Computer Networks 89 (2015) 59–77

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Opportunistic resilience embedding (ORE): Toward cost-efficient resilient virtual networks

Rodrigo R. Oliveira [a], Daniel S. Marcon [a], Leonardo R. Bays [a], Miguel C. Neves [a], Luciano P. Gaspary [a], Deep Medhi [b], Marinho P. Barcellos [a,*]

[a] Federal University of Rio Grande do Sul, Brazil
[b] University of Missouri-Kansas City, Missouri, USA

ARTICLE INFO

Article history:
Received 26 September 2014
Revised 4 June 2015
Accepted 15 July 2015
Available online 21 July 2015

Keywords:
Network virtualization
Virtual network embedding
Resilience
Survivability
Multipath embedding
Multiple paths

ABSTRACT

Network Virtualization promotes the development of new architectures and protocols by enabling the creation of multiple virtual networks on top of the same physical substrate. One of its main advantages is the use of isolation to limit the scope of attacks – that is, avoiding traffic from one virtual network to interfere with the others. However, virtual networks are still vulnerable to disruptions on the underlying network. Particularly, high capacity physical links constitute good targets since they may be important for a large number of virtual networks.

Previous work protects virtual networks by setting aside backup resources. Although effective, this kind of solution tends to be expensive, as backup resources increase the cost to infrastructure providers and usually remain idle. This paper presents ORE (opportunistic resilience embedding), a novel embedding approach for protecting virtual links against substrate network disruptions. ORE's design is two-fold: while a *proactive* strategy embeds each virtual link into multiple substrate paths in order to mitigate the initial impact of a disruption, a *reactive* one attempts to recover any capacity affected by an underlying disruption. Both strategies are modeled as optimization problems. Additionally, since the embedding problem is $\mathcal{NP}$-Hard, ORE uses a simulated annealing-based meta-heuristic to solve it efficiently. Numerical results show that ORE can provide resilience to disruptions at a lower cost.

1. Introduction

The Network Virtualization paradigm aims to simplify the development of new architectures and protocols. In particular, it allows networks to coexist over a shared physical infrastructure [1]. One of the main advantages of this paradigm is the use of isolation to limit the scope of attacks. This can be achieved by creating different, isolated virtual networks for each task or service so that traffic from one virtual network does not interfere with the others [2].

In this paradigm, the physical infrastructure (also referred to as the substrate network) may embed[1] virtual networks in an on-demand manner [3]. To guarantee isolation, a substrate network administrator may perform admission control to each virtual network. In particular, it assigns shared physical resources to each virtual node and link of a virtual network. Hereafter, the term virtual network embedding (VNE) will denote this process or the phase.

During VNE, several virtual networks are overlaid in specific regions of the substrate network [4]. From a security

* Corresponding author. Tel.: +55 51 9993 1980.
E-mail addresses: ruas.oliveira@inf.ufrgs.br (R.R. Oliveira), daniel.stefani@inf.ufrgs.br (D.S. Marcon), lrbays@inf.ufrgs.br (L.R. Bays), mcneves@inf.ufrgs.br (M.C. Neves), paschoal@inf.ufrgs.br (L.P. Gaspary), dmedhi@umkc.edu (D. Medhi), marinho@inf.ufrgs.br, marinho.barcellos@gmail.com, marinho@acm.org (M.P. Barcellos).

[1] The terms *embed*, *map* and *allocate* will be used interchangeably.

http://dx.doi.org/10.1016/j.comnet.2015.07.010
1389-1286/

perspective, this characteristic makes virtual networks more vulnerable to disruptions in the physical infrastructure. In particular, VNE tends to increase the dependency on a small set of physical resources. Hence, a failure (or a successful attack) on a physical link can affect a large number of virtual networks.

Previous research tackled this problem by setting aside extra resources as backup [5–9]. This type of strategy may not be affordable since backup resources are not used for new allocations. Therefore, these approaches can reduce the ability of the network to handle new requests.

