Virtual Machine Pair Communication Cost Reduction via Replication in Datacenters

Pedro Ledesma, Bin Tang, Mohsen Beheshti
pledesma3@toromail.csudh.edu, btang, mbeheshti@csudh.edu
Computer Science Department, California State University Dominguez Hills

California State University DOMINGUEZ HILLS

Abstract

The increasing popularity of Infrastructure as a Service (IaaS) have made the efficiency of Cloud Data Centers (CDCs) paramount as they provide the computational resources needed to run online services on the scale that modern companies require. CDCs are able to provide such services with the help of virtualization technologies which make it possible for many Virtual Machines (VMs) to run on one physical host. The VMs that constitute a service provided by the CDC may be in different physical machines. Thus minimizing the communication cost (i.e. the number of switches traversed) between a VM pair can lead to increased throughput in the CDC. We aim to reduce the total communication cost between all VM pairs in a CDC by replicating a portion of the VMs and placing them into other physical host such that there is a reduction in the total communication cost. We have developed three algorithms that we will test in this paper: Greedy, Cluster, and Exhaustive. We test our algorithms in a simulation environment written in Python using a k = 8 Fat Tree topology containing 500 to 2000 VM pairs. Our results show that exhaustive performed the best, but is the most computationally expensive.

Introduction

With the advent of Infrastructure as a Service (IoS), Cloud Data Centers are increasingly providing the computational resources needed to run online services such as social networking applications, video on demand streaming services, and online gaming platforms on the scale that modern companies require. This is done with the help of server virtualization technologies such as VMWare, Xen, and Microsoft Virtual Services which make it possible for a single physical server in the data center to host many different Operating System environments. These environments are then able to host independent processes segmented into Virtual Machines (VMs) in order to provide companies with the minimum computational resources required to run their services. The amount of data that needs to be transmitted to VMs in a data center is expected to rise, with growing demand it is vital to improve data flow or risk poor quality of service. One such way of improving quality of service is to minimize data congestion, which can be accomplished with proper virtual machine management. Virtual machine management can be further improved with the use of software defined networks and network function virtualization. In software defined cloud data centers, data storage facility runs on the internet and allows all infrastructure elements such as, networking, storage, CPU, and security, to be virtualized and delivered as a service. A software defined network means that deployment, operation, provisioning, and configuration are abstracted from hardware and implemented through software intelligence using a centralized controller that can oversee the whole network.

Problem Formulation

- Given a datacenter with a mapping of communication pairs
- Replicate select existing VMs into PMs with sufficient capacity
- Minimize total communication cost within the datacenter.

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.

