Acinonyx: Dynamic Flow Scheduling for Virtual Machine Migration in SDN-enabled Clouds

The 16th IEEE International Symposium on Parallel and Distributed Processing with Applications

Adel Nadjaran Toosi
Faculty of Information Technology
Monash University
Email: adel.n.toosi [AT] monash.edu
Homepage: http://adelnadjarantoosi.info

Rajkumar Buyya
CLOUDS Laboratory
School of Computing and Information Systems
The University of Melbourne

Outline

- Virtualization and Live VM Migration
- SDN and SND-enabled Cloud Data Centres
- Research Question
- ACINONYX
  - Dynamic Flow Scheduling for VM migration
- System Architecture
- Performance Evaluation
  - with a Prototype and a real-world testbed
- Summary and future directions

Virtualization

- Virtualization
  - Critical building block of the **operation** and **maintenance** in cloud data centres
- Virtual Machine (VM)
  - A substitute for a real machine
  - Allowing for multiple OS (isolated from one another), on the same physical machine.
- Live VM migration
  - Relocating running VMs between servers with no or minimum impact on the VM's availability.
  - hardware maintenance, load balancing and consolidation, energy saving, and disaster recovery.

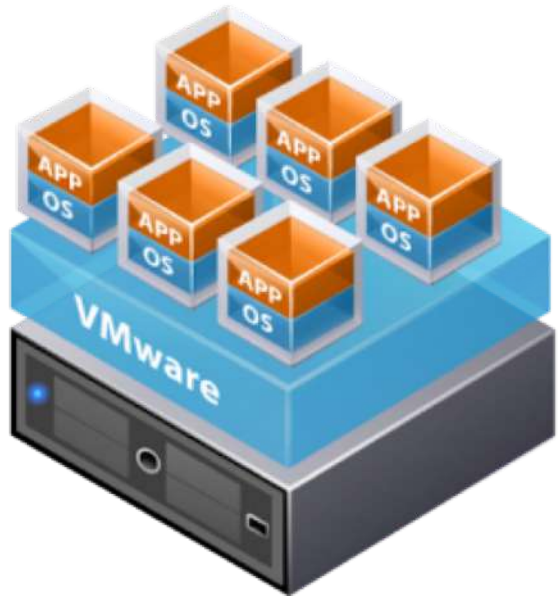

VM Migration Impacts on the Network

- Live VM migration
  - Transfer from Source to Destination
    - VM's CPU state
    - All memory pages
    - Disks
- Adverse Impacts on network
  - Elephant flows over the network links
  - Causing network congestion for other applications
- Focus of this work
  - The selection of network paths for VM migration to avoid network congestion.
  - Is possible to reduce live VM migration time and network overhead by dynamically scheduling flows in a cloud data centre?

Software-Defined Networking

- Separation of control plane from data forwarding plane
- Platform is decoupled from infrastructure
- Centralized controller, network-wide control by controller SW that performs routing and traffic engineering

