Energy-Efficient Virtual Machine Placement in Policy-Driven Data Centers

{ H. Flores, B. Tang, and M. Beheshti } California State University Dominguez Hills

Abstract

Policy-driven data centers are data centers where all traffic between the Virtual Machines (VMs) that encapsulate hosted services must first traverse a set of Middleboxes (MBs) in order to guarantee a certain level of security or performance. Unlike the VMs in the data center, once placed the MBs in a data center are difficult to move around. In this paper we develop two algorithms, MBUniAware and MBVarAware, that map VM-pairs to physical hosts with sufficient capacity such that the total number of switches every VM-pair's communications must traverse to reach each other is minimized. We prove that MBUniAware is an optimal solution while MBVarAware is a valid approximation of an NP hard problem. We demonstrate the performance of these algorithms via simulations comparing them to existing algorithms in a Fat Tree topology.

Problem Formulation

Assuming the topology has m MBs that each VM pair communication must traverse, we distribute a set of VM pairs P into physical hosts H where each VM pair is comprised of VMs v_i and v_i' and each pair communicates with frequency λ_i. We do this in a way that minimizes the total number of switches each VM pair's communications in the topology must traverse.

Algorithms

MBUniAware

1. Sort hosts in increasing order with respect to their distance to the first, H_s, and last, H_l, MB in the universal traversal sequence
2. Sort VM pairs in decreasing order with respect to their communication frequency, V_o
3. For all VM pairs in V_o, place the first VM in each pair into the first host in the ordered set H_s that can contain it
4. For all VM pairs, place the second VM in each pair into the first host in the ordered set H_l that can contain it

MBVarAware

1. For each MB in the topology create an ordered set containing each physical host in increasing order with respect to their distance from that MB
2. Sort VM pairs in decreasing order with respect to their communication frequency, V_o
3. Using the ordered sets created for each MB, find the two physical hosts closest to any MB and place the first VM pair in V_o in those two physical hosts

Acknowledgements

This research is funded in part through the National Science Foundation (NSF) under Grant No.1551221.

