

Efficient VNF Migration and Recovery in Edge Computing Environments

Yu Chen[†], Sheng Zhang^{†*}, Yanchao Zhao[‡], Zhuzhong Qian[†], Yu Liang[†], Jidong Ge[†], Sanglu Lu[†]

[†]State Key Lab. for Novel Software Technology, Nanjing University

[‡]Nanjing University of Aeronautics and Astronautics

Corresponding author: sheng@nju.edu.cn

ABSTRACT

The combination of network function virtualization (NFV) and software defined networking (SDN) contributes to making the network more efficient, and it promises to guarantee high flexibility and low cost in network operations. Meanwhile, the continuous development of mobile edge computing (MEC) makes it possible for the edge of mobile network to run cloud computing tasks. However, edge computing environments are dynamic, and VNF demands for computing resources usually fluctuate over time. Thus, when the demand of a VNF cannot be satisfied, we need to migrate the VNF so as to meet its demand and retain the network performance. Our paper focuses on efficiently migrating VNFs (MV) and recovering connections (RC) after migration, such that the migration and connection paths can meet the bandwidth requirement for data transmission. We prove that both of MV and RC are NP-complete. We present efficient algorithms to tackle them. Extensive simulations show that the proposed algorithms are efficient and effective.

1. INTRODUCTION

Mobile edge computing (MEC) [2] enables the cloud computing capabilities at the edge of the mobile network, which greatly lowers the delay in data transmission. Virtual network functions (VNFs) can be dynamically deployed [3], migrated [4], and removed.

Due to the dynamic nature of edges, the network changes from time to time. For instance, the topology of the network may change, the resource capacity of a specific machine may vary at some point, and the bandwidth between two machines may increase or decrease due to different traffic situations. In the daytime, some VNFs may require more computing resources due to heavy traffic situations, and they need to be migrated to the machines with massive resources in central network nodes; in the night, due to the decreasing demand for computing resources, some VNFs need to be migrated back to edge network nodes with fewer resources, such that the machines with massive resources can be va-

lated for sudden large tasks. Therefore, we need to efficiently complete the migration of the VNF chains, as well as the task of VNF connection recovery after migration.

In this paper, we use ‘to-be-migrated VNFs’ to represent the VNFs which we need to migrate. We consider the problem of migrating multiple VNFs to new physical machines, as the to-be-migrated VNFs’ demands for computing resources fluctuate due to different traffic situations. As a first step, to initiate a meaningful study, this paper narrows the scope of this problem to a manageable extent: we assume the network node to which each VNF is migrated is known in advance; thus, we focus on selecting the migration paths that meet the bandwidth requirements for all to-be-migrated VNFs, such that the task of migrating VNFs can be completed within the specified time. After migration, we need to recover the connection between the migrated VNFs which are adjacent along a VNF chain. To prevent the traffic congestion in the network, we need to find out the connection paths that minimize bandwidth usage.

Related work. Zhang et al. [4] discussed the benefits and challenges of VM migration and lists many open issues. Eramo et al. [1] focuses on saving the energy consumption of migration. Different from previous works, in this paper, VNF migration and connection recovery are jointly considered, and we propose effective algorithms to tackle them.

2. PROBLEM FORMULATION

Migrating VNFs (MV). We assume that there are c VNF chains to be migrated, and the j th ($j = 1, 2, \dots, c$) VNF chain has C_j to-be-migrated VNFs. Thus, the total number of to-be-migrated VNFs is denoted as $n = \sum_{j=1}^c C_j$. Besides, the i th ($i = 1, 2, \dots, n$) to-be-migrated VNF is denoted as f_i , and we migrate it from the current node $S_{M,i}$ to the destination node $D_{M,i}$, which is specified in advance. In the cloud-based edge network, there are $r_{M,i}$ paths between $S_{M,i}$ and $D_{M,i}$, and they are $R_{M,i}^1, R_{M,i}^2, \dots, R_{M,i}^{r_{M,i}}$. We need to select one of them, e.g., $R_{M,i}^k$ ($k = 1, 2, \dots, r_{M,i}$) as the migration path $P_{M,i}$ for f_i . We define $BW(R_{M,i}^k)$ as the bandwidth of $R_{M,i}^k$, which equals the minimum bandwidth of the links along $R_{M,i}^k$. We allocate $B_{M,i}$ units of bandwidth for $P_{M,i}$, and $B_{M,i}$ is no more than $BW(R_{M,i}^k)$. If multiple migration paths traverse the same link e , the total bandwidth allocated for them should not exceed B_e .

