Optical Switching and Networking 11 (2014) 154–166

Contents lists available at ScienceDirect

Optical Switching and Networking

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

Local protection: A cost efficient technique for reliable virtual infrastructure design

Hao Di [a,*], Vishal Anand [b], Dan Liao [a,c], Du Xu [a], Lemin Li [a], Gang Sun [a,c], Ninghai Bao [d]

[a] Key Lab of Optical Fiber Sensing and Communications, Ministry of Education, University of Electronic Science and Technology of China, China
[b] Department of Computer Science, The College at Brockport, State University of New York, USA
[c] Institute of Electronic and Information Engineering in Dongguan, UESTC, China
[d] School of Communication and Information Engineering, Chongqing University of Posts and Telecommunications, China

ARTICLE INFO

Available online 21 June 2013

Keywords:
Virtual infrastructure
Virtual infrastructure reliability
Multiple clusters
Local protection

ABSTRACT

Network virtualization allows multiple virtual infrastructures (VIs) to coexist on the same shared underlying substrate infrastructure, thus simplifying resource management, increasing infrastructure utilization and reducing operational costs. However, with this sharing a single server failure will affect all the VIs mapped on to it. Thus, backup resources should be reserved intelligently to provide cost effective VI reliability.

In this paper, we study the reliable VI mapping over multiple clusters of servers that are geographically distributed and interconnected over a wide-area network. Since the bandwidth cost on the wide-area network is much higher than that in a cluster, we present the local protection based reliable VI mapping algorithm (LP-RVIM*) to minimize the total reliable mapping cost by only using backup VI nodes in the same cluster to protect the primary VI nodes. The LP-RVIM* uses the additional backup cost metric to cost efficiently share backup VI nodes by primary VI nodes. Through detailed simulations we show that the LP-RVIM* algorithm can reduce the VI mapping costs.

1. Introduction

Using network virtualization technology [1,2] different networking approaches and applications can be deployed on the same shared underlying substrate network architecture. Network virtualization can help researchers develop new technologies and networking paradigms without having to redesign or by making only minimal changes to the substrate network [3]. In addition, by using the cloud computing [4] paradigm one can also deploy various applications without the expenses of maintaining the substrate.

In such a network virtualization environment, the computing and networking resources such as CPU, memory, and communication bandwidth that are required by the applications can be modeled as virtual infrastructures (VIs). Accordingly a resource request from the VI (a VI request) consists of virtual nodes (computing nodes) and virtual links (connections) between these nodes. The VI nodes and links of a VI request indicate the computing and bandwidth resource requirements of the application. To deploy an application, each VI node is assigned to a physical server in the form of virtual machine (VM), and each VI link is assigned to a substrate link or path. Thus, the VI mapping problem can be defined as: mapping the VI request onto the underlying substrate network while

* Corresponding author. Tel.: +86 13438114723.
E-mail addresses: haodi84@gmail.com, dihao@uestc.edu.cn (H. Di). vanand@brockport.edu (V. Anand), liaodan@uestc.edu.cn (D. Liao). xudu.uestc@gmail.com (D. Xu), gangsun@uestc.edu.cn (G. Sun). baonh@cqupt.edu.cn (N. Bao).

1573-4277/$ - see front matter
http://dx.doi.org/10.1016/j.osn.2013.06.004

satisfying the resource requirements of the VI request subject to the resource constraints of the underlying substrate network. The works in [5–10] study the VI mapping problem without considering fault tolerant against substrate failures. More recently reliable or survivable VI mapping is studied in [11–18], where backup computing and bandwidth resources are reserved for possible substrate failures.

