

PAM & PAL

Policy-Aware Virtual Machine Migration and Placement in Dynamic Cloud Data Centers

Hugo Flores Vincent Tran Bin Tang

Department of Computer Science
California State University Dominguez Hills,
Carson, CA 90747, USA

This work was supported by NSF Grant CNS-1911191.

IEEE INFOCOM, July 2020

Outline

- 1 MOTIVATION
- 2 SYSTEM MODELS
- 3 PAL: POLICY-AWARE VM PLACEMENT IN PADCS
- 4 PAM:POLICY-AWARE VM MIGRATION IN PADCS
- 5 PERFORMANCE EVALUATION
- 6 CONCLUSION AND FUTURE WORK

Goals of Project

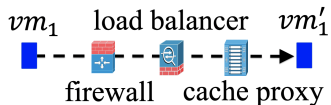
- Focused on policy-aware data centers (**PADCs**), wherein virtual machine(**VM**) traffic traverses in
 - 1 **Ordered**: Sequence of middle boxes.
 - 2 **Unordered**: Unsequence of middle boxes.
- Proposed two new VM placement and migration problems.
 - 1 **PAL**: Policy-aware virtual machine placement.
 - 2 **PAM**: Policy-aware virtual machine migration.
- Satisfied the constraint: resources capacity of PADCs.
- Addressed dynamic communicating traffic in cloud data center such as cloud chatbots¹

¹Slack, the collaboration software that moves work forward. <http://slack.com.>, and Amazon lex. <https://aws.amazon.com/lex/>.

Policy Aware Data Centers (PADCs)

--- → VM communication

■ VM 📦 PM 🌐 Switch



(a)

Figure 1: A data center policy with 1 virtual machine pair example

Policy Aware Data Centers (PADCs)

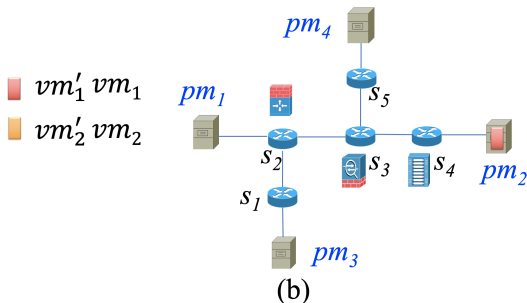


Figure 2: A PADC example with 2 virtual machine pairs

Policy Aware Data Centers (PADCs)

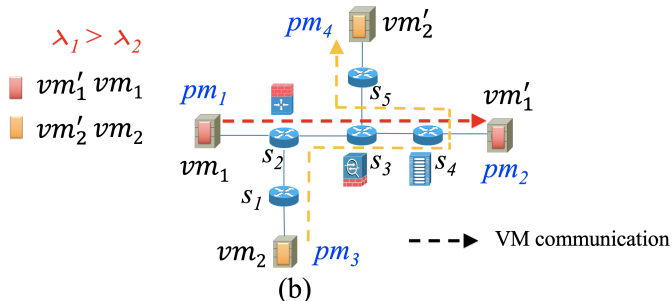


Figure 2: A PADC example with 2 virtual machine pairs

Policy Aware Data Centers (PADCs)

