Computer Networks 91 (2015) 712–737

Contents lists available at ScienceDirect

Computer Networks

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

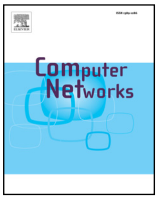

A green energy-aware hybrid virtual network embedding approach

Nizar Triki[a], Nadjia Kara[a], May El Barachi[b,*], Souad Hadjres[a]

[a] University of Quebec (ETS), Montreal, Quebec, H3C 1K3, Canada
[b] Zayed University, Khalifa City B, P.O. Box 144534, Abu Dhabi, United Arab Emirates

ARTICLE INFO

Article history:
Received 20 January 2015
Revised 17 July 2015
Accepted 18 August 2015
Available online 2 September 2015

Keywords:
Network virtualization
Virtual network embedding
Hybrid approach
Energy-awareness
Green computing

ABSTRACT

In the past few years, the concept of network virtualization has received significant attention from industry and research fora, as it represents a promising way to diversify existing networks and ensure the co-existence of heterogeneous network architectures on top of shared substrates. Virtual network embedding (VNE) is the process of dynamically mapping virtual resources (i.e. virtual nodes and links) onto physical substrate resources. VNE is the main resource allocation challenge in network virtualization and is considered as an NP-hard problem. Several centralized and distributed VNE approaches have been proposed, with the aim of satisfying different objectives ranging from QoS, to economical profit, and network survivability. More recently, emerging VNE approaches started investigating the optimization of new objectives such as energy-efficiency and networks' security. In this work, we propose a green energy-aware hybrid VNE hybrid VN embedding approach that aims at achieving energy efficiency and resource consolidation, while minimizing CO_2 emissions resulting from VNs operation. This approach consists of a hierarchical virtual networking management architecture in which control and management nodes collaborate for the splitting and embedding of sub-VNs requests to the cleanest substrate resources (i.e. the resources deployed in a sector with the smallest CO_2 emission factor) available. Three different variants of our VNE algorithms, taking into consideration different resources' selection criteria (i.e. energy source, request priority, and request location) are presented, and their performance is compared with two existing VNE algorithms based on centralized and distributed embedding approaches. The comparative performance analysis shows that our proposed approach enables a more efficient VN embedding in terms of: a reduced number of substrate resources needed, a faster request mapping time, as well as resource consolidation and reduced resource cost. Furthermore, it enables a reduction of the carbon footprint of the VNE operation, thus resulting in a more green and environmentally conscious approach to network virtualization.

1. Introduction

The network virtualization concept is gaining in popularity and several research studies have been conducted to demonstrate its potential benefits. It consists in the creation of several coexisting virtual networks (VNs) over a shared physical substrate network. Due to the potential it offers in terms of diversifying existing networks and ensuring the co-existence of heterogeneous network architectures on top of shared substrates, network virtualization is often considered as an enabler of a polymorphic Internet and a cornerstone of the future Internet architecture [1].

Enabling network virtualization requires the mapping of VN requests onto substrate networks – an operation that is referred to as virtual network embedding (VNE). VNE

---

[*] Corresponding author. Tel.: +971 56 1211 784; fax: +971 2 4434 847.
E-mail addresses: nizar.triki.1@ens.etsmtl.ca (N. Triki), Nadjia.Kara@etsmtl.ca (N. Kara), may.elbarachi@zu.ac.ae (M.E. Barachi), souad.hadjres@gmail.com (S. Hadjres).

http://dx.doi.org/10.1016/j.comnet.2015.08.016
1389-1286/

consists in assigning virtual nodes and virtual links, with specific constraints, to substrate nodes and substrate paths. The substrate path represents a single substrate link or a set of substrate links connecting two substrate nodes. The problem of VNE represents the main resource allocation challenge in network virtualization and is known as an NP-hard problem [2,4].