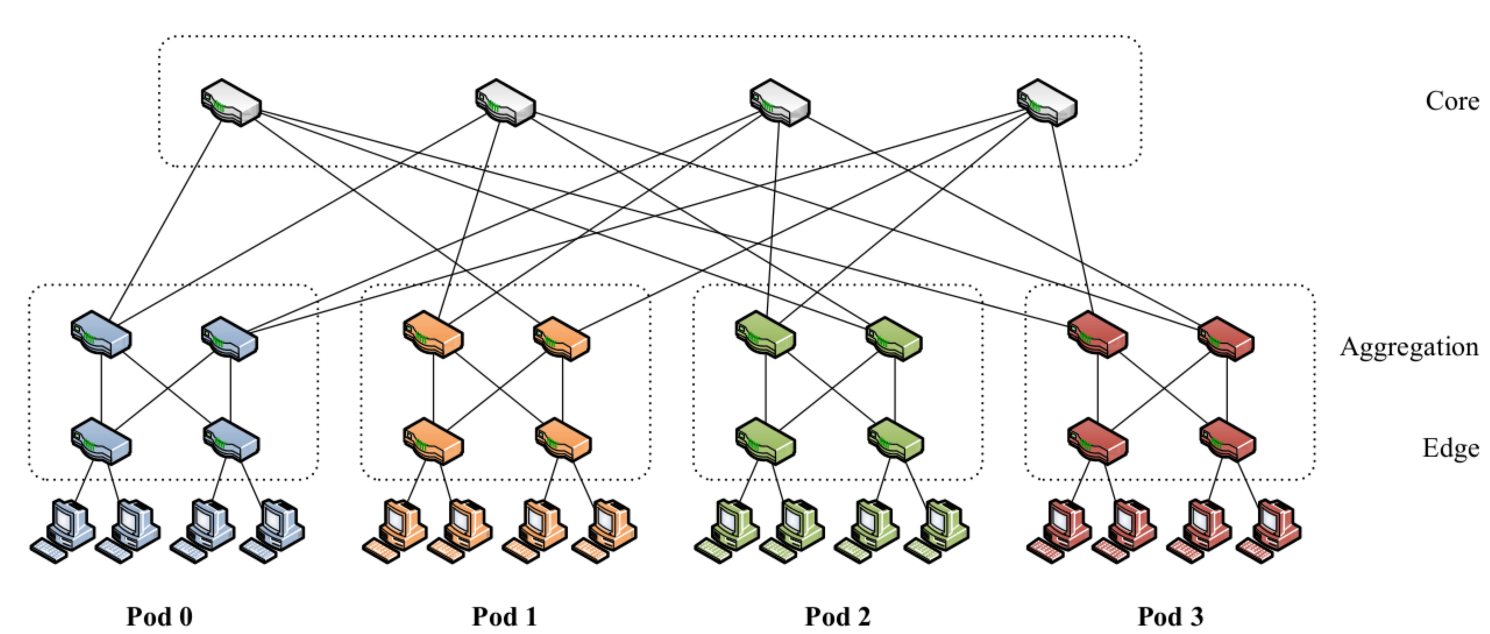

Figure 1

Algorithms and Time Complexity

Greedy Algorithm:

- List all VM pairs with starting with highest frequency
- Search all physical machines for available capacity
- Use the ordered list of VM pairs starting from highest frequency and list of physical machines available capacities to determine replication placement and reduction communication cost.
- Time complexity: *O(L log L)*

Cluster Algorithm:

- Attempt to replicate VM Pairs in the same physical machines.
- Else attempts replicate under the same edge switch
- Else attempt replicate in the same POD to reduce communication cost
- Else transverse to next and VM Pairs repeat step A until ordered list is exhausted.
- Time complexity: $O(L*M*K^3)$

Exhaustive Algorithm:

- For every space available in physical machines, it tries every combination of communication pairs
- Storing the combination resulting in the lowest cost
- Time complexity: O(M*L!)

Simulation Parameters and Analysis

- K=8 Fat Tree Topology
- 2 Total Runs per Simulation Type (Slow)
- Uniform VM Size
- Simulations Alter Storage Space (2 * L + F * L)/128
  - F = 0.50
  - F = 0.25