We start to migrate the n to-be-migrated VNFs at the same time, and all the migration tasks need to be finished within the specific time T . We assume the amount of the to-be-migrated data of f_i is d_i . Remember that the bandwidth

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICPP'19, Aug 5–8, 2019, Kyoto, Japan.

Copyright 2019 ACM 978-1-4503-2193-8/13/07 ...\$15.00.

we allocate for f_i along $P_{M,i}$ is $B_{M,i}$, so the time required for f_i 's migration is $d_i/B_{M,i}$. All to-be-migrated VNFs need to be migrated to the destinations within time T , so we should select an appropriate migration path for each VNF. There is no motivation for each VNF to finish migration in the time less than T , because it will occupy extra bandwidth that might be utilized by other VNFs. Thus, the bandwidth requirement of $P_{M,i}$ is equal to d_i/T , i.e., $B_{M,i} = d_i/T$.

Recovering Connections (RC). After migration, we need to recover the connection path between each migrated VNF and its neighbor VNFs along the VNF chain. We assume that there are m connection paths to be recovered in total, and they are each selected between node $SC_{,1}$ and $DC_{,1}$, node $SC_{,2}$ and $DC_{,2}$, ..., node $SC_{,m}$ and $DC_{,m}$. In the cloud-based edge network, there are $r_{C,i}$ ($i = 1, 2, \dots, m$) paths between $SC_{,i}$ and $DC_{,i}$, and they are $R_{C,i}^1, R_{C,i}^2, \dots, R_{C,i}^{r_{C,i}}$. Similar to MV, we should select one of them, e.g., $R_{C,i}^k$ ($k = 1, 2, \dots, r_{C,i}$) as the recovered connection path $PC_{,i}$ between node $SC_{,i}$ and $DC_{,i}$. The bandwidth requirement of the connection path between node $SC_{,i}$ and $DC_{,i}$ is $BC_{,i}$. Thus, we need to allocate $BC_{,i}$ units of bandwidth for $PC_{,i}$, and $BC_{,i}$ is no more than $BW(R_{C,i}^k)$. If multiple recovered connection paths traverse the same link e , the total bandwidth allocated for them should not exceed B_e .

The total bandwidth allocated for the connection paths traversing the link e_i is denoted as t_{e_i} , and we define the bandwidth occupancy rate after recovering VNF connection as $\alpha = \max_{i=1}^{|E|} \frac{t_{e_i}}{B_{e_i}}$. It is easy to see α can reflect the traffic condition of the network after migration. Our objective is to minimize α and prevent the traffic congestion.

NP-completeness result. By reducing the Widest Pair of Paths Problem to RC and reducing the Widest Pair of Disjoint Paths Decoupled Problem to MV, we can both of MV and RC are NP-complete.

3. SOLUTION

Heuristics for MV. We need to select n migration paths satisfying bandwidth requirements for all to-be-migrated VNFs in MV. To ensure that as many VNFs as possible can be migrated successfully, we migrate VNFs in the non-decreasing order of the bandwidth requirements and first select the shortest or the minimum wide-edge migration path, such that each migration path will use fewer links and the bandwidth of each link can be made the most use of on the whole.