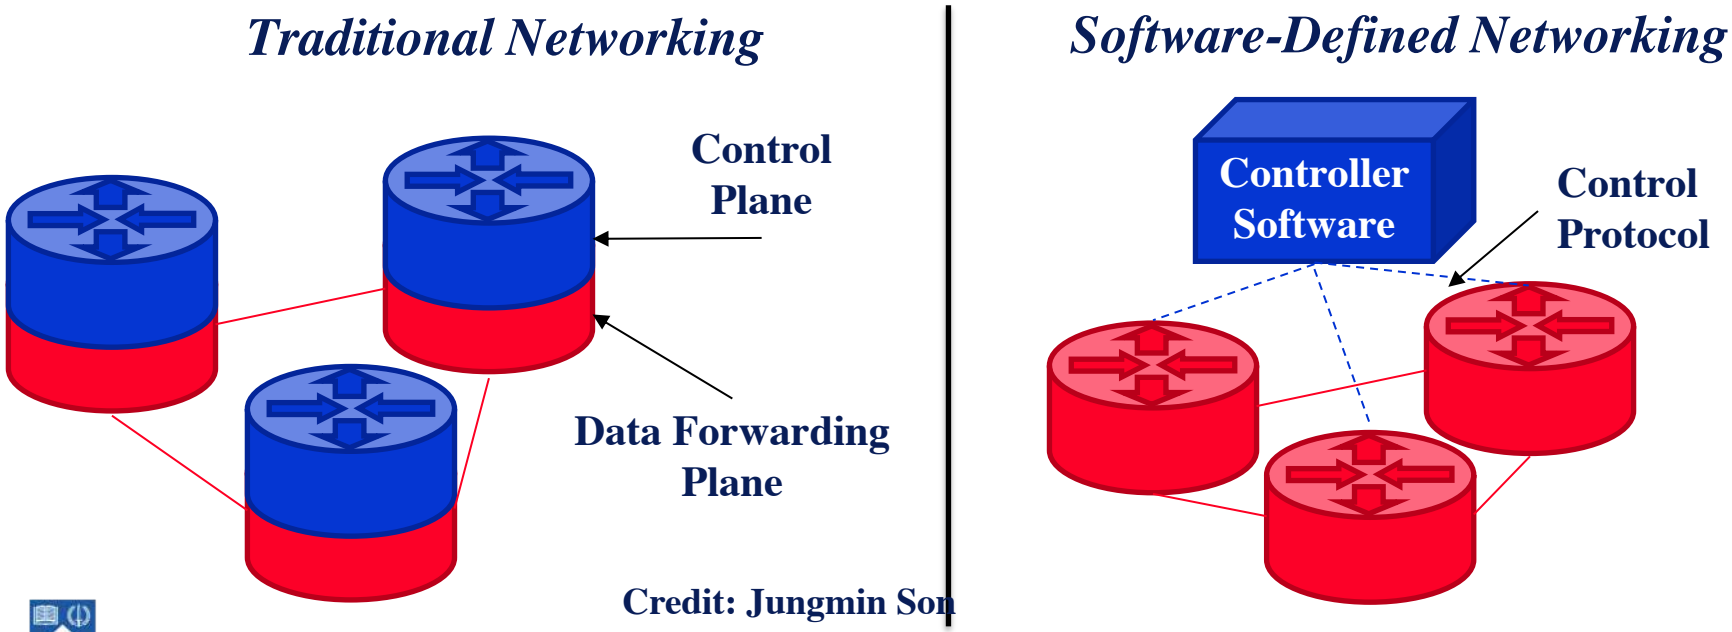

SDN Benefits

- Programmable network
- Enables dynamic configuration of networking
- Open opportunities for innovation

Data Centre Network

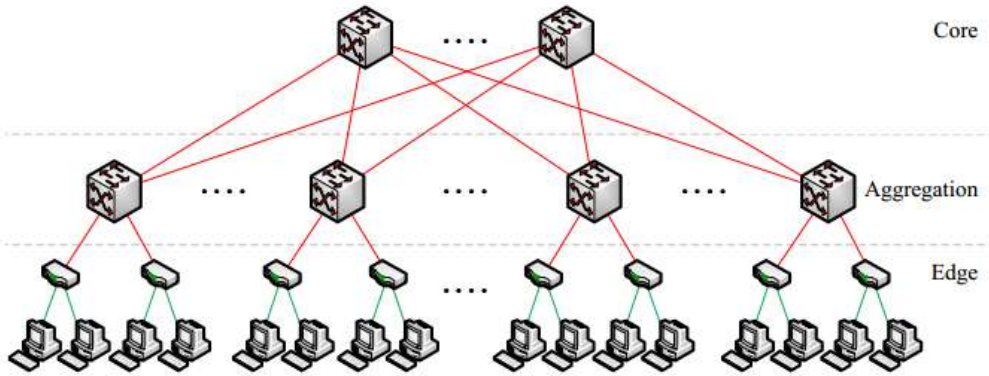

- Network topology
  - Fat tree, VL2, PortLand, and Bcube
  - Specialized hardware (middleboxes) implementing networking functions
  - NAT, load balancing, WAN optimization, firewall…
- Specialized communication protocols for top tiers
- Communication patterns between hosts change frequently

Dynamic Flow Scheduling for Virtual Machine Migration

- "Is it possible to reduce **live VM migration time** and **overhead** by dynamically scheduling flows in a cloud data center with multiple paths available between a given pair of physical hosts?"