Figure 2

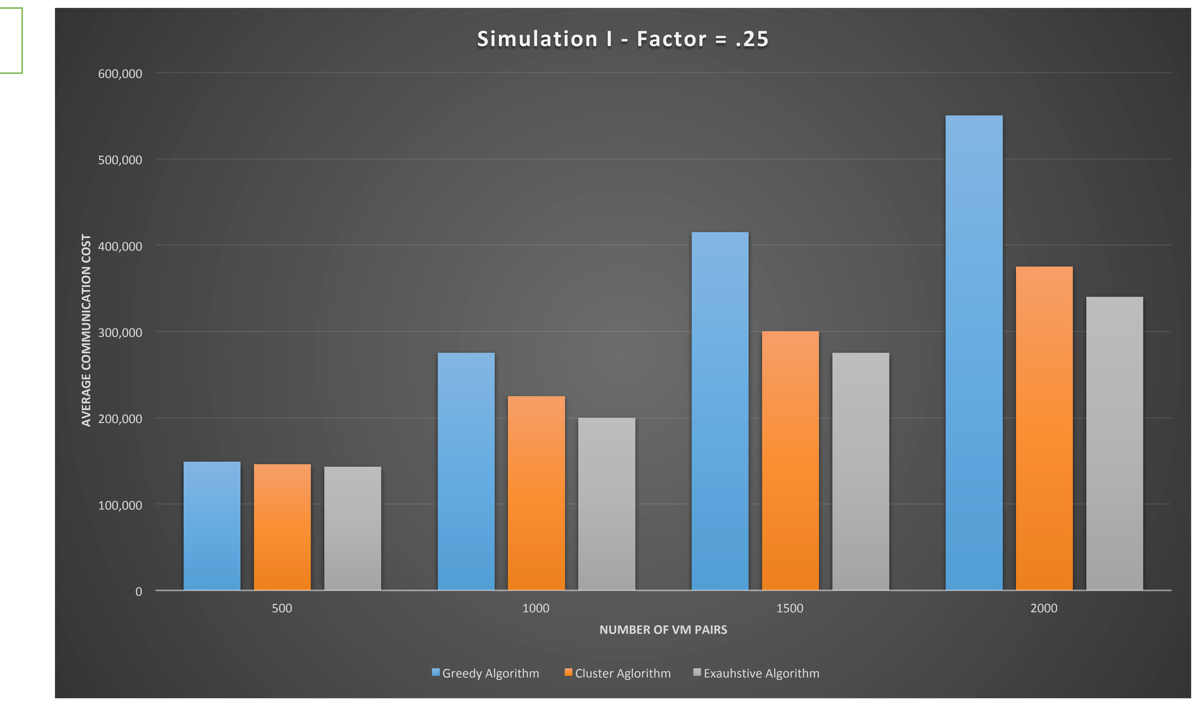

Figure 3

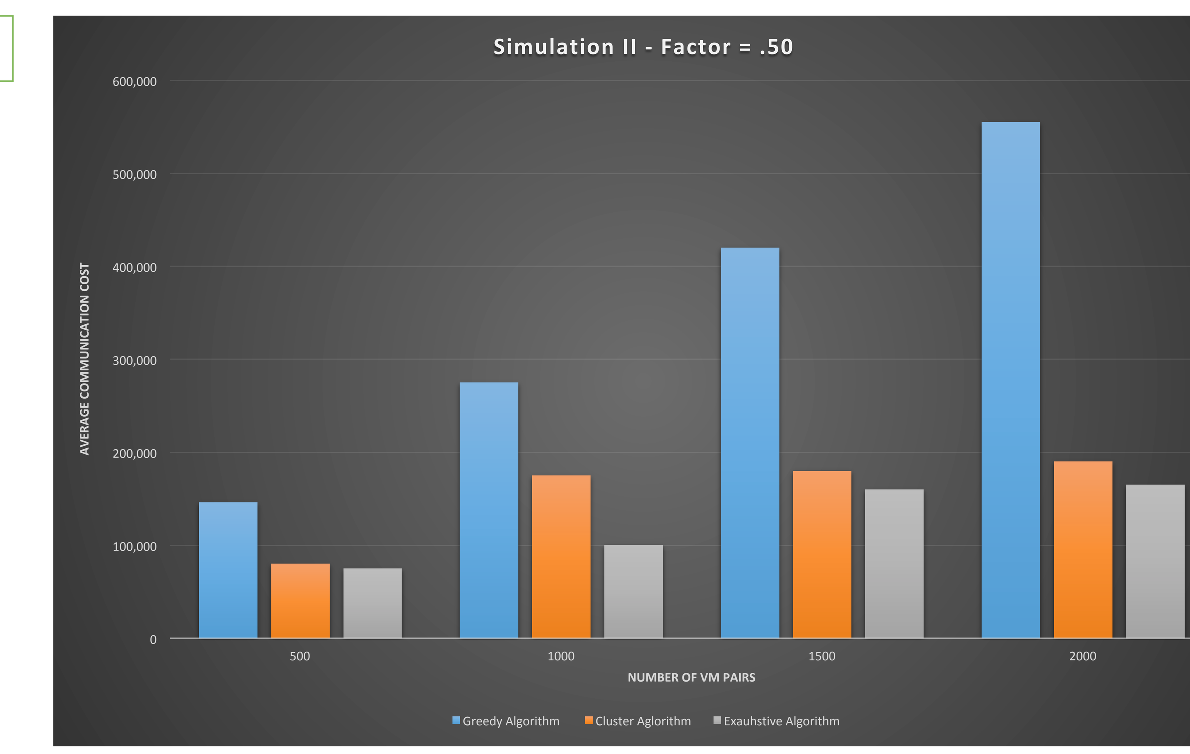

- The results shown above show the average number of hope after replication, for each algorithm.
- You can see that the greedy algorithm uses the most hops. Then the cluster algorithm uses the second most amount of hops. Finally the exhaustive algorithms uses the least amount of hops.

Conclusions

- We formulate three virtual machine replication algorithms, in SDN-enabled data centers
- Extensive simulations show that exhaustive performs the least amount of hops, but is the most computationally expensive.

Acknowledgements

This research is funded in part through the National Science Foundation (NSF) under Grant No.1649271.

CAHSI

NSF