Elastic Virtual Network Function Placement

CloudNet 2015

M. GHAZNAVI, A. KHAN, N. SHAHRIAR, KH. ALSUBHI, R. AHMED, R. BOUTABA

DAVID R. CHERITON SCHOOL OF COMPUTER SCIENCE

UNIVERSITY OF WATERLOO

Outline

Introduction

State of the Art

Problem: Elastic Virtual Network Function Placement

Solution: Simple Lazy Facility Location

Evaluation

Conclusion

Introduction

MIDDLE-BOXES

NETWORK FUNCTION VIRTUALIZATION

VNF SERVICES IN CLOUD

Middle-Boxes

"any intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host" [1]

Expensive hardware

Hard to deploy

Hard to modify

Hard to scale

Provision for peak-load

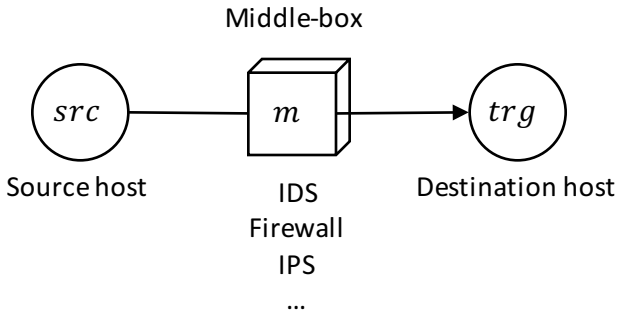

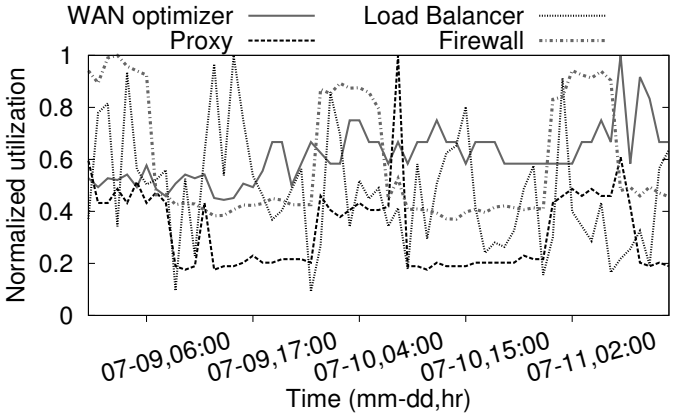

Middle-box utilization peak at different times [2]

[1] CARPENTER, B., AND BRIM, S. Middleboxes: Taxonomy and Issues. RFC 3234, https://tools.ietf.org/rfc/ rfc3234.txt, 2002.

[2] V. Sekar, N. Egi, S. Ratnasamy, M. K. Reiter, and G. Shi. Design and implementation of a consolidated middlebox architecture. In Proceedings of NSDI 12, 2012.

Network Function Virtualization

Virtualization (Softwarization) of middle-boxes

Software middle-boxes are called Virtual Network Function (VNF)

NFV *"involves the implementation of network functions in software that can run on a range of industry standard server hardware, and that can be moved to, or instantiated in, various locations in the network as required, without the need for installation of new equipment."*[1]

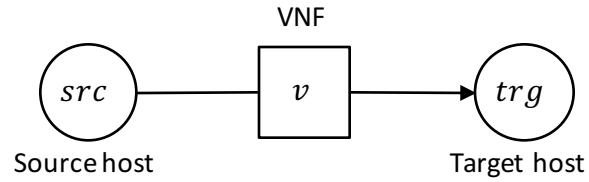

[1] "Network Functions Virtualization". ISG web portal: https://portal.etsi.org/nfv/nfv_white_paper.pdf

Network Function Virtualization

MIDDLE-BOXES	VIRTUAL NETWORK FUNCTIONS
Expensive hardware	Low-cost software
Hard to deploy	Easy to deploy
Hard to modify	Easy to modify
Hard to scale	Easy to scale
Provision for peak-load	Scale resources on demand

VNF Services in Cloud

Offered by cloud providers

- IBM Bluemix
- Microsoft Azure
- Amazon EC2

Services

- Riverbed STEELHEAD WAN optimizer [1]
- McAfee Next Generation firewall [2]
- Virtual LoadMaster load balancer [3]

[1] http://media-cms.riverbed.com/documents/Spec+Sheet+-+Steelhead+Family+-+05.06.2015.pdf
[2] https://kc.mcafee.com/resources/sites/MCAFEE/content/live/PRODUCT_DOCUMENTATION/25000/PD25151/en_US/NGFW_57_HW_Requirements.pdf
[3] http://kemptechnologies.com/files/downloads/documentation/Datasheets/VLM-AWS.pdf

VNF Services in Cloud

WHAT CLOUD PROVIDER SHOULD SUPPORT

Pay per use

- Clients pay only for real used resources

Elasticity

- Scale resources on demand
  - Upon arrival or departure of service request
  - Variation of workload of admitted service request

CHALLENGES OF CLOUD PROVIDER

Minimizing Costs:

- Trade-off between Host & Bandwidth Resources

Elasticity

- Which mechanisms to apply
- Elasticity benefit vs. its overhead

VNF Services in Cloud

Where to place VNF instances?

Which request must be assigned to which VNF instance?

