Computer Networks 93 (2015) 460–479

Contents lists available at ScienceDirect

Computer Networks

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

Energy-aware node and link reconfiguration for virtualized network environments

Ebrahim Ghazisaeedi*, Changcheng Huang

Department of Systems and Computer Engineering, Carleton University, Ottawa, Ontario, Canada

ARTICLE INFO

Article history:
Received 5 March 2015
Revised 1 August 2015
Accepted 19 September 2015
Available online 19 October 2015

Keywords:
Virtualized network environment
Virtualized data center
Energy efficiency
Network reconfiguration
Network optimization

ABSTRACT

Energy consumption in information and communications technology (ICT) accounts for a large portion of the total energy consumed in industrial countries, which is rapidly increasing. Virtualized network environments (VNEs) have recently emerged as a solution in this technology to address the challenges of future Internet. VNEs also play a fundamental role toward virtualizing data centers. Consequently, it is essential to develop novel techniques for reducing the energy consumption of VNEs. In this paper, we discuss multiple energy-saving solutions that locally optimize the node and link energy consumption of VNEs during the off-peak period by reconfiguring the mapping of already allocated virtual nodes and links. The proposed reconfigurations enable the providers to adjust the level of the reconfiguration and accordingly control possible traffic disruptions. An integer linear program (ILP) is formulated for each solution. Because the defined ILPs are $\mathcal{NP}$-hard, a novel heuristic algorithm is also proposed. The proposed energy-saving methods are evaluated over random VNE scenarios. The results confirm that the solutions save notable amounts of energy in physical nodes and links during the off-peak period.

1. Introduction

Several reports from different information and communications technology (ICT) organizations from around the world have confirmed the increasing energy demands in this technology, which is a major concern for future Internet. In the case where no green technology will be deployed in communication networks, the Global e-Sustainability Initiative (GeSI) predicts an energy consumption of 35.8 TWh for European telecom operators in 2020 [1].

Recently, virtualization has emerged for communication networks in ICT. It shares the resources in the network environment [2]. A virtualized network environment (VNE) supports the coexistence of multiple virtual networks (VNs) over a single substrate network [3]. Virtual nodes and links are mapped onto physical nodes and links, respectively. Network virtualization decouples the functionality of the current network architecture into infrastructure providers (InPs) and service providers (SPs). In addition, virtual machines (VMs) traditionally virtualize physical servers' resources. VNs together with VMs underpin virtualized data centers (VDCs). Traditional data centers are moving toward virtualized data centers to address cloud computing limitations regarding network performance, security, and manageability [4]. Hence, network virtualization has been regarded as a promising technology to flexibly share the resources, and therefore, the corresponding solutions to energy savings in this type of network become essential.

In fact, virtual networks' traffic loads change over time. Virtual networks might be highly utilized during a period of time (peak time, e.g., day hours), whereas they are under-utilized during another notable period of time (off-peak time, e.g., night hours). Virtual networks' processing demands are likely related to their traffic loads. Traffic

* Corresponding author. Tel.: +1 6138188373.
E-mail addresses: eghazisaeedi@sce.carleton.ca (E. Ghazisaeedi), huang@sce.carleton.ca (C. Huang).

http://dx.doi.org/10.1016/j.comnet.2015.09.041
1389-1286/

variations in virtual networks correspondingly change the utilization of the substrate network. The reports for 40 North American and 25 European network providers reveal a 60% difference between peak and minimum off-peak traffic rates over their substrate networks [5]. However, today's substrate networks are provisioned to support VNs' peak time traffic demands, with some additional over-provisioning accommodating unexpected traffic rates [5]. The substrate network's elements are always switched on, neglecting the traffic behavior.