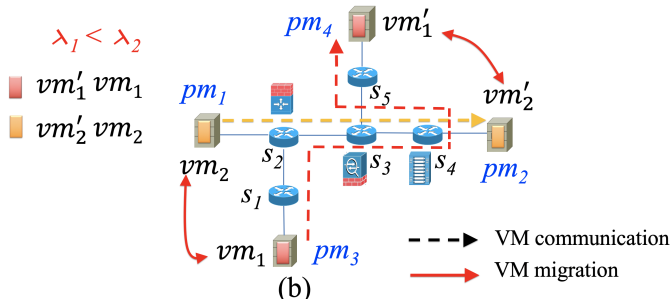


Figure 2: A PADC example with 2 virtual machine pairs

Fat Tree

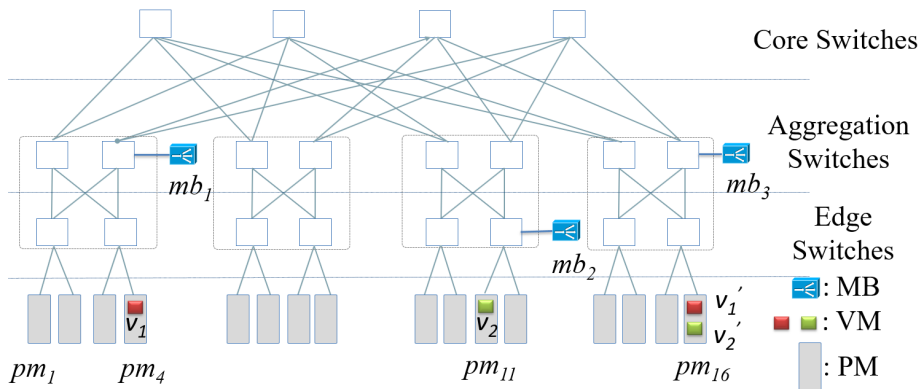


Figure 3: A PADC ($k=4$) with 16 PMs, 3 MBs, and two VM pairs

Ordered Policy: $vm_1 \rightarrow vm'_1$

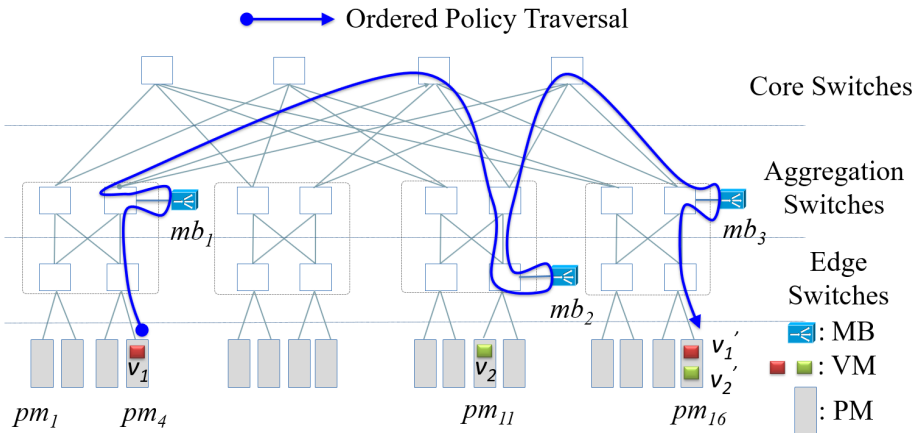


Figure 4: A PADC ($k=4$) with 16 PMs, 3 MBs, and two VM pairs

Unordered Policy: $vm_2 \rightarrow vm'_2$

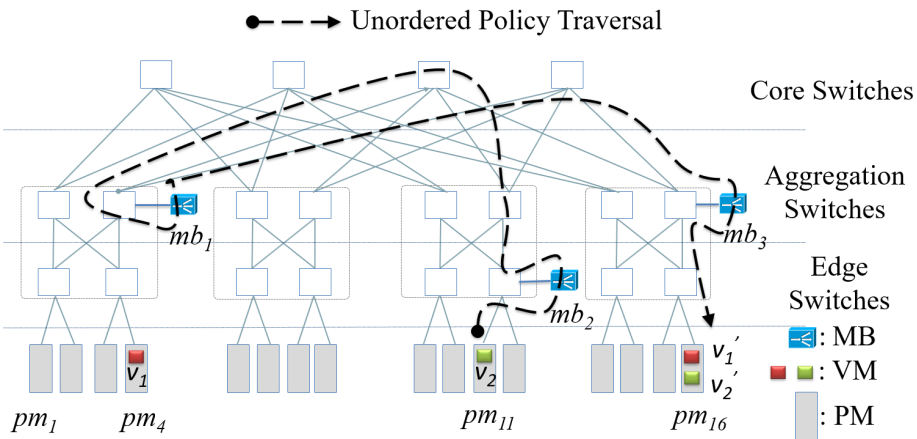


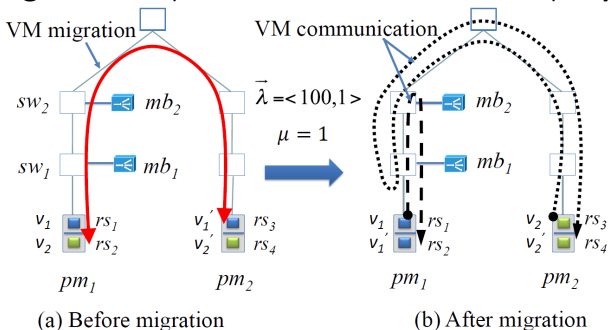
Figure 5: A PADC ($k=4$) with 16 PMs, 3 MBs, and two VM pairs

