Energy-Efficient Middlebox Placement in SDN-Enabled Data Centers

Daniel Hernandez, Alexander Ing, Bin Tang, Mohsen Beheshti
danhernandez, aing1@toromail.csudh.edu, btang,mbeheshti@csudh.edu
Computer Science Department, California State University Dominguez Hills

Abstract

Data center policies require that VM traffic in data centers traverse a sequence of specified middleboxes, for the purposes of security and performance. In cloud data center, the integration of Software Defined Network (SDN) and Network Function Virtualization (NFV) has been recently proposed to enables efficient placement of software-based middblebox in commercial off the-shelf switches. There are a few studies that address such middlebox placement problem, however, none of them are from an algorithmic angle. In this poster, we first formulate this problem formally. We then propose two time efficient heuristic algorithms viz. Random and Greedy to solve the problem. We show via extensive simulations that Greedy outperforms Random in energy consumptions in all different network parameters.

Introduction

A middlebox, also called network appliance, is a computer networking device that transforms, inspects, filters, or otherwise manipulates traffic for purposes other than packet forwarding. Traditional middlebox hardware is widely deployed in enterprise networks to improve network security and performance. The widespread deployment of middleboxes has resulted in some challenges and criticism due to poor interaction with higher layer protocols . To address above challenges, the integration of Software Defined Network (SDN) and Network Function Virtualization (NFV) has been recently proposed to enables efficient placement of software-based middleblox in commercial off the-shelf switches.

Fat-Tree Data Center

A k-ary fat-tree is shown in Fig. 1 with k = 4, where k is the number of ports of each switch. There are three layers of switches: edge switch, aggregation switch and core switch A fat-tree built with k-port switches supports $k^3/4$ physical machines. In the small data center of Fig. 1, there are 16 physical machines.