Today's data centers contain thousands of computers/servers to support a myriad of cloud services that may require significant computing and bandwidth requirements. In this paper, we study reliable VI mapping over *multiple* clusters of servers (i.e., data centers [19,20]) with the aim of minimizing the total mapping cost. The clusters are geographically distributed and interconnected over a wide-area network [21] with the VI requests mapped over multiple clusters [9]. Different from the study of reliable VI mapping in a *single* cluster [17,18], the bandwidth cost on the wide-area network is much higher than that in a cluster. In [17,18] backup VI nodes are added in the virtualization layer to guarantee a desired reliability, and then backup VI links connecting the neighbors of each protected primary VI node and any one of the backup VI nodes protecting the primary VI node are also added. These backup VI links support the communication between each protected primary VI node and its neighbors after recovering/replacing the protected VI node by a backup VI node. Sharing backup VI nodes by the primary VI nodes can lead to the reduction of the required backup computing resources. In [17,18] all the protected primary VI nodes are protected together, i.e., share the same backup VI nodes, so as to reduce the required number of backup VI nodes as much as possible. The bandwidth within a cluster is quite cheap, and in [17] this bandwidth cost is not included in the simulation. However, if the backup computing resources are shared across multiple server clusters (i.e., the protected VI nodes and the corresponding backup VI nodes are not in the same cluster), the backup connections will pass through the wide-area network, where the backup bandwidth cost is considerable and must be taken into consideration. Thus, in this study we use the backup computing resources in the same cluster to protect the mapped primary VI nodes so as to avoid the high backup bandwidth cost. To further reduce the cost of backup computing resources we share the backup computing resources among the primary VIs in each cluster.

In this study, we assume that each physical server in a cluster can fail with a known failure rate or probability, e.g., 1/MTBF (mean time between failures). The VI requests are mapped onto the substrate so as to guarantee a desired reliability. Similar to [17,18], the reliability of a VI request is defined as the probability of all VI nodes remaining in operation. A cluster contains servers that are located in multiple racks, which are then interconnected by switches and routers [19,20]. Since servers in different racks can be connected to different switches and power supplies, we assume server failure probabilities in different racks of a cluster or in different clusters to be independent. We design the local protection based reliable VI mapping algorithm (LP-RVIM*) to map the VI requests onto the substrate to guarantee the desired reliabilities, while minimizing the total (computing and bandwidth) mapping cost. In LP-RVIM*, the primary VI nodes mapped in a cluster are protected by the shared backup VI nodes in the same cluster to avoid the high backup bandwidth cost on the wide-area network. Since multiple VI requests are mapped on the same substrate, some primary VI nodes are in the same clusters, which can be protected together to reduce the backup computing cost. We use the additional backup cost metric while mapping each primary VI node, and choose the mapping that reduces this metric and promotes cost efficient sharing of backup computing resources in a cluster.

The rest of the paper is organized as follows. Section 2 describes the related works. Section 3 describes reliable VI mapping over multiple clusters using local protection. Section 4 formulates the reliable VI mapping problem as a mixed integer programming (MIP) problem. Section 5 describes the LP-RVIM* algorithm. Section 6 presents the simulation results, and Section 7 concludes the paper.

2. Related work

The VI mapping problem is how to map VI nodes and links of VI requests onto the shared substrate while satisfying the resource demands of the VI requests subject to the resource constraints of the substrate. There has been some work on non-survivable VI mapping [5–10]. In these studies, VI survivability or reliability is not considered; the strategies developed focus on minimizing the cost of computing and bandwidth resources allocated for VI requests (i.e., mapping cost) and/or maximizing the number of coexisting VIs.

In a network virtualization environment multiple VIs are mapped on the same substrate, and hence several VIs may be affected by the failure of a single device or equipment in the substrate. Thus, it is important that VIs should be fault tolerant against substrate failures. The survivable or reliable VI mapping problem is more complex than survivability in single layer networks [22–24] due to the fact that there may be numerous ways to map a VI request including its VI nodes and VI links onto the substrate network, and accordingly many ways to recover (remap) a mapped VI request based on the nature of the failure and mapping. Furthermore, since multiple VI nodes and links from different VI requests may be mapped onto the same substrate components increases the problem complexity. Thus survivability or reliability guarantee for VI requests must be coordinated with the VI mapping problem, i.e., both virtualization layer and substrate layer should be considered. The survivable VI mapping problem has similarities with approaches for optical layer survivability in WDM networks [25–27], and hence some of the lightpath protection approaches can be applied for protecting VI links. However, in the study of optical layer survivability only backup bandwidths are reserved to protect lightpaths, and the survivability for computing resources on physical servers is not included. While considering VI survivability, the bandwidth resources on substrate links and the computing resources on physical