Network Model

- An undirected general graph $G(V, E)$.
 - $V = V_p \cup V_s$
 - $V_p = \{pm_1, pm_2, \dots, pm_{|V_p|}\}$.
 - V_s is a set of switches.
 - E is a set of edges, each connecting either one switch to another switch or a switch to a physical machine.
- A set of n middle boxes, $M = \{mb_1, mb_2, \dots, mb_n\}$ with mb_j installed at switch $sw_j \in V_s$.
- ℓ pairs of communicating VMs $P = \{(v_1, v'_1), (v_2, v'_2), \dots, (v_\ell, v'_\ell)\}$ in which v_i is source and v'_i is destination.
- $\vec{\lambda} = \langle \lambda_1, \lambda_2, \dots, \lambda_\ell \rangle$.
 - λ_i is the traffic rate or transmission rate between (v_i, v'_i) .
 - Change over time.

A $k = 2$ data Center with $n = 2, \ell = 2$

Figure 5: Computation the total cost in ordered policy



Before Migration	After Migration
$C_t(m) = \underbrace{6 \cdot 100}_{(v_1, v_1')} + \underbrace{6 \cdot 1}_{(v_2, v_2')} = 606$	$C_t(m) = \underbrace{4 \cdot 100}_{(v_1, v_1')} + \underbrace{10 \cdot 1}_{(v_2, v_2')} + \underbrace{(6 + 6) \cdot 1}_{\text{migration cost}} = 422$

Problem formulation of PAL for Ordered Policy

The total communication cost of all the l VM pairs under a VM placement function p as $C_c(p)$

$$C_c(p) = \underbrace{\sum_{i=1}^l \lambda_i \cdot \sum_{j=1}^{n-1} c(sw_j, sw_{j+1})}_{MBs} + \underbrace{\sum_{i=1}^l \lambda_i \cdot (c(p(v_i), sw_1) + c(sw_n, p(v'_i)))}_{c_{in}(rs), c_e(rs)}$$

The objective of PAL is to find a VM placement p to minimize $C_c(p)$ while satisfying resource constraint of PMs: $|\{v \in \mathbb{V} | p(v) = i\}| \leq rc_i, \forall i \in \mathbb{V}_p$.

PAL Algorithm for Ordered Policy

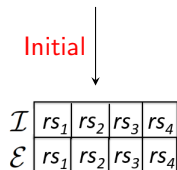
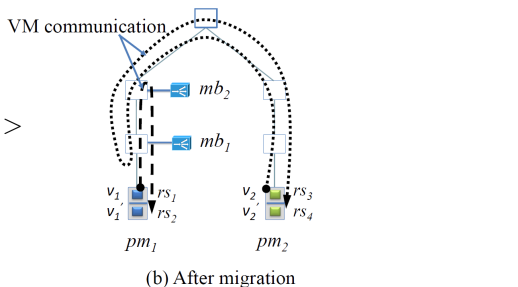
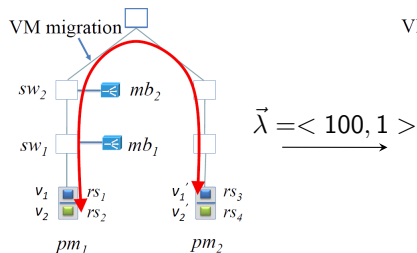
Algorithm 1

- **Input:** A PADC with ordered policy $(mb_1, mb_2, \dots, mb_n)$, VM pairs P , $V_p = \{pm_i\}$, resource capacity rc_i .
- **Output:** A placement p and the total communication cost $C_c(p)$
- **Algorithm(summary)**
 - 1 Sort resource slots in ascending order of their ingress (and egress) costs
 - 2 Find optimal resource slots for (v_k, v'_k)
 - 3 Descending frequencies communication of virtual machine pair
 - 4 Place VM pairs and calculate cost
- **Return:** A placement p and total communication cost.

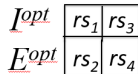
Theorem

PAL Algorithm for Ordered Policy finds the VM placement that minimizes total communication cost for the ℓ VM pairs.

PAL Algorithm for Ordered Policy



PAL Ordered Alg.