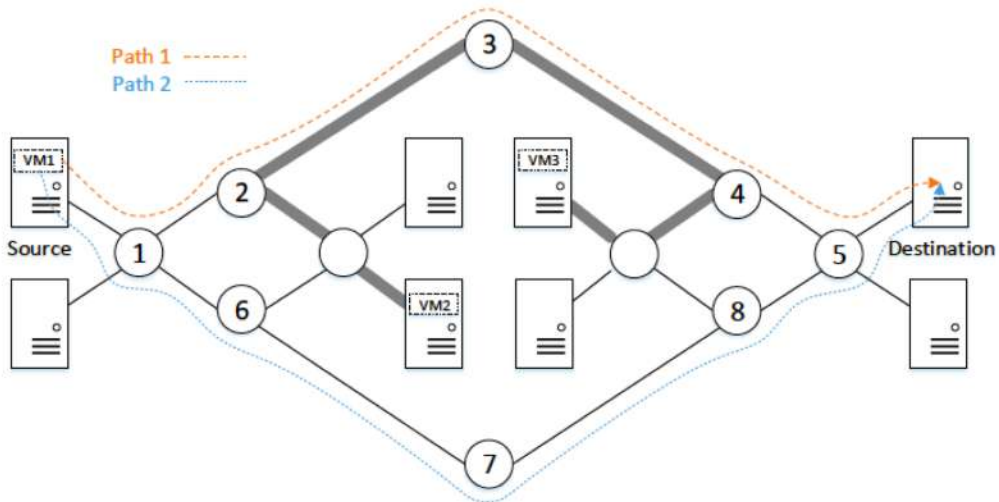

ACINONYX: Dynamic Flow Scheduling

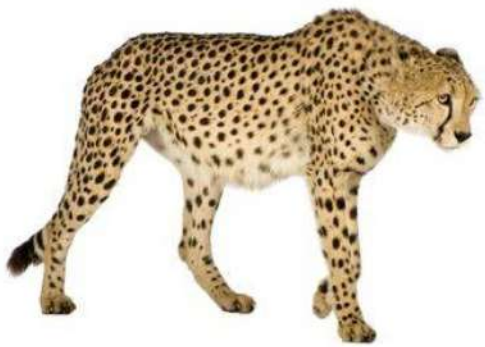

Algorithm 1 Acinonyx Dynamic Flow Scheduling

Input: s, d

```
1:  while MIGRATION-IS-IN-PROGRESS() do
2:      G ← GET-TOPOLOGY()
3:      P ← FIND-SHORTEST-PATHS(G, s, d)
4:      min ← +∞
5:      for p in P do
6:          max ← 0
7:          for link in p do
8:              (b, f) ← GET-BYTE-RATE(link)
9:              r ← b − f
10:             if r > max then
11:                 max ← r
12:             end if
13:         end for
14:         if max ≤ min then
15:             path ← p
16:             min ← max
17:         end if
18:     end for
19:     mbr ← MAXBYTERATE(currentPath)
20:     if mbr − min > mbr × α then
21:         PUSH-FLOWS(path)
22:     end if
23:     SLEEP(β)
24: end while
```