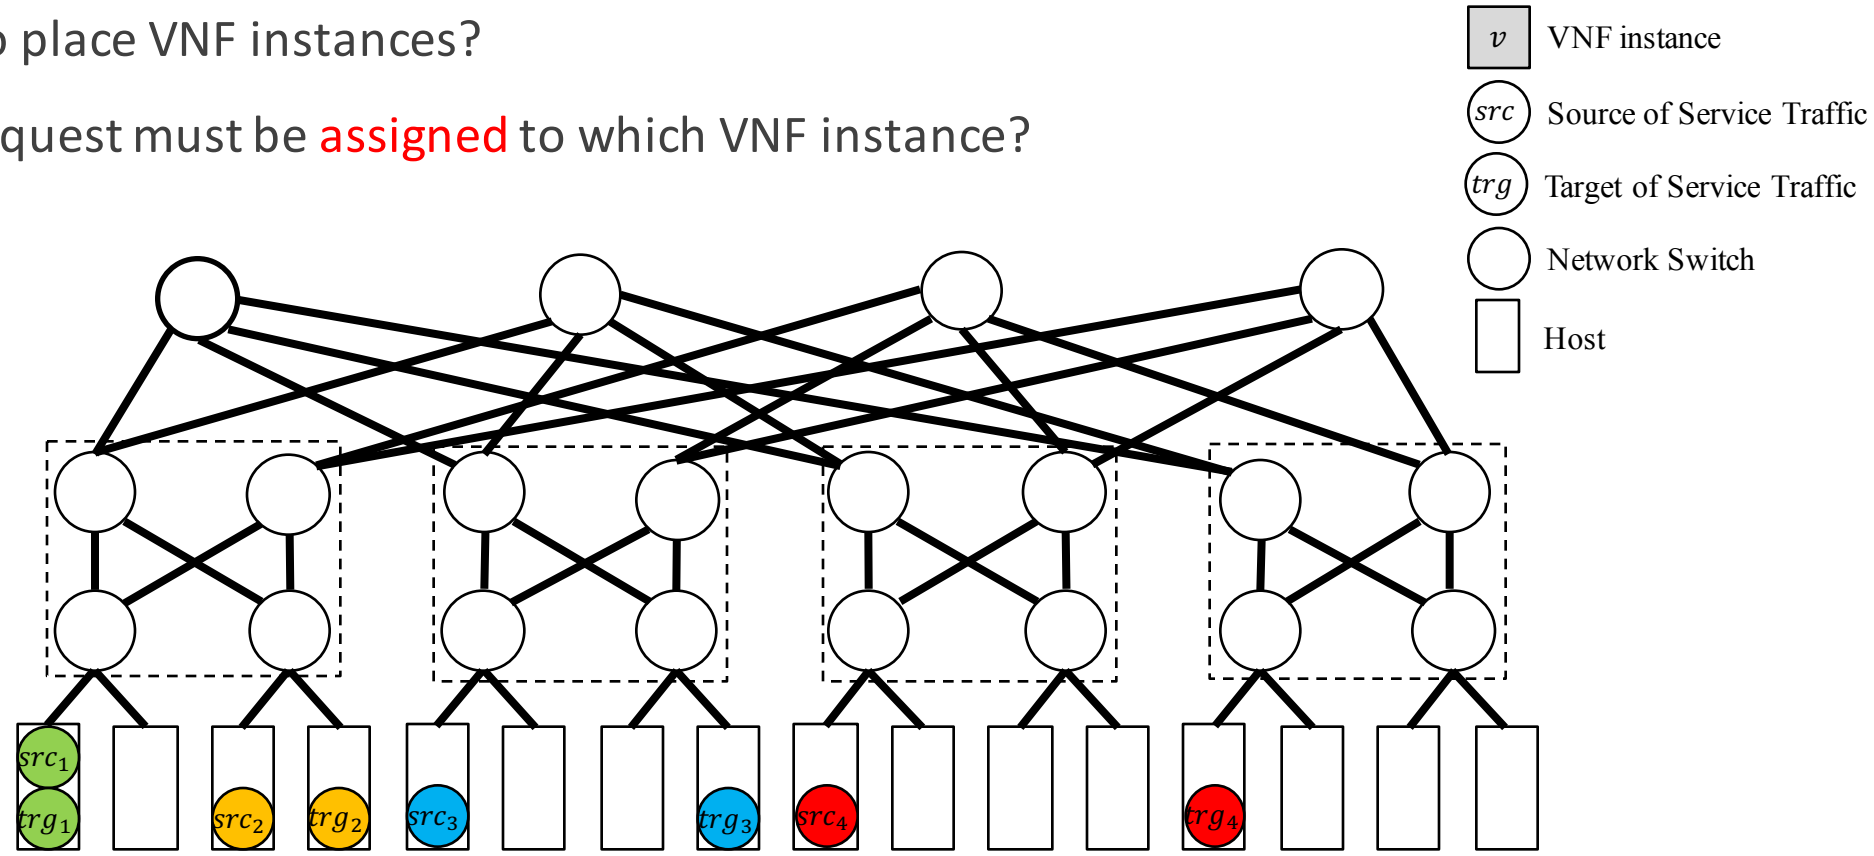

VNF Services in Cloud

A solution can be

- v_1 serves the first and second service traffics
- v_2 serves the third and forth service traffics

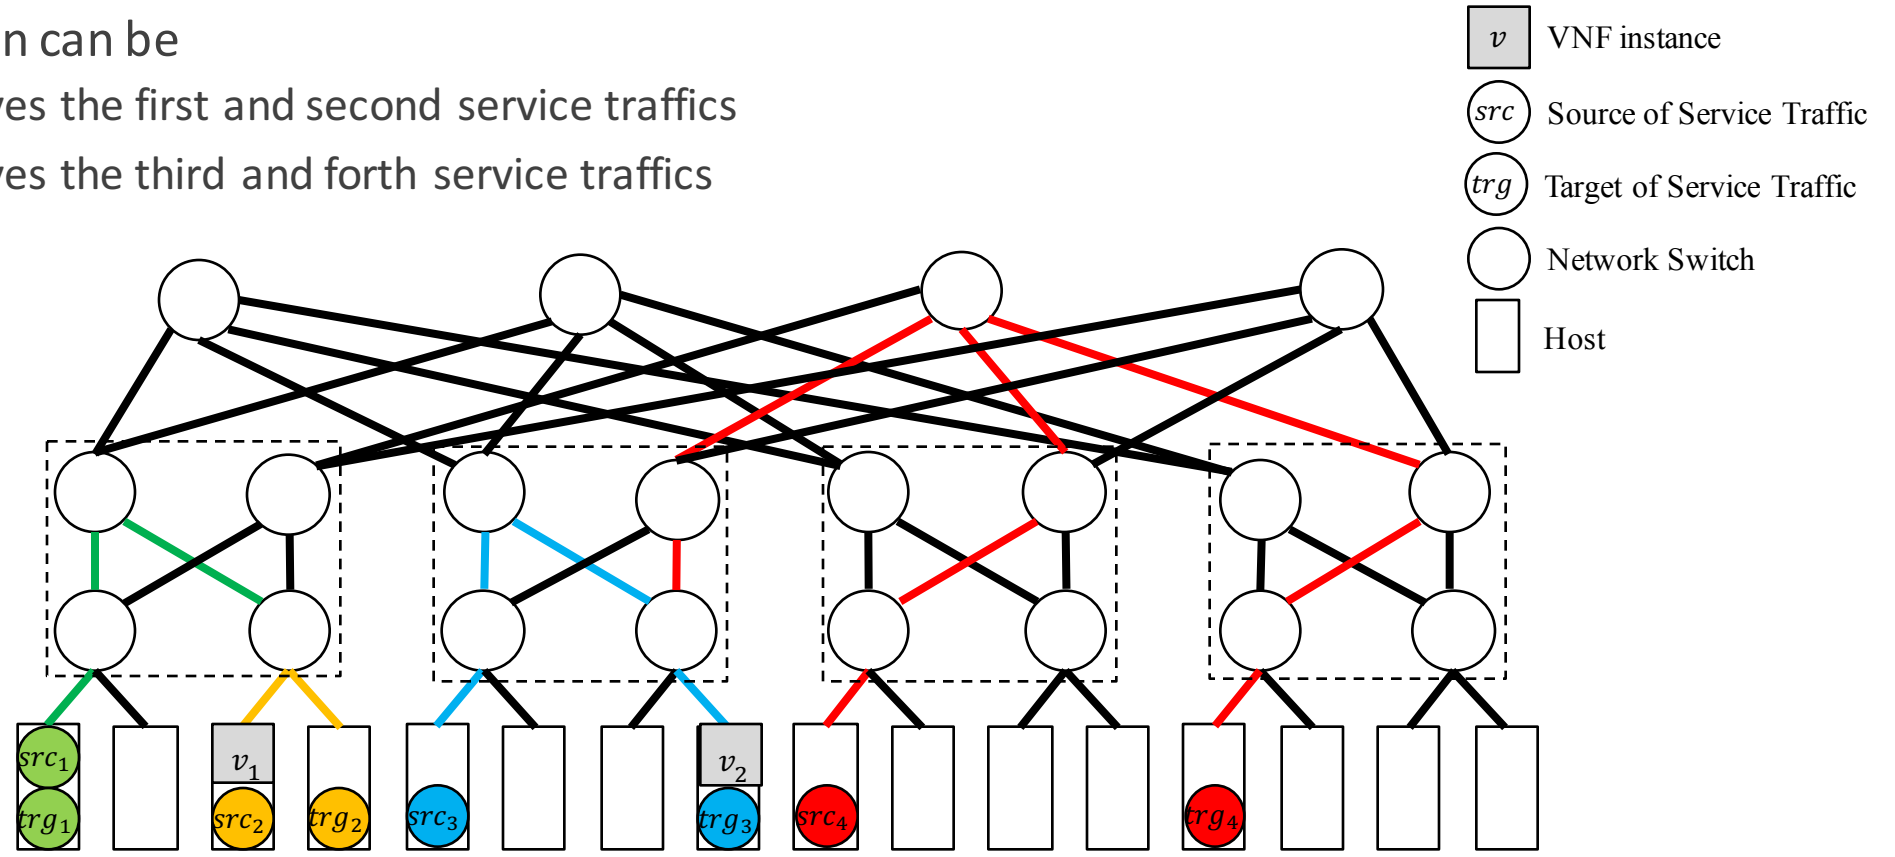

State of the Art

COMPARISON OF STATE OF THE ART

Comparison of State of the Art

Paper	Host Res. Cost	Bandwidth Res. Cost	Elasticity
Elastic Virtual Network Function Placement (EVNFP)	✔	✔	✔
Elasticity in Cloud [1, 2, 3]	✔	✘	✔
Dynamic VM Placement [2, 4]	✔	✘	✔
Network Aware VM Placement [5, 6, 7]	✔	✔	✘
Virtual DPI Placement [8]	✔	✔	✘

