

LB-MAP: LOAD-BALANCED MIDDLEBOX ASSIGNMENT IN POLICY-DRIVEN
DATA CENTER

A Thesis

Presented

to the Faculty of

California State University Dominguez Hills

In Partial Fulfillment

of the Requirements for the Degree

Master of Science

in

Computer Science

by

Manar Alqarni

Summer 2017

THESIS: LB-MAP: LOAD-BALANCED MIDDLEBOX ASSIGNMENT IN POLICY-
DRIVEN DATA CENTER

AUTHOR: MANAR ALQARNI

APPROVED:

Bin Tang, PhD
Thesis Committee Chair

Jianchao (Jack) Han, PhD
Committee Member

Mohsen Beheshti, PhD
Committee Member

I dedicate my thesis to my mother, father, my sister Manahil, and to my brothers especially my brother Mazen who supported me in every single step.

ACKNOWLEDGEMENTS

I would like to thank Dr. Bin Tang for his support. This thesis would not be possible without his guidance. I would also like to thank Dr. Jack Han and Dr. Mohsen Beheshti for their valuable advice.

TABLE OF CONTENTS

APPROVAL PAGE.....	ii
DEDICATION.....	iii
ACKNOWLEDGEMENTS.....	iv
TABLE OF CONTENTS.....	v
LIST OF FIGURES.....	vii
LIST OF TABLES.....	viii
ABSTRACT.....	ix

CHAPTER	PAGE
1. INTRODUCTION.....	1
2. RELATED WORK.....	5
3. LOAD BALANCED MIDDLEBOX ASSIGNMENT PROBLEM(LB-MAP).....	9
Network Model.....	9
Middlebox Model.....	9
Energy Model.....	11
Problem Formulation of LB-MAP.....	12
4. MINIMUM COST FLOW (MCF) OPTIMAL ALGORITHM	15
Minimum Cost Flow Problem (MCF).....	15
Transforming a Data Center Network to a Flow Network.....	17
5. THREE HEURISTIC ALGORITHMS	21
VM-Based Algorithm.....	21
MB-Based Algorithm	22
VM-MB-Based Algorithm.....	24
6. PERFORMANCE EVALUATION.....	25
Simulation Setting.....	25
Effect of Number of VM Pairs l	26
Effect of Number of MB Instances m	27
Comparison in Large Data Centers.....	28

CHAPTER	PAGE
Comparison Under Skewed Energy Model.....	29
7. CONCLUSION AND FUTURE WORK	31

REFERENCES.....33

LIST OF FIGURES

	PAGE
1. A k-ary fat tree with $k = 4$ and 16 physical machines (PMs).....	7
2. Load-Balanced VM communication with $\kappa = 1$	13
3. LB-MAP is equivalent to minimum cost flow problem.....	17
4. Varying number of VM pairs under uniform energy model.....	27
5. Varying number of MB instances under uniform energy model.	28
6. Varying number of VM pairs under uniform energy model.....	29
7. Varying number of MB instances under uniform energy model.....	29
8. Varying number of VM pairs under skewed energy model.	29
9. Varying number of MB instances under skewed energy model.....	29
10. Varying number of VM pairs under skewed energy model.	30
11. Varying number of MB instances under skewed energy model.	30

LIST OF TABLES

PAGE

1. NOTATION SUMMARY.....	9
--------------------------	---

ABSTRACT

In this thesis, the Load-Balanced Middlebox Assignment Problem (LB-MAP) is studied, which minimizes the communication energy cost of VM pairs while satisfying their policy requirement as well as the capacity constraint of the switches that the middleboxes placed upon. The thesis shows that LB-MAP is equivalent to the classic minimum cost flow problem (MCF), which can be solved optimally and efficiently. A suite of efficient heuristic algorithms is also designed based on different criteria viz. VM-Based, MB-Based, and VM+MB-Based. Via extensive simulations, it has been showed that all the heuristic algorithms perform close to the optimal minimum cost flow algorithm, while VM+MB-Based performs best among all the heuristic algorithms. This thesis is the first work that addresses the energy cost minimization for VM communications as well as load-balancing for middleboxes in policy-driven data centers.

CHAPTER 1

INTRODUCTION

Middleboxes, also known as “network appliances” or “network functions (NFs)”, are intermediary computer networking devices that transform, inspect, filter, or otherwise manipulate network traffic for purposes other than packet forwarding. Middleboxes are widely deployed in enterprise networks such as data centers and play an increasingly important role in improving the networks’ security (e.g., firewalls and intrusion detection systems (IDSs)) and performance (e.g., load balancers) as well as reducing the bandwidth cost (e.g., WAN optimizers).

A recent study (Sherry et al., 2012) shows that the number of middleboxes is on par with the number of routers in enterprise networks. Traditional middleboxes are diverse, stateful systems supporting narrow specialized network functions. As such, they are mainly proprietary and purpose-built hardware, which is typically closed and expensive. Their deployment and operation represent a significant part of the network capital and operational expenditure as well as costing large amount of space and power consumption. This greatly impedes the equipment upgrades and service addition in any networks, inducing the so-called network ossification problem (Li & Chen, 2015).

Network Function Virtualization (NFV) has been recently proposed (Guerzoni et al., 2012; Mijumbi et al., 2015) to alleviate this problem. NFV is a network virtualization technology that virtualizes middleboxes (or network functions) into building blocks that create communication services. It allows network operators and service providers to implement middleboxes in software instead of running on purpose-built hardware. This dramatically improves the management efficiency of middleboxes as they can be moved to or instantiate in various locations in the network without the need for new equipment installation and network operator involvement. Software Defined Networking (SDN)

(Feamster et al., 2015; Lee et al., 2014) being a complementary technology to NFV, further alleviates the network ossification problem by moving management functions out of the hardware and placing them in a centralized controller. Such centralized management, coupled with standardized protocol between the controller and network functions, enables dynamic and flexible configuration and placement of middleboxes inside networks.

The major enabler and beneficiary of NFV and SDN are cloud data centers, for two reasons. First, virtualization has become the core building block of modern data centers. The hardware resources such as CPU cycles, memory, and bandwidth are divided into smaller isolated computing units, known as virtual machines (VMs), which can be rented to tenants in a pay-as-you-go manner. Thus data centers provide an ideal enabling platform to implement and experiment the concept of NFV and SDN. Second, cloud data centers deploy a variety of middleboxes to protect, manage and improve the performance of the applications and services (Joseph et al., 2008). As a result, the introduction of NFV and SDN has greatly improved the middlebox management in cloud data centers (Hasan et al., 2012; Chiang et al., 2013; Ghadiyal et al., 2012).