Methods & Approach

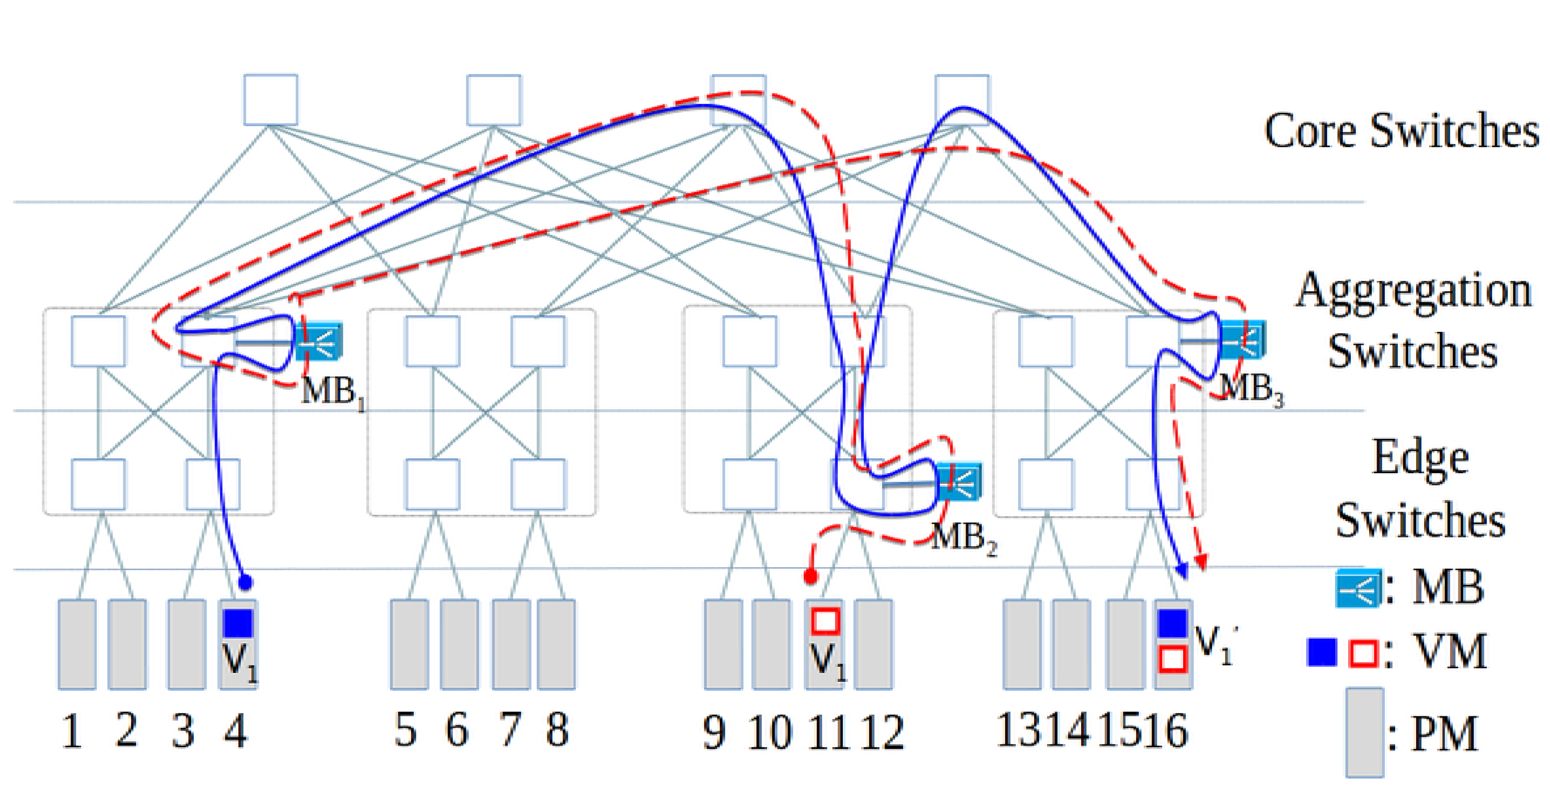

Figure 1: A $k = 4$ Fat Tree topology with two pairs of VMs and one MB set. VM pair one must traverse the set in order while VM pair two can traverse the set of MBs in any order.

General Simulation Parameters:

1. Test in Fat tree topologies with k set to 4, 8, 16
2. Set the number of VM pairs to 100, 500, 1000, 1500, and 2000
3. The number of MBs is set to 3
4. Run simulation 20 times to get 95% confidence interval

MBUniAware Simulation Parameters:

1. MB set is generated randomly and must be traversed in a set order
2. MBUniAware performance compared to random distribution of VMs and traffic- aware VM placement algorithm

MBVarAware Simulation Parameters:

1. MB set is generated randomly and can be traversed in any order
2. MBVarAware performance compared to Host-MB-Host minimum spanning tree algorithm