System Architecture

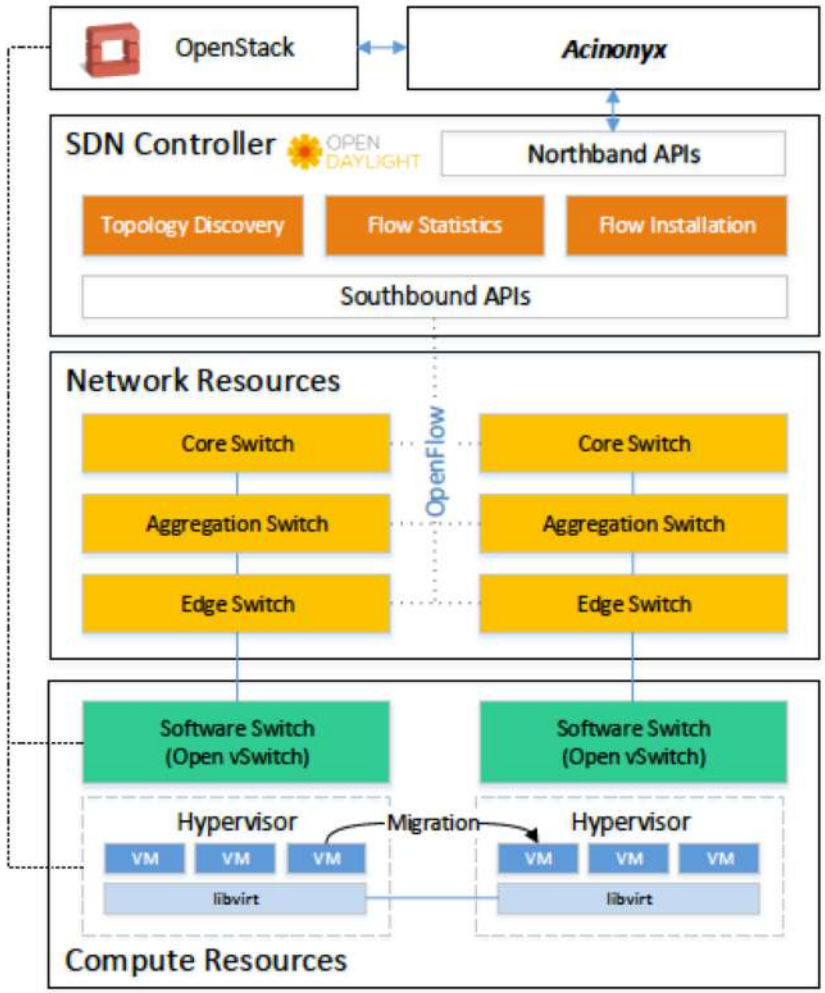

Performance Evaluation

- Real-world testbed for SDN-enabled cloud computing
- 8 heterogeneous servers
- **Fat-tree** topology
- **Raspberry Pis** (Pi 3 Model B) with Open vSwitch (OVS) as Switch
- **OpenStack** (Ocata release) as cloud platform
- **OpenDaylight** (ODL) as SDN controller

Hardware

Machine	CPU	Cores	Memory	Storage
3 x IBM X3500 M4	Intel(R) Xeon(R) E5-2620 @ 2.00GHz	12	64GB (4 x 16GB DDR3 1333MHz)	2.9TB
4 x IBM X3200 M3	Intel(R) Xeon(R) X3460 @ 2.80GHz	4	16GB (4 x 4GB DDR3 1333MHz)	199GB
2 x Dell OptiPlex 990	Intel(R) Core(TM) i7-2600 @ 3.40GHz	4	8GB (2 x 4GB DDR3 1333MHz)	399GB

USB 2.0 to 100Mbps Ethernet adapters

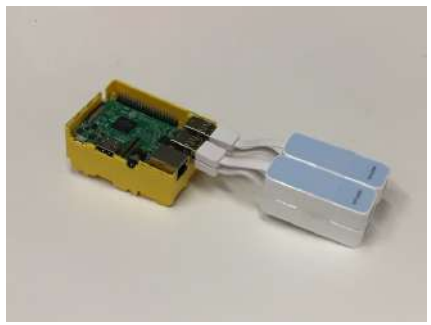

Raspberry Pis (Pi 3 MODEL B)