In order to provide security and performance guarantee in data centers, network policies are established to demand traffic to traverse a sequence of specified middleboxes. For example, each VM communicating pair in data center could go through an intrusion detection system and a load balancer in that order so that malicious traffic can be filtered then trusted traffic be diverted to avoid network congestion. In fact, due to the diverse user applications demands, network policies have become an inseparable part of the Service Level Agreement (SLA) of data centers, which sets the expectations and commitments between the cloud data center consumers and cloud service providers. Those policies become an important measurement of the efficiency and efficacy of any data centers; we refer to such data centers as policy-driven data centers.

However, there exists a dilemma in the middlebox management of policy-driven data centers. On one side, the middleboxes have limited packet hardware resources such as CPU, memory, or accelerators. Besides, the SDN switches upon which middleboxes are installed use Ternary Content Addressable Memory (TCAM), which is expensive in both cost as well as power consumption, making the amount of forwarding rules available at each switch very limited. On the other side, middleboxes involve complex and extensive processing to capture application-semantics by using deep packet inspection. As such, overload is one of the dominant reasons for middlebox failures (Hasan et al., 2012), causing packet loss, traffic delay as well as wasting energy consumption. How to well balance the workload of middleboxes while minimize the VM communication energy cost is an important problem in policy-driven data centers.

Fortunately, due to the software implementation of middleboxes brought by NFV and SDN, it is possible to replicate and place multiple copies of the same middlebox inside the policy-driven data centers for the purpose of load-balancing and fault tolerance (Carpio et al., 2017). Each copy is referred to as an instance of the middlebox. In this thesis, it considers that there are multiple instances of one type of middlebox in the policy driven data center. To satisfy the policy requirement of security or performance, each VM communication pair must traverse one of the middlebox instances while different VM pairs can traverse different instances.

In particular, in this thesis, we study that given a set of VM pairs and a set of middlebox software instances in the data center, how to assign a middlebox instance for each VM pair to traverse, in order to minimize the total energy cost of the VM pairs while satisfying the resource constraint of each middlebox instances. The problem is refer as load-balanced middlebox assignment problem (LB-MAP). LB-MAP is formulated formally and prove that LB-MAP in data center network is equivalent to the well-known minimum cost flow problem (MCF) in a transformed flow network. MCF can be solved

optimally and efficiently (Ahuja et al., 1993). Therefore the LBMAP can also be solved optimally and efficiently. A suite of efficient heuristic algorithms is also designed viz. VM-Based, MB-Based, and VM+MB-Based. Via extensive simulations, the results show that all the heuristic algorithms perform close to the optimal minimum cost flow algorithm, while VM+MB-Based performs best among all the heuristic algorithms. To the best of knowledge, the main work is the first that collectively addresses the energy cost minimization for VM communications as well as load-balancing for middleboxes in policy-driven data centers.

CHAPTER 2

RELATED WORK

Joseph et al. (2008) was one of the first architectural work discussing middlebox management in data centers. In particular, it proposed and designed a policy-aware new layer-2 switching layer in data centers that explicitly forwards different types of traffic through different sequences of middleboxes. They showed that their approach traverses middleboxes correctly as well as efficiently. Qazi et al. (2013) further addressed key system design and algorithmic challenges for aforesaid policy enforcement layer. They proposed efficient data plane support for policy composition, unified switch and middlebox resource management, and automatically dealing with dynamic packet modifications. Sekar et al. (2012) took a different perspective and presented a new architecture for middlebox deployments. They proposed to consolidate individual middleboxes as well as their management to multiplex hardware resources and reuse processing modules across different applications.

On the system side, Sherry et al. (2012) showed that outsourcing middlebox processing to the cloud relieves enterprises of major problems caused by today's enterprise middlebox infrastructure. Recently Gember et al. (2012) realized a software defined middlebox networking framework by representing, manipulating, and knowledgeably controlling middlebox state. Zhang et al. (2013) presented a framework for SDN-enabled services that dynamically route traffic through any sequence of middleboxes. They also proposed an algorithm to select the best locations for placing to optimize the performance.

Middlebox management has strong theoretic roots as well. Liu et al. (2015) studied middlebox placement problem, in which given network information and policy specifications, it attempts to determine the optimal locations to place the middleboxes so

that either end-to-end delay or the bandwidth consumption is optimized. They showed this problem is NP-hard and proposed two heuristic algorithms. Li et al. (2013) studied the policy-aware cloud application embedding problem and designed simple and online primal-dual algorithms. Cui et al. (2016) proposed a suite of synergistic schemes to jointly consolidate network policies and virtual machines. They studied dynamic virtual machine consolidation and dynamic network policy (re)allocation to meet both efficient DC resource management and middleboxes traversal requirements.

Among above research, Qazi et al. (2013) was the only work that specifically addressed load-balancing issue of middleboxes in data centers. It formulated an online integer linear program (ILP) to minimize the maximum middlebox load across the network. However, as is stated in the thesis, the load balancing ILP might take a long time for a large network. In contrast, our LB-MAP has different goals and use different solution techniques. We aim to minimize the total energy cost of all the VM pairs in data centers while satisfying the capacities of middleboxes. This thesis proposes a time-efficient minimum cost flow optimal solution and a suite of time-efficient heuristic algorithms. Recently, Tu et al. (2015) introduced a programmable middlebox that distributes data center traffic more evenly in order to enhance bandwidth utilization and reduce traffic delay. The SDN middlebox controller collects traffic distribution and server loads information and performs load balancing accordingly. However, it did not consider the capacity constraint of each middlebox and did not aim to minimize the VM communication energy, therefore is different from LB-MAP.

Data Center Topology




Figure 1. A k -ary fat tree with $k = 4$ and 16 physical machines (PMs). There are two communicating VM pairs: (v_1, v'_1) and (v_2, v'_2) , and two middlebox instances MB1 and MB2. The capacity of each MB $\kappa = 2$. The minimum communication takes place as follows: (v_1, v'_1) traverses MB1 (colored blue, with cost of 3) while (v_2, v'_2) , traverses MB1 too (colored red, with cost of 5), resulting in minimum total cost of 8 under uniform energy model.