Total Communication cost: $100 \cdot 4 + 1 \cdot 10 = 410$

Problem Formulation of PAL for Unordered Policy

Total Communication Cost: $C_c(p, \vec{\pi}) = \sum_{i=1}^{\ell} c_i^{p, \pi^i}$, where

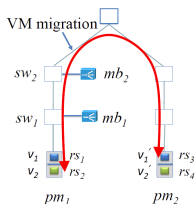
- Given p and π^i , denote (v_i, v'_i) 's communication cost as

$$c_i^{p, \pi^i} = \lambda_i \cdot c(p(v_i, sw(\pi^i(1)))) \\ + \lambda_i \cdot \sum_{j=1}^{n-1} c(sw(\pi^i(j)), sw(\pi^i(j+1))) + \lambda_i \cdot (sw(\pi^i(n)), p(v'_i))$$

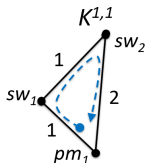
- MB traversal function π^i : a permutation function indicating the j^{th} MB that (v_i, v'_i) visits is $mb_{\pi^i j}$

PAL Algorithm for Unordered Policy

The algorithm achieves 2-approximation when $\ell = 1$

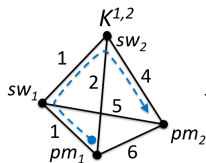


(a) Before migration



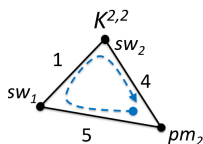
(a)

$$X = (1, 1, 4)$$



(b)

$$X = (1, 2, 6)$$



(c)

$$X = (2, 2, 10)$$

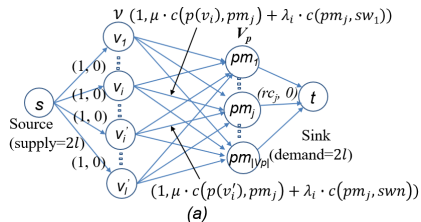
- (v_1, v'_1) is placed at pm_1 in (a), and (v_2, v'_2) at pm_2 in (c).
- Total cost: 410

Problem formulation of PAM for Ordered Policy

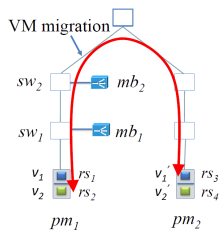
Let $C_t(m)$ be the total migration and communication cost of all pairs after m . Then total cost is calculated as follow:

$$\begin{aligned}
 C_t(m) = & \overbrace{\sum_{i=1}^{\ell} \lambda_i \cdot \sum_{j=1}^{n-1} c(sw_j, sw_{j+1})}^{\text{communication cost between middle boxes}} \\
 & + \underbrace{\sum_{i=1}^{\ell} (\mu \cdot c(p(v_i), m(v_i)) + \lambda_i \cdot c(m(v_i), sw_1))}_{\text{migration and communication } v_i \rightarrow \text{first mbs}} \\
 & + \underbrace{\sum_{i=1}^{\ell} (\mu \cdot c(p(v'_i), m(v'_i)) + \lambda_i \cdot c(sw_n, m(v'_i)))}_{\text{migration and communication } v'_i \rightarrow \text{last mbs}}
 \end{aligned}$$

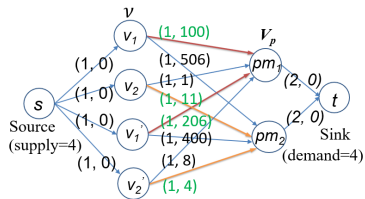
Transforming PADC to FLOW Network for $k = 2$



$$C_t = 100 + 11 + 206 + 4 + \underbrace{101}_{sw_1 \rightarrow sw_2} = 422$$



(a) Before migration



(b)

PAM is equivalent to minimum cost flow problem.

