Chair of Communication Networks
Department of Electrical and Computer Engineering
Technical University of Munich

*Flex*Nets: Evaluating Flexibility in Softwarized Communication Networks

Wolfgang Kellerer

Technical University of Munich

Zagreb, June 29, 2017

ConTEL 2017

IEEE

This work is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement No 647158 – **FlexNets (2015 – 2020)**.

European Research Council
Established by the European Commission

Introduction

- Networking today: new requirements from vertical industries, dynamically changing user behavior, and global digitalization

- Less (explicitly) addressed: ***flexibility*** and hence ***adaptation***

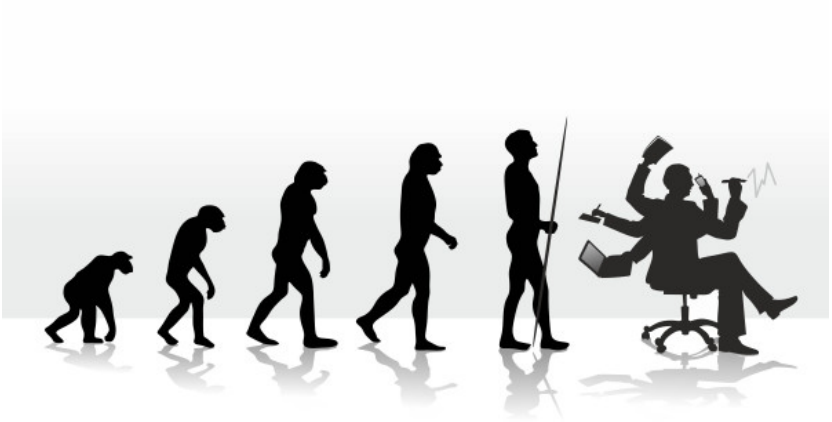

Image source: http://www.paleoplan.com

- In this talk, I will …
  … present our definition of a measure for network flexibility …
  … give concrete examples of how to apply …
  … raise more questions

FlexNets
2015 - 2020

erc
European Research Council
Established by the European Commission

Outline of this Talk

- On flexibility in softwarized networks
- Proposal for a *flexibility measure*
- Use Cases
  - The Function Placement Problem
  - Dynamic Controller Placement
  - HyperFlex: a flexible SDN Hypervisor solution

The Internet

… is able to adapt its resources

… *somehow* (best-effort, TCP elasticity, BGP, OSPF)

early-days simplicity

→ complex and ossified network system

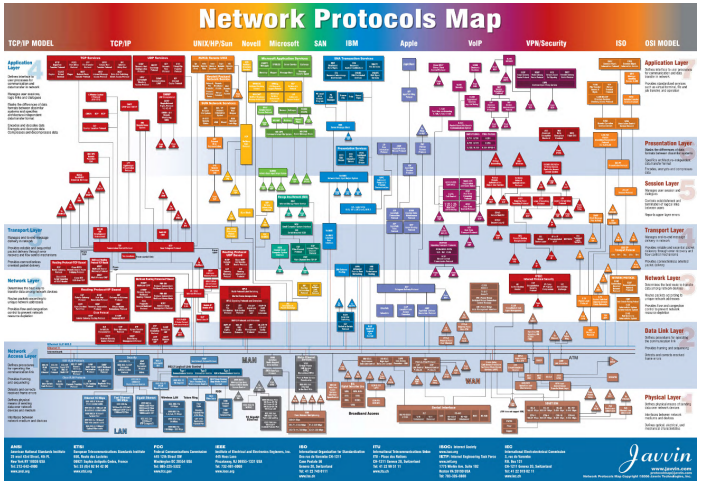

source: SFB MAKI

very slow adaptation to new requirements

→ reaction to dynamic changes hardly possible

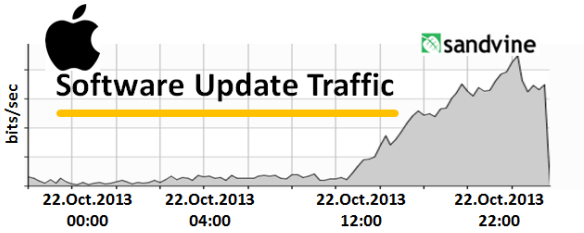

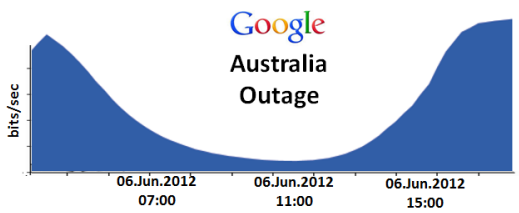

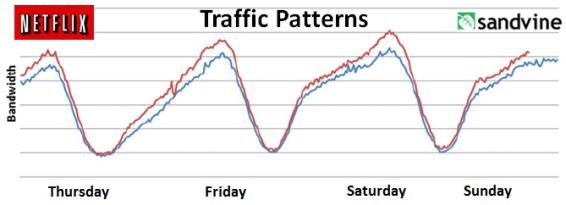

New concepts such as …

Network Virtualization (NV), Software Defined Networking (SDN) and Network Function Virtualization (NFV)

…*promise* to create and adapt networks and functions on demand in *software*

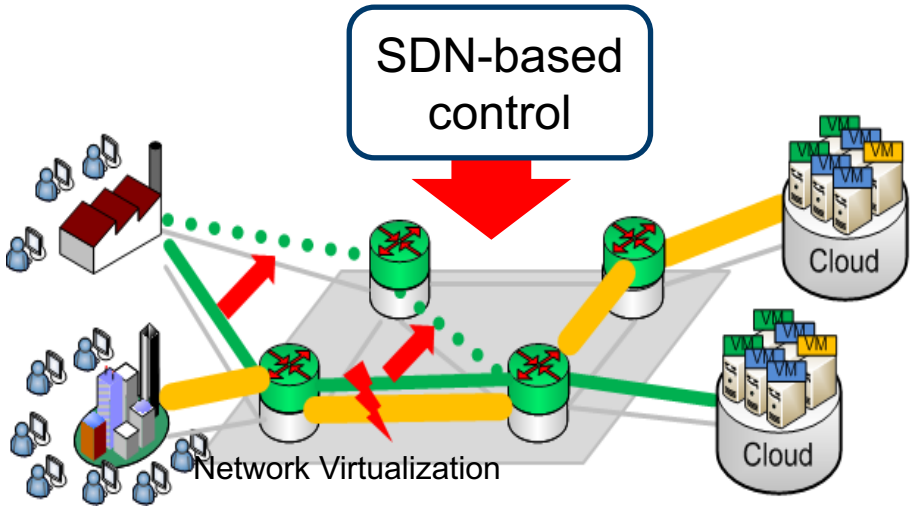

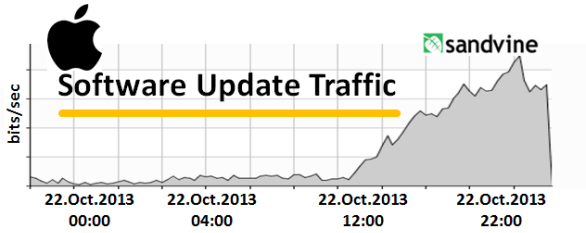

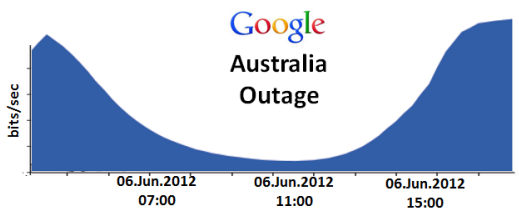

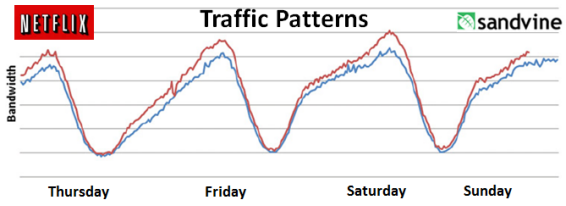

New concepts such as …

Network Virtualization (NV), Software Defined Networking (SDN) and Network Function Virtualization (NFV)

…*promise* to create and adapt networks and functions on demand in *software*

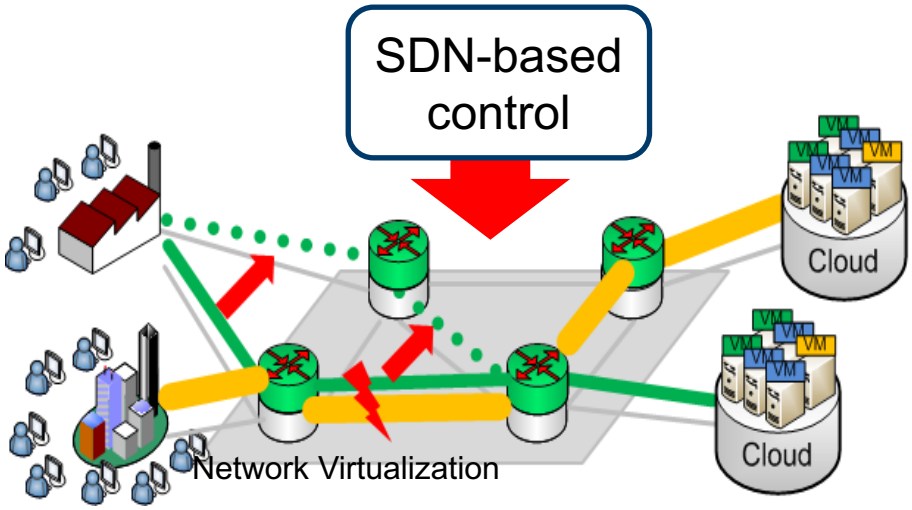

→ ***Softwarized Networks***

All problems solved?

- Are we fully flexible already?
- How far can we go? What is the right network design?

We need
- a **fundamental understanding** of how to provide flexibility
- a **quantitative measure** for flexibility pro and contra certain designs

For networks, **flexibility** = ability to *support new requests* to change design requirements (traffic pattern, latencies,…)

This work is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 program grant agreement No 647158 – **FlexNets (2015 – 2020)**.

FlexNets
2015 - 2020

European Research Council

A simple measure

For networks, **flexibility** = ability to *support new requests* to change design requirements (traffic pattern, latencies,…)
via adaptation of resources (topology, capacity, ...) if needed

$$\varphi(S) = \frac{|supported\ new\ requests|}{|total\ number\ of\ given\ new\ requests|}$$

- fraction of the number of **new requests** that can be supported of all given requests

- $\varphi(S) \in [0,1]$ „percentage“

What is missing?
The time aspect of flexibility

"Heatposter" by Source. Licensed under Fair use via Wikipedia – http://en.wikipedia.org/wiki/File:Heatposter.jpg#/media/File:Heatposter.jpg

What Robert de Niro says on ***flexibility***

in HEAT (1995) as Neil McCauley:
"Don't get attached to anything you can't walk out on in 30 seconds flat if you feel the heat around the corner."

Not only the number of options, but the time matters for ***flexibility***!

Flexibility Measure – proposed definition

$$\varphi_T \quad (S \mid \text{state } i) = \frac{|supported\ new\ requests\ within\ T|}{|total\ number\ of\ given\ new\ requests|}$$

- fraction of the number of new requests that can be supported in a time interval T of all given new requests

$$\varphi^{aspect}_{T->\infty} \quad (S) = \frac{|supported\ new\ requests|}{|all\ given\ new\ requests|}$$

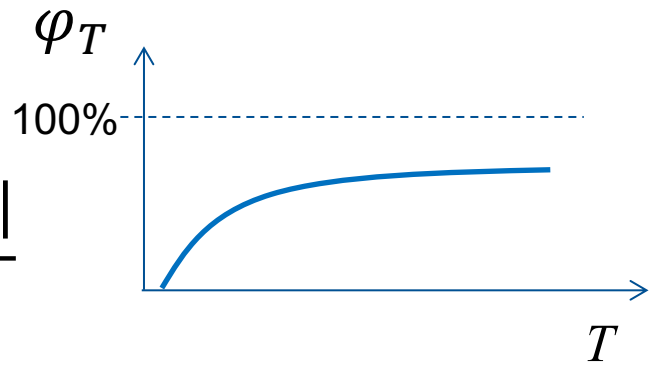

A simple illustration (1)

- New request to an SDN-network: Controller Capacity (cc) is increased
- Can such new request be supported?
  e.g. by migrating the controller to a node with higher capacity (NC)
- BUT: migration time cannot exceed "1 hop" (T)