Smallest path Heuristic (SH) algorithm. SH starts from the migration path $P_{M,i}$ with the smallest bandwidth requirement. SH utilizes the Dijkstra's Algorithm to calculate the shortest path sp between $S_{M,i}$ and $D_{M,i}$.

Minimum Wide-edge Heuristic (MWH) algorithm. MWH differs from SH only in that instead of using Dijkstra we use a similar algorithm that finds the minimum wide-edge path between two given vertices.

Heuristics for RC. To minimize the bandwidth occupancy rate and uniformly use all links' bandwidths, we select the migration paths that have large bandwidth capacity and traverse few links.

Widest Heuristic (WH) algorithm. WH starts from the path with the smallest bandwidth requirement, and checks whether there is any path that satisfies its demand.

Modified Shortest Heuristic (MSH) algorithm. Different from WH, MSH uses the Dijkstra's algorithm to find a suitable for a given vertices pair.

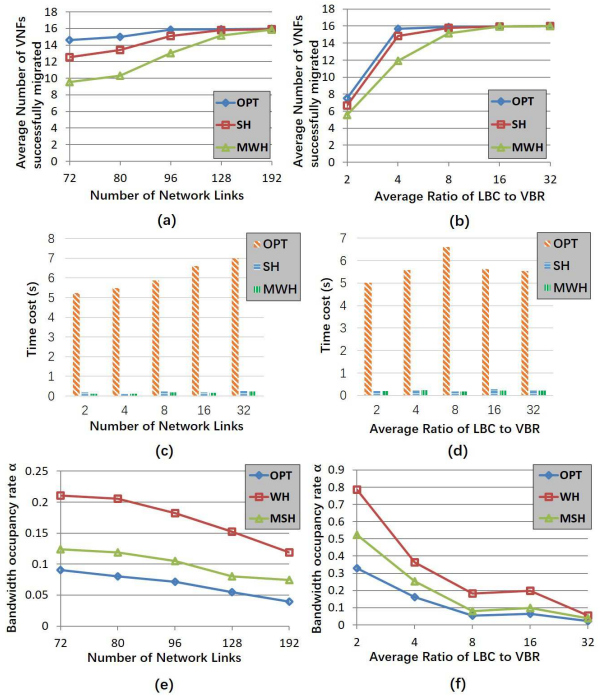


Figure 1: Simulation results

4. EVALUATION

We simulate the network using the Python package *NetworkX*. In general, the simulation results demonstrate the efficiency and effectiveness of the proposed algorithms, which are nearly as good as the optimal solution. For example, in Fig. 1(a)(b), we can see that the performance of SH and MWH is close to the optimal solution when the number of links is more than twice that of network nodes, i.e., 128. Besides, when the ratio of LBC to VBR in the network is larger than 8, SH and MWH also perform nearly as well as OPT. From Fig. 1(c)(d), we learn that the running time cost of SH and MWH is much less than that of OPT. Thus, when there are a number of links in the network and the ratio of LBC to VBR is not very low, our algorithms for MV can achieve good results.

Acknowledgements

This work was supported in part by NSFC (61872175), NSF of Jiangsu Province (BK20181252), and Collaborative Innovation Center of Novel Software Technology and Industrialization. Sheng Zhang is the corresponding author.

5. REFERENCES

- [1] V. Eramo, M. Ammar, and F. G. Lavacca. Migration energy aware reconfigurations of virtual network function instances in nvf architectures. *IEEE Access*, 5:4927–4938, 2017.
- [2] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies. The case for vm-based cloudlets in mobile computing. *IEEE Pervasive Computing*, 8(4):14–23, Oct 2009.
- [3] S. Wang, M. Zafer, and K. K. Leung. Online placement of multi-component applications in edge computing environments. *IEEE Access*, 5:2514–2533, 2017.
- [4] F. Zhang, G. Liu, X. Fu, and R. Yahyapour. A survey on virtual machine migration: Challenges, techniques and open issues. *IEEE Communications Surveys Tutorials*, pages 1–1, 2018.