[1] Z. Gong, X. Gu, and J. Wilkes. Press: Predictive elastic resource scaling for cloud systems. In *IEEE CNSM, 2010*
[2] U. Sharma, P. Shenoy, S. Sahu, and A. Shaikh. A cost-aware elasticity provisioning system for the cloud. In *IEEE ICDCS 2011*.
[3] Z. Shen, S. Subbiah, X. Gu, and J. Wilkes. Cloudscale: Elastic resource scaling for multi-tenant cloud systems. In *ACM SoCC, 2011*.
[4] A. Verma, P. Ahuja, and A. Neogi. pmapper: Power and migration cost aware application placement in virtualized systems. In *ACM/IFIP/USENIX Middleware, 2008*.
[5] O.Biranetal.A stable network-aware vm placement for cloud systems. In *CCGRID*, pages 498–506, 2012.
[6] V. Mann, A. Kumar, P. Dutta, and S. Kalyanaraman. Vmflow: Leveraging vm mobility to reduce network power costs in data centers. In *IFIP NETWORKING, 2011*.
[7] X. Meng, V. Pappas, and L. Zhang. Improving the scalability of data center networks with traffic-aware virtual machine placement. In *IEEE INFOCOM, 2010*.
[8] M. Bouet, J. Leguay, and V. Conan. Cost-based placement of vdpi functions in nfv infrastructures. In *NetSoft*, 2015.

Problem: Elastic Virtual Network Function Placement (EVNFP)

SCOPE AND ASSUMPTIONS

CONFLICTING OBJECTIVES

ELASTICITY MECHANISMS AND OVERHEAD

Scope and Assumptions

SCOPE

Single cloud provider

Single data-center

Centralized optimization

ASSUMPTIONS

One VNF instance-type

Multi-tenancy

Elasticity Mechanisms

- Horizontal Scaling
- Migration of VNF instances
- Reassignment of workload

Conflicting Objectives

Minimizing the bandwidth cost, and

Minimizing the number of installed VNFs

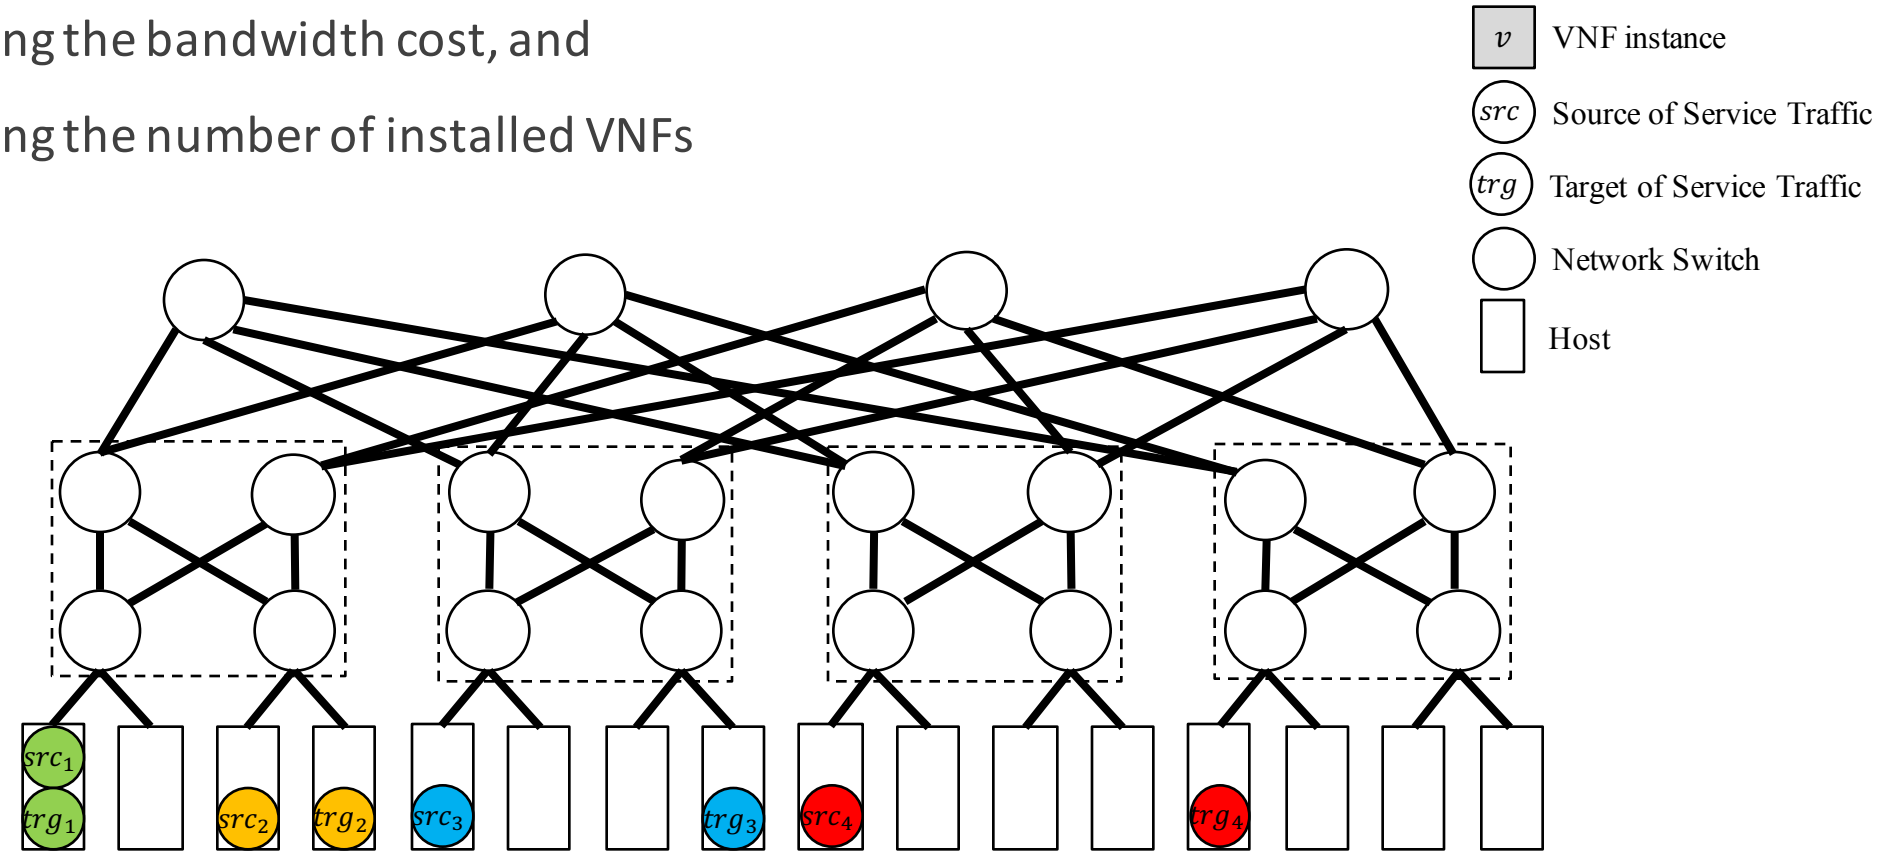

Conflicting Objectives

Minimizing the bandwidth cost:

- 12 Unit of Bandwidth over 12 Links
- 4 VNF instances

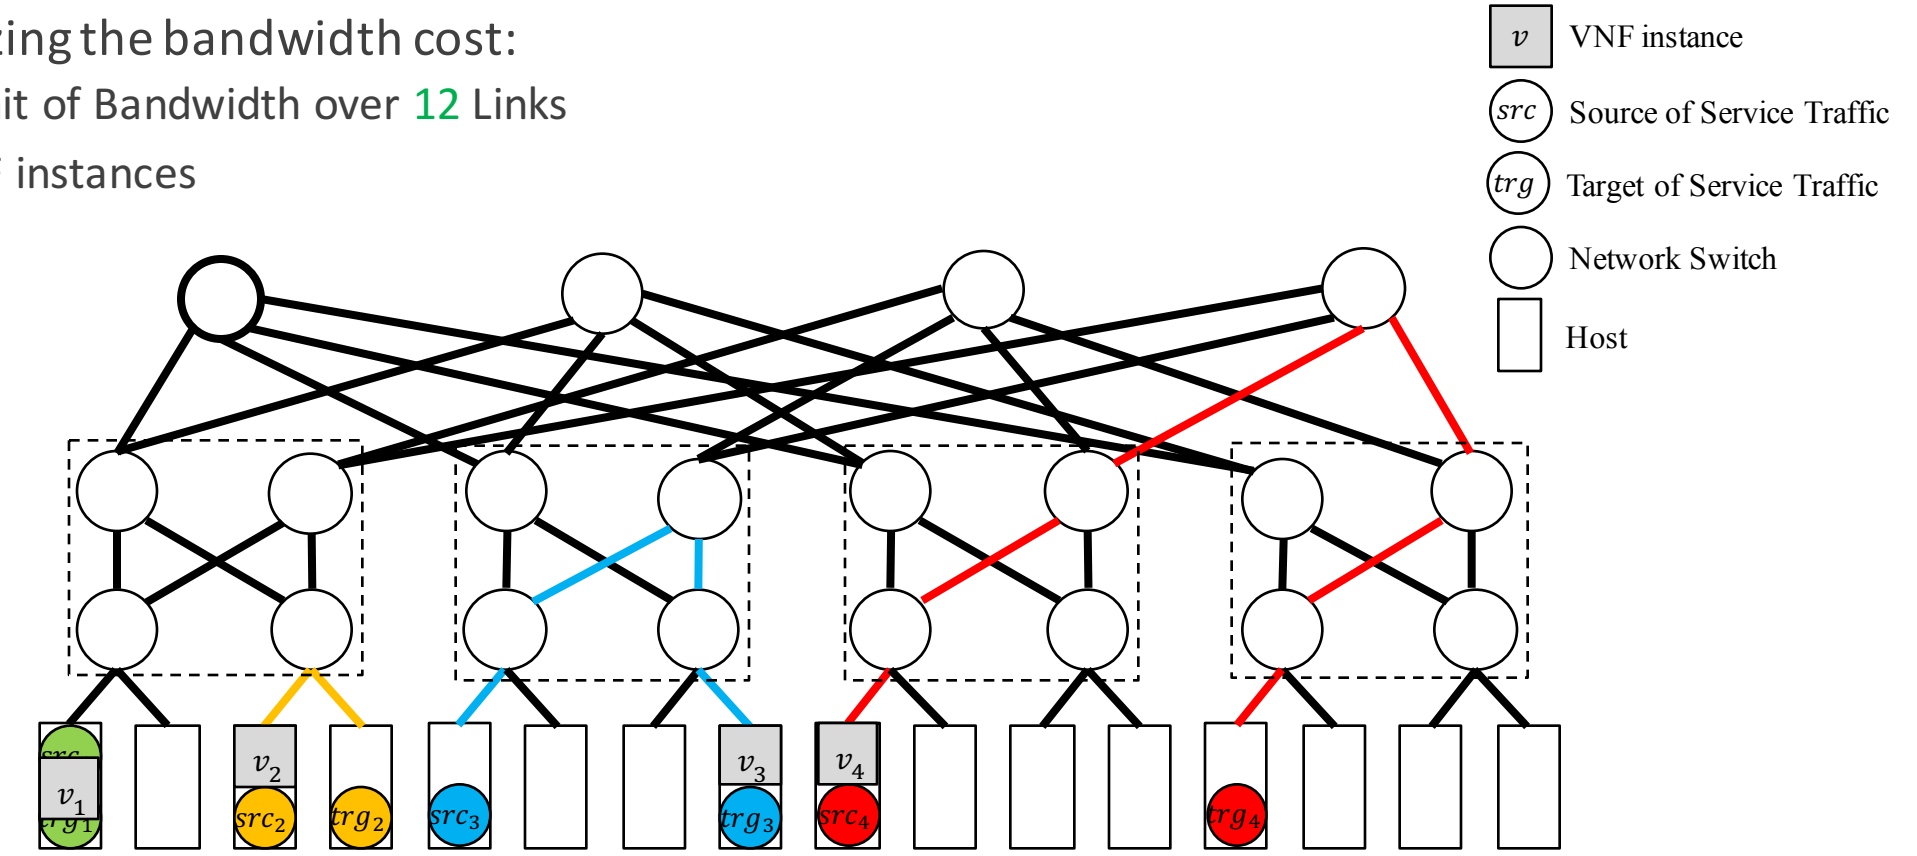

Conflicting Objectives

Minimizing the number of installed VNFs

- 1 VNF instance
- 34 Unit of Bandwidth over 20 Links

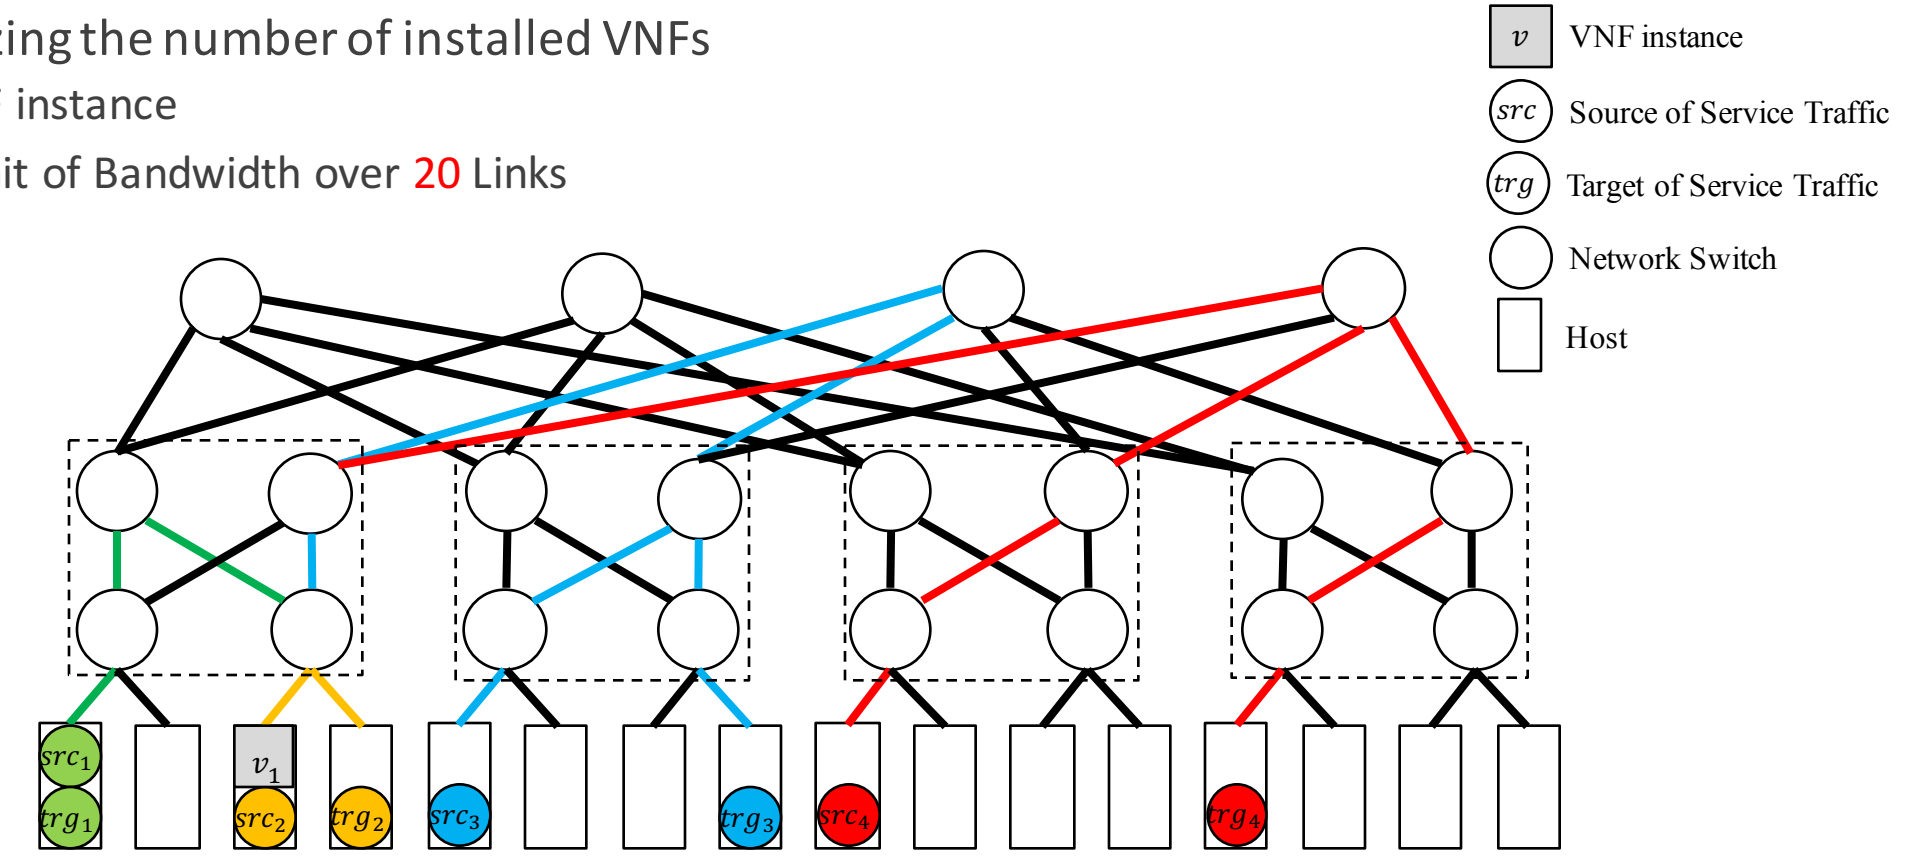

Elasticity Mechanisms and Overhead

MECHANISMS

Horizontal Scaling of VNF instance