max. migration time T = **1 hop**

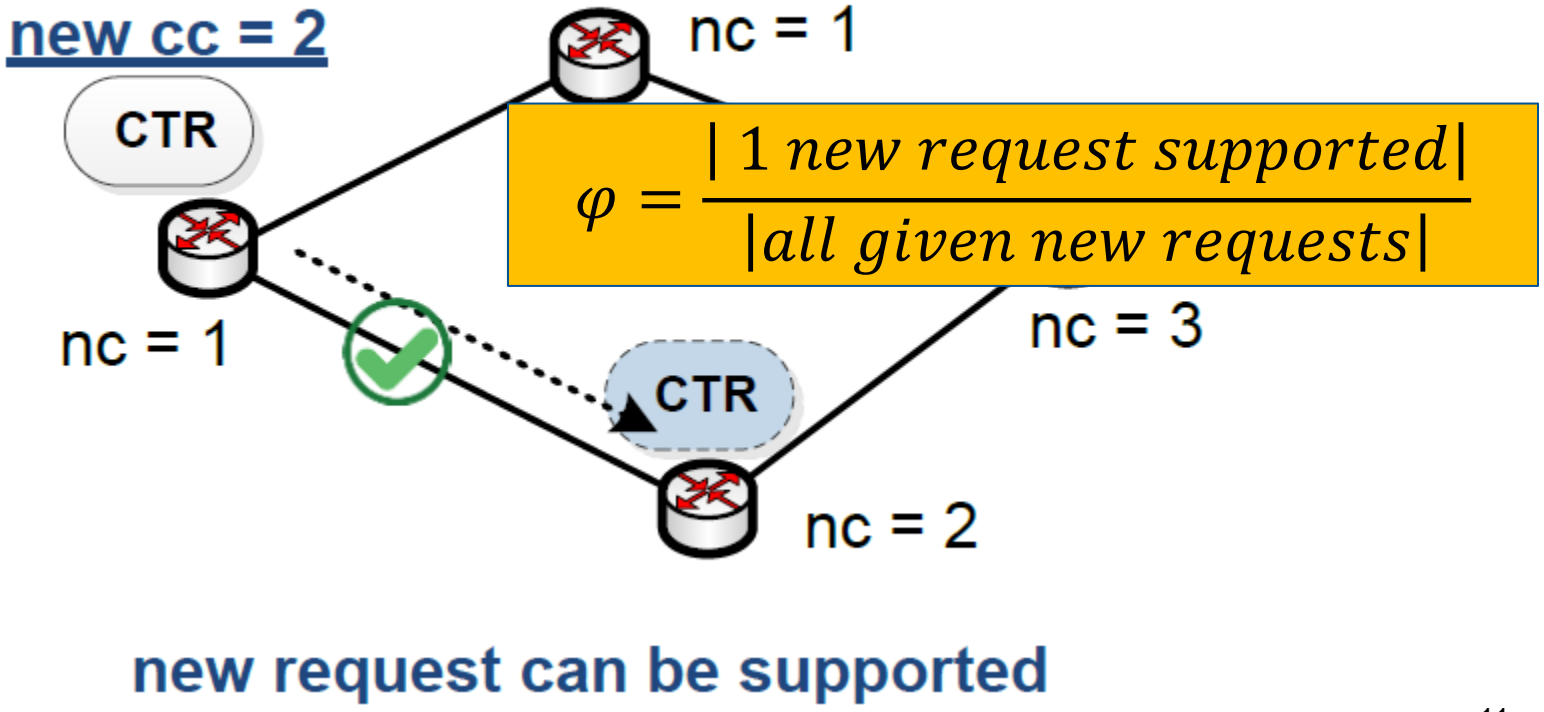

new request can be supported

A simple illustration (2): more requests

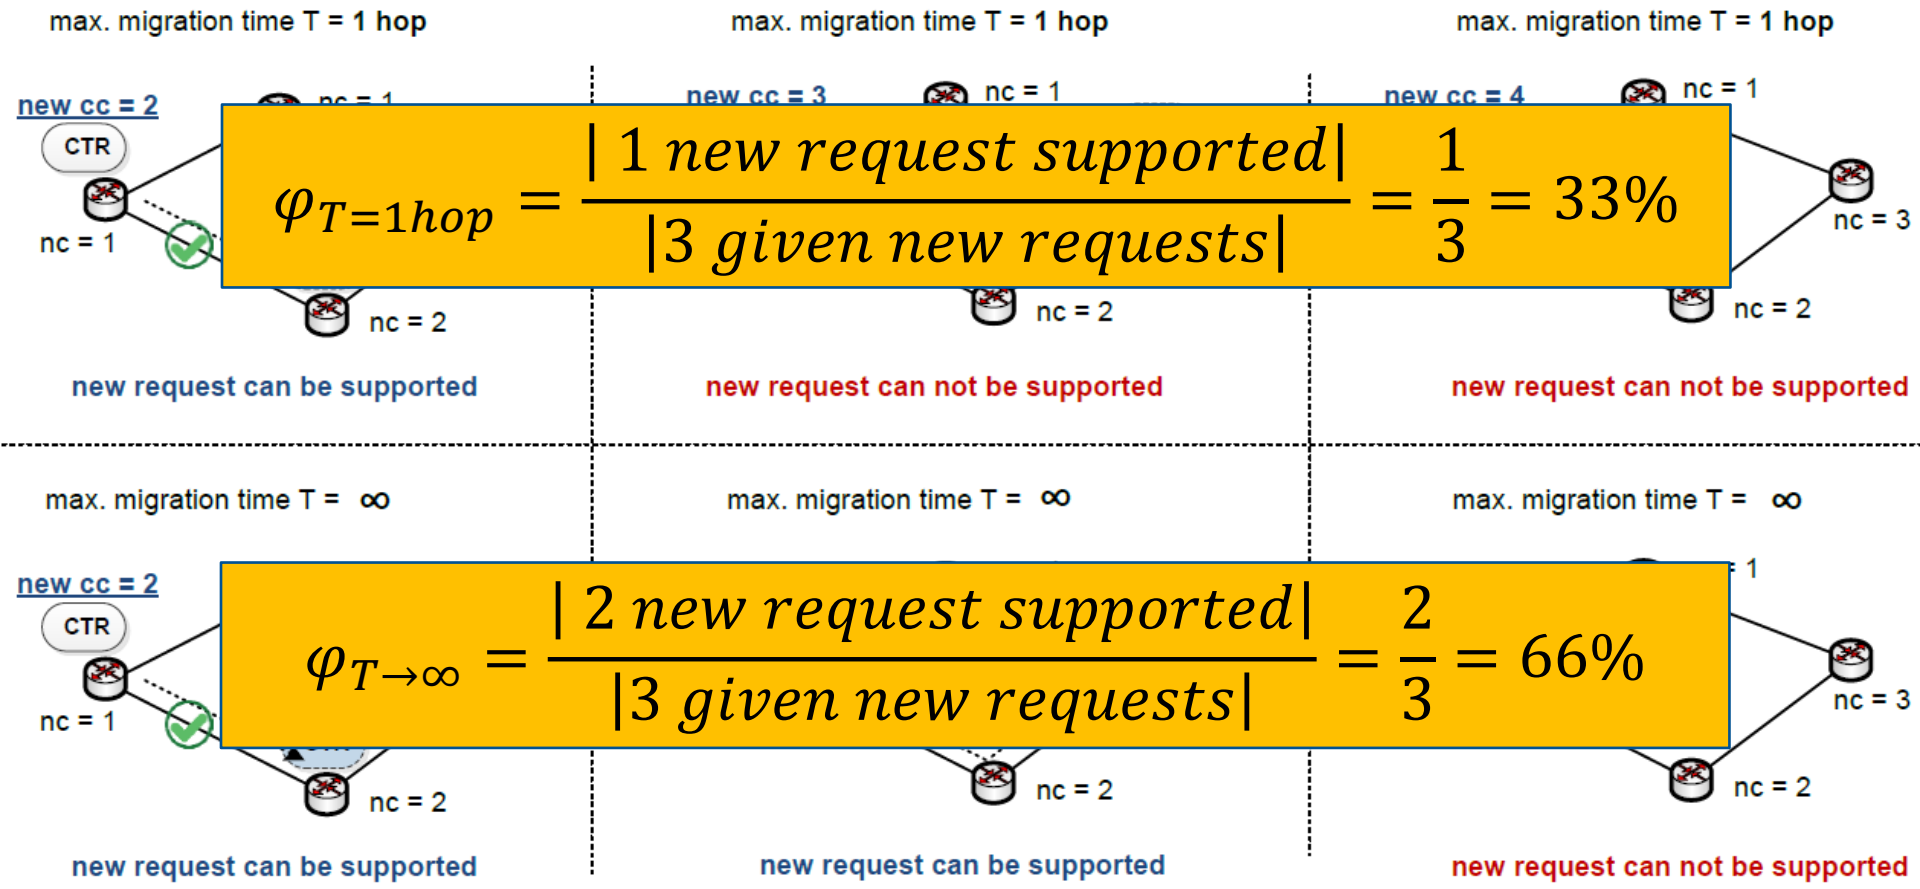

$$\varphi_{T=1hop} = \frac{|\,1\ new\ request\ supported|}{|3\ given\ new\ requests|} = \frac{1}{3} = 33\%$$

$$\varphi_{T\to\infty} = \frac{|\,2\ new\ request\ supported|}{|3\ given\ new\ requests|} = \frac{2}{3} = 66\%$$

Flexibility a new measure? - Yes

no single quality indicator for a ***Quality of Flexibilty (QoF)***

- similar to QoS
- to be regarded by case (requirements, design goals, system)

we propose: ***flexibility aspects*** [1, 2]

- similar as we do with QoS (rate, delay, throughput, jitter,…)
- shall allow us to quantitatively compare two different system designs

[1] W. Kellerer, A. Basta, A. Blenk, Using a Flexibility Measure for Network Design Space Analysis of SDN and NFV, SWFAN'16, IEEE INFOCOM Workshop, April 2016.
[2] W. Kellerer, A. Basta, A. Blenk, Flexibility of Networks: a new measure for network design space analysis?. arXiv preprint arXiv:1512.03770, 2015.

Flexibility Aspect example 1: Flow steering and reconfiguration

Parameters (for change requests):

- number of flows,
- granularity (forwarding, duplicating,…),
- time to change

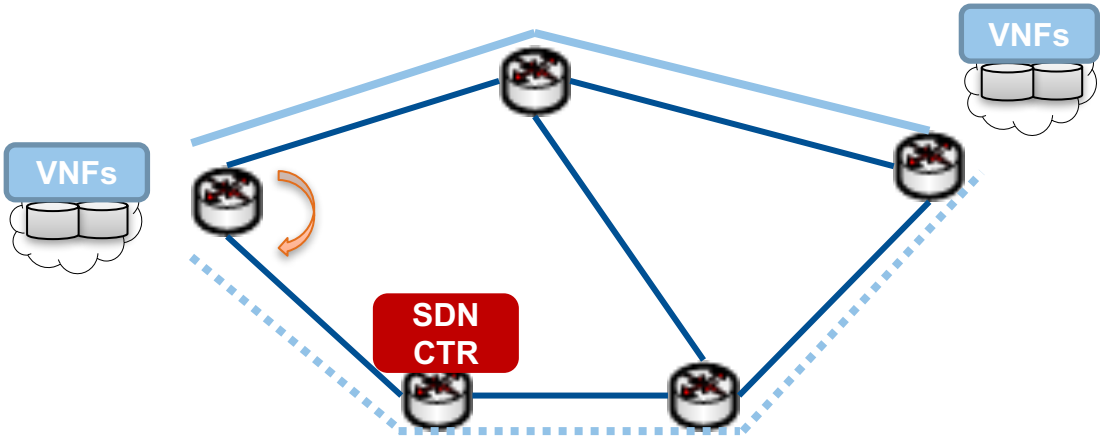

Flexibility Aspect example 2: Function Placement

Parameters:

- set of possible locations,
- number of supported requirements (latency, …),
- time of placement (static, dynamic)

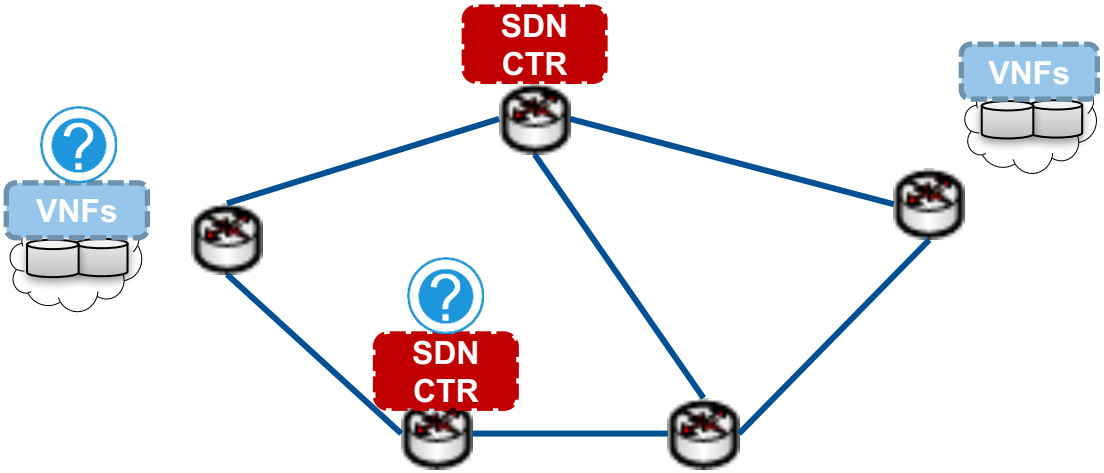

Use Case 1: *The Function Placement Problem*

- NFV = virtualize & move function (= everything) to DC

Example: mobile core network functions

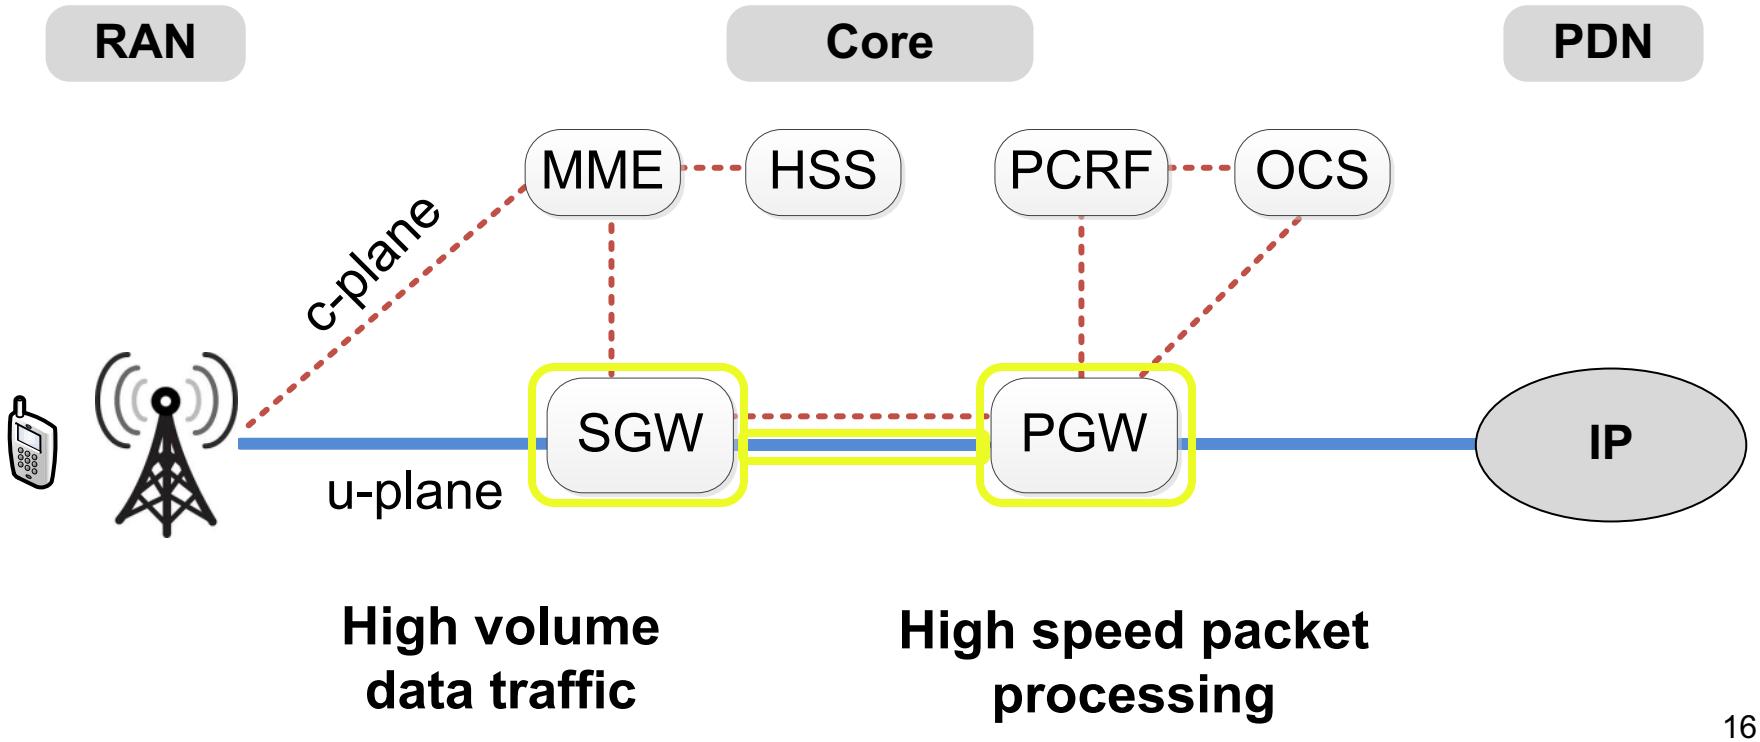

Function Realization based on NFV

- Virtualization of GW functions [1] → NFV

data-plane latency?

depends on the DC placement

network load?

traffic transported to DC (longer path → cost)

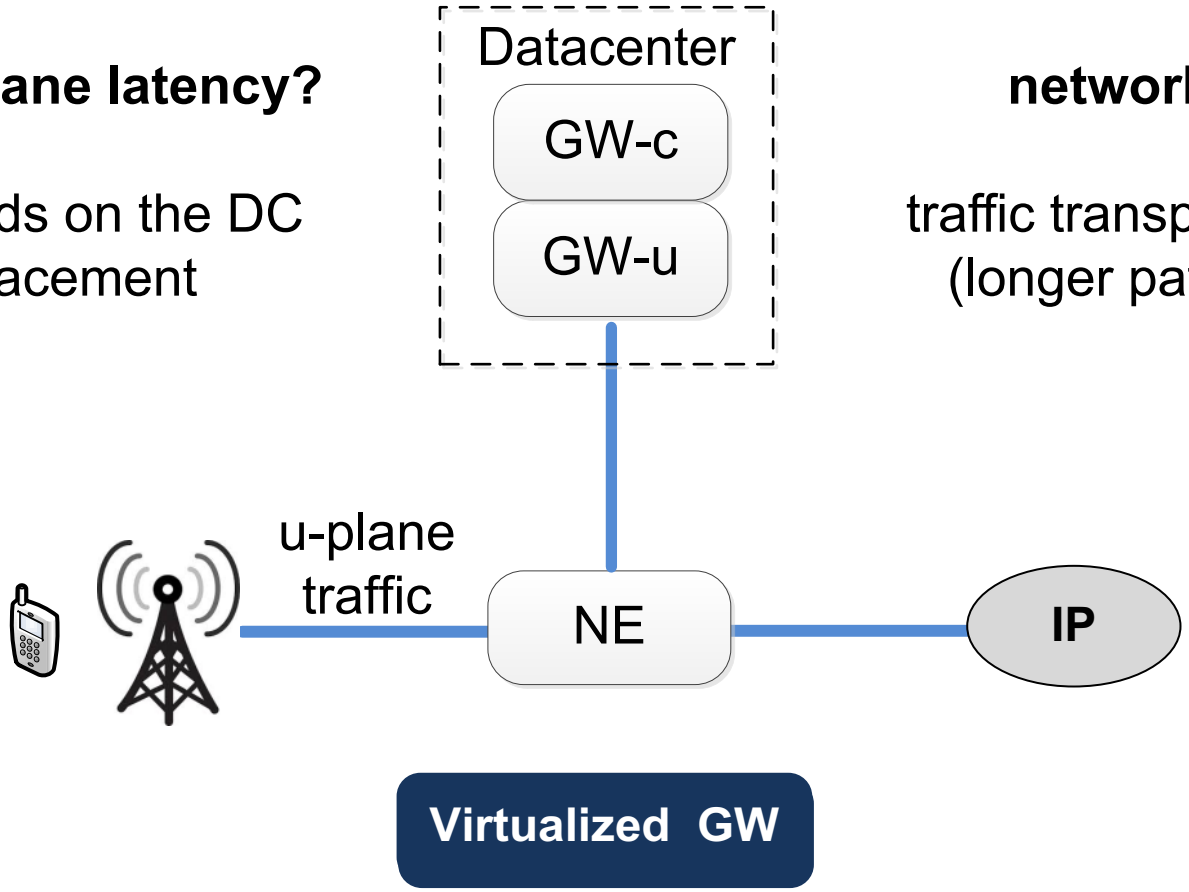

[3] A. Basta et al., A Virtual SDN-enabled EPC Architecture : a case study for S-/P-Gateways functions, SDN4FNS 2013.