In this paper, we propose a novel embedding approach that optimistically improves the resilience of virtual networks without expending additional resources. Our approach is composed of two complementary optimization strategies, one **proactive** and the other **reactive**. The proactive strategy attempts to mitigate the initial impact of an attack by embedding each virtual link into *multiple paths*, thus preventing the virtual links from losing all their capacity. The reactive strategy aims at partially or fully recovering any capacity lost due to failures or compromised by attacks. To achieve this goal, the set of unaffected paths (if such paths exist) is used to *opportunistically recover* the compromised capacity. The design of opportunistic resilience embedding (ORE) is inspired by known restoration techniques (i.e., those applied to traditional networks for improving survivability).

The first strategy requires solving the VNE problem. In a previous work [10], a mixed-integer programming (MIP) formulation of a variation of this strategy was presented. Yet, since the VNE problem is NP-Hard [11–14], MIP models cannot scale to solve large problems. Hence, to achieve solutions within a reasonable timeframe, we use a simulated annealing-based metaheuristic. It runs efficiently and provides near-optimal solutions, converging toward a global optimum. The second strategy requires solving a max-flow problem over a set of pre-computed paths (i.e., those computed by the first strategy). We present a linear programming formulation to solve this problem.

The main contributions of this paper are summarized as follows:

- We present the design and implementation of a novel virtual network embedding algorithm, ORE. In contrast to the previous work, ORE does not make use of dedicated backup resources. Instead, ORE attempts to attain efficient resource utilization to the substrate network and resilience to the virtual links together. To our knowledge, ORE is the first approach to integrate embedding and restoration for virtual networks and analyze its advantages;
- ORE's strategies are modeled as optimizations problems. The proactive strategy is implemented with a simulated annealing-based algorithm, which achieves near-optimal solutions at a reasonable computing time. The reactive strategy is modeled as a linear programming (LP) problem, which is solved optimally and efficiently;
- Through simulation studies, we measured benefits of our approach. Based on our experimental studies, we found that ORE can reduce both bandwidth loss and severity of disruptions by increasing the number of paths used per virtual link. Further, when compared to backup-oriented schemes, ORE achieves a higher acceptance rate.

The rest of the paper is organized as follows: Section 2 introduces preliminary concepts and assumptions that will be used throughout the paper. ORE is described in Sections 3 and 4; Section 3 provides and overview of the ORE's design, while Section 4 describes the implementation of both the proactive and the reactive strategies. Section 5 discusses numerical results, Section 6 reviews related work, and Section 7 concludes this paper.

2. Preliminaries

2.1. Virtual network embedding problem

The virtual network embedding (VNE) problem emerges from the interaction between *virtual networks* and the underlying *substrate network*. In short, it consists of an efficient allocation of virtual network topologies on top of the substrate network. Efficiency comes from the goal of maximization or minimization of an objective as defined by the owner of the substrate network, subject to satisfying demands imposed by an SLA agreed upon between the owners of the virtual networks and the substrate network provider. Depending on the requirements and the purpose of the virtual networks, the problem is to be solved either in real-time, or offline ahead of time. If virtual network requests arrive *on-demand* without any prior knowledge, they are to be allocated in real-time or near real-time; on the other hand, if demand is known beforehand for a future request, the problem can be solved offline ahead of time so that virtual networks can be allocated.

In general, each virtual network is composed of a virtual topology (interconnected by virtual routers and virtual links) and a variety of requirements. The substrate network is composed of substrate routers, which support a virtualization technology, interconnected by substrate links. Both substrate routers and substrate links have attributes and constraints (e.g., router CPU capacity, link bandwidth capacity) that dictate which features are available to the overlay virtual networks [15]. Therefore, virtual networks are free to implement potentially arbitrary services, as long as the substrate network is capable of offering the necessary features.

In order to match virtual requirements with substrate capabilities and achieve optimal utilization of resources, the substrate network administrator invokes an allocation process. This process consists of selecting feasible substrate devices and assigning shares of their resources to the overlaid virtual networks [13,16,17]. Thus, each virtual node is mapped to a substrate node that supports all requirements; likewise, each virtual link is mapped to a substrate link or a continuous, loop-free substrate path that supports all requirements. Additionally, the endpoints of the substrate link (or path) selected to embed a virtual link should be the substrate nodes embedding the endpoints of the virtual link.

Fig. 1 exemplifies possible mappings of two virtual networks into a substrate network. Note that the example is not exhaustive in terms of all possible mappings; rather, it is meant to illustrate a VNE scenario. For simplicity, this figure represents resources used by slices on the substrate