In this thesis, the focus is on fat-tree networks (Al-Fares et al., 2008) as they are widely adopted in data centers to interconnect commodity Ethernet switches. Fat tree is a variation of three-stage Clos networks (Clos, 1953), which is rearrangeably nonblocking with 1:1 oversubscription ratio (Al-Fares et al., 2008). Rearrangeably non-blocking means all the bandwidth available to the end hosts can always be saturated for any communication patterns. Over-subscription is the ratio of the worst-case achievable aggregate bandwidth among the end hosts to the total bisection bandwidth of a particular communication topology. An oversubscription of 1:1 indicates that all hosts may potentially communicate with arbitrary other hosts at the full bandwidth of their network interface. 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 from bottom to top. Core switches handles huge amount of traffic across the

entire data center, therefore consuming lots of energy power. In contrast, aggregate switches and edge switches transmit less amount of traffic therefore consume less power. The lower two layers are separated into k pods. Pods are modular units of compute, storage, and networking resources that are designed together as a unit in data center, each containing $k/2$ aggregation switches and $k/2$ edge switches, while forming a complete bipartite graph in between. In particular, each edge switch is directly connected to $k/2$ physical machines (PMs); and each of its remaining $k/2$ ports is connected to each of the $k/2$ aggregation switches from the same pod. There are $\frac{k^2}{4}$ k -port core switches, each of which is connected to each of k pods. In general, a fat-tree built with k -port switches supports $\frac{k^3}{4}$ PMs. Figure 1 shows a data center of 16 PMs.

CHAPTER 3

LOAD BALANCED MIDDLEBOX ASSIGNMENT PROBLEM (LB-MAP)

Network Model A data center is modeled as an undirected general graph $G(V, E)$. $V = V_p \cup V_s$ includes the set of PMs V_p and the set of (edge, aggregate, and core) switches V_s . E is the set of edges, each connecting either one switch to another switch or a switch to a PM, as shown in Figure 1. In the data center network, there are l communicating VM pairs $P = \{(v_1, v'_1), (v_2, v'_2), \dots, (v_\ell, v'_\ell)\}$. v_i and v'_i ($1 \leq i \leq \ell$) are referred to as the source VM and the destination VM, respectively. Each VM v is located in one of the PMs, denoted as $S(\mathcal{V})$. In particular, the PMs $S(v_i)$ and $S(v'_i)$ are referred to as the source PM and destination PM of (v_i, v'_i) , respectively. A PM can store multiple VMs, and can be source and destination PM simultaneously. In Figure 1, there are two communicating VM pairs: (v_1, v'_1) and (v_2, v'_2) . Table I shows all the notations used in this thesis.

TABLE I
NOTATION SUMMARY

Notation	Explanation
V_p	The set of physical machines (PMs) in the data center
V_s	The set of switches in the data center
P	The set of l VM communication pairs of (v_i, v'_i) , $1 \leq i \leq \ell$
M	The set of m middlebox instances, $mb_j, 1 \leq j \leq m$
$S(\mathcal{V})$	The PM where VM \mathcal{V} is stored
κ	The capacity of each middlebox instance
r_e, r_a, r_c	The energy consumption on edge, access, and core switch
$c(i, j)$	The energy cost between PM (or switch) i and j
$c_{i,j}$	The energy cost when VM pair (v_i, v'_i) traverses mb_j
C_p	The total energy cost for a MB assignment function p

Middlebox Model According to Gill et al. (2011), among all the network devices in data center, load balancers have the highest failure probability. This is due to high number of software faults (such as software bugs and configuration errors) and hardware faults

related to application-specific integrated circuit (ASIC) and memory. It is assumed that there is one type of middlebox in the data center network (e.x., load balancer). However, there are multiple copies of the same middlebox type inside the cloud data centers (Jukan F et al., 2017). The more general case wherein there are multiple middlebox types and each has multiple instances, and each VM pair needs to traverse a sequence of middleboxes, will be discussed in Future Work Section.

In particular, we assume there are m software instances of this middlebox $M = \{mb_1, mb_2, \dots, mb_m\}$, which have already been placed inside the data center network. Each instance is placed in one of the switches. Let's assume that mb_j ($1 \leq j \leq m$) is located at switch $sw(j) \in V_s$. The policy specifies that for each communicating VM pair (v_i, v'_i) , it must traverse one of the instances. However, due to the capacity limit of the middlebox, at most κ VM pairs can be served by any middlebox instance. In Figure 1, there are two load balancer instances: MB_1 and MB_2 , the capacity of each MB is $\kappa = 2$.

Bump-Off-The-Wire Design Traditional middlebox appliances are deployed into the data center network using an inline “bump-in-the-wire” design. In this configuration, dedicated middlebox hardware is plugged in the physical network data path, processing all the traffic passing through it. There are two drawbacks of this design. First, such on-path deployment of middleboxes forces all traffic on a network path to traverse the same sequence of middleboxes. This is unnecessary since different applications may have different requirements, different application traffic may therefore need to traverse different middleboxes. Second, when there are multiple instances of the same type of middlebox, it is a waste of processing time and hardware resources when the traffic must pass through more than one of them. Instead, we adopt the “bump-off-the-wire” design proposed by Joseph et al. (2008), which is further improved by Zhang et al. (2013). In their designs,

middleboxes are taken off the physical network path and implemented as software modules or VMs that are installed on a PM plugged into each switch. Therefore, network traffic is explicitly forwarded to the middleboxes. Due to its low latency links and minimal performance overhead, the data center network is very suitable for such explicit traffic redirection.

Energy Model The power consumption of any VM pair communication is measured by focusing on the switches the communication goes through. In particular, r_e , r_a , and r_c are used to denote the power consumption on an edge, aggregate, and core switches respectively, when it transmits a VM communication. In this thesis, two energy consumption models are considered that are currently adopted in cloud data center research.

Uniform Energy Model In this model, the energy consumption of VM communication is measured as the minimum number of switches it traverses (Meng et al., 2010). That is, it costs the same amount of energy by going through either a core, aggregation, or edge switch: $r_e = r_a = r_c$. For example, if $r_e = r_a = r_c = 1$, the power consumption between v_1 and v'_1 and between v_2 and v'_2 in Figure 1 are 3 and 5, respectively.

Skewed Energy Model This model is based on the fact that the core switches handle more traffic therefore usually consume more energy power than aggregate switches, which consume more energy power than edge switches (Boru et al., 2015). Accordingly, we set $r_e < r_a < r_c$. For example, if $r_e = 1$, $r_a = 5$, and $r_c = 10$, the power consumption between v_1 and v'_1 and between v_2 and v'_2 in Figure 1 are 7 and 22, respectively.