Function Realization based on SDN: *move functions back*

- Decomposition of GW functions [1] via SDN

data-plane latency?

additional latency
is avoided

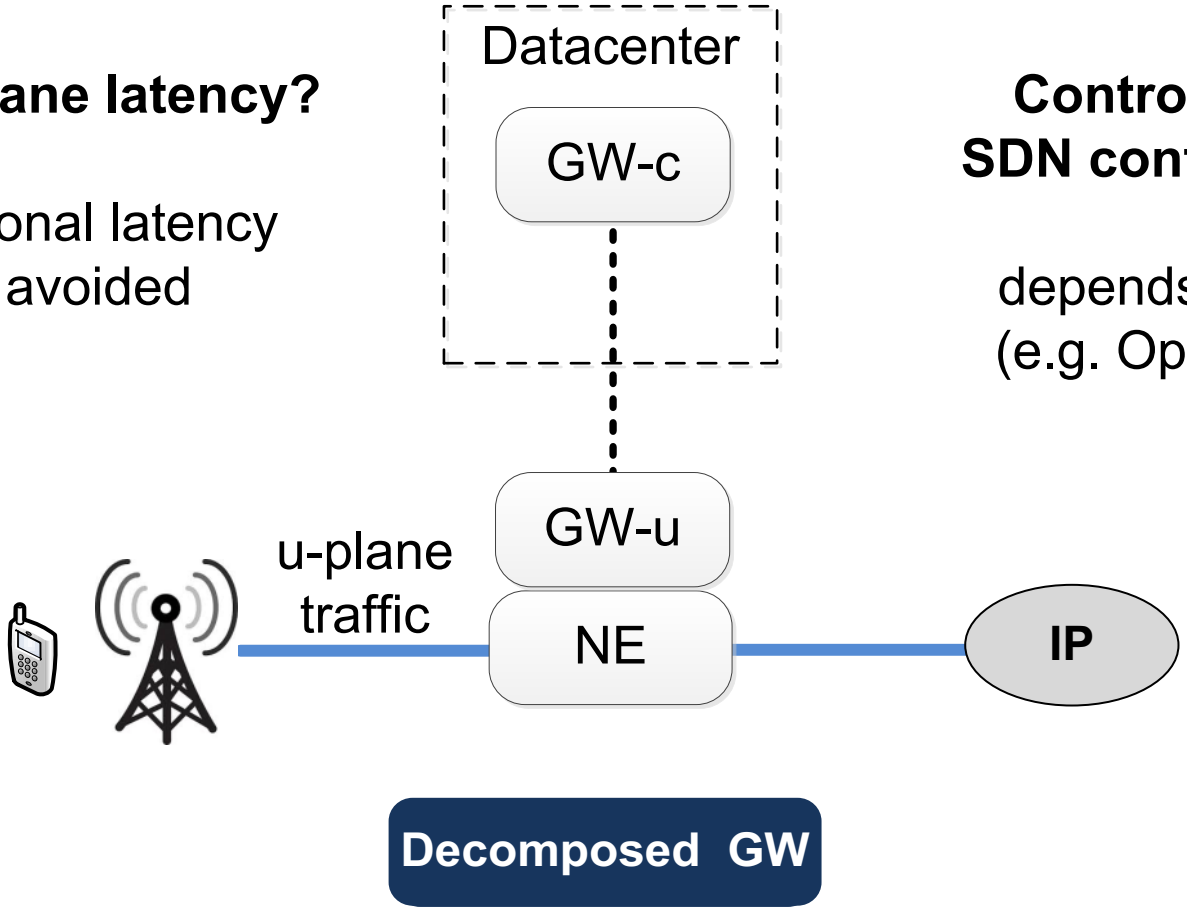

Control load?
SDN control load!

depends on API
(e.g. OpenFlow)

[3] A. Basta et al., A Virtual SDN-enabled EPC Architecture : a case study for S-/P-Gateways functions, SDN4FNS 2013.

Interdependencies → Function chains (mixed design)

- Propagation latency depends on function chain = path SGW - PGW

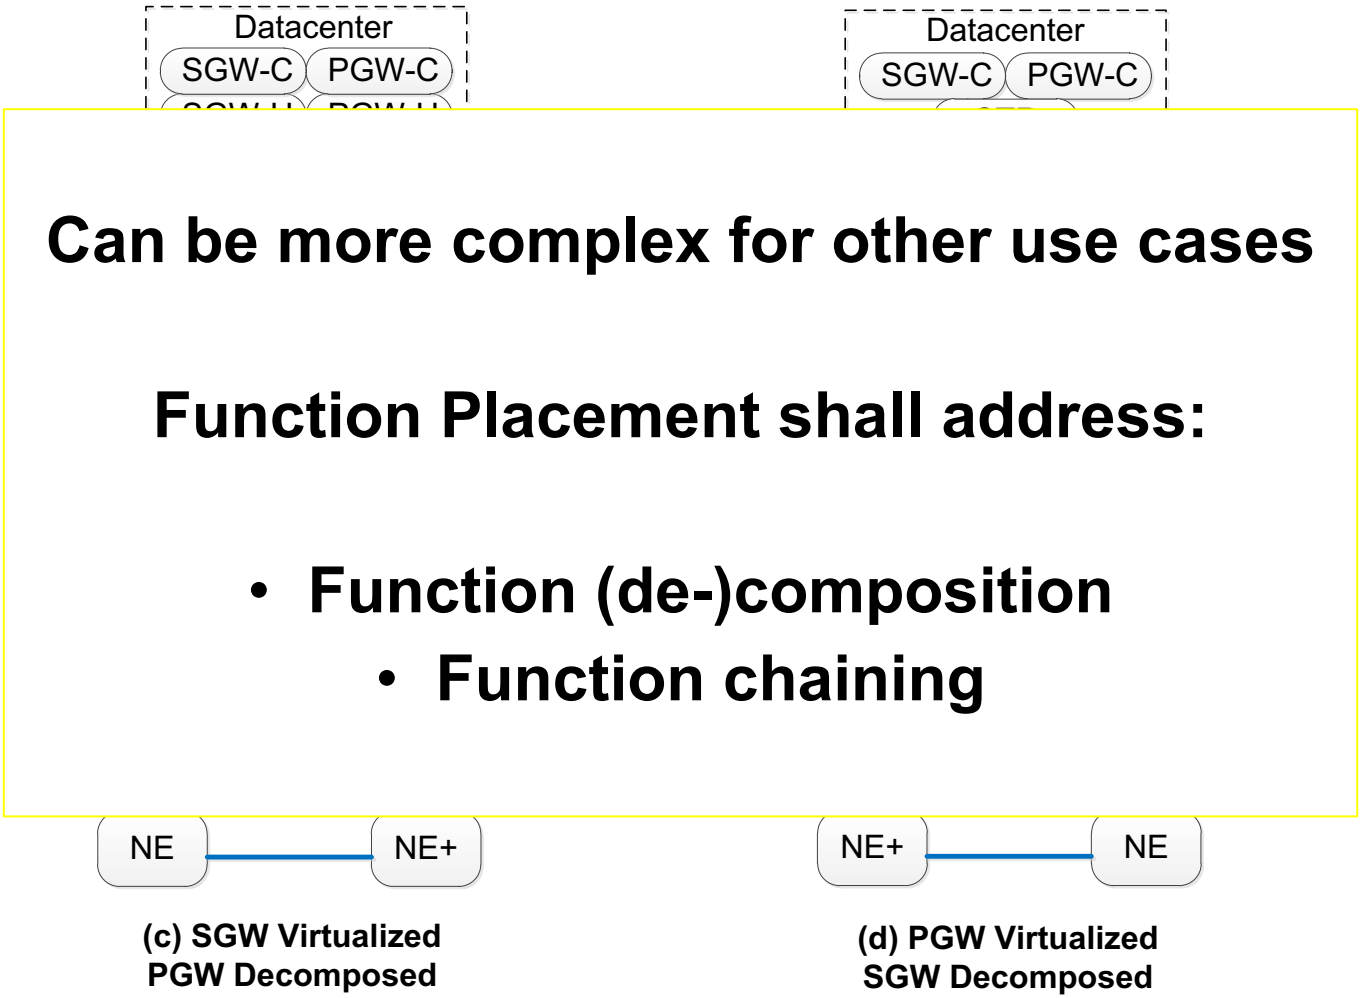

Some Evaluation Studies

- Virtualize all GWs? decompose all? mixed deployment?
  - Which GWs should be virtualized? decomposed? DC(s) placement?
    - minimize core load $
    - satisfy data-plane latency ✓

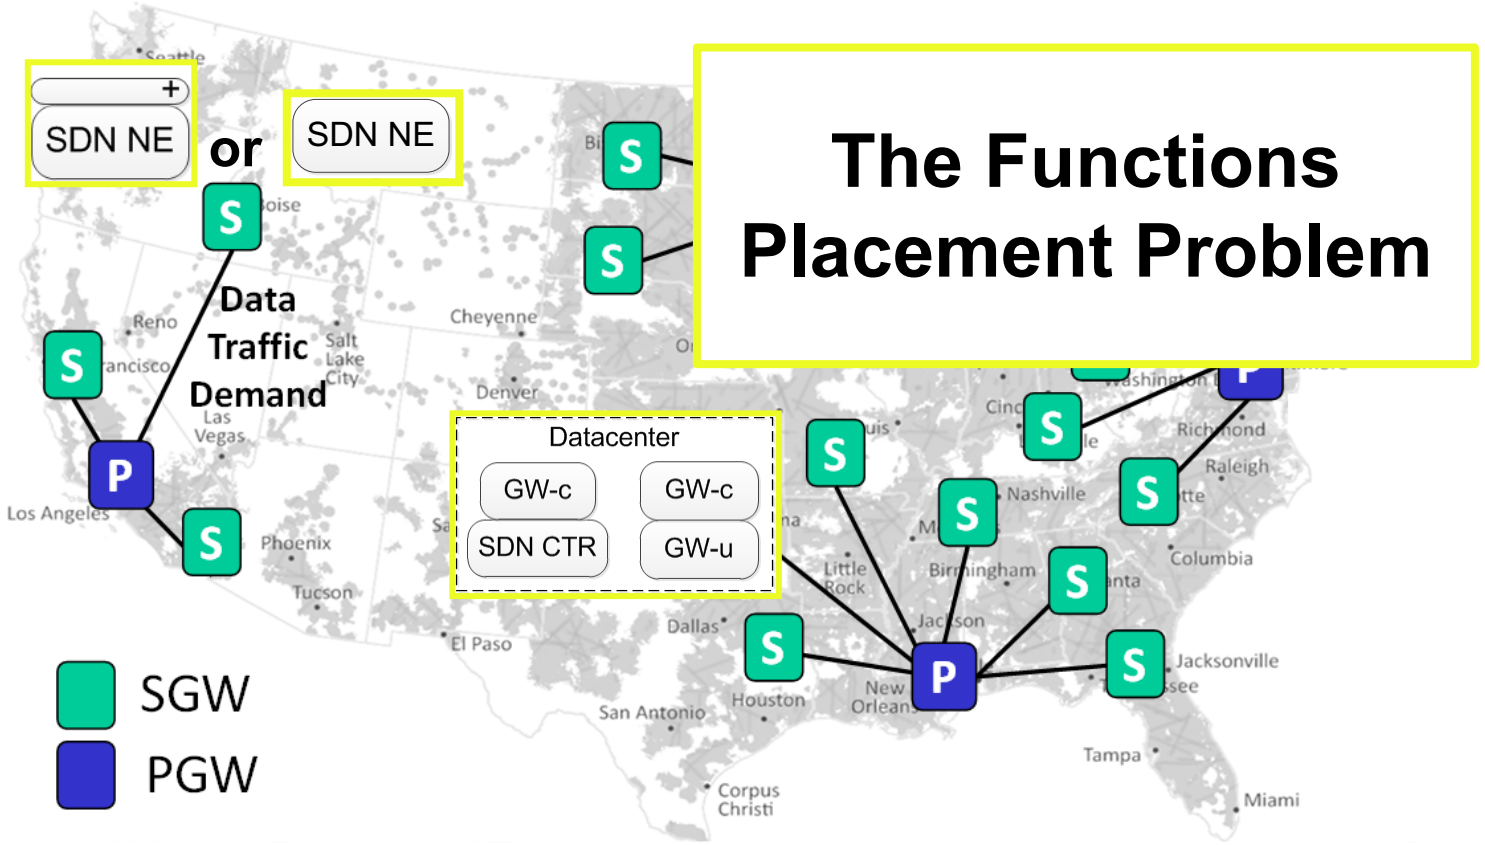

The Functions Placement Problem

[4] A. Basta, W. Kellerer, M. Hoffmann, H. Morper, K. Hoffmann, Applying NFV and SDN to LTE Mobile Core Gateways; The Functions Placement Problem, AllThingsCellular14, Workshop ACM SICGOMM, Chicago, IL, USA, August 2014

Flexibility Analysis of Function Placement

Use Case 1

3 design choices (= systems) to compare [5]:

(1) SDN design

(2) NFV design

(3) mixed SDN/NFV design

Parameter in focus:

- Flexibility to support different latency requirements for
  - control plane latency and data plane latency

    e.g.: {5, 10, 15,…, 45, 50} ms

all requests:
10 x10 =100

[1] W. Kellerer, A. Basta, A. Blenk,
Using a Flexibility Measure for Network Design Space Analysis of SDN and NFV, SWFAN'16,
IEEE INFOCOM Workshop, April 2016.

Design Choices

Legacy LTE core design:
Gateways (GW) as dedicated middleboxes

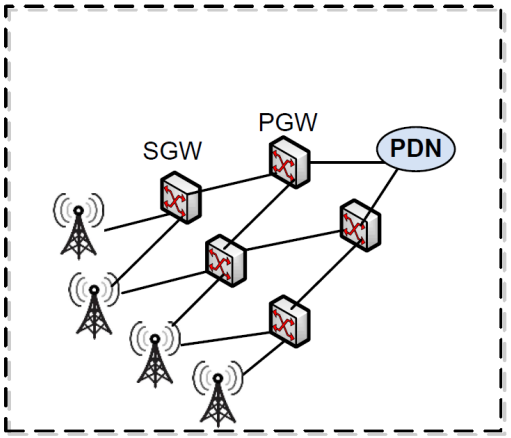

(a) Current LTE Core GW Architecture

(1) SDN design:
separation of control and data plane for GWs

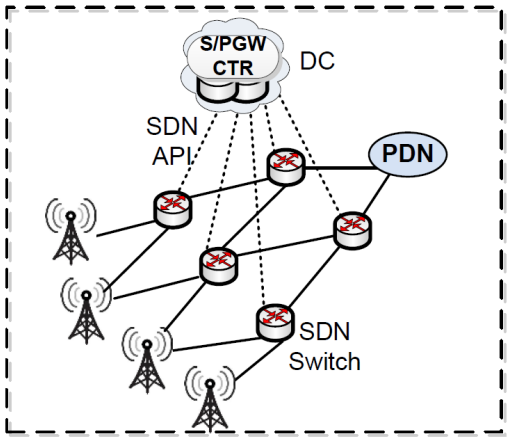

(b) SDN Core GW Architecture

only control to cloud

(2) NFV design:
all functions (data and control) run in a cloud

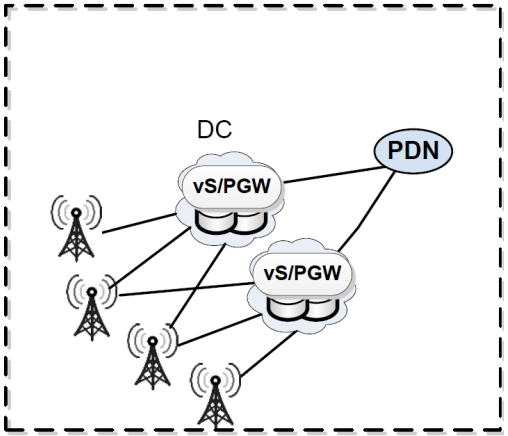

(c) NFV Core GW Architecture

control and data to cloud

(3)
mixed SDN/NFV design:

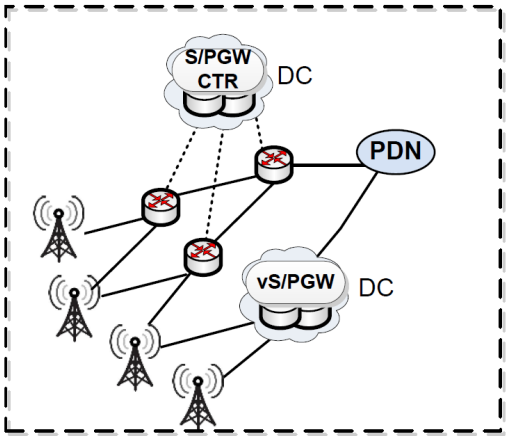

d) Mixed SDN and NFV Core GW Architecture

Flexibility measure and evaluation setup

Flexibility measure:

$$\varphi^{placement}\ (design.x) = \frac{\left(\sum_i \sum_j feasibleSol_{i,j} \cdot w_{i,j}\right)}{\sum_i \sum_j w_{i,j}}$$

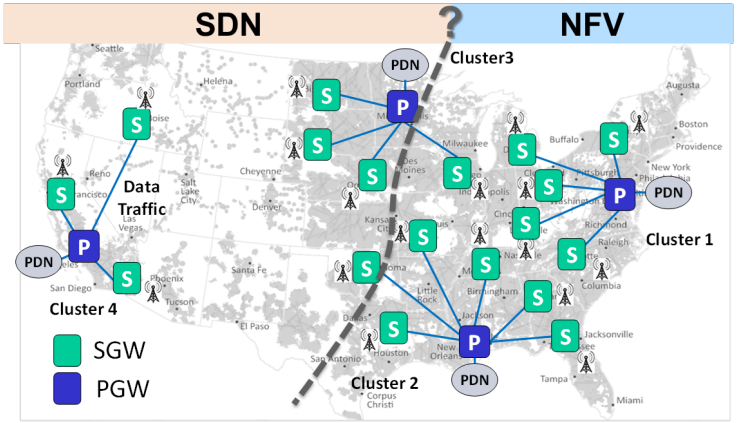

Function placement problem formulated as a MILP [6]

- SGW and PGW (VNF) placement
- constraints on data and control plane latency
- weights

$$w_{i,j} = \frac{\alpha}{dataLatency_i} + \frac{\beta}{controlLatency_j}$$

[4] A. Basta, W. Kellerer, M. Hoffmann, H. J. Morper, K. Hoffmann, Applying NFV and SDN to LTE mobile core gateways, the functions placement problem, All things cellular Workshop ACM SIGCOMM, Chicago, August, 2014.