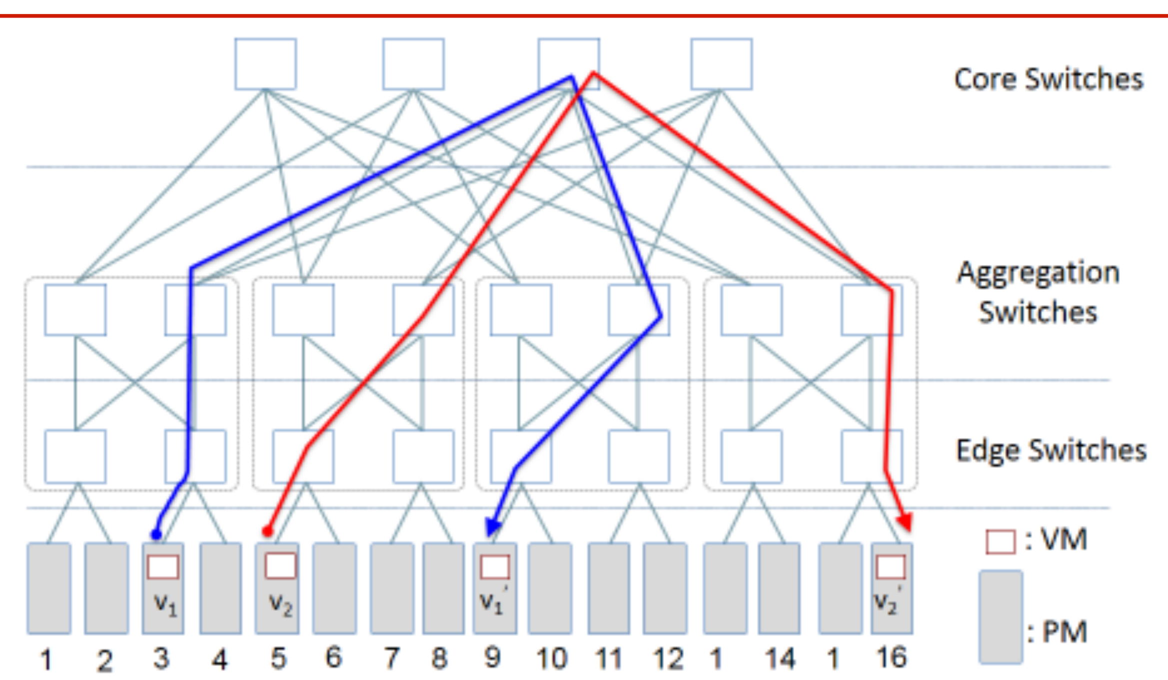

Figure 1

Problem Formulation

- Given a data center graph $G(V,E)$ (See Figure 1)
- There are m MBs $M=\{mb_1, mb_2, \ldots, mb_m\}$ to be placed inside the data center
- A set of p communicating node pairs P, each pair (s,t) in P needs to traverse $mb_1, mb_2, \ldots, mb_m$ in that order
- The cost for $p = (s,t)$ is $c(p) = d(s, mb_1) + d(mb_1, mb_2) + \ldots + d(mb_{m-1}, mb_m) + d(mb_m, t)$
- Goal: where to place the m MBs in data centter such that the total cost of all p pairs is minimized

Algorithms and Time Complexity

Random Placement Algorithm:

- Randomly place the m MBs inside the data center
- Time complexity: $|V|^3 + p$

Greedy Placement Algorithm:

- Takes place in m rounds
- In round i, it places mb_i at a node that minimizes the total communication cost in that round
- Time complexity: $|V|^3 + |V|^m p$

Simulation Parameters and Analysis

- *k = 8, 16, and 32*
- *Number of servers = 128, 1024, 8192*
- *p = 100, 200, 300, 400, and 500*
- *m = 1, 3, 5, 7*
- Greedy outperforms Random in all sets of parameters
- When k is fixed, with the increase of p, cost increases
- When p is fixed, with the increase of k, cost remains the same. The network diameter (maximum number of hops between any two nodes) does not increase with increase of k.
- When p is fixed, with the increase of m, performance difference between Random and Greedy increase. This indicates Greedy works better for large number of middleboxes.