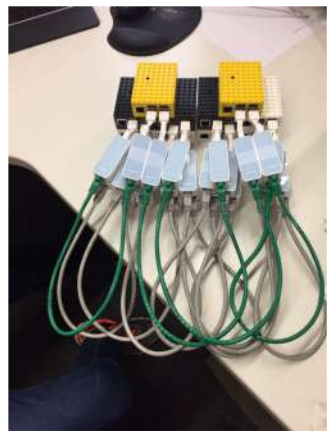

Testbed

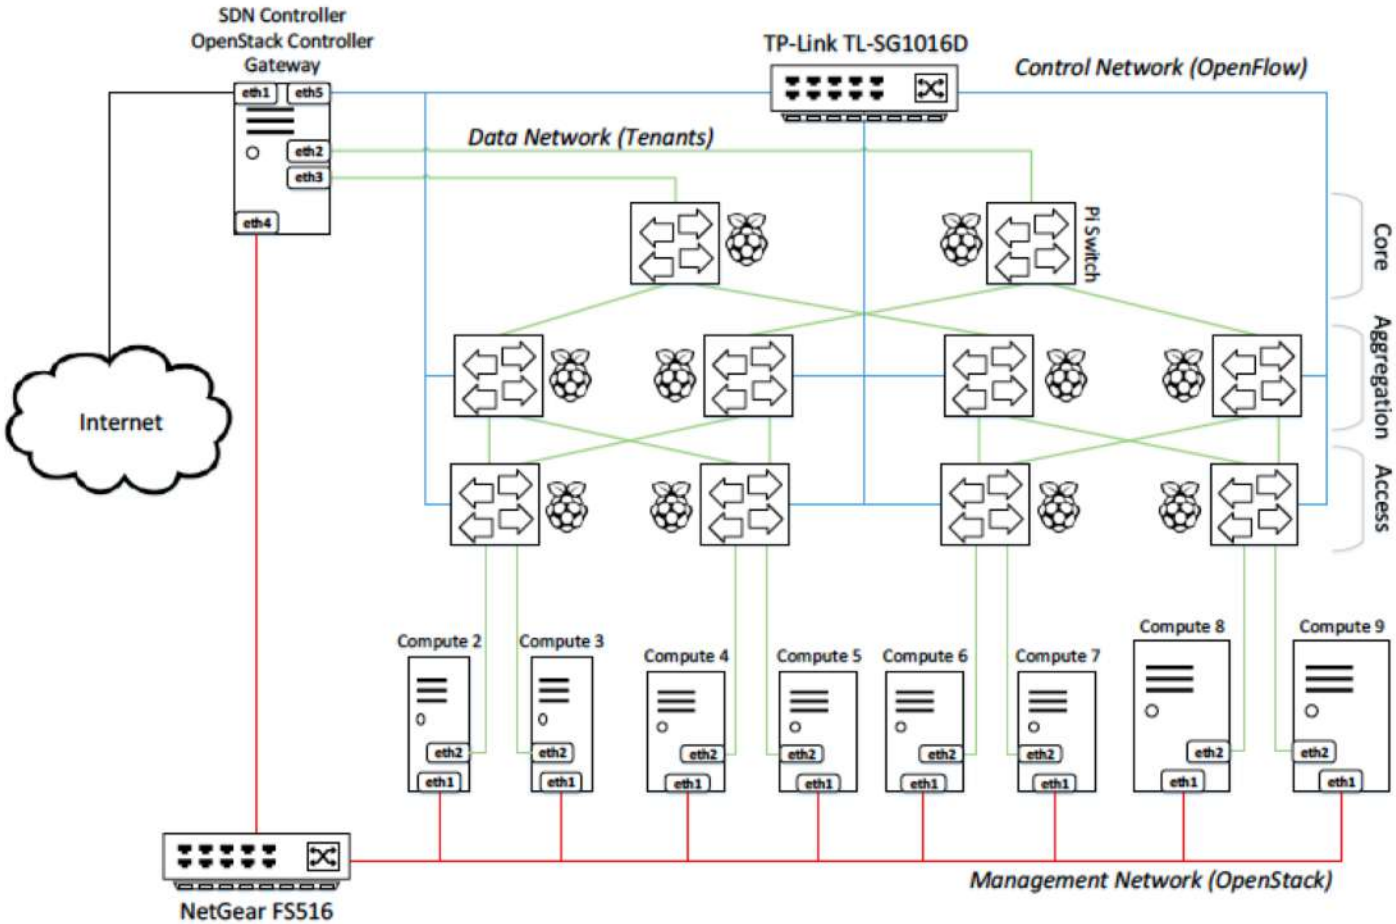
SDN Controller
OpenStack Controller
Gateway
eth1
eth5
eth2
eth3
eth4
TP-Link TL-SG1016D
Control Network (OpenFlow)
Data Network (Tenants)
Pi Switch
Core
Aggregation
Access
Internet
Compute 2
Compute 3
Compute 4
Compute 5
Compute 6
Compute 7
Compute 8
Compute 9
eth2
eth1
NetGear FS516
Management Network (OpenStack)