Results [5]

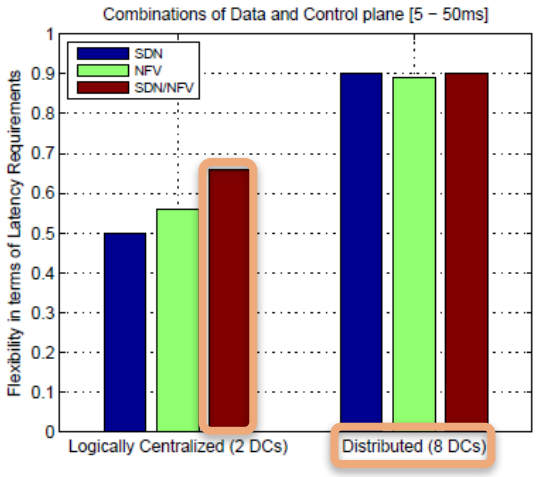

(a) equal weights for data and control latencies

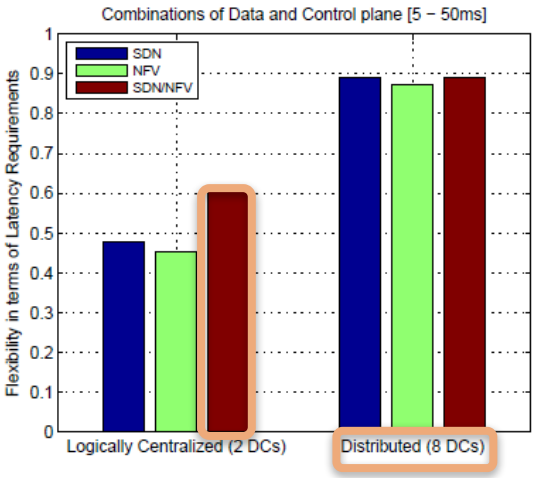

(b) weights biased by data latency

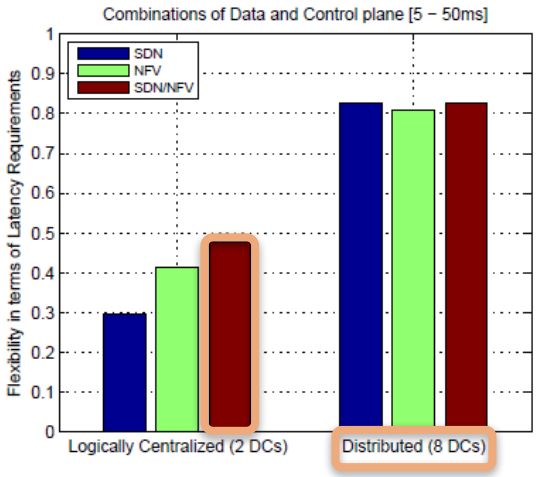

(c) weights biased by control latency

With respect to the support of latency requirements in function placement:

- mixed SDN/NFV is more flexible for a logically centralized data center infrastructure
- for distributed data centers all three design choices are equally flexible

[1] W. Kellerer, A. Basta, A. Blenk, Using a Flexibility Measure for Network Design Space Analysis of SDN and NFV, SWFAN'16, IEEE INFOCOM Workshop, April 2016.

Nothing is for free: Cost of Flexibility

What are the costs of a design for flexibility?

- in terms of signaling overhead, number of data centers,…

Possible relationship (to be confirmed):

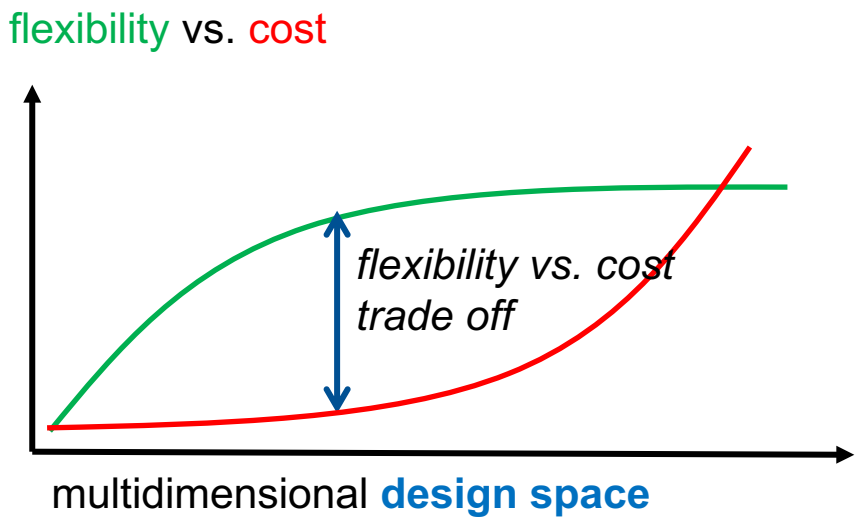

Use Case 2: Dynamic Controller Placement Problem

- place 1 ..n SDN controllers for time varying traffic input
  → controller migration/reconfiguration

- Evaluation parameters
  - Abilene network topology (11 nodes, 14 links)
  - new requests: 100 different flow profile requests over time (random)
  - N = 1,…, 4 controllers (design choices for comparison)
  - Algorithm finds optimal controller placement and flow to controller assignment
    optimization goal: minimize avg. flow setup time (**performance**)
  - How many controllers can be migrated (incl. control plane update) in time T?
    (success ratio → **Flexibility**)
  - Migrations and reconfigurations → **Cost**

[5] M. He, A. Basta, A. Blenk, W. Kellerer, *How Flexible is Dynamic SDN Control Plane?,*
IEEE INFOCOM Workshop, SWFAN, Atlanta, USA, May 2017.

[6] M. He, A. Basta, A. Blenk, W. Kellerer, *Modeling Flow Setup Time for Controller Placement in SDN: Evaluation for Dynamic Flows,*
IEEE International Conference on Communications (ICC), Paris, France, May 2017.

Simulation Results

Use Case

Flexibility **Performance** **Cost**

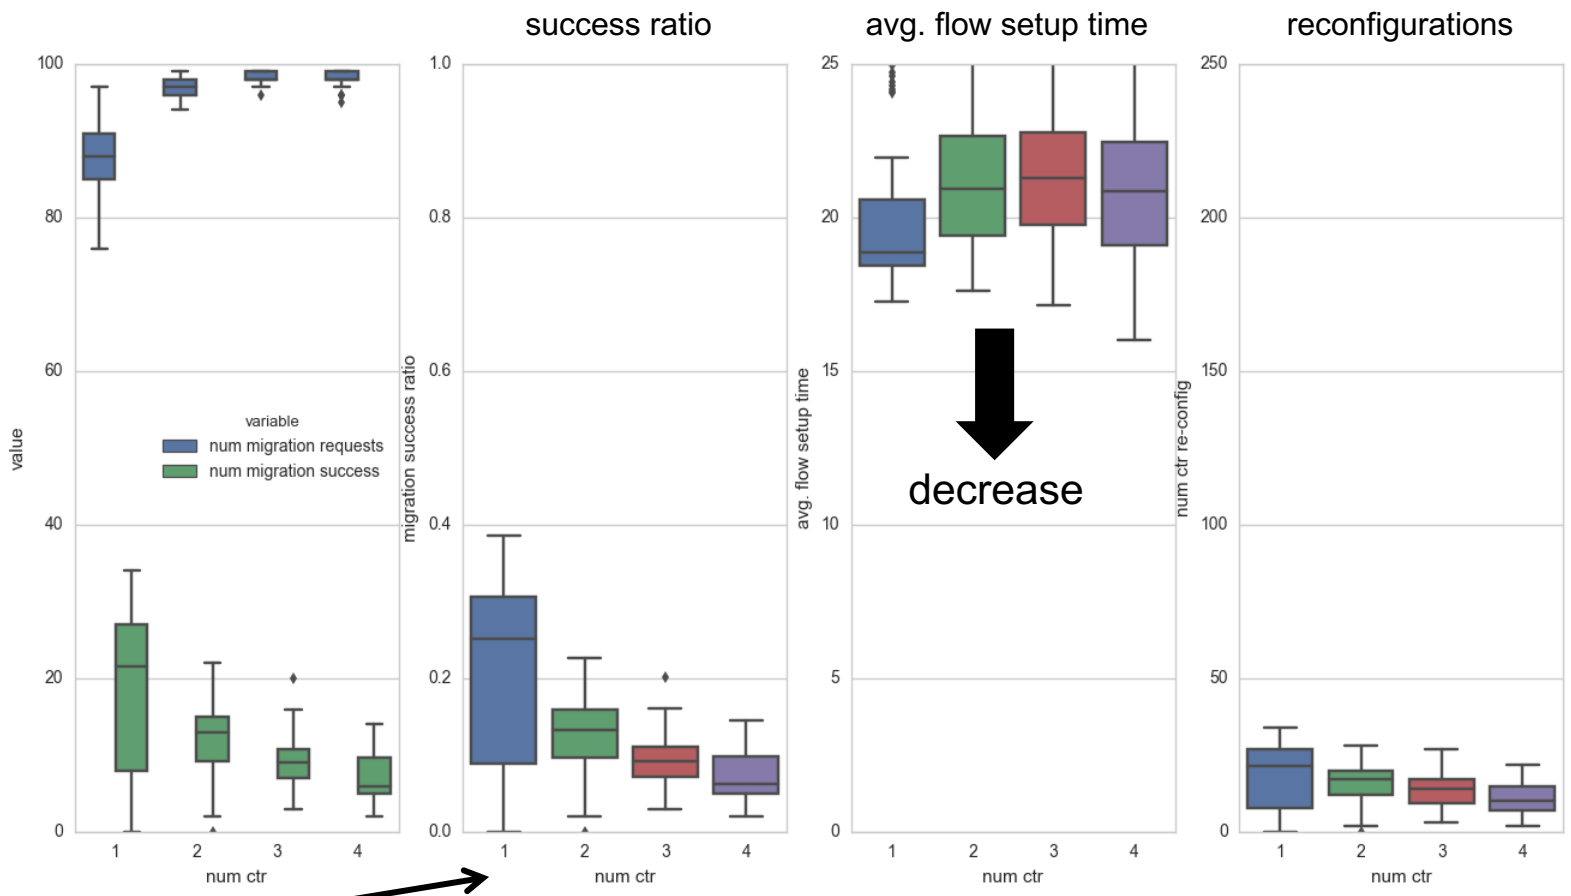

Number of controllers N

migration time threshold = 803 ms

T is very short (800 ms is transmission delay of 1 controller)

Simulation Results

Use Case

Flexibility **Performance** **Cost**

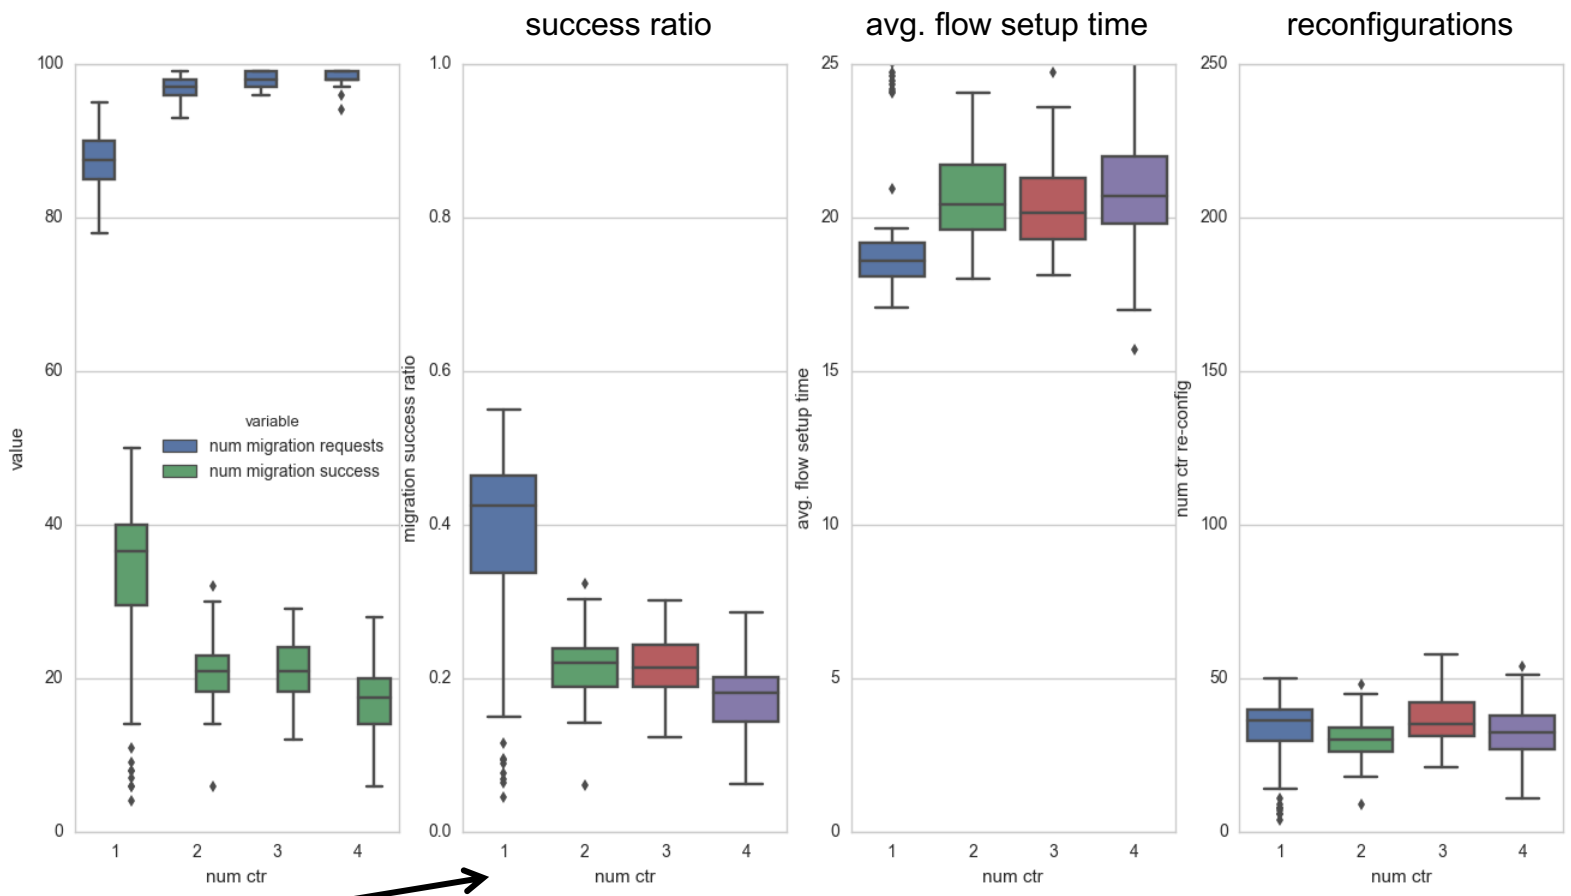

Number of controllers N

migration time threshold = 804 ms

T is very short (800 ms is transmission delay)