Graph transformation and MCF results for PADC

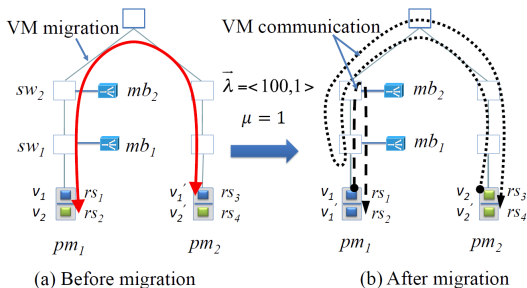
Problem formulation of PAL for Unordered Policy

Let $\vec{\pi} = \langle \pi^1, \pi^2, \dots, \pi^\ell \rangle$ and $C_t(m, \vec{\pi})$ denote the total cost of all the VM pairs with m and $\vec{\pi}$. Then,

$$\begin{aligned}
 C_t(m, \vec{\pi}) = & \overbrace{\sum_{i=1}^{\ell} (\mu \cdot c(p(v_i), m(v_i)) + \mu \cdot c(p(v'_i), m(v'_i)))}^{\text{total migration cost}} \\
 & + \sum_{i=1}^{\ell} \lambda_i \cdot \\
 & \underbrace{\left(\sum_{j=1}^{n-1} c(\text{sw}(\pi^i(j)), \text{sw}(\pi^i(j+1))) + c(m(v_i), \text{sw}(\pi^i(1))) + c(\text{sw}(\pi^i(n)), m(v')) \right)}_{\text{total communication cost}}
 \end{aligned}$$

PAM Algorithm for Unordered Policy

- 1 Compute the path route between two physical machine.
- 2 Find PM pair for VM pair, capacity of each physical machine is checked.
- 3 Update migration scheme and total cost.



Migrate $v_1' \rightarrow pm_1 : 406$, $v_2 \rightarrow pm_2 : 16$, $C_t(m) = 422$

Simulation Results

- Fat-tree PADCs, $k = 8$ with 128 PMs and $k = 16$ with 1024 PMs.
- The Traffic rates of VM pairs are in range of $[0, 1000]$.
- Traffic rates²: 25% in $[0, 300)$, 70% in $[300, 700]$, and 5% in $(700, 1000]$.
- VM Pairs³: 80% are placed into the PMs under the same edge switches, 20% under different edge switches.
- Each data point in an average of 20 runs with 95% confidence interval.
- In each run a new set of VM pairs are placed or migrated in the PADC.

Table 1: COMPARING PAM AND PAL ALGORITHMS

	Ordered Policy	Unordered Policy	Existing Work
PAL	Optimal (Algo. 1)	Approx-PAL (Algo. 3)	TrafficAware
PAM	MCF	Approx-PAM (Algo. 4)	PLAN

²Facebook Data Centers(Section 5.1, A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren. Inside the social network's (datacenter) network. In Proc. of ACM SIGCOMM 2015.)

³Cisco Design Guide, Cisco virtualized multi-tenant data center, version 2.0 compact pod design guide. <http://hyperurl.co/hpj2xt>.

PLAN & Traffic Aware


1 PLAN⁴ - Greedy algorithm

- Works in rounds.
- Each round, computes that which VM is migrated to which PM with available resources, such that the utility of this migration is the maximum among all the VMs that have not been migrated.
- Continues until all the VMs are migrated, or no more VM migration gives any positive utility.

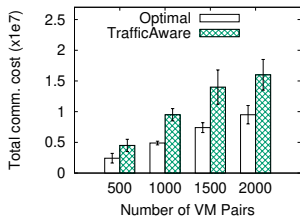
2 Traffic Aware⁵

- Only assigns VMs to the same PMs or PMs in close proximity, and does not consider the proximity of the PMs to the MBs.
- In ordered policy, it places VM pairs (in non-ascending order of their traffic rates) to the PMs that are closest to the ingress switch.
- In unordered policy, TrafficAware always places VM pairs in the same PM if possible.

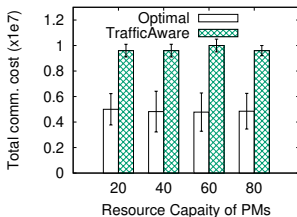
⁴L. Cui, F. P. Tso, D. P. Pezaros, W. Jia, and W. Zhao. Plan: Joint policy- and network-aware vm management for cloud data centers. *IEEE Transactions on Parallel and Distributed Systems*, 28(4):1163–1175, 2017

⁵X. Meng, V. Pappas, and L. Zhang. Improving the scalability of data center networks with traffic-aware virtual machine placement. In *Proc. of IEEE INFOCOM 2010*. 