Cloud providers could determine the off-peak time period of the substrate network and traffic/processing demands of each VN in that period through information given by VNs' customers or using network traffic prediction techniques, e.g., [6,7], that estimate the future traffic by examining the current traffic state. During the off-peak period, it is possible to reduce the VNE's energy consumption by reconfiguring the mapping of the already embedded VNs according to their decreased demands. In this context, virtual networks are accepted and embedded onto the substrate network by a normal (not energy efficient) VNE embedding process to accommodate the peak traffic behavior. The cloud provider runs the reconfiguration technique during normal network operations on networks transitioning from the peak period to the off-peak period to save energy in the off-peak period. This is typically called elasticity in cloud terminology, and it can be performed in three ways: proactive cyclic scaling, proactive event-based scaling, and auto-scaling based on demand. However, when the traffic load changes from a peak level to an off-peak level, some traffic flows that last in both time periods might suffer from traffic disruptions imposed by applying the reconfiguration [5]. Moreover, reconfiguring the mapping of embedded VNs may require additional signaling traffic that is necessary for notifying all the involved routers [8]. This may introduce a significant work load for the signaling controller, particularly when the reconfiguration attempts to make changes to a large number of nodes at the same time. Consequently, it may not be a good practice to reconfigure the mapping of every embedded virtual node/link.

The reconfiguration process requires four steps. First, it learns the peak and off-peak periods and the associated demands. Second, it calculates the off-peak mapping. Third, at the beginning of each off-peak period, it reconfigures the mapping of all the affected VNs based on the result from the second step. Fourth, at the beginning of each peak period, it restores the original mapping.

The major cost of the reconfiguration consists of three components: collecting traffic information with each VN, calculating the off-peak mapping, and conducting the reconfiguration operation. Because the third item is supposed to be completed as fast as possible, its energy cost is insignificant. The first and second items depend on how often they are performed. For example, if all the traffic of all VNs are roughly periodic, the first and second items only need to be performed once whenever a new VN is added or a VN is leaving.

In this paper, we discuss multiple novel energy-saving solutions that optimize a VNE's node and link energy consumption during the off-peak time according to two defined power models. We approach the problem by reconfiguring the mapping for some already embedded virtual nodes and links. We assume that the substrate and virtual networks are homogeneous. All the substrate and virtual nodes are assumed to be switches/routers to reflect the network environment. This is the case in the majority of existing related research studies.

In this regard, we first formulated a mixed integer linear program (MILP) for off-peak node and link energy optimization through local node and link reconfiguration. A stress rate is defined for substrate nodes. Accordingly, the MILP *may* set less stressed substrate nodes and their adjacent substrate links into sleep mode for the off-peak time. We re-map a virtual link *if and only if* we sleep the substrate node that hosts its source/sink virtual node or at least one substrate node that relays its traffic over its embedded path. The energy consumption of a re-mapped virtual link is also minimized.

Because reconfiguring the mapping for both virtual nodes and virtual links may be expensive and cause interruptions to normal network operations, we derive another methodology that optimizes the node and link energy consumption of a VNE during the off-peak time by reconfiguring the mapping for only some of the virtual links. We do not reconfigure the mapping of the embedded virtual nodes in this case. We call the latter problem off-peak node and link energy optimization by local link reconfiguration. We define a different stress rate for intermediate substrate nodes, which do not host a virtual node and only relay the traffic. Accordingly, a solution is proposed to optimize the total energy consumption of the intermediate substrate nodes and substrate links during the off-peak time. This method *might* set the less stressed intermediate substrate nodes and their respective substrate links into sleep mode for the off-peak time. We re-map a virtual link *if and only if* we sleep at least one intermediate substrate node over its embedded path. The energy consumption of a re-mapped virtual link is also minimized. A MILP for splittable traffic and a binary integer linear program (BILP) for non-splittable traffic are formulated for this strategy.

These methods enable providers to change level of the reconfiguration by adjusting the stress rate's threshold. Therefore, they can control the possible traffic interruptions of the reconfiguration. Clearly, there is a trade-off between the energy-savings level and the possible traffic interruptions.

Because the formulated optimization programs are $\mathcal{NP}$-hard, we also propose a heuristic algorithm for off-peak node and link energy optimization through local link reconfiguration. The simulation results confirm that the heuristic algorithm can achieve points that are close to the optimum points set by the formulated optimization program, while it is considerably faster. The proposed heuristic is scalable to large network sizes.

The remainder of this paper is organized as follows. Most of the related works in the literature and our contributions in this paper are discussed in Section 2. Two power models for physical nodes and links are studied in Section 3. Multiple optimization programs are defined and formulated as ILPs in Section 4, and the proposed heuristic is discussed in Section 5. To evaluate the performances of the ILPs and the heuristic algorithm, multiple random scenarios have been simulated over randomly generated VNEs, and the results are analyzed in Section 6. The paper will conclude in Section 7.