- Installing a new VNF instance
- Removing an existing VNF instance

Migration of a VNF instance

Reassignment of workload to another VNF instance

OVERHEAD

Migration overhead

Reassignment overhead

Elasticity Mechanisms and Overhead

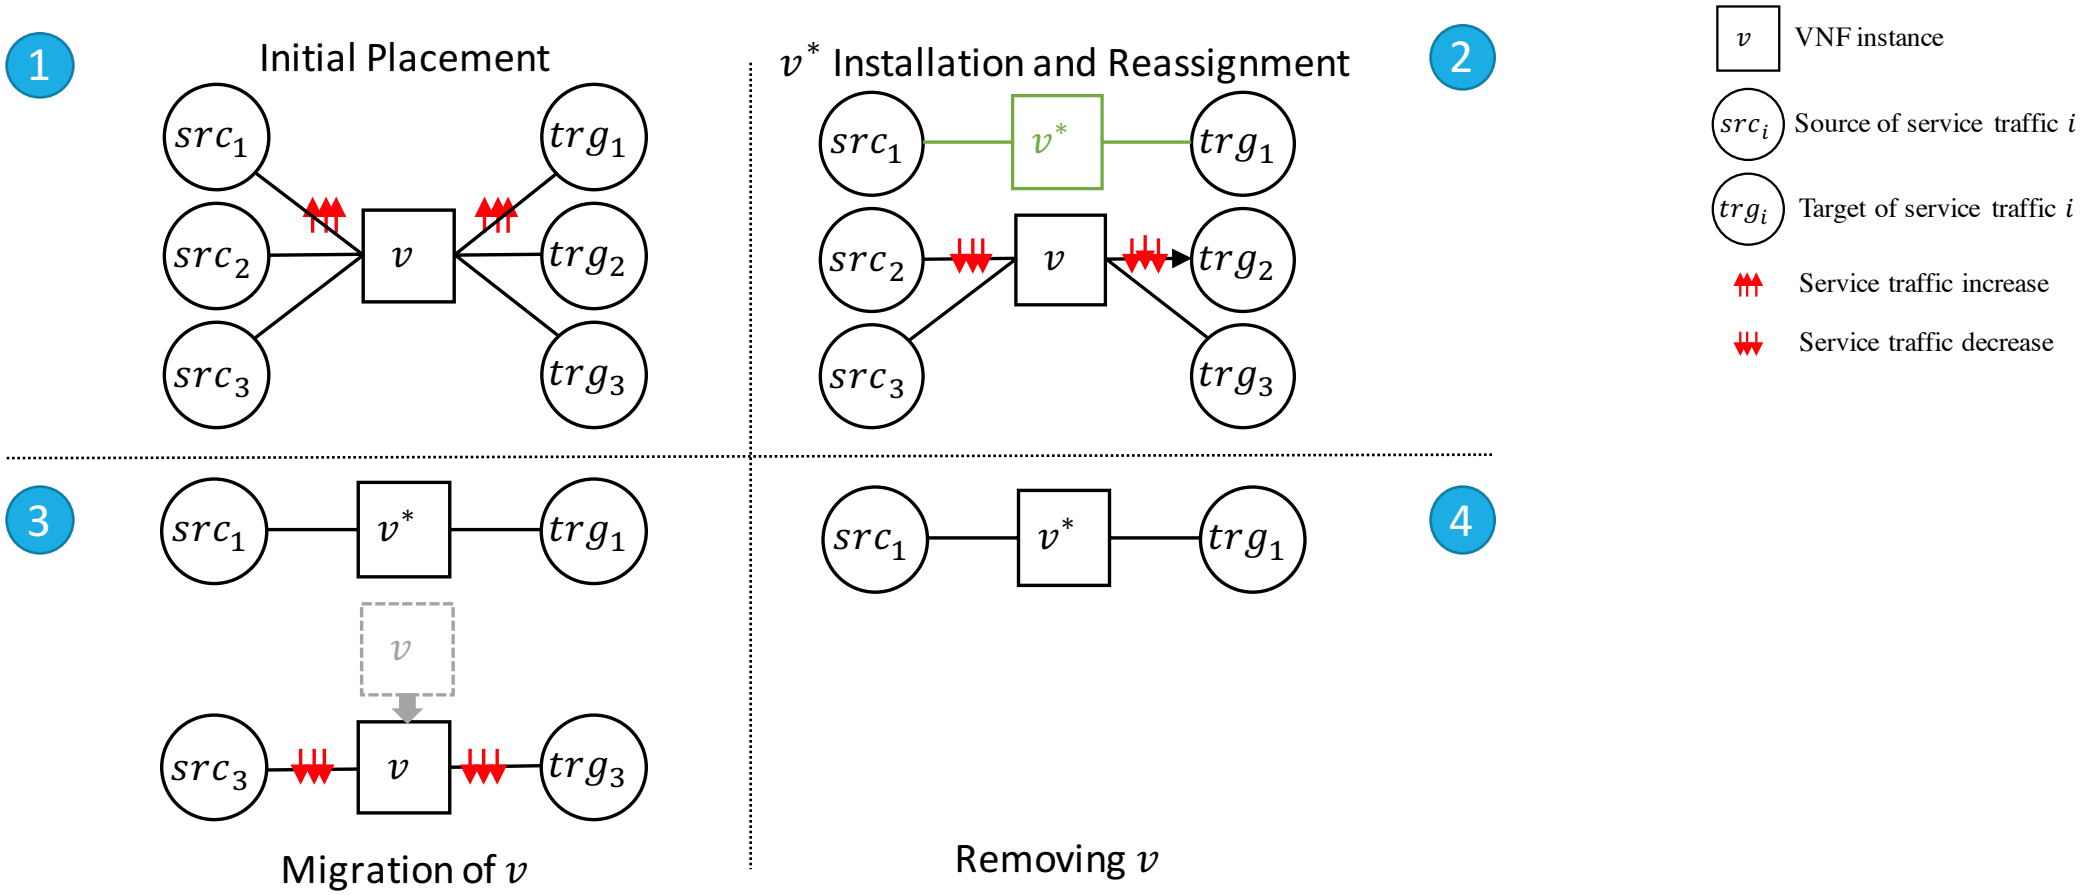

Solution: Simple Lazy Facility Location(SLFL)