Comparing with TrafficAware



(a) Varying ℓ , $n = 3$, $rc = 40$

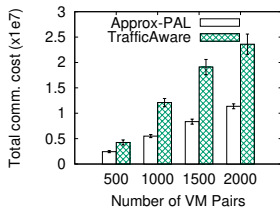


(b) Varying rc , $\ell = 1000$, $n = 3$

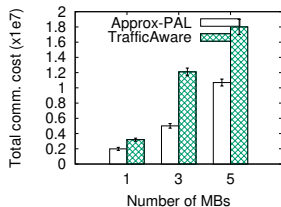
Figure 6: Comparing with TrafficAware in ordered policy, $k = 16$

- In (a), varies the number of VM pairs ℓ and shows that Optimal yields 46-49% less costs than TrafficAware.
- In (b), varies resource capacities of PMs rc and shows that Optimal outperforms TrafficAware by around 48%.

Comparing with TrafficAware (Cont.)



(a) Varying ℓ , $n = 3$, $rc = 40$

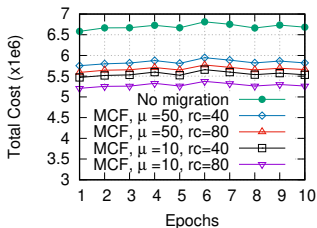


(b) Varying n , $\ell = 1000$, $rc = 40$

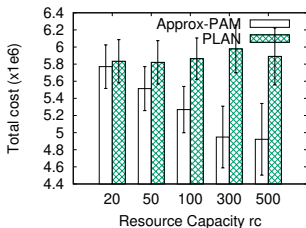
Figure 7: Comparing with TrafficAware in unordered policy, $k = 16$

- Varies ℓ as well as number of MBs n and shows that Approx-PAL outperforms TrafficAware by 37-58% in all scenarios.

Effect of VM Migrations



(a) Effects of VM migration

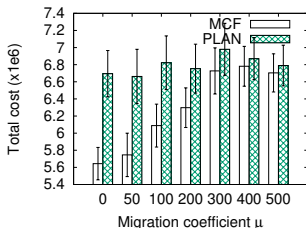


(b) Comparing with PLAN

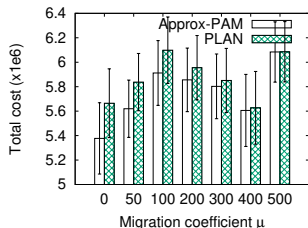
Figure 8: $k = 8, \ell = 1000, m = 3$

- VM migration reduces the total costs by up to 25% ($\mu = 10$) and $rc = 80$ compared to without migration.

Comparing with PLAN



(a) Ordered Policy



(b) Unordered Policy

Figure 9: Comparing with PLAN, $k = 8, \ell = 1000, n = 3, rc = 40$

- In (a), ordered policy, the MCF outperforms are the PLAN by around 20% when μ small. With the increase of μ , PLAN and MCF start to perform close due to high migration cost.
- In (b), unordered policy, Approx-PAM outperforms PLAN slightly for the entire range of μ .

Comparing with PLAN (cont.)

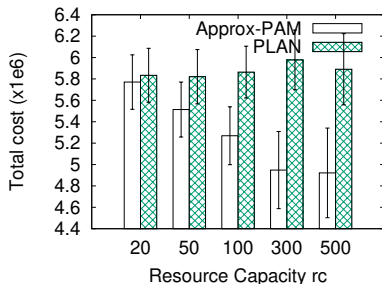


Figure 10: Comparing with PLAN, $k = 8$, $\ell = 1000$, $n = 3$, $rc = 40$

- Compares Approx-PAM and PLAN by varying rc while fixing μ as 20, and shows that Approx-PAM outperforms PLAN for the entire range of rc .
- With the increase of the rc , the performance difference between Approx-PAM and Greedy gets larger, around 30%.

Conclusion

- VM migration is an effective technique to alleviate dynamic VM traffic in PADCs.
- Working together, PAL and PAM place and then migrate VMs in the event of dynamic traffic fluctuation, achieving optimal and near-optimal network resource management for a PADC's lifetime.
- Uncovered a suite of new policy-aware problems and designed optimal, approximation, and heuristic algorithms.

Future Work

- Study if the optimality and approximability of our algorithms still hold when VMs have different resource demands.
- Study how VNF migration mitigates diverse and dynamic traffic and design a holistic VNF+VM migration scheme to achieve ultimate resource optimization in PADCs.