Figure 2

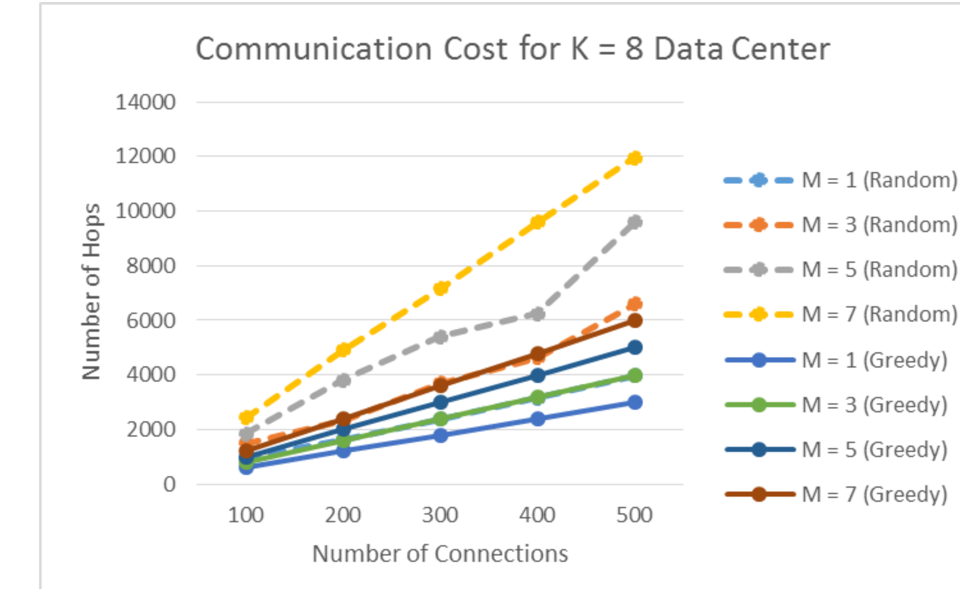

Figure 3

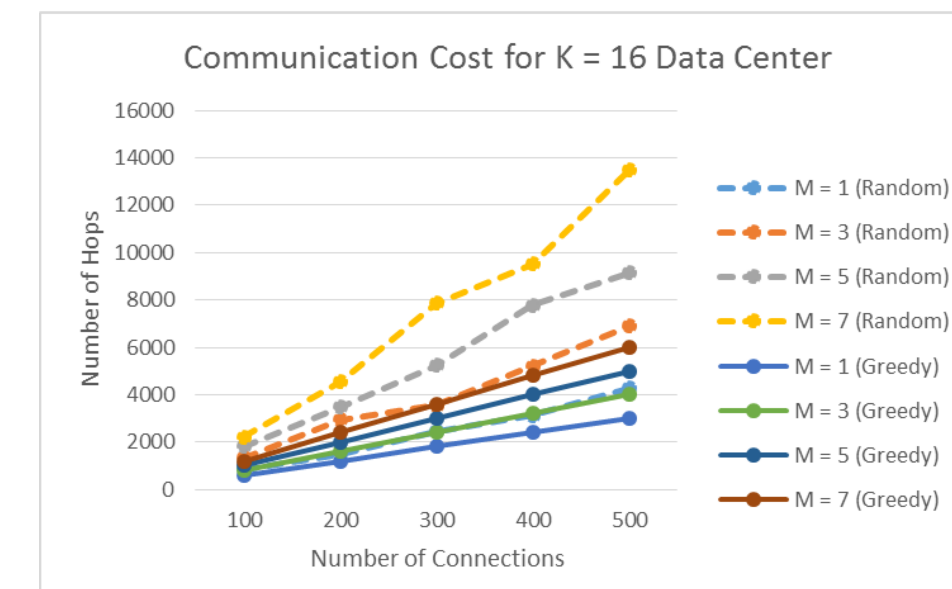

Figure 4

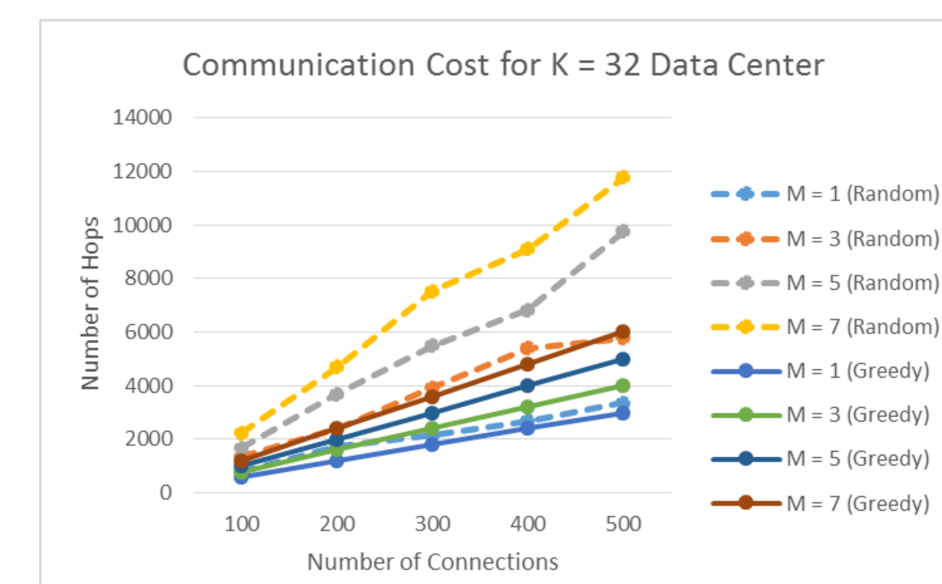

Conclusions

- We formulate energy-efficient middlebox placement problem and designed two time-efficient algorithms, in SDN-enabled data centers
- Extensive simulations show that Greedy always outperforms Random in all different network scenarios.

Acknowledgements

This research is funded through the National Science Foundation (NSF) under grant #HRD-1302873 and the Chancellor's Office of the California State University. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the Chancellor's Office of the CSU.