More Photos

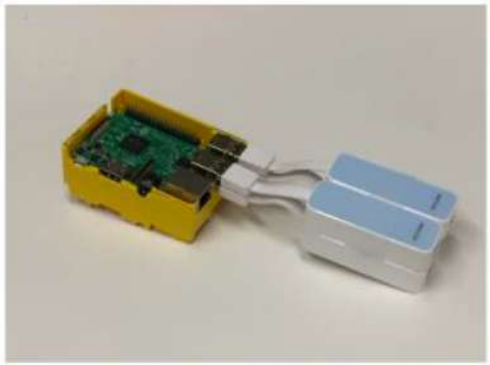

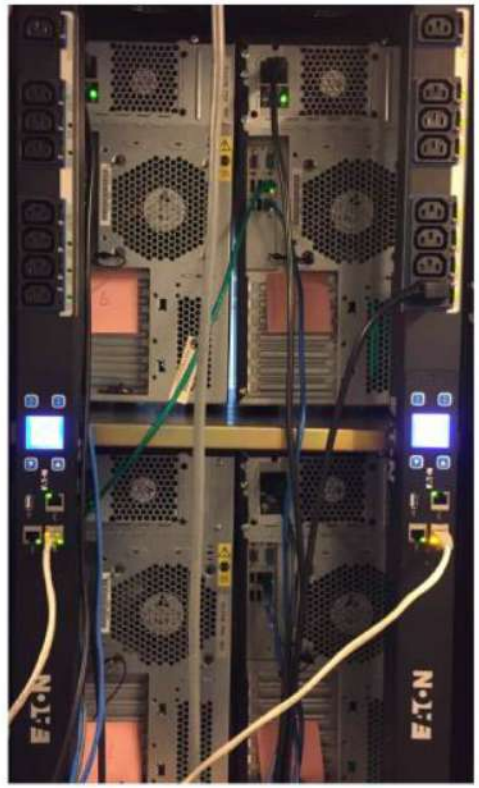

Traffic Generation and Routing

- Iperf3 tool

Conn-x	Time	Source	Destination	Length	BW	Path
Conn-1	0	Compute-5	Compute-4	60s	10	8-5-1-3-7
Conn-2	25	Internet-1	Compute-2	120s	30	0-2-6
Conn-3	40	Compute-4	Compute-5	120s	60	7-3-1-5-8
Conn-4	125	Compute-8	Compute-2	120s	40	9-4-0-2-6
Conn-5	180	Internet-1	Compute-6	30s	50	0-4-8
Conn-6	185	Compute-5	Compute-8	60s	20	8-5-9
Conn-7	200	Compute-2	Compute-8	90s	40	6-2-0-4-9

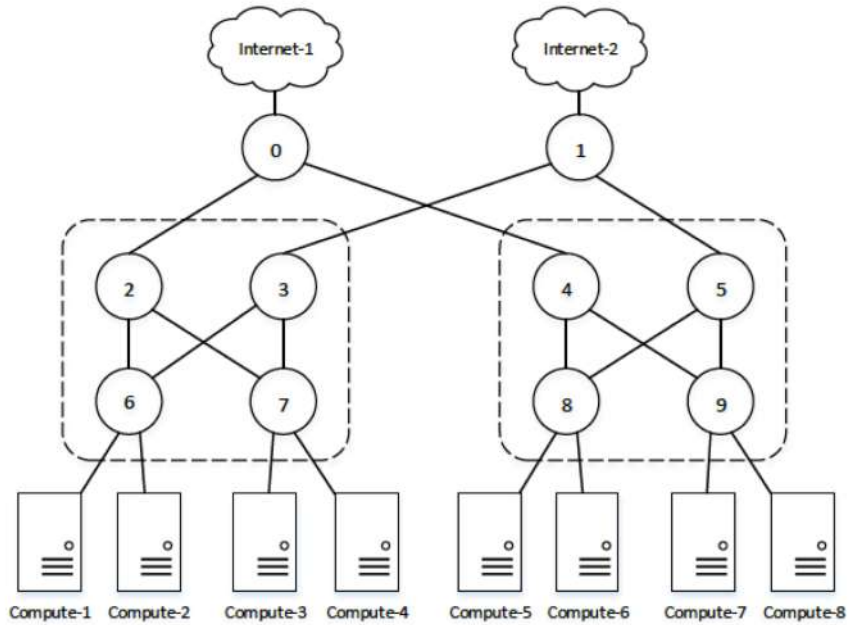

Experimental Results (EXP-1)

- A live VM migration in OpenStack
  - m1.small VM with Ubuntu-16.04 image
  - from Compute-1 to Compute-7
- Results:

Metric	Static Routing	Acinonyx
Migration Time (s)	287	**256**
Average Throughput (Mbs)	32.0	**34.4**

