*] K. Thulasiraman *et al.*, “Logical topology augmentation for guaranteed survivability under multiple failures in ip-over-wdm optical networks,” *Optical Switching and Networking*, vol. 7, no. 4, pp. 206–214, 2010.

Novelty

Explore an alternate survivability model, *CoViNE*, requiring significantly less backup resources than traditional survivability approaches in SVNE literature.

Solution: augment a VN with minimal number of virtual links while preserving the topological structure of the VN.

The first condition for surviving k SLink failures is that a VN must be k + 1 edge connected, augment.

two ways: i) VLinks can be augmented between arbitrary pair of VNodes to ensure k + 1 edge connectivity, changes topology

ii) the other way is to augment only parallel VLinks between adjacent VNodes

k-protected component: *set of parallel VLinks augmented in such a way that simultaneous removal of* k *arbitrary VLinks will not partition.*

Definition 1. ***k-protected component:*** *A k-protected component of a graph $\bar{G}$ is a multi-graph $\hat{G}_k = (\hat{V}_k, \hat{E}_k)$, where $\hat{V}_k \subseteq \bar{V}$, $\hat{E}_k = \bar{E}_k \cup \tilde{E}_k$, $\bar{E}_k \subseteq \bar{E}$, $\tilde{E}_k \subseteq \tilde{E}$ and $\tilde{E}_k$ is a set of parallel VLinks augmented in such a way that simultaneous removal of k arbitrary VLinks in $\hat{G}_k$ will not partition $\hat{G}_k$.*

Definition 2. ***Conflicting VLinks:*** *Two VLinks are considered as conflicting if they must be embedded on edge-disjoint substrate paths in order to ensure $k+1$ edge connectivity.*

Definition 3. ***Conflicting set:*** *A conflicting set of a VLink $(\hat{u}, \hat{v})$, denoted by $\chi^{\hat{u}\hat{v}}$, is the set of VLinks in $\hat{E}$ those are conflicting with $(\hat{u}, \hat{v})$. A Conflicting set of a VN $\hat{G} = (\hat{V}, \hat{E})$, denoted by $\chi^{\hat{G}}$, is defined as $\chi^{\hat{G}} = \bigcup_{\forall(\hat{u},\hat{v})\in\hat{E}} \chi^{\hat{u}\hat{v}}$.*