Simulation Results

Use Case

Flexibility **Performance** **Cost**

success ratio avg. flow setup time reconfigurations

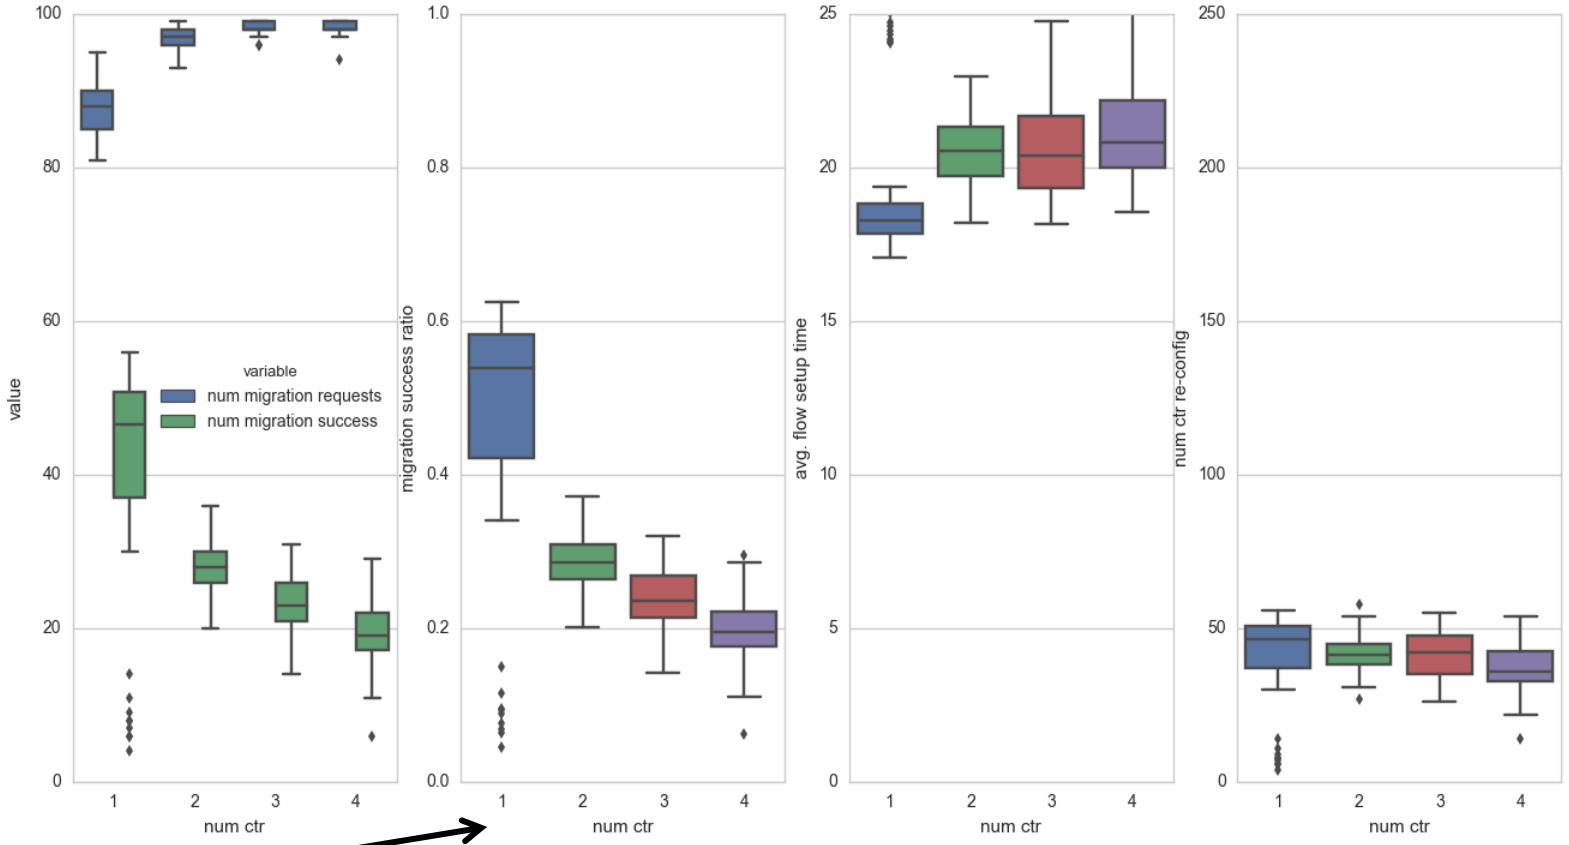

Number of controllers N

migration time threshold = 805 ms

Simulation Results

Use Case

Flexibility **Performance** **Cost**

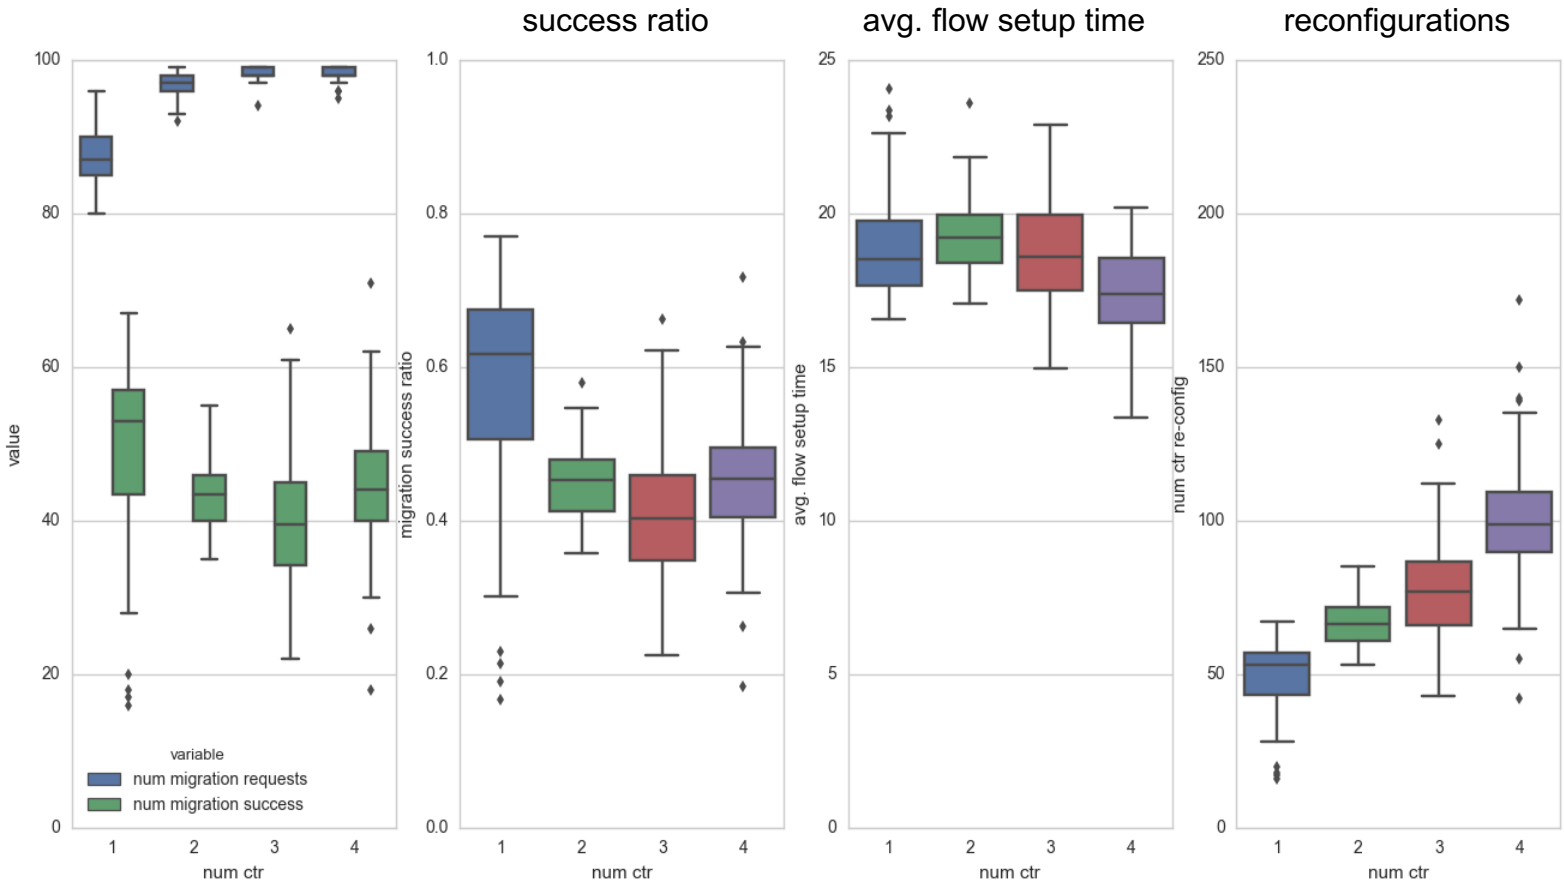

migration time threshold = 806 ms

1 controller has highest flexibility at low cost

But: performance is not good (flow setup time)