EXAMPLE 1 In Figure 1, the capacity of each MB $\kappa = 2$. Each VM pair needs to traverse one of the instances. To minimize the communication cost, the (v_1, v'_1) traverses

MB_1 (colored blue, with cost of 3) while (v_2, v'_2) traverses MB_2 too (colored red, with cost of 5), resulting in minimum total communication cost of 8 under uniform energy model.

Problem Formulation of LB-MAP Let $c(i, j)$ denote the minimum energy consumption between PM (or switch) i and j . Let $c_{i,j}$ be the minimum power consumption for VM pair (v_i, v'_i) when it is assigned to middlebox instance mb_j . Then,

$$c_{i,j} = c(S(v_i), sw(j)) + c(sw(j), S(v'_i)).$$

Now we define the load balanced middlebox assignment function as $p : P \rightarrow M$, signifying that VM pair $(v_i, v'_i) \in P$ is assigned to middlebox instance $p(i) \in M$. Given any middlebox assignment function p , the power consumption for VM pair (v_i, v'_i) is then

$$c_{i,p(i)} = c(S(v_i), sw(p(i))) + c(sw(p(i)), S(v'_i)).$$

Denote the total energy consumption of all the l VM pairs with middlebox assignment p as C^p . Then

$$C^p = \sum_{i=1}^l c_{i,p(i)} = \sum_{i=1}^l \left(c \left(S(v_i), sw(p(i)) \right) + c \left(sw(p(i)), S(v'_i) \right) \right)$$

Let p_{min} be an assignment that yields the minimum total energy consumption among all the middlebox assignments P , i.e., $C^{p_{min}} \leq C^p, \forall p \in P$. The objective of LB-MAP is to find such a p_{min} under the constraint that at most κ VM pairs can be served by any middlebox instance:

$$|\{1 \leq i \leq l | p(i) = j\}| \leq \kappa, \forall j, 1 \leq j \leq m.$$




Figure 2. Load-Balanced VM communication with $\kappa = 1$. The minimum energy communication is then: (v_1, v_1') traverses MB1 (shown in blue color, with cost of 3) while (v_2, v_2') traversing MB2 (shown in red color, with cost of 9), resulting in total communication cost of 12 under uniform energy model.

EXAMPLE 2 If the capacity of each MB is $\kappa = 1$, the two VM pairs in Figure 1 can not traverse MB1 simultaneously. Figure 2 shows a load-balanced VM pair communication. (v_1, v_1') traverses MB1 (shown in blue color, with cost of 3) while (v_2, v_2') traversing MB2 (shown in red color, with cost of 9), resulting in total communication cost of 12 under uniform energy model.

CHAPTER 4
MINIMUM COST FLOW (MCF) OPTIMAL ALGORITHM

In this chapter, it shows that the LB-MAP is equivalent to minimum cost flow problem, which can be solved efficiently and optimally (Ahuja et al., 1993). We first present the minimum cost flow problem with its well-known algorithms, and then transform the data center network to a flow network to show the equivalency.

Minimum Cost Flow Problem (MCF) MCF (Ahuja et al., 1993) is a classic optimization problem that finds the cheapest possible way of sending a certain amount of flow through a flow network, considering that each edge in the flow network has an associated capacity and cost. It is formally defined as below

Let $G = (V, E)$ be a directed graph. The capacity of edge $(u, v) \in E$, denoted by $c(u, v)$, represents the maximum amount of flow that can pass through this edge. The cost of edge $(u, v) \in E$, denoted by $d(u, v)$, represents the cost when one amount of flow that can pass through an edge. Besides, there is a source node $s \in V$ with b amount of supply, and a sink node $t \in V$ with b amount of demand. A flow on an edge $(u, v) \in E$, denoted by $f(u, v)$, is a mapping $f : E \rightarrow \mathbb{R}^+$, subject to the following two constraints:

1) Capacity constraint: $f(u, v) \leq c(u, v), \forall (u, v) \in E$. That is, the flow of an edge cannot exceed its capacity.

2) Flow conservation constraint: $\sum_{u \in V} f(u, v) = \sum_{u \in V} f(v, u)$, for each $v \in V \setminus \{s, t\}$.

That is, the sum of the flows entering a node must equal the sum of the flows exiting a node, except for the source and the sink nodes. The net flow out of source node s is

$\sum_{u \in V} (f(s, u) - f(u, s)) = b$; the net flow into sink node t is

$\sum_{u \in V} (f(t, u) - f(u, t)) = b$.

The goal of MCF is to find a flow function f such that the total cost of the flow in the network is minimized. That is,

$$\min \sum_{(u,v) \in E} (d(u,v) \cdot f(u,v)) .$$

MCF Algorithms

MCF can be solved efficiently by many combinatorial algorithms (Ahuja et al., 1993). There are pseudo-polynomial algorithms including cycle-canceling (Morton, 1967), successive shortest path (Edmonds & Karp, 1972), and out-of-kilter algorithm (Fukerson, 1961), polynomial algorithms including cost- and capacity scaling (Edmonds & Karp, 1972), (Orlin, 1988), (Goldberg and Tarjan, 1990), and network simplex algorithm (Orlin, 1997). Here, we adopt the scaling push-relabel algorithm proposed by Goldberg (Goldberg, 1997), which works well over a wide range of problem classes. For any flow network, the algorithm has the time complexity of $O(a^2 \cdot b \cdot \log(a \cdot c))$, where a , b , and c are the number of nodes, number of edges, and maximum edge capacity in the flow network, respectively.