Results

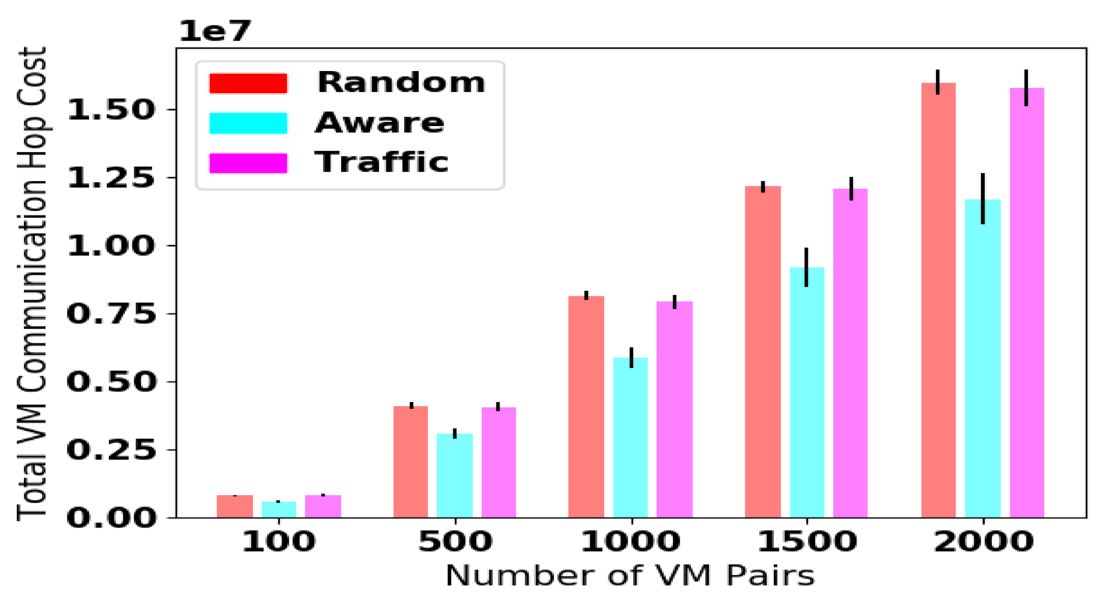

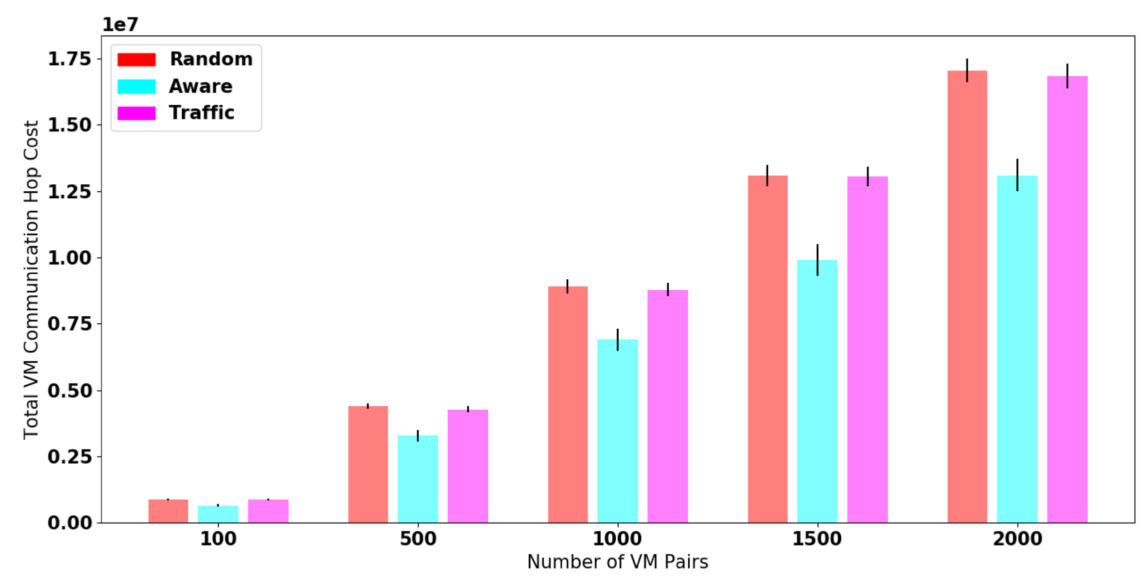

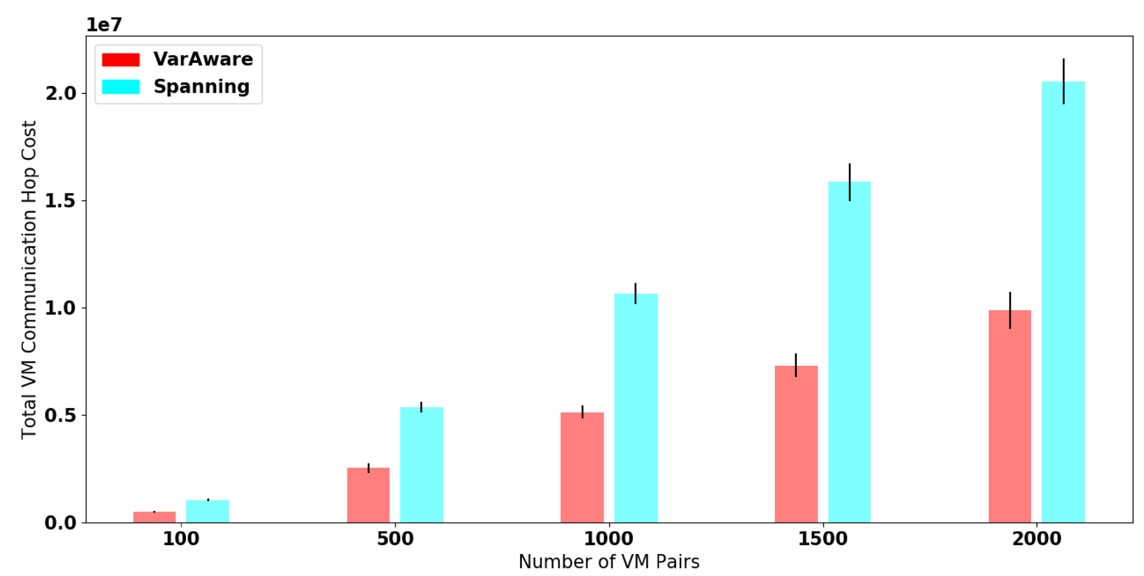

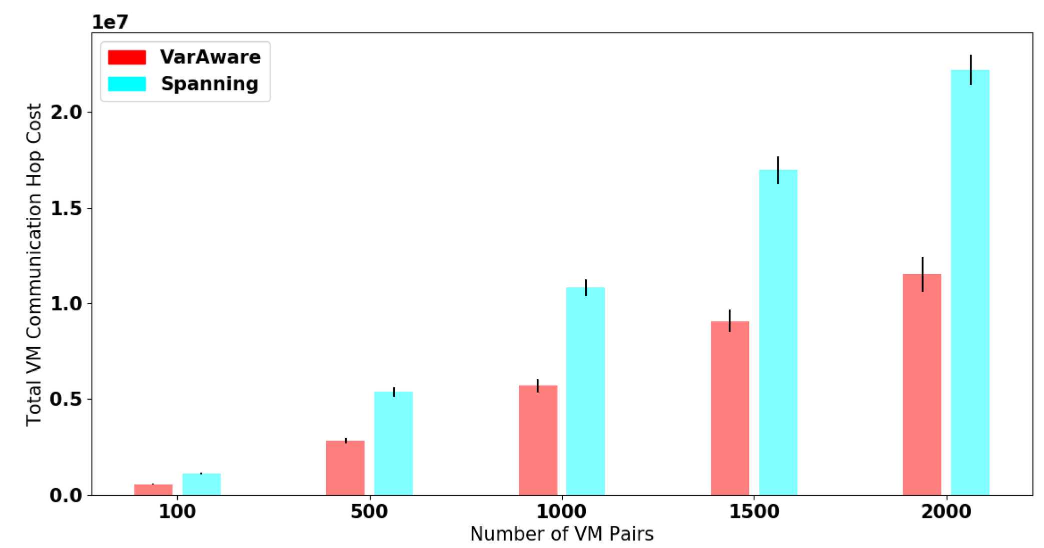

Conclusions

1. In the case where each VM pair's communications must traverse the MB set in a given order, MBUniAware reduces the total number of switches traversed by an average of 31% when compared to a random distribution of VM pairs and 32% when compared to the traffic aware algorithm
2. In the case where each VM pair's communications must traverse the MB set in any order, MBVarAware reduces the total number of switches traversed by an average of 114% when compared to the the minimum spanning tree algorithm