Simulation Results

Use Case

Flexibility **Performance** **Cost**

success ratio avg. flow setup time reconfigurations

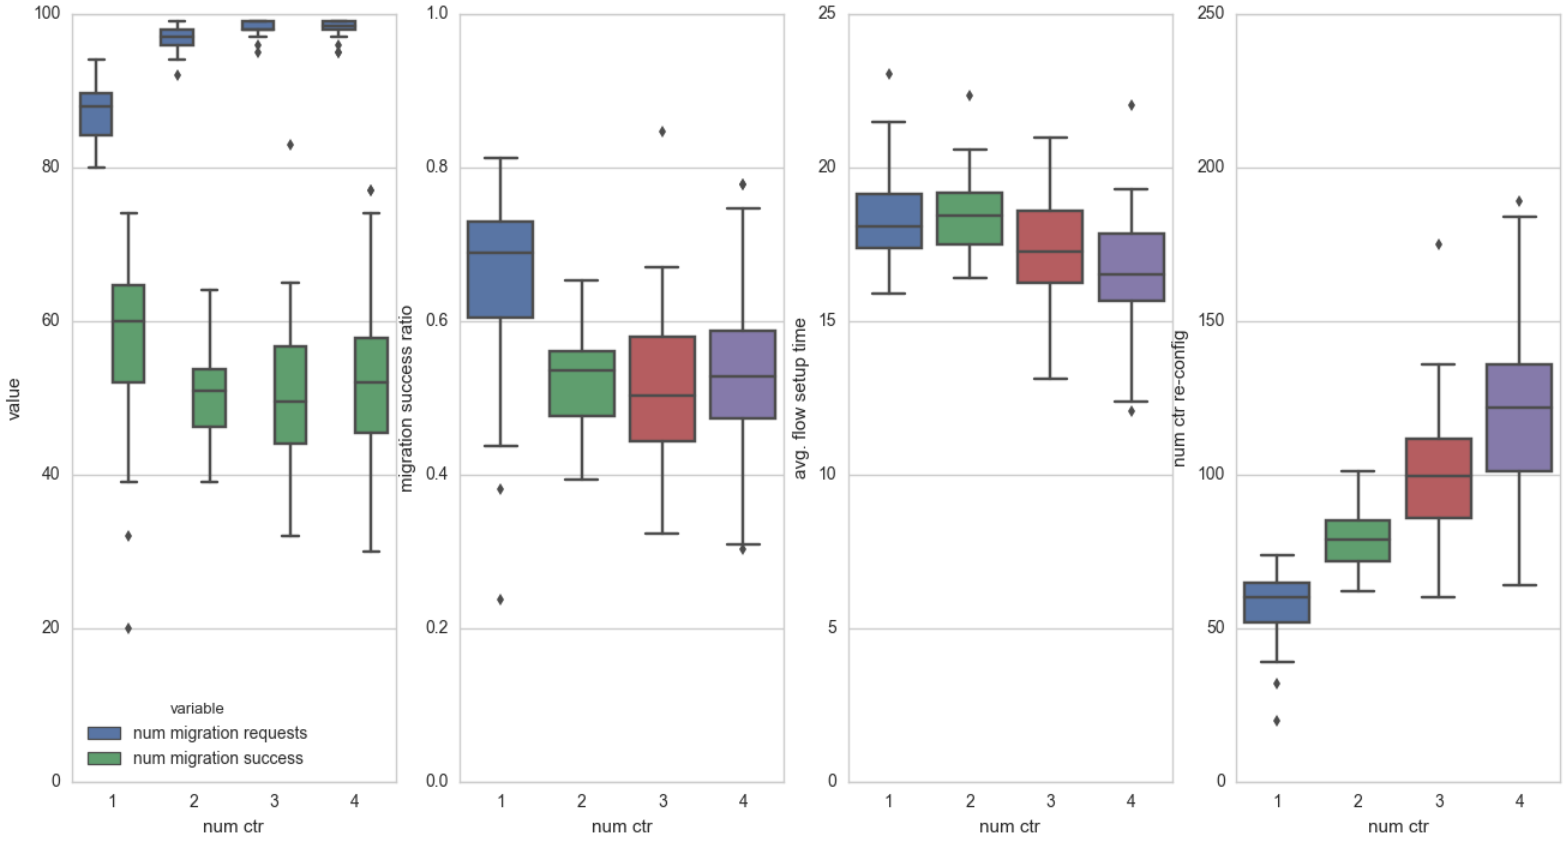

migration time threshold = 807 ms

Simulation Results

Use Case

Flexibility **Performance** **Cost**

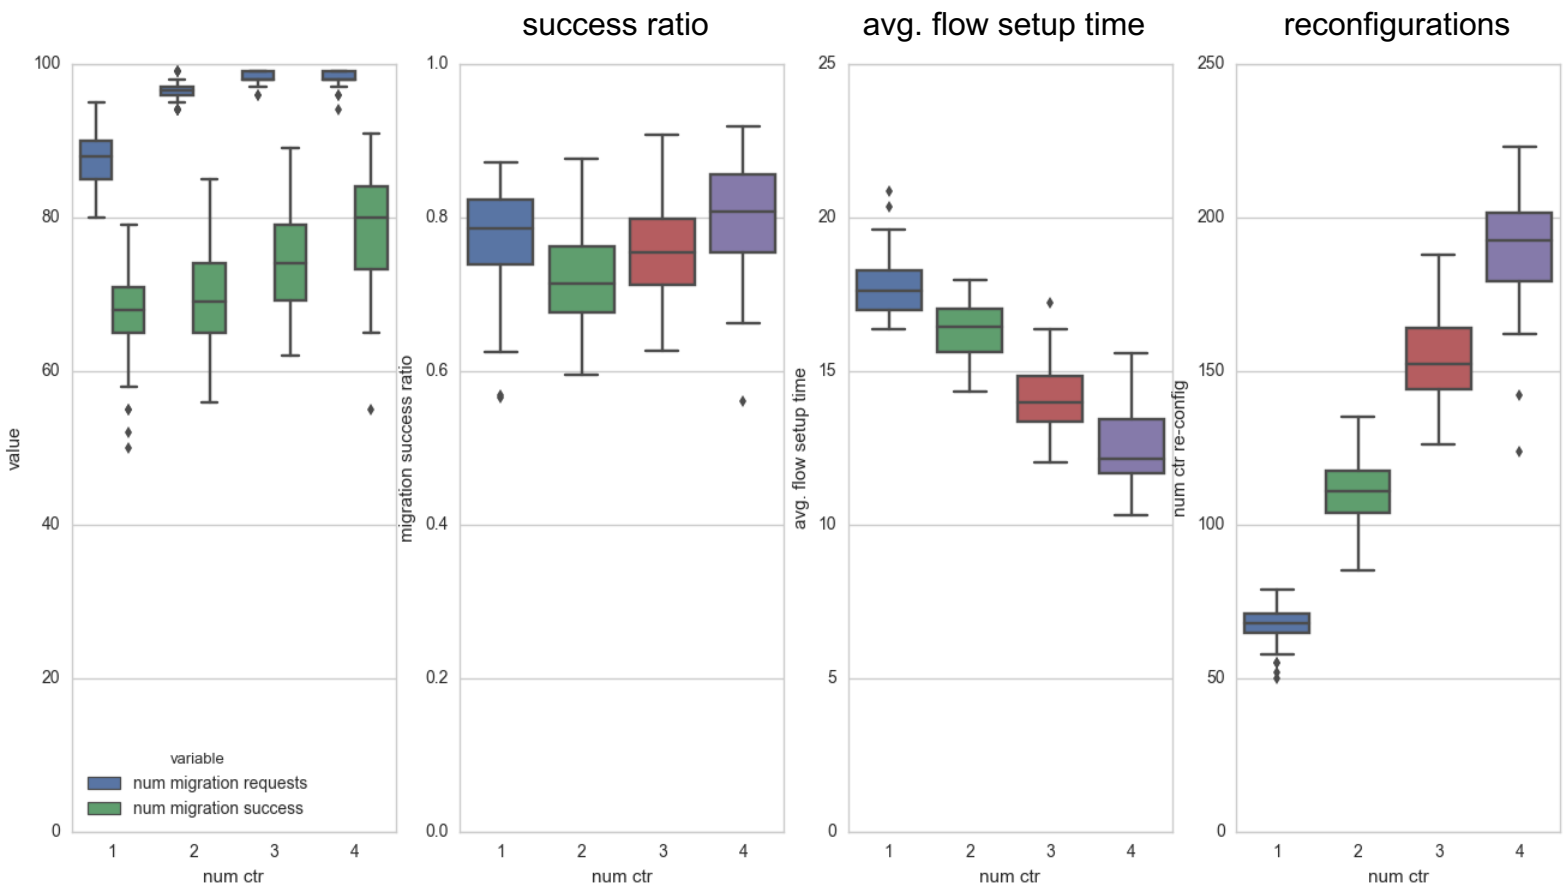

migration time threshold = 808 ms

Simulation Results

Use Case

Flexibility **Performance** **Cost**

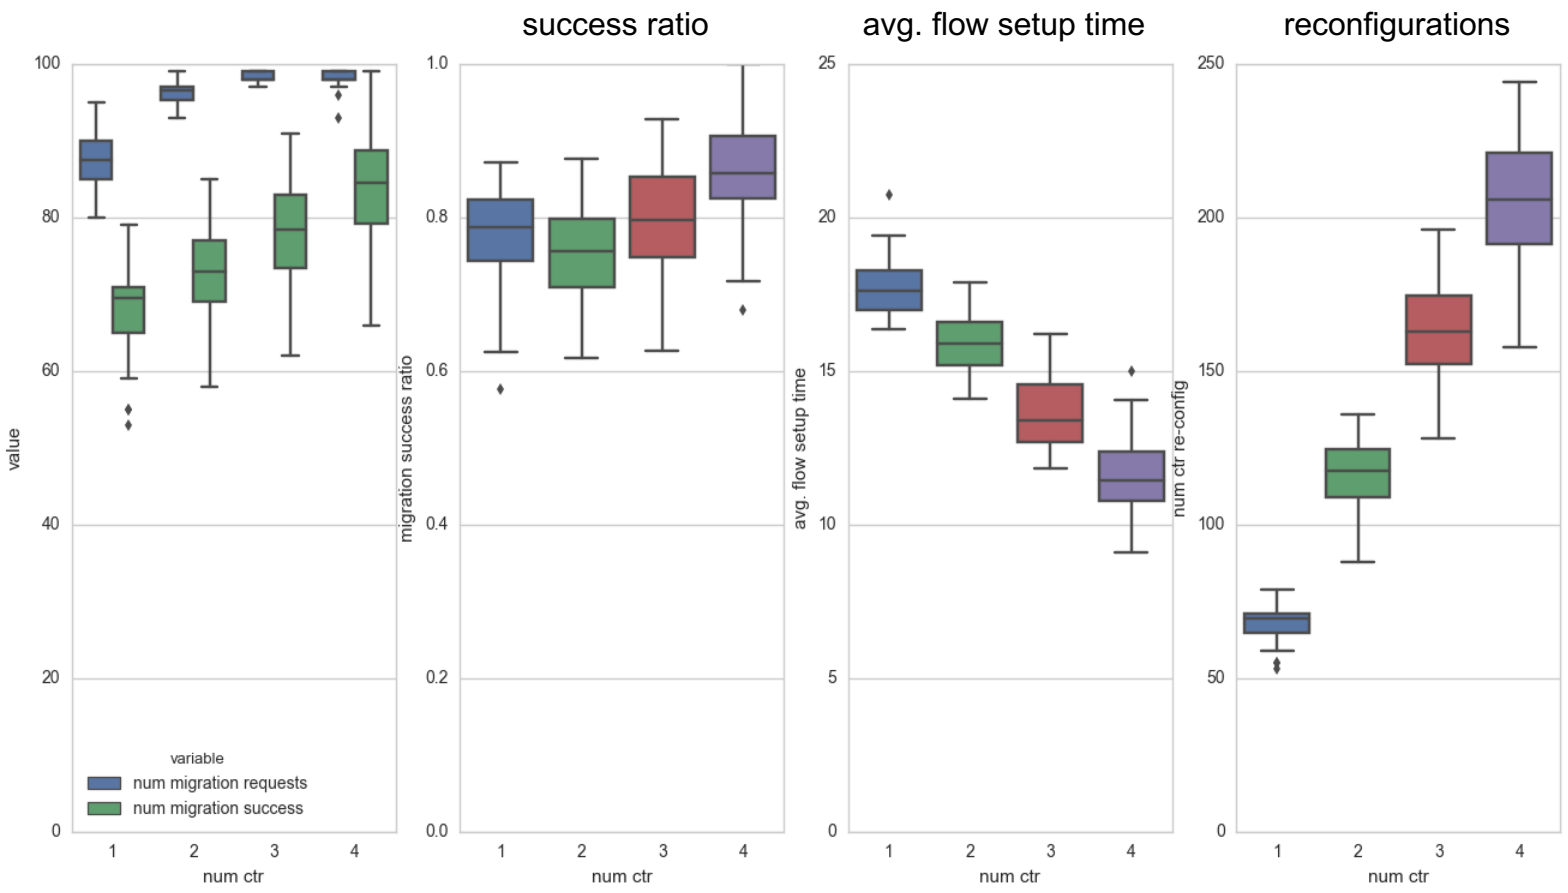

migration time threshold = 809 ms

Simulation Results

Use Case

Flexibility **Performance** **Cost**

success ratio avg. flow setup time reconfigurations

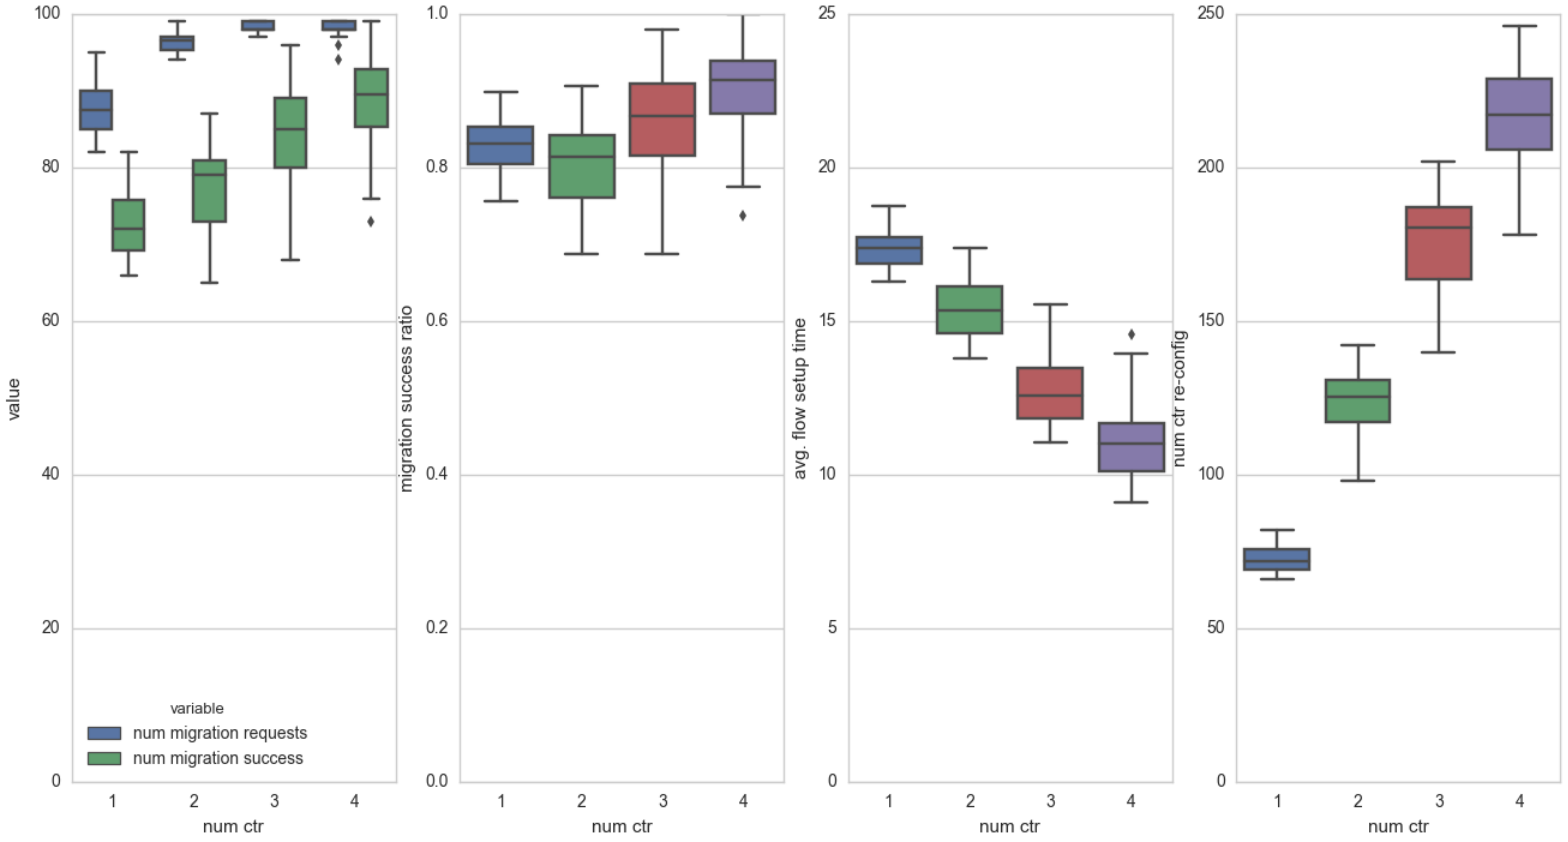

migration time threshold = 810 ms

T is moderate: more controllers → higher flexibility at higher cost

Simulation Results

Use Case

Flexibility **Performance** **Cost**

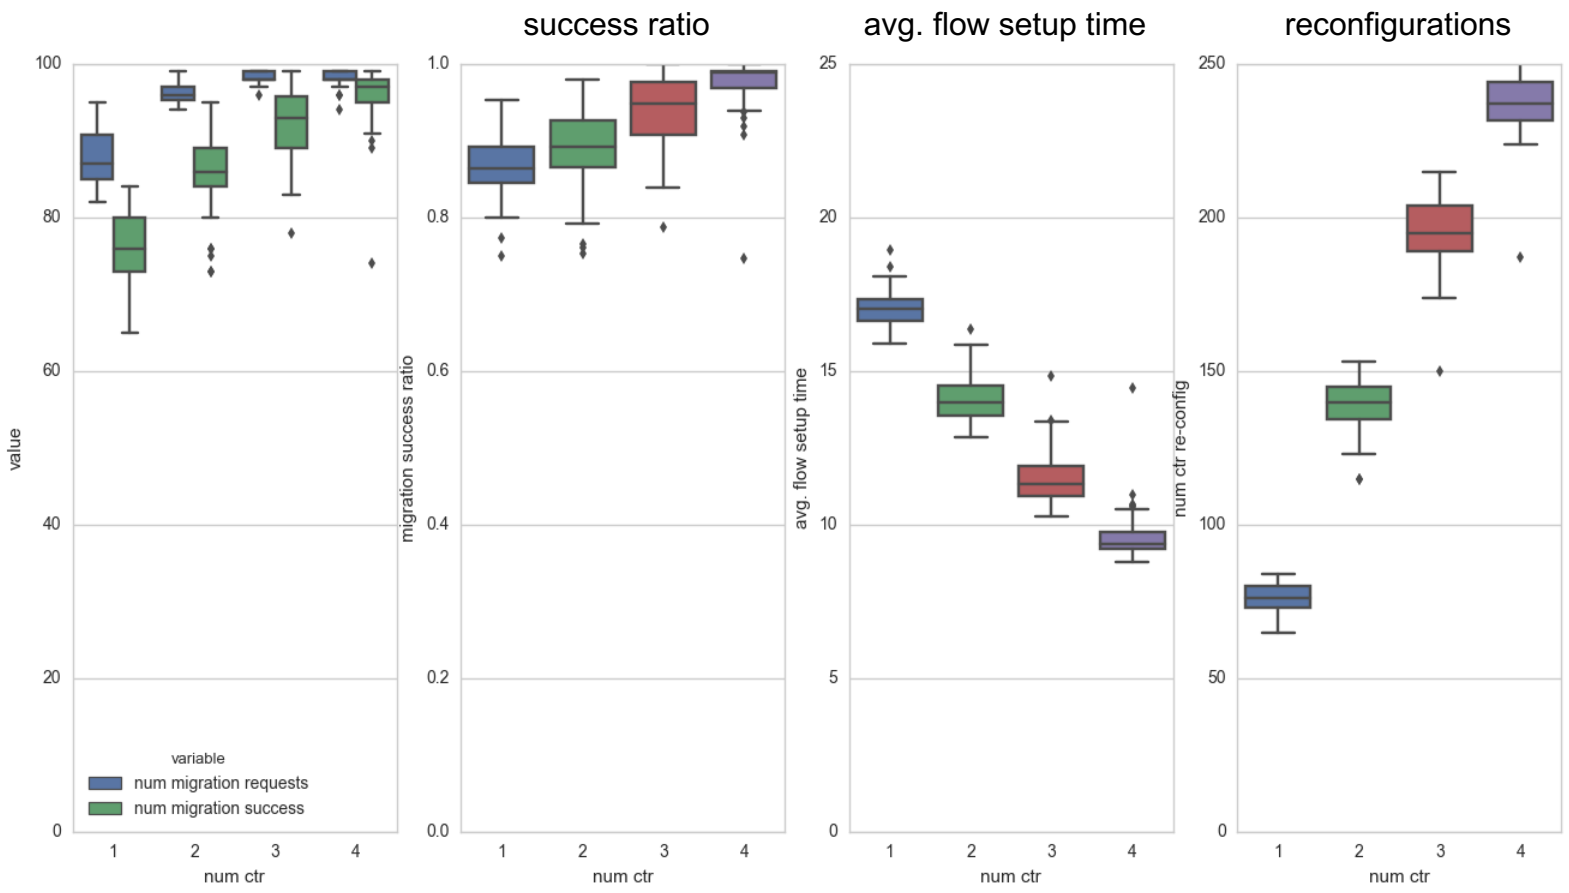

migration time threshold = 811 ms

T is moderate: more controllers → higher flexibility at higher cost

Simulation Results

Use Case

Flexibility **Performance** **Cost**

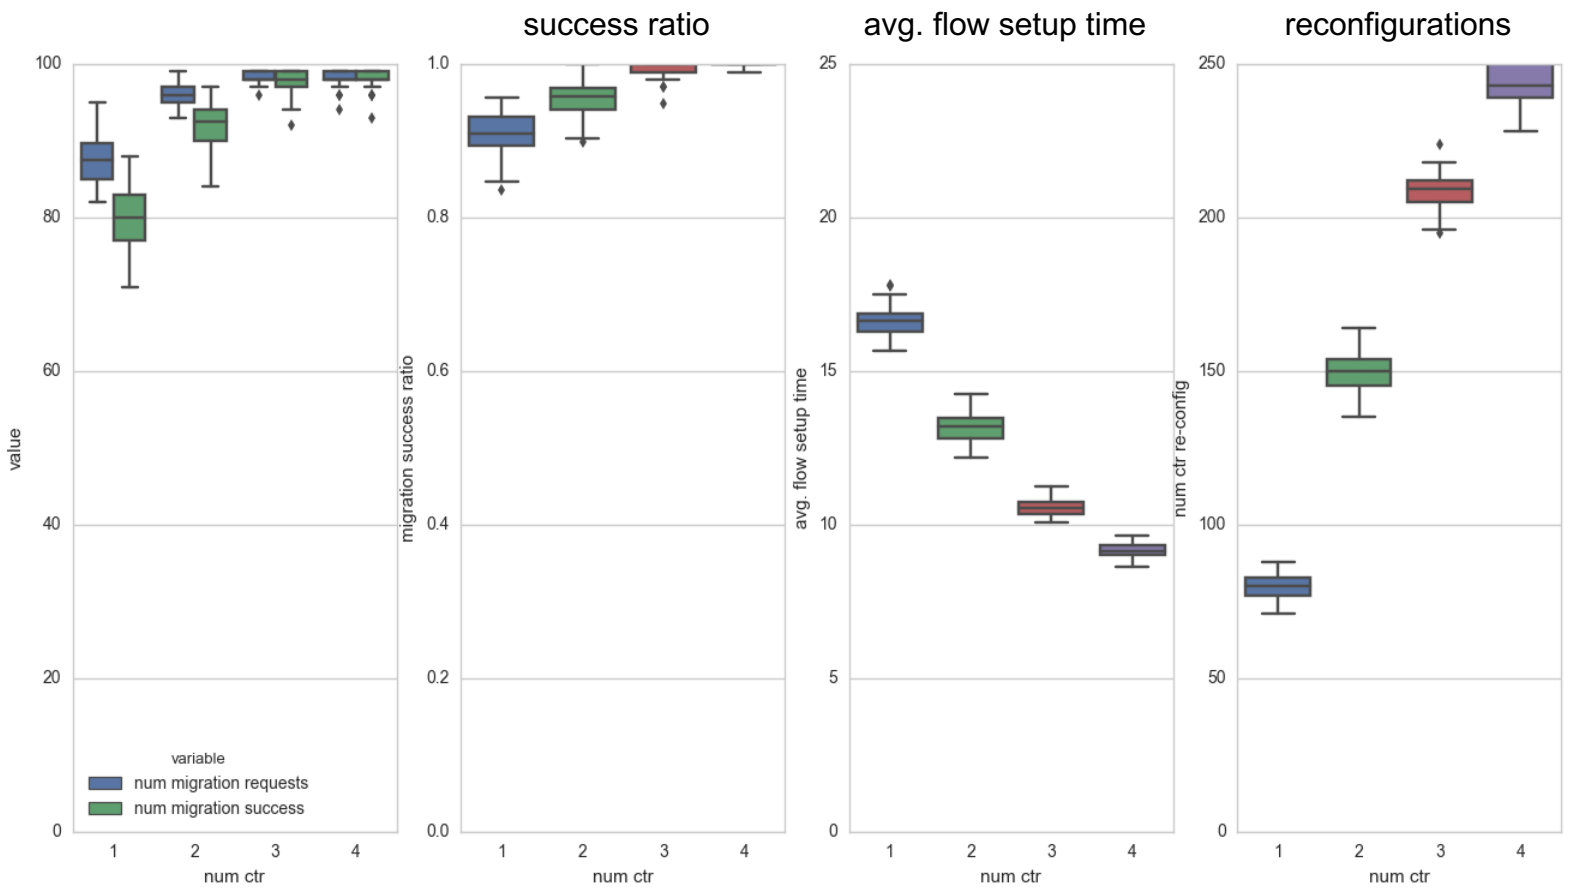

migration time threshold = 812 ms

T is moderate: more controllers → higher flexibility at higher cost

Interpretation

Use Case

Flexibility **Performance** **Cost**

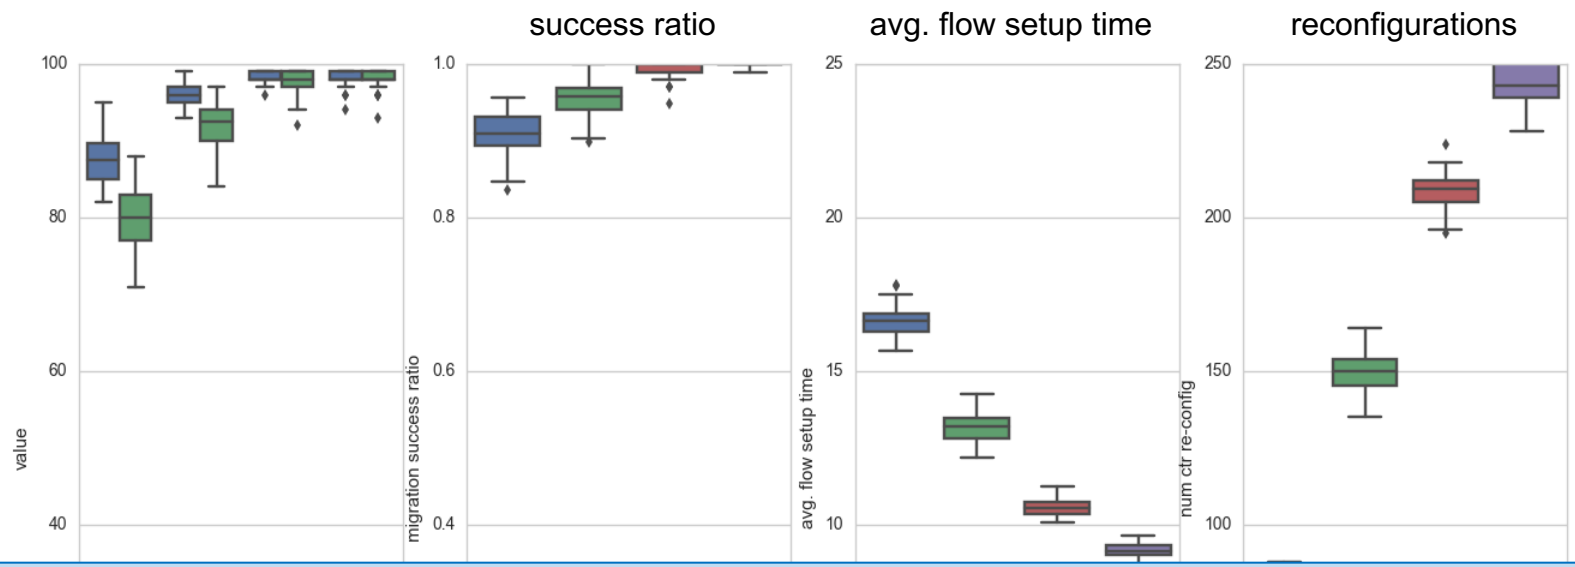

- Some cases: 1 controller is more flexible (short T)
- T considerable for adaptation: more controllers → more flexible
- There is a cap in gain – cost is rising

migration time threshold = 812 ms

Summary (from [5])

Flexibility

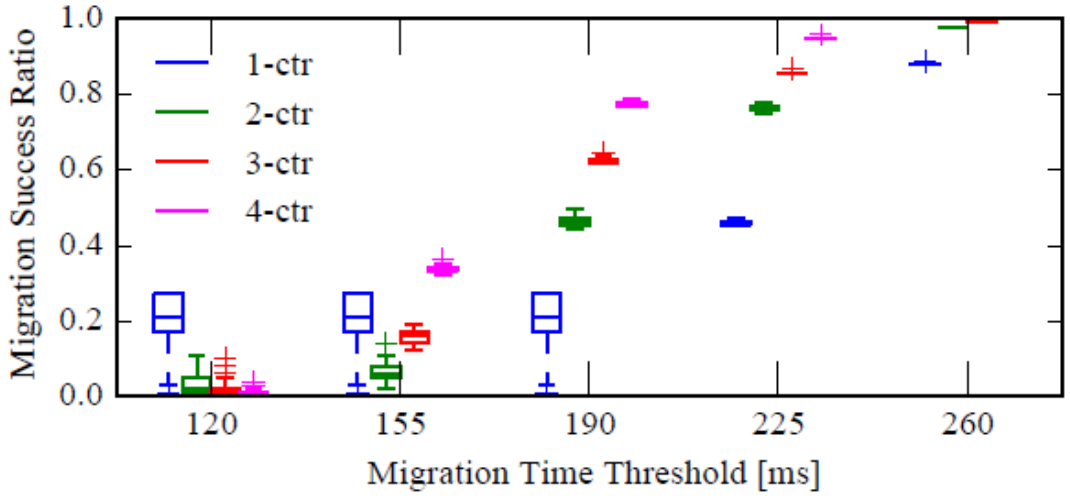

(a) Abilene, Flexibility Measure

Performance

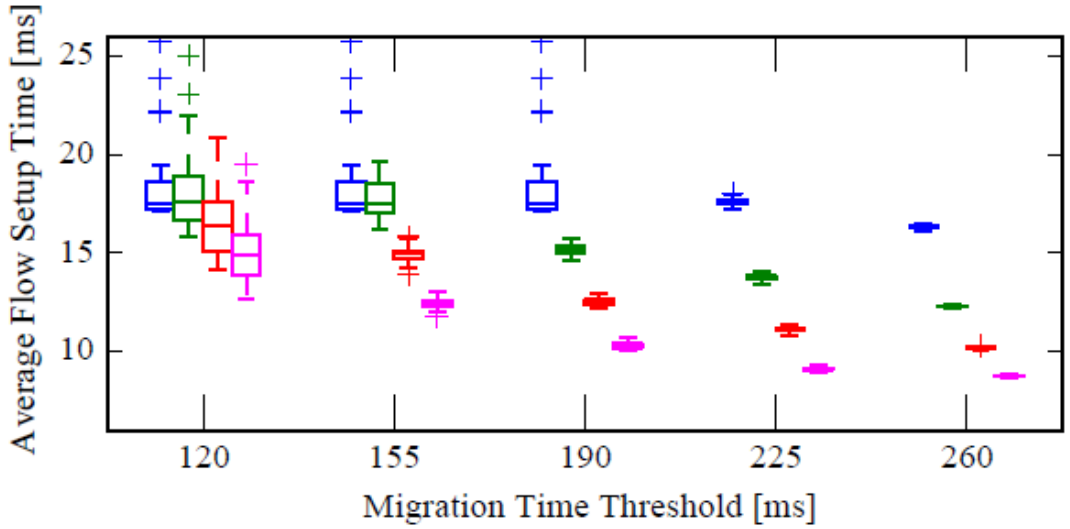

[5] M. He, A. Basta, A. Blenk, W. Kellerer, *How Flexible is Dynamic SDN Control Plane?,* IEEE INFOCOM Workshop, SWFAN, Atlanta, USA, May 2017.