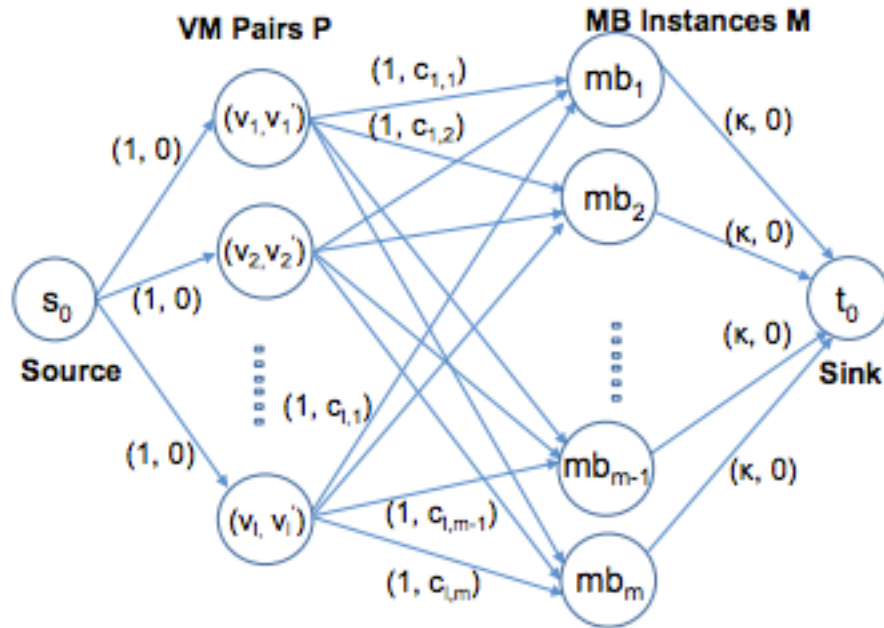


Figure 3. LB-MAP is equivalent to minimum cost flow problem. In each parenthesis on the edge, the first value is the capacity of the edge and the second the cost of the edge.

Transforming a Data Center Network to a Flow Network Next, the data center network $G(V, E)$ is transformed (Figure 1) into a new flow network $G'(V', E')$ (Figure 3). LB-MAP in $G(V, E)$ is equivalent to MCF in $G'(V', E')$. The transformation consists of the following five steps:

Step 1. $V' = \{s_0\} \cup \{t_0\} \cup P \cup M$, where s' is the new source node and t' is the new sink node in the flow network.

Step 2. $E' = \{(t_0, (v_i, v_i')) : (v_i, v_i') \in P\} \cup \{(v_i, v_i'), mb_j) : (v_i, v_i') \in P, mb_j \in M\} \cup \{(mb_j, t_0) : mb_j \in M\}$. Note that it is a complete bipartite graph between P and M .

Step 3. For each edge $(s_0, (v_i, v'_i))$, set its capacity as 1 and its cost 0. For each edge (mb_j, t_0) , set its capacity as κ and its cost 0.

Step 4. For each edge $((v_i, v'_i), mb_j), (v_i, v'_i) \in P, mb_j \in M$, set its capacity as 1 and its cost as $c_{i,j}$, the minimum energy cost of VM pair (v_i, v'_i) when it is assigned to middlebox instance mb_j .

Step 5. Set the supply at s_0 and the demand at t_0 as 1. The technique used above is similar to those in (Tang et al.), (Khani et al., 2015).

Next, the generated flow network (Figure 3) is passed to the MCF algorithm discussed above, which outputs the load-balanced middlebox assignment. That is, for each VM pair, the MCF outputs its assigned middlebox instance without violating the capacity constraint of each middlebox.

Time Complexity of LB-MAP Algorithm

The algorithm consists of two parts: the graph transformation and the MCF algorithm. For the transformation, calculating $c(i, j)$ and $c_{i,j}$ each takes $O(|V|^3)$. Constructing $G'(V', E')$ takes $|V|^3 + l + l \cdot m + m$, which is $O(l \cdot m + |V|^3)$. Therefore it takes $O(l \cdot m + |V|^3)$ for transformation. For the minimum cost flow, the scaling push-relabel algorithm (Goldberg, 1997) that adopted has the time complexity of $O(a^2 \cdot b \cdot \log(a \cdot c))$, where a, b , and c are the number of nodes, number of edges, and maximum edge capacity in the flow network. From the transformation, $|V'| = l + m + 2$ and $|E'| = l + m + l \cdot m$. Therefore the time complexity of the minimum cost flow is $O((l + m)^2 \cdot l \cdot m \cdot \log((l + m) \cdot \kappa))$, where κ is the largest number of VM

communications each middlebox instance can accommodate. The time complexity of the LB-MAP algorithm is thus $O(|V|^3 + (l + m)^2 \cdot l \cdot m \cdot \log((l + m) \cdot \kappa))$

Theorem 1: LB-MAP is equivalent to minimum cost flow problem.

Proof: By applying minimum cost flow algorithm upon above flow network, it achieves that a) each of the l VM communication pairs is assigned to exactly one middlebox instance while b) satisfying the capacity constraints of middleboxes, and c) achieving the minimum energy consumption for all the l VM pairs. First, with above transformation, sending l amount of flow from s_0 to t_0 ensures that each of the l VM pairs be assigned to one middlebox instance. In particular, since the amount of supply at s_0 is l (Step 5), and the capacity of each edge $(s_0, (v_i, v'_i))$ ($1 \leq i \leq l$) is one (Step 3), a valid flow of l amount from s_0 to t_0 must consists of one amount on edge $(s_0, (v_1, v'_1))$, one amount on $(s_0, (v_2, v'_2))$, ..., and one amount on $(s_0, (v_l, v'_l))$. Now, since the capacity on each edge $((v_i, v'_i), mb_j)$ is one (Step 4), according to flow conservation, then one amount of flow must come out of any edge (v_i, v'_i) and go into exactly one of the middlebox instance mb_j . This results in that each VM pair is assigned to exactly one middlebox instance. Note that the capacity of each edge $((v_i, v'_i), mb_j)$ could be set as any positive integers as it does not change above analysis. Next, the VM pair assignment does not violate the capacity constraint of middleboxes. Since the edge capacity of edge (mb_j, t_0) is κ (Step 3), it ensures that no more than κ amount of flow coming out of each node $mb_j \in M$. This guarantees that each middlebox instance mb_j will not accommodate more than κ VM pairs, satisfying the capacity constraint of middleboxes. Finally, as for the cost, note that the edge cost of $((v_i, v'_i), mb_j)$ is $c_{i,j}$, the minimum energy consumption of VM pair

(v_i, v'_i) when it is assigned to middlebox mb_j , while all other edges have cost zero. This indicates that only the VM communication cost is considered in minimum cost flow. Minimum cost flow algorithm gives the minimum cost of sending l amount of flow from s_0 to t_0 , showing that the corresponding VM communication cost obtained is indeed minimum.