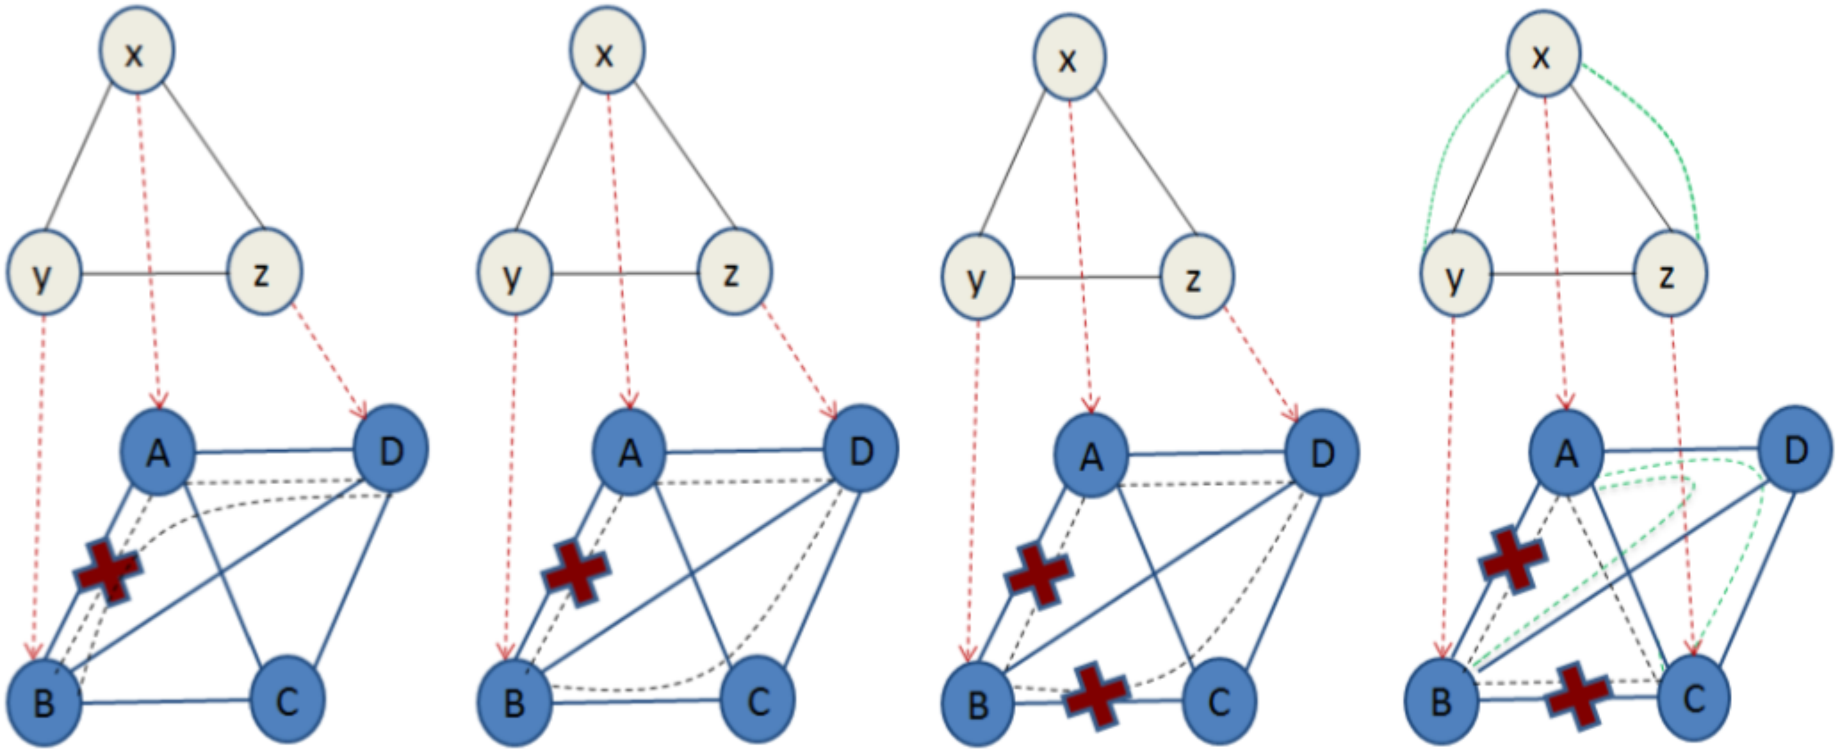

(a) Single Failure (b) Double Failure

Fig. 1. CoViNE examples

Problem Formulation

A. Conflicting Set Computation

Assume: VNs are k-protected

To remain connected after k SLink failures, embedding algo must ensure k + 1 edge connectivity between SNodes for every pair of VNodes of the k-protected VN.

Achieved if VLinks of every edge-cut in G are embedded on at least k + 1 edge-disjoint paths in G. Exponential number of edge-cuts and there are combinatorial number of ways of choosing k + 1 conflicting VLinks from an edge-cut, the number of possibilities for computing a conflicting set is enormous.

optimal conflicting set is one that ensures k + 1 edge connectivity of the embedding while minimizing disjoint path requirement in the embedding.

Find minimum number of partitions of the VLinks such that the VLinks in a partition are not conflicting with each other. Since the VLinks in a partition do not impose any disjointness constraint, minimizing the number of partitions will yield optimal conflicting set.

Computing the optimal conflicting set of a VN is NP-complete.(reduced to *Minimum vertex coloring* problem)

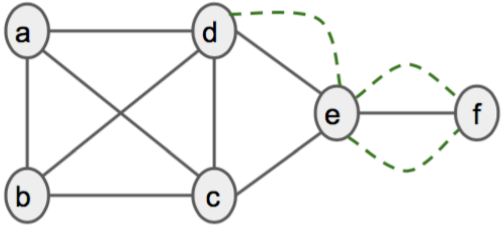

Fig. 2. The VN with only solid edges is the input VN, $\bar{G}$. The VN with both solid edges ($\bar{E}$) and dashed edges ($\tilde{E}$) is the 2-protected VN, $\hat{G}$. Any subgraph of $\hat{G}$ having 3 edge connectivity is $\hat{G}_2$.

For example, in Fig. 2, VNodes a and b will remain connected in presence of 2 SLink failures if the VLinks on paths $P_1^{ab} = (a, b)$, $P_2^{ab} = \{(a, d), (d, c), (c, b)\}$, and $P_3^{ab} = \{(a, c), (c, e), (e, d), (d, b)\}$ are mapped to disjoint SN paths. Hence, $\chi^{ab} = P_2^{ab} \cup P_3^{ab}$.

For less computation: First, ensure connectivity in Gˆ by ensuring connectivity in a minimum spanning tree. For VN in Fig. 2, k+1 edge-disjoint path computations are required for the VLinks = {(a,b),(a,c),(c,d),(d,e),(e,f)} instead of all the 12 VLinks in Gˆ. This method yields smaller conflicting set χab = pab ∪ pab, where pab = {(a, c), (c, b)}, and pab = {(a, d), (d, b)} in Fig. 2. 23

B. VLink Augmentation

augment a given VN with parallel VLinks in order to make it a k-protected VN.

Now, the challenge here is to minimize the number of augmented parallel VLinks. We use Menger's Theorem to find the pair of VNodes with less than k + 1 edge connectivity and add parallel VLinks as needed.