from fundamental research to practice:

an implementation solution for flexibility

Designing for Flexibility: Network Slicing

- Why do we need network virtualization “slicing“?

- NGMN 5G white paper [7]
  - logical virtual mobile network slices
  - reliable and on-demand slices

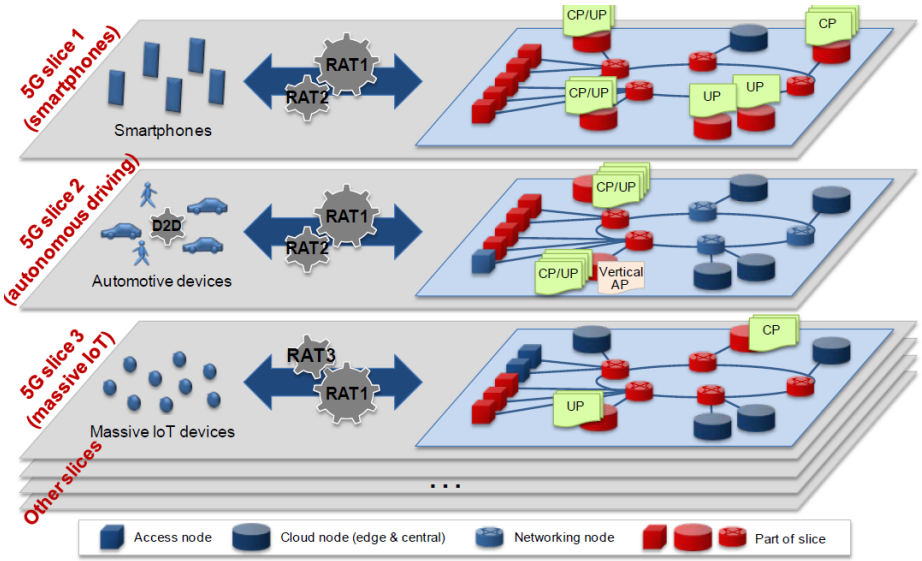

Source: NGMN 5G white paper

- METIS 5G system concept and technology roadmap [8]
  - application and service differentiation
  - logical virtual mobile network slices
  - heterogenous and dynamic slices

[7] 5G Initiative Team, NGMN 5G White Paper, 2015, https://www.ngmn.org/uploads/media/NGMN-5G-White-Paper-V1-0.pdf
[8] Mobile and wireless communications Enablers for the Twenty twenty Information Society (METIS), Final report on architecture (Deliverable D6.4), 2015, https://www.metis2020.com/wpcontent/uploads/deliverables/METIS-D6.4-v2.pdf

5G Slicing: SDN virtualization

- Why do we need SDN virtualization “slicing“ in 5G?
  - Bring your own controller
  - Full flexibility and programmability

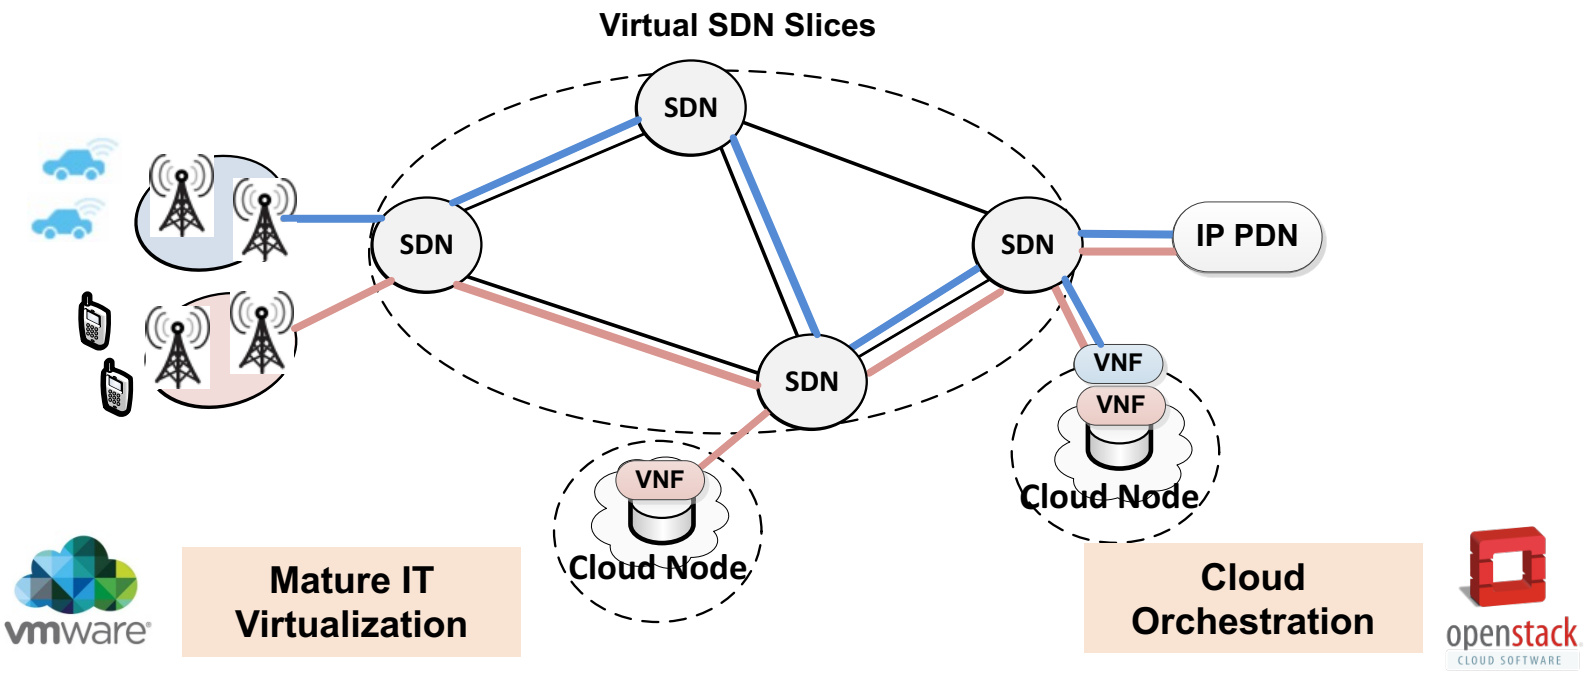

SDN Virtualization Overview

- How to achieve slicing for SDN networks?
  - SDN virtualization layer, i.e., SDN hypervisors
  - e.g. FlowVisor [9], OpenVirteX [10]

- What should an SDN hypervisor do?

- Virtual SDN abstraction
- Control plane translation
- Data and control slice isolation

- … in a most flexible way

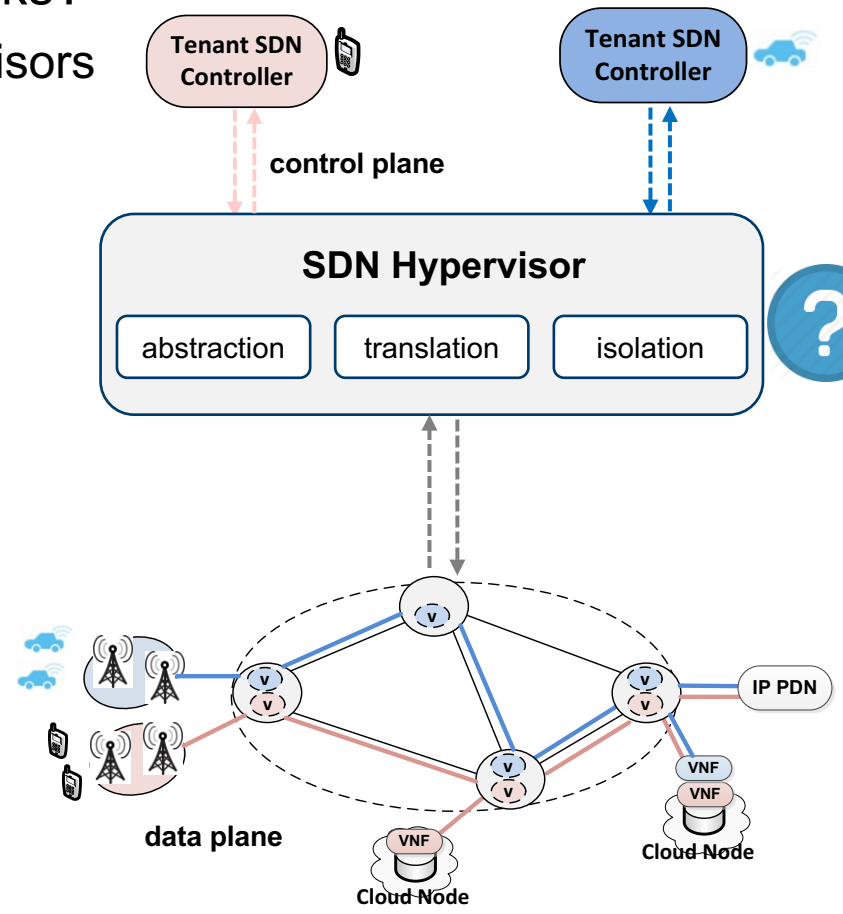

[9] R. Sherwood et al., Carving research slices out of your production networks with OpenFlow, ACM CCR, 2010

[10] A. Al-Shabibi et al, OpenVirteX: A network hypervisor, Open Networking Summit, 2014

State-of-the-art Limitations [11]

- SDN Slices
  - focus on data plane slices
  - control performance impacts the data plane performance in SDN!