IDEA

SLFL: SIMPLE LAZY FACILITY LOCATION

Idea

Arrival and departure of a request, or workload variation alter the locality

SLFL locally optimizes the placement of VNF instances in a greedy manner

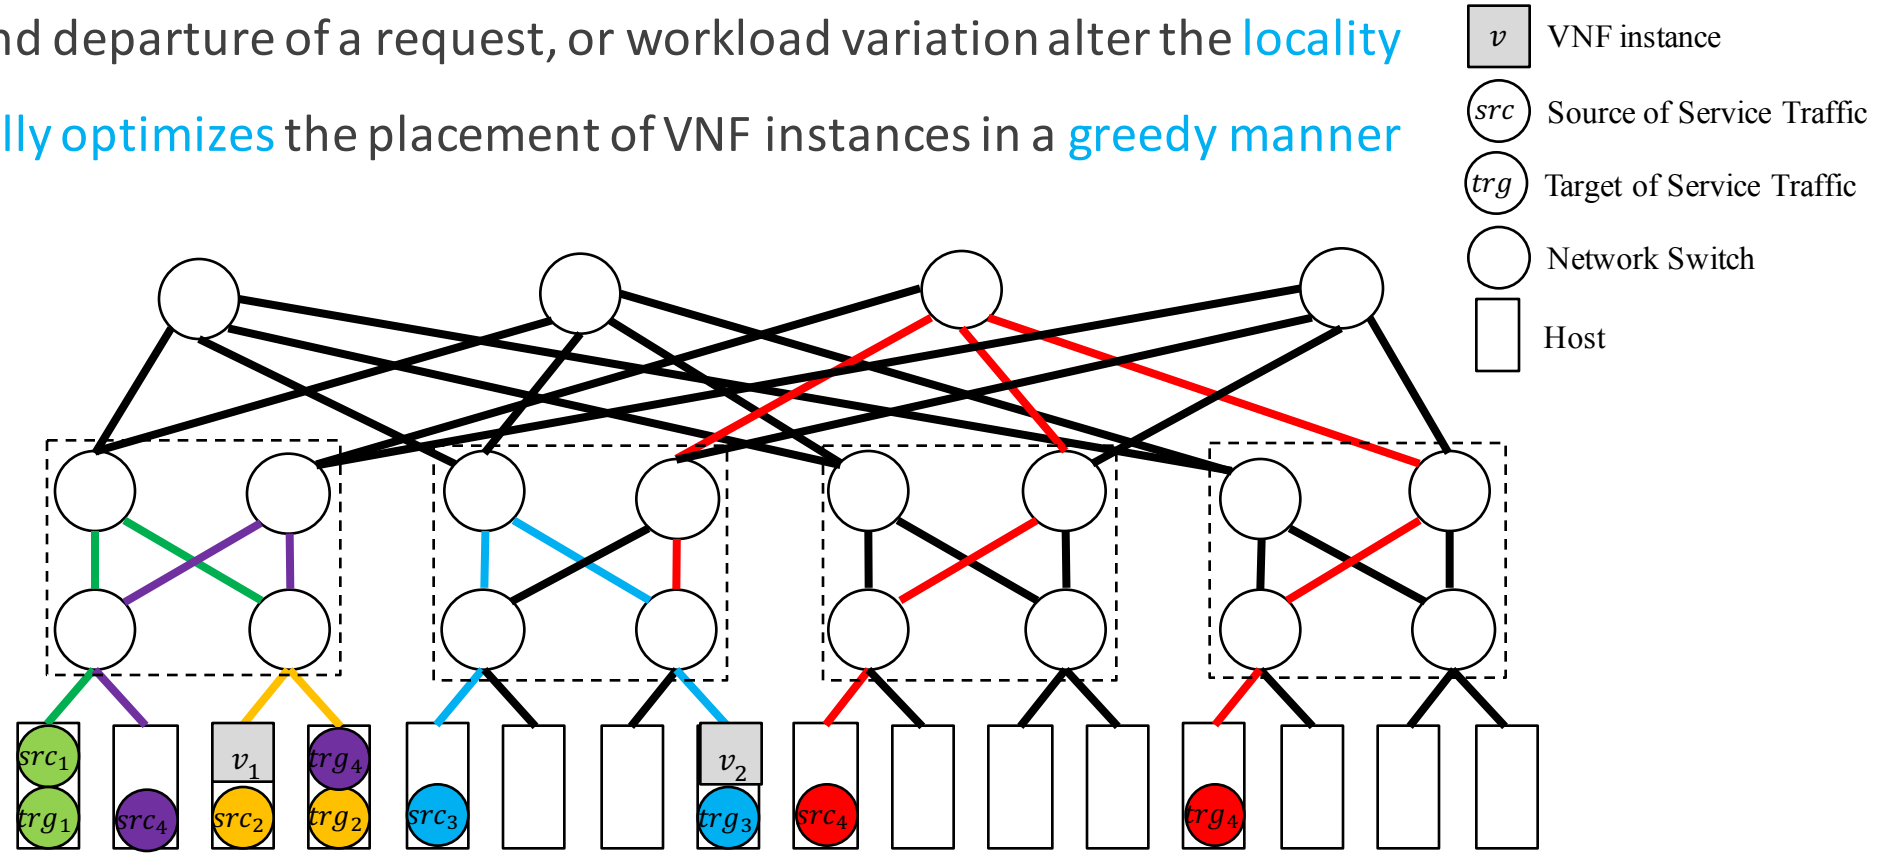

SLFL: Simple Lazy Facility Location

UPON REQUEST ARRIVAL OR WORKLOAD INCREASE

Installation potential

- Installing a VNF instance
- Set of reassignments
- *The difference of operational cost before and after installing the VNF instance and reassignments*

Migration potential

- Migration of a VNF instance
- *The difference of operational cost before and after migration of the VNF instance*

UPON REQUEST DEPARTURE OR WORKLOAD DECREASE

Removing potential

- Removing a VNF instance
- Set of reassignments
- *The difference of operational cost before and after removing the VNF instance and reassignments*

Emigration potential

- Migration of a VNF instance
- *The difference of operational cost before and after migration of the VNF instance*

SLFL: Simple Lazy Facility Location

UPON REQUEST ARRIVAL OR WORKLOAD INCREASE

Apply the best action among:

- Installing a VNF instance
  - Considering the installation potential
- Migrating a VNF instance
  - Considering the migration potential of the VNF instance
- Assign to one of existing VNFs
  - Considering bandwidth cost

UPON REQUEST DEPARTURE OR WORKLOAD DECREASE

Apply the best action among:

- Removing a VNF instance
  - Considering the installation potential
- Migrating a VNF instance
  - Considering the emigration potential of the VNF instance

Evaluation

EXPERIMENTAL SETUP AND OBJECTIVES

ACCEPTANCE RATIO AND OPERATIONAL COST

RESOURCE UTILIZATION

Experimental Setup and Objectives

EXPERIMENTAL SETUP

Network
- Fat-tree of 99 nodes
- 54 hosts with 8 Core CPU
- 1 GB full bisection bandwidth

VNF
- Bro IDS [2]: 80 Mbps, 1 vCPU, 1GB of memory

Requests
- 20,000 requests
- Arrival: Poisson distribution
- Duration: Exponential distribution

OBJECTIVES

Evaluating
- The acceptance ratio
- Operational cost
  - Balancing bandwidth and host resource costs
- Resource Utilization
  - Balancing bandwidth and host resource utilization ?