Embedding

Node mapping constraint: only certain locations.

Link Capacity constraint.

Disjointness Constraints: To ensure the desired survivability of CoViNE, a VLink $(\hat{u}, \hat{v}) \in \hat{E}$ should never share an SLink with it's conflicting VLinks in $\chi_{\hat{u}\hat{v}}$ in their mappings.

objective is to minimize the bandwidth provisioning cost over all the SLinks used by the mappings of all the VLinks of a VN.

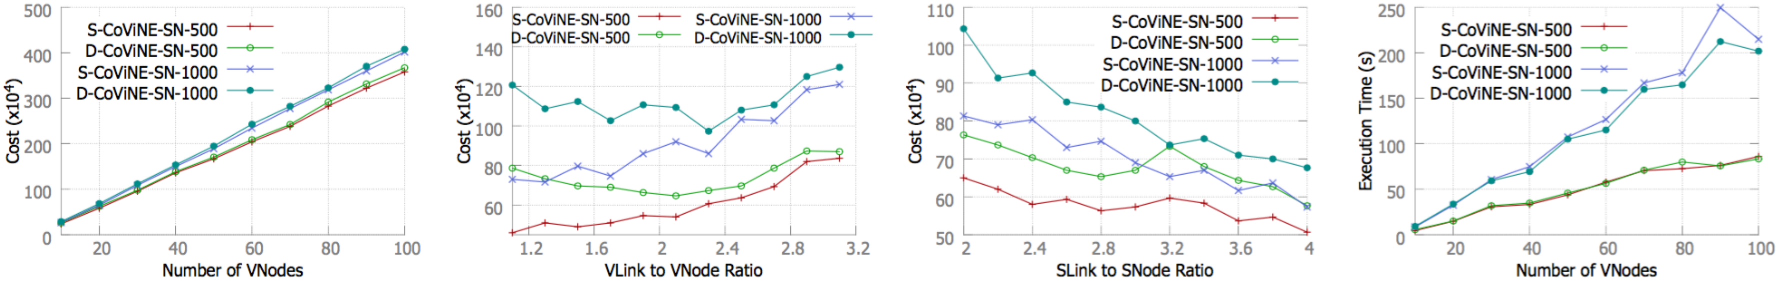

(a) Cost Vs. VN Size (b) Cost Vs. VN LNR (c) Cost Vs. SN LNR (d) Time Vs. VN Size

Fig. 4. Large Scale Performance

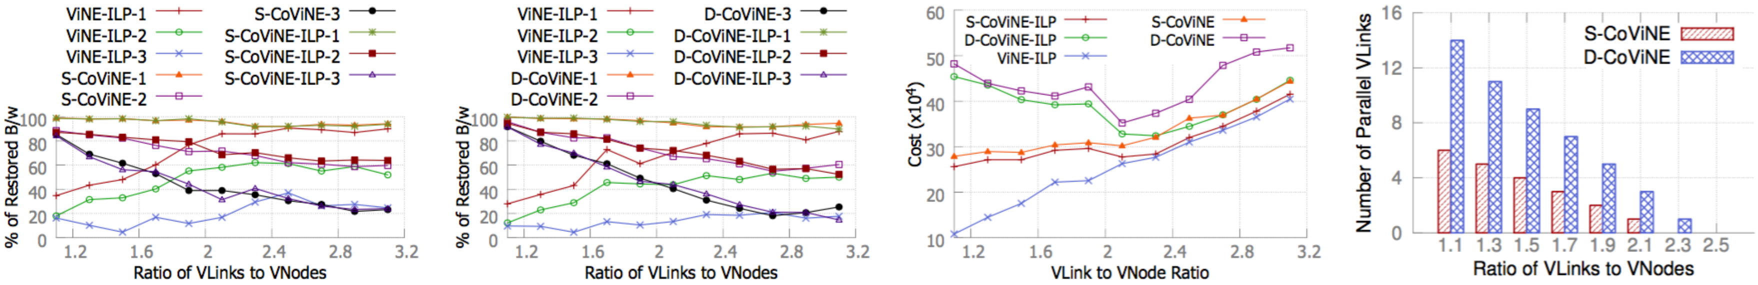

(a) Single Failure (b) Double Failure (c) Overhead (d) Parallel VLink Augmention

Fig. 5. Impact of Failure

Notation	Failures	Disjointness	Embedding
S-CoViNE	Single	Algorithm 1	Algorithm 2
D-CoViNE	Double	Algorithm 1	Algorithm 2
S-CoViNE-ILP	Single	Algorithm 1	§ VI
D-CoViNE-ILP	Double	Algorithm 1	§ VI
S-Cutset-ILP [10]	Single	Optimal Cut-set	ILP
ViNE-ILP [24]	None	None	MCUF ILP[1]

UCDAVIS