CHAPTER 5

THREE HEURISTIC ALGORITHMS

This chapter proposes three polynomial-time greedy algorithms for comparison purpose. Each greedy algorithm takes place in rounds, and is based on a different criteria to choose middlebox for each VM pair. VM-Based Algorithm. The VM-Based algorithm works as follows. For each VM pair, it is assigned to an MB instance such that it gives the minimum energy consumption for this VM pair among all the MB instances, while satisfying this MB instance's capacity. Finding all the minimum energy consumption paths between all pair of PMs takes $O(|V|^3)$. Assigning each VM pair to a MB instance takes $O(l \cdot m)$. Therefore the time complexity of VM-Based Algorithm is $O(|V|^3 + l \cdot m)$.

VM-Based Algorithm

Input: A data center $G(V, E)$ with l VM pairs and m MBs

Output: Total energy cost C for all the l VM pairs.

Notations

i : the index for VM pairs

j : the index for middlebox instances

$load(j) = 0$: the current load of mb_j

c_{min}^i : the minimum energy cost for VM pair (v_i, v'_i)

j^* : middlebox mb_{j^*} is assigned to (v_i, v'_i)

1. $C = 0$;
2. **for** ($i = 1$ to l)

3. $c_{min}^i := infinite$;
4. **for** ($j = 1$ to m)
5. **if** ($c(i, j) \leq c_{min}^i$ and $load(j) < \kappa$)
6. $c_{min}^i = c(i, j)$;
7. $j^* = j$;
8. **end if**;
9. **end for**;
10. $load(j^*)++$;
11. $C = C + c_{min}^i$;
12. **end for**;
13. **RETURN** C .

MB-Based Algorithm

In this algorithm, for each MB instance, it is assigned κ VM pairs among all the VM pairs that give the minimum energy consumption when going through that MB instance. The running time of this algorithm is $|V|^3 + m \cdot (l + l \cdot lgl + \kappa)$, which is $O(|V|^3 + m \cdot l \cdot lgl)$.

Algorithm 2: MB-Based Algorithm.

Input: A data center $G(V, E)$ with l VM pairs and m MBs

Output: Total energy cost C for all the l VM pairs.

Notations:

i : the index for VM pairs

j : the index for middlebox instances

X_j : the set of VM pairs assigned to mb_j

$assigned[i]$: true if (v_i, v'_i) is assigned, false if not

1. $C = 0$;
2. **for** ($i = 1$ to l)
 3. $assigned[i] = \text{false}$;
 4. **end for**;
 5. **for** ($j = 1$ to m)
 6. $X_j = \varnothing$;
 7. **for** ($i = 1$ to l)
 8. **if** ($assigned[i] == \text{false}$)
 9. $X_j = \{(i, c(i, j))\} \cup X_j$;
 10. **end if**;
 11. **end for**;
 12. Sort X_j in the non-descending order of $c(i, j)$;
 13. $X_j = \{(x_1, (x_1, c(x_1, j))), (x_2, (x_2, c(x_2, j))), \dots\}$,
 14. where $c(x_1, j) \leq c(x_2, j) \leq c(x_3, j) \dots$;
 15. **for** ($k = 1$ to κ)
 16. $C = C + c(x_k, j)$;
 17. **end for**;
 18. **end for**;
 19. **RETURN** C .

VM-MB-Based Algorithm

In this algorithm, in each round, it checks which VM pair is assigned to which MB instance, such that when that VM pair traverses that MB instance, it yields the minimum energy consumption among all the unassigned VM pairs and all the MB instances in that round. The time complexity of this algorithm is $O(|V|^3 + m \cdot l)$.

Algorithm 3: VM-MB-Based Algorithm.

Input: A data center $G(V, E)$ with l VM pairs and m MBs

Output: Total energy cost C for all the l VM pairs. Notations:

i : the index for VM pairs

j : the index for middlebox instances

$load(j) = 0$: the current load of mb_j

j^* : middlebox mb_j is assigned in each round,

c_{min} : the minimum energy obtained in each round

1. $C = 0$;
2. **for** ($i = 1$ to l)
3. $assigned[i] = false$;
4. **end for**;
5. **while** (there are still unassigned VM pairs)
6. $c_{min} = infinite$;
7. **for** ($i = 1$ to l)
8. **if** ($assigned[i] == false$)
9. **for** ($j = 1$ to m)
10. **if** ($load(j) < \kappa$ and $c(i, j) \leq c_{min}$)


```
11.           $c_{min} = c(i, j);$ 
12.           $j^* = j;$ 
13.          end if;
14.          end for;
15.          end if;
16.          end for;
17.          load( $j^*$ )++;
18.           $C = C + c_{min};$ 
19. end while;
20. RETURN C.
```

CHAPTER 6

PERFORMANCE EVALUATION

Simulation Setting In this chapter, the performances of the four LB-MAP algorithms are compared: Minimum Cost Flow Algorithm (referred to as **MCF**), VM-Based Algorithm (referred to as **VM**), MB-Based (referred to as **MB**), and VM+MB-Based Algorithm (referred to as **VM+MB**). Two different sizes of data centers are created: $k = 8$, a small data center with 128 PMs and $k = 16$, a relatively large size data center with 1024 PMs. The source and destination VMs of each VM pair are randomly placed on the PMs and the MB instances are randomly placed on the switches. In all the simulation plots, each data point is an average of 10 runs, and the error bars indicate 95% of confidence interval. For fair comparison, for each simulation run, the four algorithms are run on the same network set up with the same initial placement of VM communication pairs and MB instances. In all the simulations, unless otherwise mentioned, the number of VM pairs is varying in the data center l from 100, 200, ..., to 500, and number of MB instances m from 1 to 2 to 3. In each case, the middlebox capacity $\kappa = \lfloor \frac{l}{m} \rfloor$. Throughout the simulation, $r_e = r_a = r_c = 1$ is set for the uniform energy model and $r_e = 1, r_a = 5$, and $r_c = 10$ for the skewed energy model. Recall that r_e, r_a , and r_c denote the power consumption on the edge, aggregate, and core switches respectively.

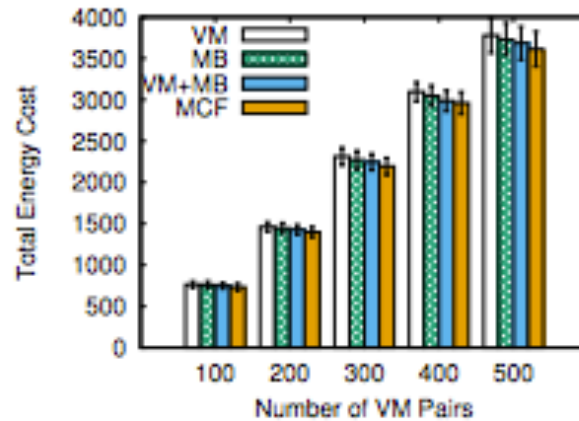


Figure 4. Varying number of VM pairs under uniform energy model. Here, number of MBs $m = 3$, number of PMs in data center is 128.