Several heuristic algorithms have been proposed to address this issue in a single domain [2–7] or across multiple network domains [8,9]. The main objective is to optimize the use of the substrate resources by providing efficient algorithms for VNs embedding. Most of these proposals use greedy methods to embed virtual nodes onto substrate nodes without any correlation between the node assignment and link assignment phases. Some of them map the VN request to the substrate network as a whole unit [5,7,10], while others split it into a number of sub-virtual networks (sub-VNs) [2,4]. On one hand, splitting the VN request into a set of sub-VNs eases the integration of large topologies with specific constraints in the substrate network. On the other hand, it enables efficient substrate resource usage and allows the substrate to accept more VN requests. In [2,3], the VN requests are divided into a number of star topologies, while in [4] only the splitting of virtual links over a set of substrate paths is supported. Moreover, some proposals carried out the VN embedding algorithm in a centralized or distributed fashion [2–10]. The centralized VN mapping requires a central coordinator to maintain global information about the substrate network. Gathering this information imposes a great number of signaling messages on this central entity that becomes a bottleneck when the number of VN requests increases. This algorithm limits scalability and increases complexity in highly dynamic or large scale virtual network embedding environments. The distributed VN embedding algorithm requires the exchange of a significant number of messages between the substrate nodes to support the VN embedding. This algorithm is not only a non-optimal solution because of the lack of a global view of the entire substrate network, but also imposes a large communication overhead on the substrate network.

From VNE optimization objectives point of view, energy efficiency is now considered by several VNE approaches [11–16]. These approaches rely on resource consolidation in order to group dispersed resources on a small number of substrate nodes and switch off the unused nodes, thus saving energy. While these solutions present interesting approaches for energy saving, this comes with a price to pay in terms of high VNRs' rejection ratio. Furthermore, all these solutions focus on the economic aspect of energy consumption (i.e. revenue maximization through energy cost reduction) but neglect the ecological aspect associated with energy utilization (i.e. CO_2 emissions resulting from the electricity generation). With the current shift towards sustainable and renewable energy solutions, there is a need to consider the type of energy source used to power substrate resources as a parameter in the VNE process.

In this paper, we propose a green energy-aware hybrid VN embedding approach that balances the scalability and load imposed on the substrate network, while minimizing CO_2 emissions resulting from VNs' operation. In this approach, we use VNs request splitting along with a hierarchical VNs embedding strategy that maintains a global view of the entire substrate network, in order to guarantee efficient resource management as well as low CO_2 emissions from VNs. Each VN request is split into a set of small sub-VNs topologies according to the virtual nodes and links constraints (e.g. CPU, delay, and packet loss constraints). Moreover, the substrate resources are divided into sectors according to the available type of energy source – i.e. renewable sources such as those coming from water, biomass, and wind, or non-renewable fossil fuel energy sources. Each sector is associated with an emission factor (obtained from Environment Canada [17]), which quantifies the CO_2 emissions resulting from the energy usage, and thus gives an indication about the energy consumption of different substrate resources. The algorithm embeds sub-VNs to the cleanest resource (i.e. a resource deployed in a sector with the smallest emission factor) available in the substrate network while satisfying the virtual nodes and virtual links constraints. If the sub-VNs constraints could not be satisfied, the algorithm identifies a sector with the closest emission factor and embeds the whole or a part of the sub-VNs set on it.

Three variants of our energy-aware hybrid VN embedding algorithm are proposed: (1) the basic hybrid algorithm that was described above; (2) a priority-based hybrid algorithm that takes into consideration VN requests' priority as an additional constraint in the embedding process; and (3) a location and priority based hybrid algorithm that takes into consideration both VNs' location and priority, as parameters in the VN embedding process. All three approaches are simulated and their performance is analyzed with respect to several performance metrics, including: the average node and link utilization; the time to map VN requests to physical resources; the CO_2 emission cost; as well as the revenue and the cost of physical resources. This performance analysis gives interesting insights on the effect of introducing various parameters and constraints as part on the VNE process, on its overall performance.

The benefits of our proposed VN embedding approach are of five folds.

- First, the proposed hybrid approach leads to a resource efficient VNE process which outperforms centralized and distributed VNE approaches by achieving a reduction of the number of physical nodes and links required for VNR mapping, and thus leading to a reduced cost of utilized physical resources.
- Second, by using energy efficiency and CO_2 emission reduction as objectives, our hybrid approach leads to resource consolidation (i.e. hosting as many virtual instances as possible in one substrate resource) and thus energy savings.
- Third, it allows faster VN request mapping when compared to existing approaches, thus a more efficient VN embedding operation.
- Fourth, it enables the introduction of various parameters and constraints (e.g. link delay and packet loss constraints, energy source used, request priority, and request location) as part of the VNE, thus leading to a finer grained selection/embedding process.
- Finally, our approach leads to a reduction in the carbon footprint of the resulting VNs, thus representing a more