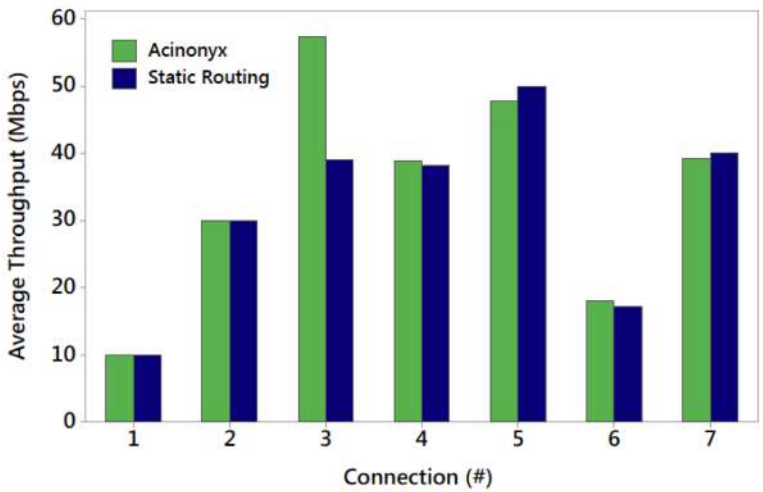

The impact of the Sleeping Interval

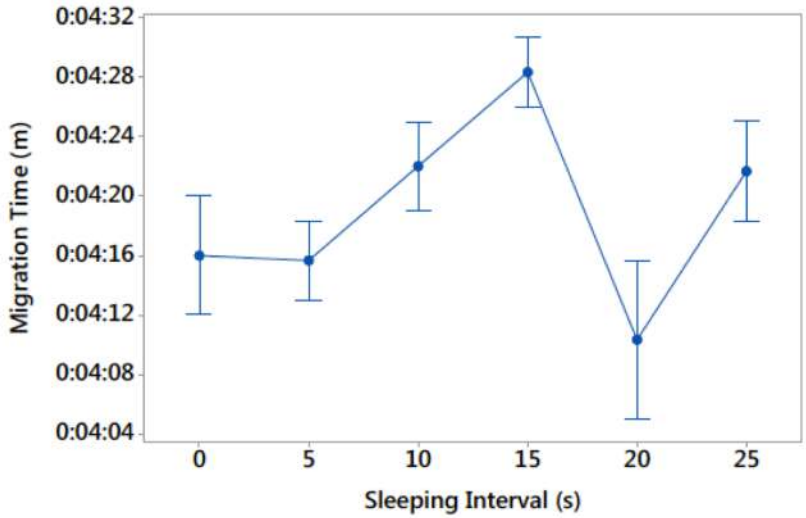

(a)

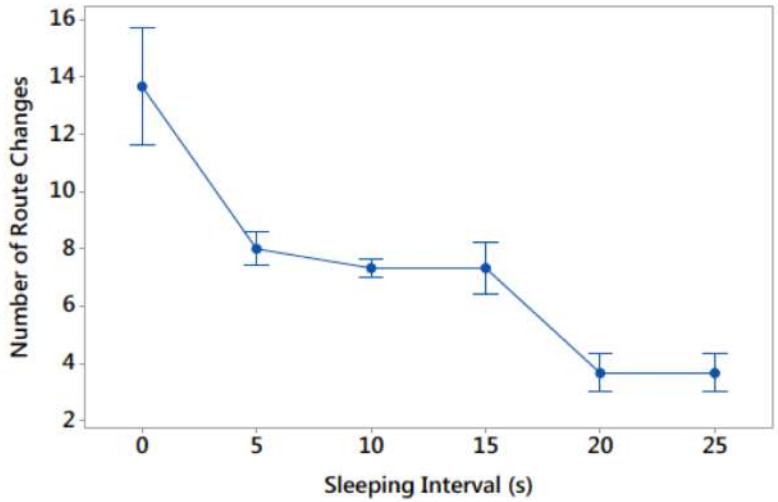

(b)

The impact of Switching Ratio Factor

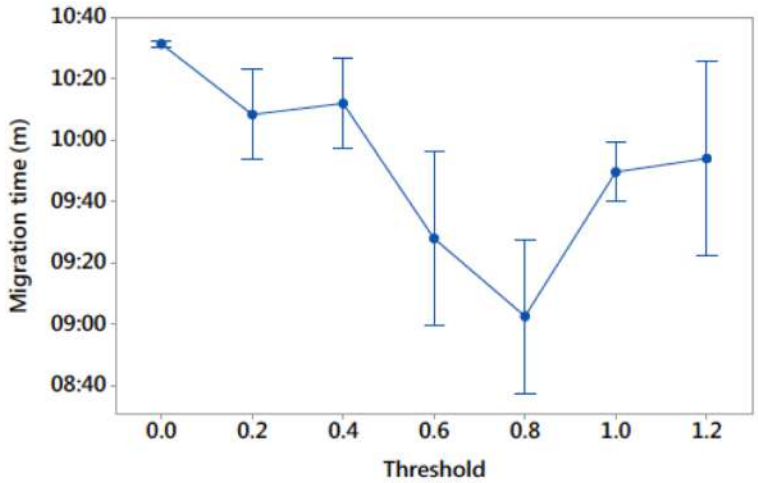

(a)