Effect of Number of VM Pairs l . First, the effect of the number of VM pairs is investigated on the total energy cost, shown in Figure 4. The number of MB instances in the data center is set as 3 and vary the number of VM pairs from 100, 200, ..., to 500. There are several observations. First, for each algorithm, with the increase of number of VM pairs, the total energy cost increases as well. This is obvious since more VM pairs in general incurs more energy consumption since all the VM pairs are randomly placed. Second, the performance order of the four algorithms are (from worst to best) is VMBased, MB-Based, VM+MB-Based, and MCF. This indirectly supports the theoretical proof that the minimum cost flow algorithm is optimal, performing better than the other three heuristic algorithms. With the increase of number of VM pairs, the performance difference between MCF and the other three seems to be more evident.

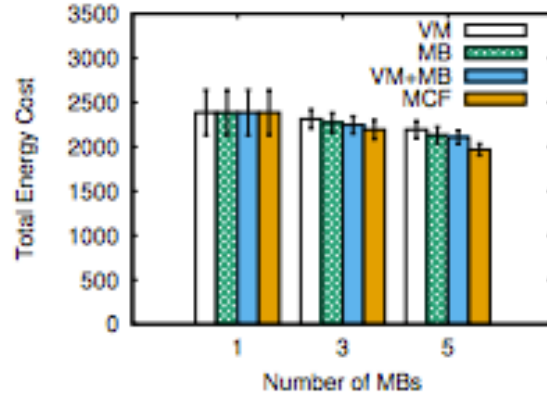


Figure 5. Varying number of MB instances under uniform energy model. Here, number of VM pairs $l = 300$, number of PMs in data center is 128.

Effect of Number of MB Instances m The effect of the number of MB instances is investigated on the total energy cost, shown in Figure 5. The number of VM pairs in the data center is set as 300 and vary the number of MB instances as 1, 2, 3. There are several observations. First, for each algorithm, with the increase of number of MB instances, the total energy cost decreases. This is because more MB instances give each VM pairs more options to choose an energy-efficient path between source and destination VMs, incurring less energy consumption. Second, when there is only one middlebox, it is observed that all the four algorithms perform the same. In this case, all the VM pairs must go through the same middlebox for policy requirement independent of the algorithms, therefore costing the same amount of energy for all four algorithms. Finally, there is still observe that MCF performs the best, yielding the minimum amount of energy consumption. Among the three heuristic algorithms, VM+MB-based performs better than MBbased, which performs better than VM-based.




Figure 6. Varying number of VM pairs under uniform energy model. Here, number of MBs $m = 3$, number of PMs in data center is 1024.




Figure 8. Varying number of VM pairs under skewed energy model. Here, number of MBs $m = 3$, number of PMs in data center is 128.




Figure 7. Varying number of MB instances under uniform energy model. Here, number of VM pairs $l = 300$, number of PMs in data center is 1024.




Figure 9. Varying number of MB instances under skewed energy model. Here, number of VM pairs $l = 300$, number of PMs in data center is 128.

Comparison in Large Data Centers Figure 6 and Figure 7 show the performance comparison in a large data center of 1024 PMs, by varying number of VM pairs and number of middleboxes, respectively. There is the same general observations as in a smaller data center shown in Figure 4 and Figure 5, however, with one notable difference. Figure 7 shows that for all the three heuristic algorithms, when number of MB instances increases from one to three, their energy costs actually increase. This is counter-intuitive at first, since more middleboxes help with reducing the energy cost. However, this

increase is due to the random initial placement of the VM pairs and middleboxes, under which it is possible that a placement with two MBs does not help as much as a placement of one MB for different set of randomly placed VM pairs.




Figure 10. Varying number of VM pairs under skewed energy model. Here, number of MBs $m = 3$, number of PMs in data center is 1024.




Figure 11. Varying number of MB instances under skewed energy model. Here, number of VM pairs $l = 300$, number of PMs in data center is 1024.

Comparison Under Skewed Energy Model Finally, how the skewed energy model affects the energy costs of the four algorithms is investigated. Figure 8 and Figure 9 show the performance comparison in a small data center of 128 PMs, while Figure 10 and Figure 11 show the performance comparison in a large data center of 1024 PMs. MCF performs best, yielding the minimum amount of energy in all three algorithms. VM+MB-Based performs better than the other two heuristic algorithms. This is mainly because in each round, it selects an assigned VM pair and a MB that that when this pair is assigned to this MB, it yields the smallest energy cost among all the unassigned VM pairs and MB combination. Whereas for the other two algorithms (VM- Based and MB-Based), it only focuses on one of them without checking thoroughly all the possibilities.

CHAPTER 7

CONCLUSION AND FUTURE WORK

LB-MAP: Load-Balanced Middlebox Assignment Problem was studied in policy-driven data centers. The goal of LBMAP is to minimize the energy cost of all the communicating virtual machine pairs who must traverse a middlebox for policy requirement, while taking into account of the limited capacity of the middlebox. Much research has been done to study how to place middleboxes inside data centers for cost efficient VM communication. However, not much research has focused on load balance of middleboxes. We formulated LB-MAP formally and proved that LB-MAP is equivalent to the well-known minimum cost flow problem (MCF), which can be solved optimally and efficiently. A suite of efficient heuristic algorithms was also designed based on different criteria viz. VM-Based, MB-Based, and VM+MB-Based. Via extensive simulations, it showed that all the heuristic algorithms perform close to the optimal minimum cost flow algorithm, while VM+MB-Based performs best among all the heuristic algorithms. To the best of our knowledge, this is the first work that addresses the energy cost minimization for VM communications as well as load-balancing for middleboxes in policy-driven data centers. In our current setup, we assume that there is only one middlebox type such as load balancers. In the future, it will consider a more general problem wherein multiple types of middleboxes (e.x., load balancers, firewalls, and IDSs) exist, each having multiple instances. To satisfy policy requirement, each VM pair needs to traverse a sequence of middlebox instances of different types in a particular order. In this so called service chain scenario, tackling middlebox loadbalancing as well as minimizing VM pair communication cost become more challenging problems.