Comparison to
- Random Placement
- First-Fit Placement

Acceptance Ratio and Operational Cost

ACCEPTANCE RATIO

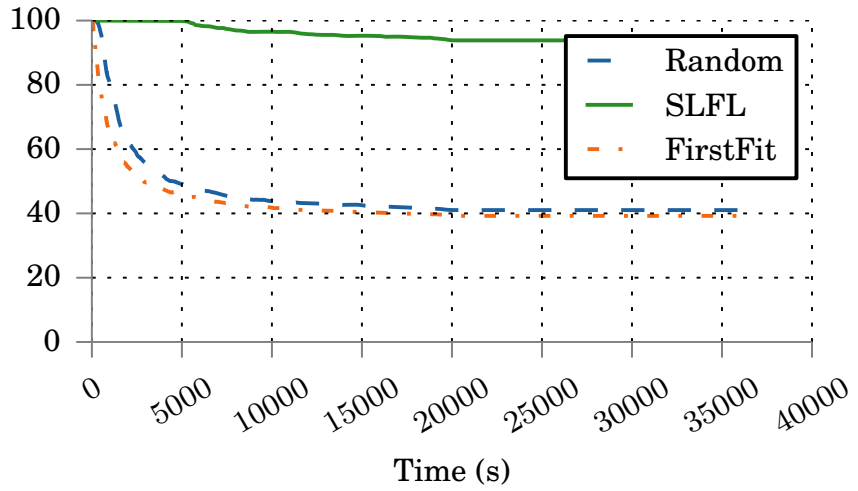

SLFL accepts ~2× workload vs basic algorithms

SLFL	97% acceptance ratio
Random	48% acceptance ratio
FirstFit	45% acceptance ratio

TOTAL OPERATIONAL COST

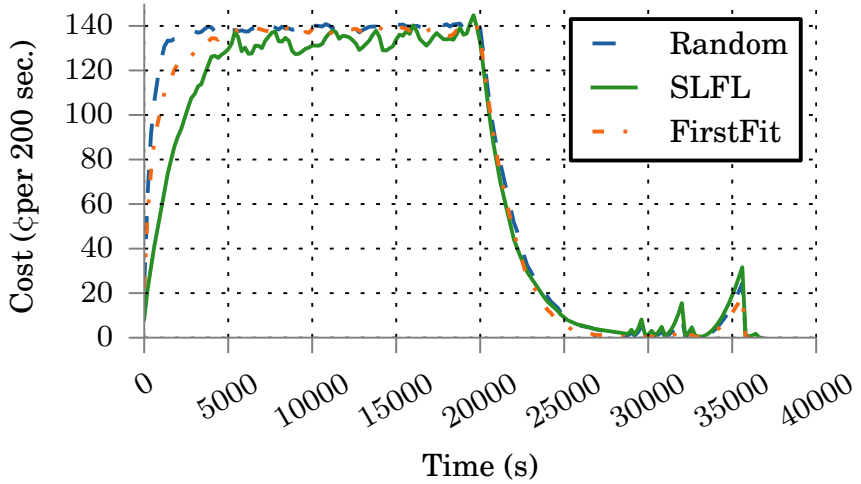

SLFL accepts ~2× workload with less cost

- 9% operational cost less than Random
- 4% operational cost less than FirstFit

Resource Utilization

BANDWIDTH RESOURCE UTILIZATION

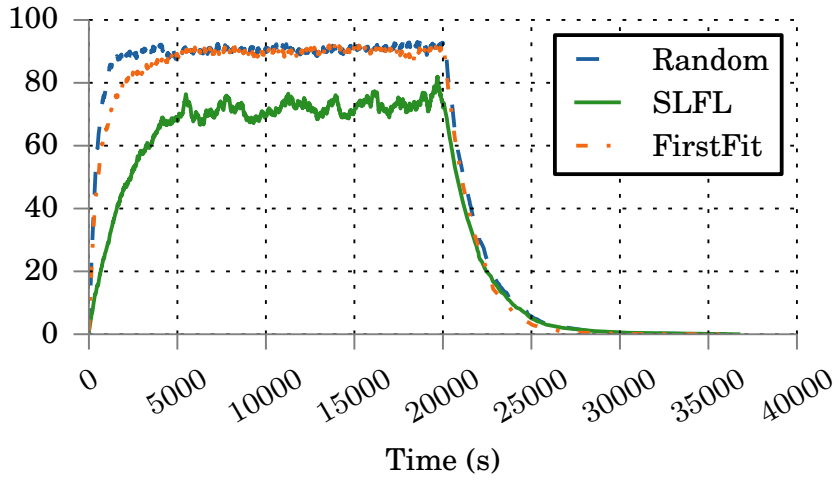

HOST RESOURCE UTILIZATION

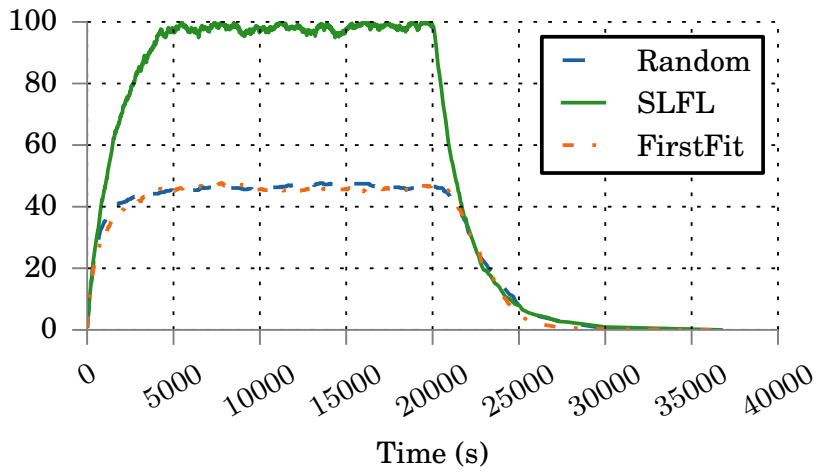

82% Utilization of bandwidth resources
91% Utilization of host resources

Conclusion

SUMMARY

Summary

Elastic Virtual Network Function Problem

- Bandwidth and host resources cost trade-off
- Elasticity Overhead

Simple Lazy Facility Location

- Balancing the bandwidth and host resource cost trade-off
- Carefully selecting the correct elasticity mechanisms
- Optimizing the elasticity overhead
- Accepting ~2× workload vs basic algorithms