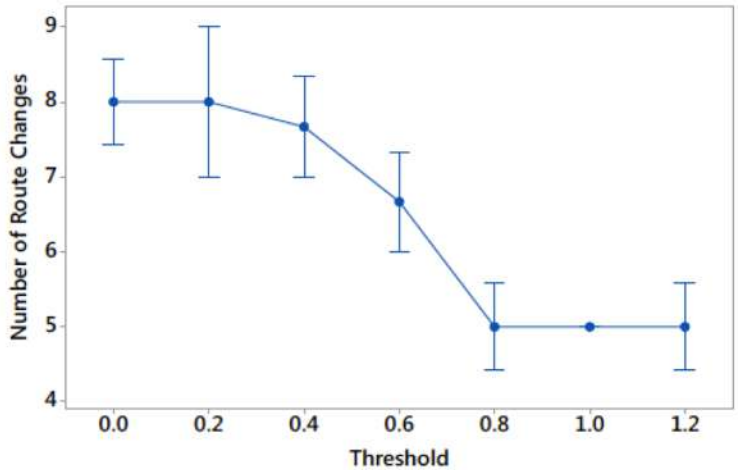

(b)

Multiple Migrations

- m1.small VMs,
  - One from Compute-2 to Compute-7 and
  - One from Compute-3 to Compute-8.

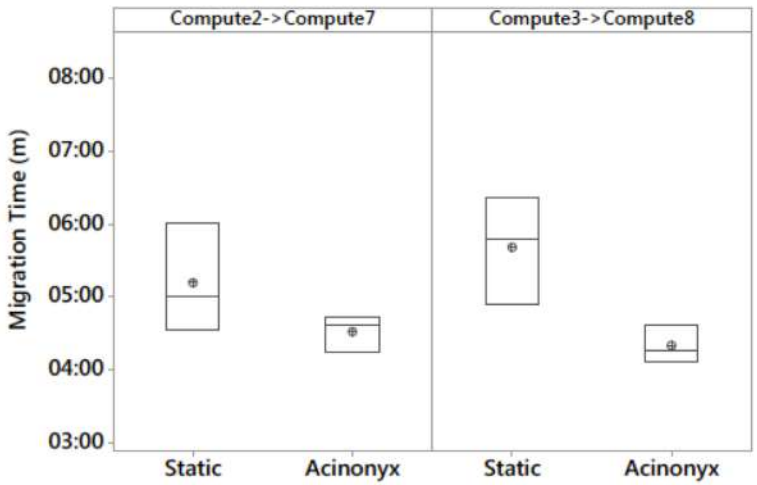

(a) No Background Traffic.

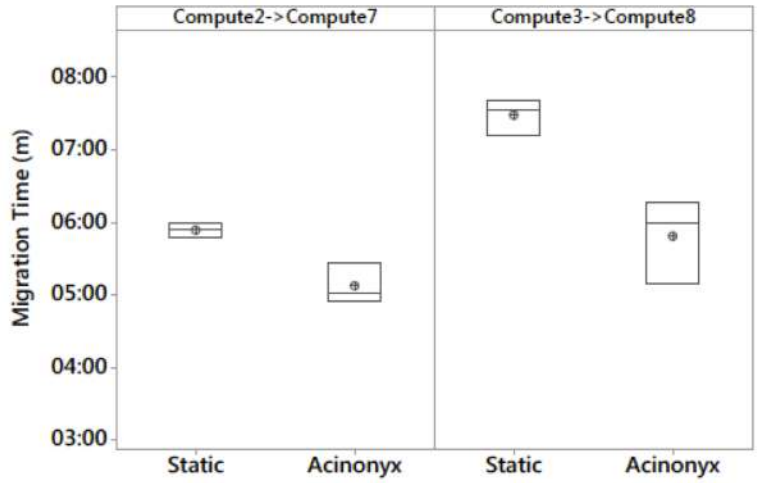

(b) With Background Traffic

- Live VM migration is frequently used in cloud data centers
- Proposed a dynamic flow scheduling for live VM migration in SDN-enabled cloud data centers
- Showed live VM migration time can be reduced up to 10% compared to existing static routing
- Demonstrated the feasibility by building a working prototype over a practical testbed.

Future work

- To explore the impact *Acynonix* on the flows of applications running in the migrating VM.
- To extend our approach for efficient flow scheduling of combined multiple migrations.
- Study the performance under the deployment of real-world applications exhibiting various network traffic characteristics
  - stream processing, data analytics, web applications, scientific workflows

THANK YOU

Questions?