REFERENCES

- Ahuja, R. K., Magnanti, T. L., & Orlin, J. B. (1993). *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc.
- Akyildiz, Ian F, Lee, A., Wang, P., Luo, M., and Chou, W.(2014). A roadmap for traffic engineering in sdn-openflow networks. *Computer Networks (Elsevier)*, 71:1–30.
- Al-Fares, M., Loukissas, A. and Vahdat, A.(2008) A scalable, commodity data center network architecture. *SIGCOMM Comput. Commun. Rev.*, 38(4) 63–74.
- Boru, D., Kliazovich, D., Granelli, F., Bouvry, P. and Zomaya, A. Y. (2015). Energy-efficient data replication in cloud computing datacenters. *Cluster Computing*, 18(1): 385–402.
- Clos, C. (1953). A study of non-blocking switching networks. *Bell System Technical Journal*, 32(2).
- Cui, L., Cziva, R., Tso, F. P. and Pezaros, D. P.(2016). Synergistic policy and virtual machine consolidation in cloud data centers. *In Proc. of IEEE INFOCOM*.
- Cui, L.and Tso, F. P. (2015). Joint virtual machine and network policy consolidation in cloud data centers. *In Proc. of IEEE CloudNet*.
- Cui, L., Tso, F. P., Pezaros, D. P., and Jia, W. (2015). Plan: A policy-aware vm management scheme for cloud data centres. *In Proc. of IEEE/ACM UCC*.
- Cui, L., Tso, F. P., Pezaros, D. P, Jia, W., and Zhao, W. (2015). Policy-aware virtual machine management in data center networks. *In Proc. of IEEE ICDCS*.
- Edmonds, J., and Karp, R.M.(1972). Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of ACM*, 19:248–264.
- Guerzoni et al, R. (2012.). Network functions virtualisation: an introduction, benefits, enablers, challenges and call for action, introductory white paper. *In SDN and OpenFlow World Congress*.

- Jukan, A., Carpio, F., Dhahri, S. (2017). Vnf placement with replication for load balancing in nfv networks. *In IEEE ICC*.
- Feamster, N., Rexford, J., and Zegura, E. (2013). The road to sdn. *Queue*, 11(12):20–40.
- Fukerson, D.R. (1961). An out-of-kilter method for minimal-cost flow problems. *Journal of the Society for Industrial and Applied Mathematics*, 9(1):18– 27.
- Gember, A., Prabhu, P., Ghadiyali, Z., and Akella, A. (2012). Toward software defined middlebox networking. *In Proc. of HotNets-XI*.
- Gill, P., Jain, N., and Nagappan, N. (2011). Understanding network failures in data centers: Measurement, analysis, and implications. *SIGCOMM Comput. Commun. Rev.*, 41(4):350–361.
- Goldberg, A. V. (1997). An efficient implementation of a scaling minimum-cost flow algorithm. *J. Algorithms*, 22:1–29.
- Goldberg, A. V., and Tarjan, R. E. (1990). Finding minimum-cost circulations by successive approximation. *Math. Oper. Res.*, 15(3):430–466.
- Joseph, D. A., Tavakoli, A., and Stoica, I. (2008). A policy-aware switching layer for data centers. *In Proc. of the ACM SIGCOMM*.
- Khani, P., Tang, B., Han, J., and Beheshti, M. (2015). Dao-r: Integrating data aggregation and offloading in sensor networks via data replication. *In Proceedings of IEEE GLOBECOM*.
- Li, L. E., Liaghat, V., Zhao, H., Hajiaghayi, M., Li, D., Wilfong, G., Yang, Y. R., and Guo, C. (2013). Pace: Policy-aware application cloud embedding. *In Proc. of IEEE INFOCOM*.
- Li, Y., and Chen, M. (2015). Software-defined network function virtualization: A survey. *IEEE Access*, 3.

- Liu, J., Li, Y., Zhang, Y., Su, L., and Jin, D. (2015). Improve service chaining performance with optimized middlebox placement. *IEEE Transactions on Services Computing*.
- McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S., and Turner, J. (2008). Openflow: Enabling innovation in campus networks. *SIGCOMM Comput. Commun. Rev.*, 38(2):69–74.
- Meng, X., Pappas, V., and Zhang, L. (2010). Improving the scalability of data center networks with traffic-aware virtual machine placement. In Proc. of IEEE INFOCOM.
- Mijumbi, R., Serrat, J., Gorricho, J. L., Bouten, N., Turck, F. D., and Boutaba, R. (2015). Network function virtualization: State-of-the-art and research challenges. *IEEE Communications Surveys and Tutorials*, 18(1).
- Morton, K. (1967). A primal method for minimal cost flows with applications to the assignment and transportation problems. *Management Science*, 14:205–220.
- Orlin, J. B. (1988). A faster strongly polynomial minimum cost flow algorithm. In Proc. of the 20th ACM Symposium on the Theory of Computing, (pp. 377–387).
- Orlin, J. B. (1997). A polynomial time primal network simplex algorithm for minimum cost flows. *Mathematical Programming*, 79:109–129.
- Qazi, Z. A., Tu, C.-C., Chiang, L., Miao, R., Sekar, V., and Yu, M. (2013). Simplefying middlebox policy enforcement using sdn. In Proc. of the ACM SIGCOMM.
- Sekar, V., Egi, N., Ratnasamy, S., Reiter, M. K., and Shi, G. (2012). Design and implementation of a consolidated middlebox architecture. In Proc. of NSDI.
- Sherry, J., Hasan, S., Scott, C., Krishnamurthy, A., Ratnasamy, S., and Sekar, V. (2012). Making middleboxes someone else’s problem: Network processing as a cloud service. In Proc. eedings of the ACM SIGCOMM.

Tang, B., Jaggi, N., and Takahashi, M. Achieving data k-availability in intermittently connected sensor networks. *In Proceedings of the International Conference on Computer Communications and Networks (ICCCN'14)*.

Tu, R., Wang, X., Zhao, J., Yang, Y., Shi, L., and Wolf T. (2015). Design of a loadbalancing middlebox based on sdn for data centers. *In IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPs)*.

Zhang, Y., Beheshti, N., Beliveau, L., Lefebvre, G., Manghirmalani, R., Mishra, R., Patney, R., Shirazipour, M., Subrahmaniam, R., Truchan, C., and Tatipamula M. (2013.). Steering: A software-defined networking for inline service chaining. *In Proc. of the IEEE ICNP*.