QA

Acceptance Ratio and Resource Utilization

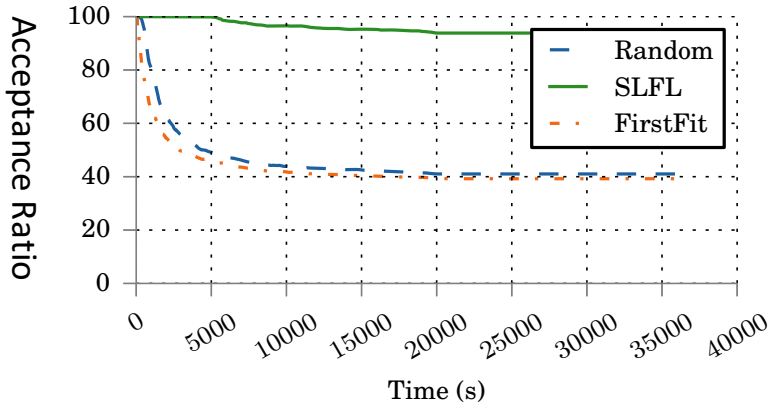

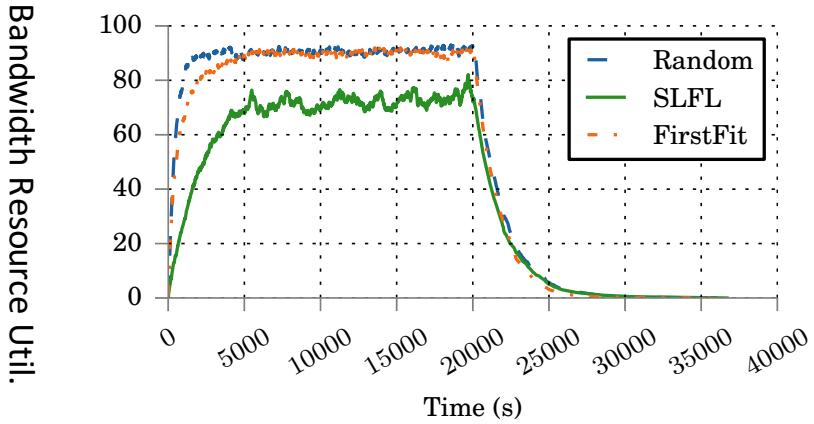

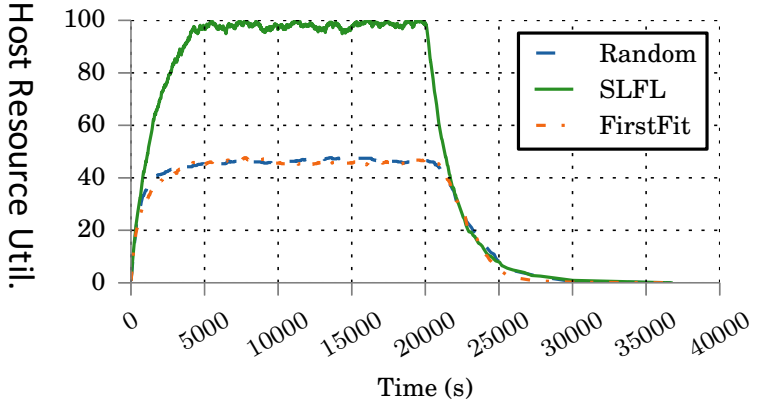

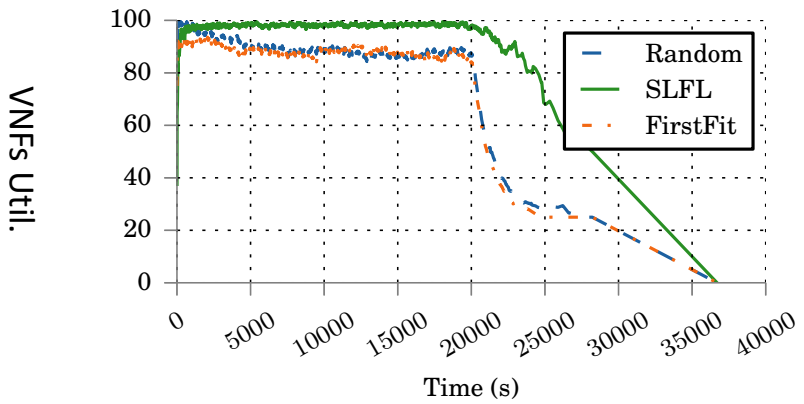

Operational Cost

Total Operational Cost

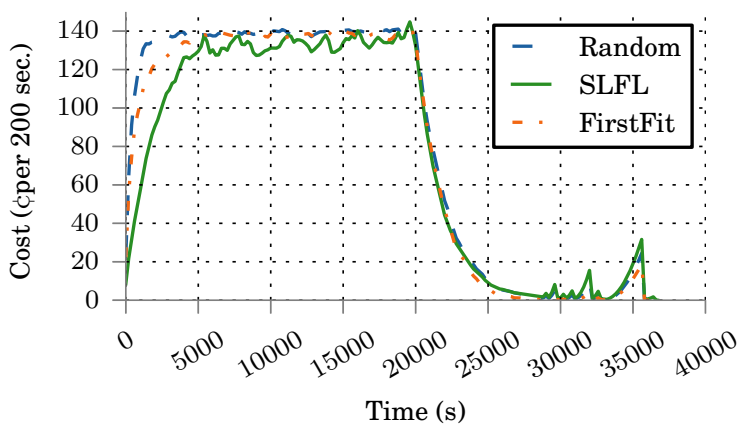

Bandwidth Resource Cost

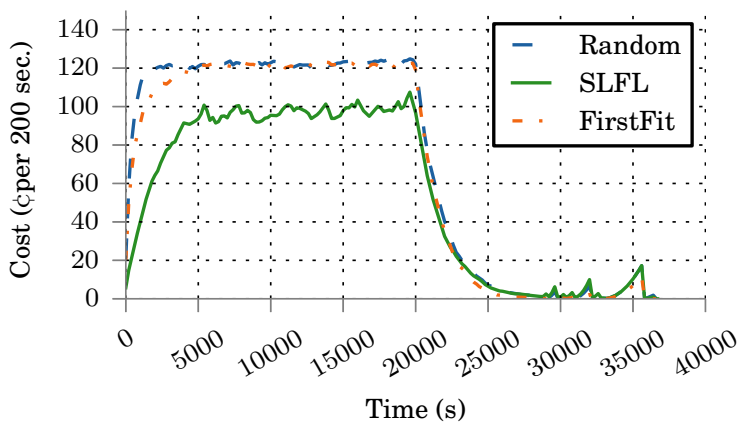

Host Resource Cost

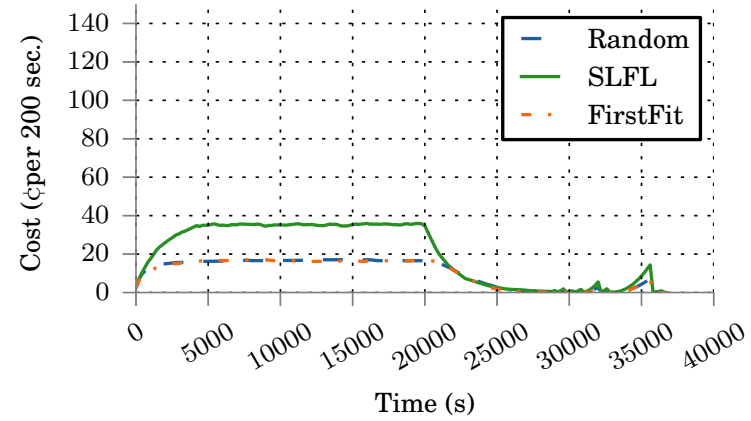

Elasticity Overhead Cost

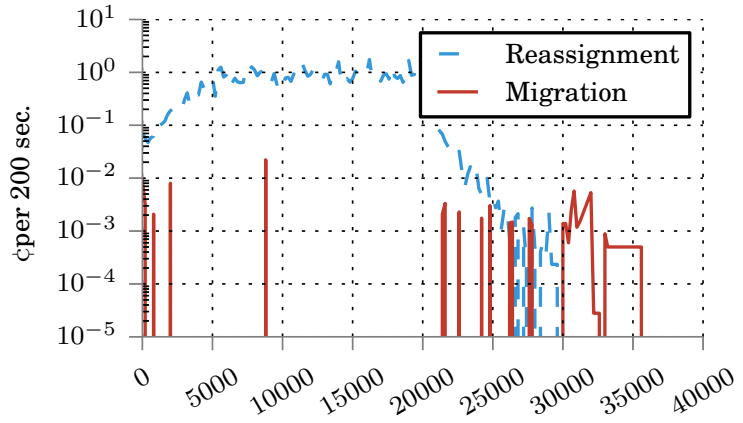

Assumptions-Horizontal Scaling

Why horizontal scaling and ignoring vertical scaling

- On the fly vertical resource scaling is not supported in most cases
- Might require system reboot
  - SLA violation

Assumptions-One VNF instance-type

		Scenario	One small flavor	Multiple flavors
Resource Consumption	Host Res.	~	- Worse	+ Better
	Bandwidth Res.	~	+ Better	- Worse
Elasticity	Installation	In a same machine	+ Better	- Worse
	Removal	In a same machine	+ Better	- Worse
	Migration overhead	~	+ Better	- Worse
	Reassign. overhead	~	= Equal	= Equal