- Management
  - automated slice request is not addressed
  - admission control interfaces are missing

- Deployment
  - no mechanisms to change the deployment on run time
  - e.g., automate adding or removing of a hypervisor instance

[11] A. Blenk, A. Basta, M. Reisslein, W. Kellerer, Survey on Network Virtualization Hypervisors for Software Defined Networking, IEEE Communications Surveys & Tutorials, vol. 18, no. 1, pp. 655-685, January 2016.

Our apporach: HyperFlex Features

- Admission Control [12]
  - automated request of virtual SDN slices
  - guarantees for data and control plane performance
  - run time update to slice
  - embedding of virtual links on the physical network

(a) Tenant View

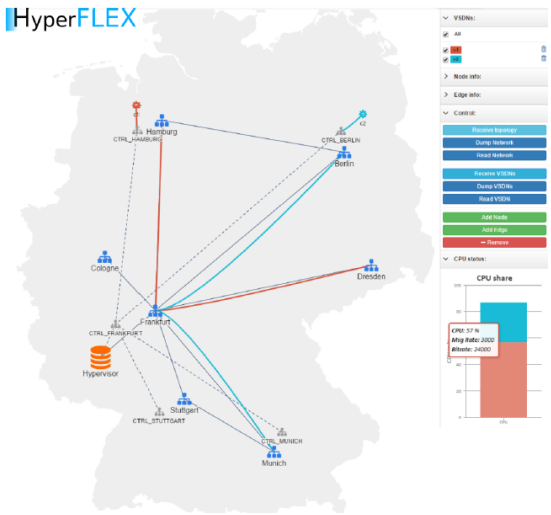

(b) HyperFlex View

[12] A. Blenk, A. Basta, J. Zerwas, M. Reisslein, W. Kellerer, Control Plane Latency with SDN Network Hypervisors: Cost of Virtualization IEEE Transactions on Network and Service Management, September 2016

SENDATE PLANETS (funded by the BMBF under Project ID 16KIS0473)
ERC Grant FlexNets (funded by the EC under grant agreement No 647158)

HyperFlex Features

- Performance Monitoring [12]
  - monitor the performance of the running hypervisors, e.g., CPU
  - monitor the performance of the SDN slices
    - control plane latency
    - control plane loss rate

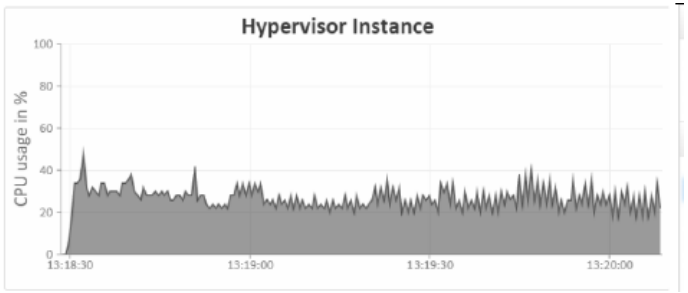

(a) Hypervisor performance

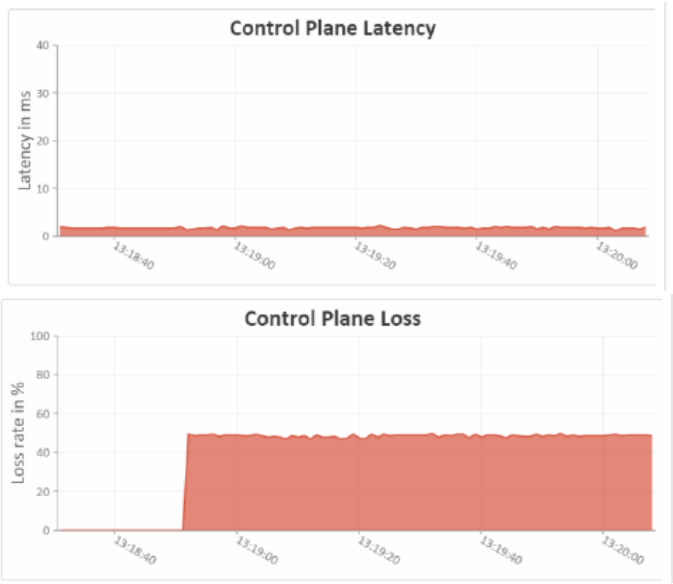

(b) Tenant control performance

[12] A. Blenk, A. Basta, J. Zerwas, M. Reisslein, W. Kellerer, Control Plane Latency with SDN Network Hypervisors: Cost of Virtualization
IEEE Transactions on Network and Service Management, September 2016

HyperFlex Features

- Dynamic Deployment “Orchestration” [12]
  - cope with the slice dynamics, e.g., new requirements, time-varying traffic, …
  - transparent to tenants, i.e., no interruption and no control loss
  - optimal placement of SDN hypervisors

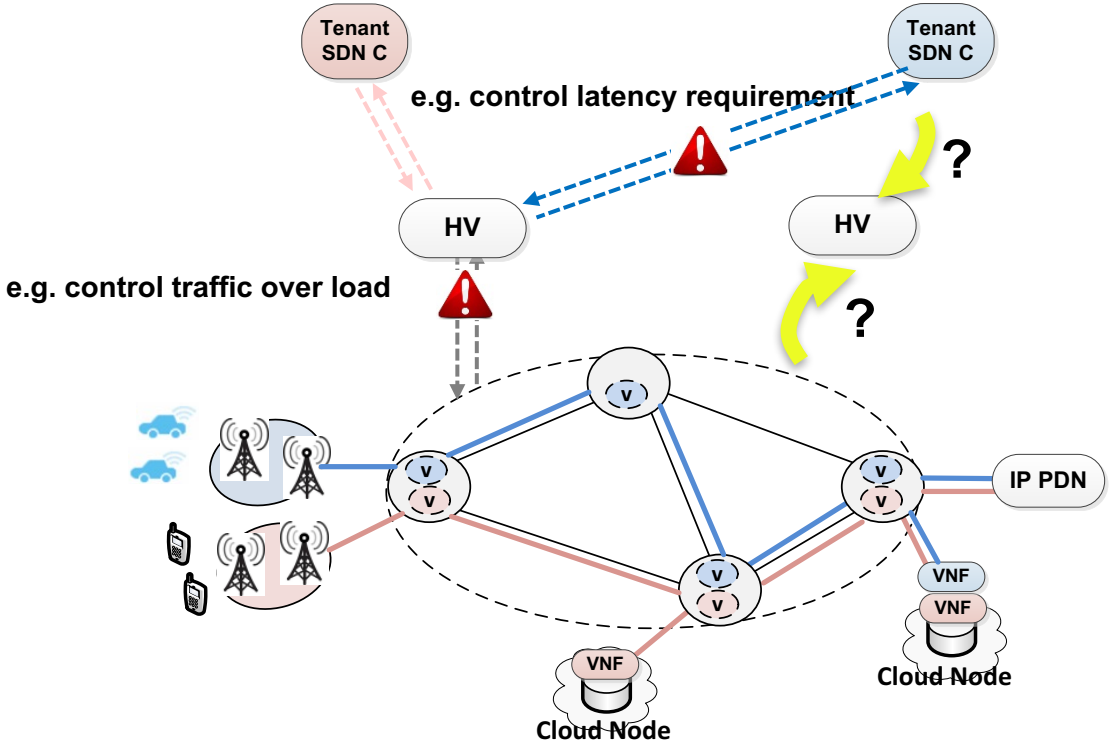

[12] A. Blenk, A. Basta, J. Zerwas, M. Reisslein, W. Kellerer, Control Plane Latency with SDN Network Hypervisors: Cost of Virtualizatio
IEEE Transactions on Network and Service Management, September 2016

Testbed@TUM: Flexible Application-to-Cloud Softwarized 5G Networks

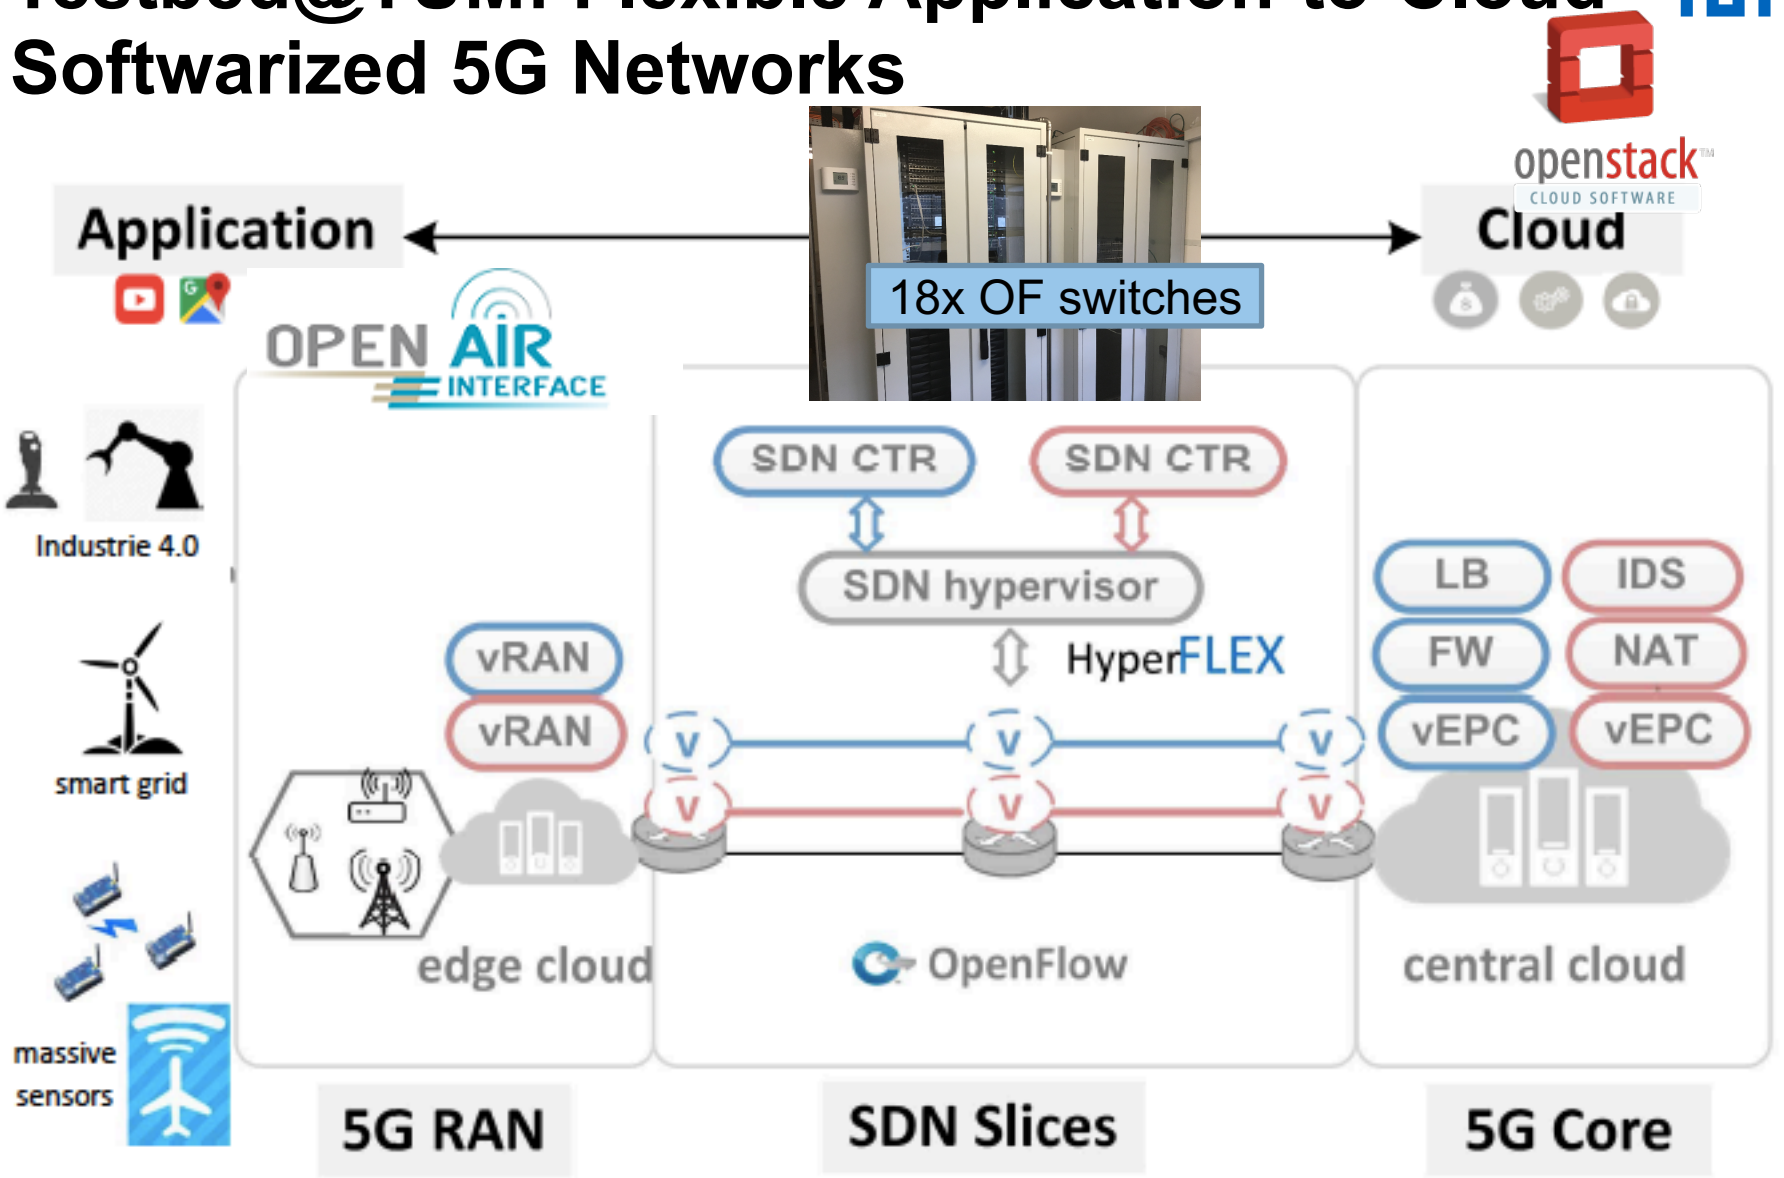

Conclusion

Key Takeaways

- Network research is faced with new requirements from emerging networked industries
- These include **flexibility**

- Network softwarization (SDN, NV, NFV) can be used

- Need for
  - a **measure** to analyse flexibility
  - new **flexible concepts** (e.g. HyperFlex)

References for further reading (1)

- M. He, A. Basta, A. Blenk, W. Kellerer, *How Flexible is Dynamic SDN Control Plane?,* IEEE INFOCOM Workshop, SWFAN, Atlanta, USA, May 2017.

- M. He, A. Basta, A. Blenk, W. Kellerer, *Modeling Flow Setup Time for Controller Placement in SDN: Evaluation for Dynamic Flows,* IEEE International Conference on Communications (ICC), Paris, France, May 2017.

- W. Kellerer, A. Basta, A. Blenk, *Using a Flexibility Measure for Network Design Space Analysis of SDN and NFV,* IEEE INFOCOM Workshop, SWFAN, San Francisco, USA, April 2016.

- A. Basta, W. Kellerer, M. Hoffmann, H. Morper, K. Hoffmann, *Applying NFV and SDN to LTE Mobile Core Gateways; The Functions Placement Problem*, AllThingsCellular14, Workshop ACM SICGOMM, Chicago, IL, USA, August 2014.

- A. Basta, A. Blenk, M. Hoffmann, H. Morper, K. Hoffmann, W. Kellerer, *SDN and NFV Dynamic Operation of LTE EPC Gateways for Time-varying Traffic Patterns,* 6th International Conference on Mobile Networks and Management (MONAMI), Würzburg, Germany, September 2014.

- W. Kellerer, A. Basta, A. Blenk, *Flexibility of Networks: a new measure for network design space analysis?,* arXive report, December 2015. http://www.lkn.ei.tum.de/forschung/publikationen/dateien/Kellerer2015FlexibilityofNetworks:a.pdf

References for further reading (2)

- A. Basta et al., A Virtual SDN-enabled EPC Architecture : a case study for S-/P-Gateways functions, SDN4FNS 2013.

- A. Blenk, A. Basta, J. Zerwas, M. Reisslein, W. Kellerer, Control Plane Latency with SDN Network Hypervisors: Cost of Virtualization,
  IEEE Transactions on Network and Service Management, September 2016

- A. Blenk, A. Basta, M. Reisslein, W. Kellerer, Survey on Network Virtualization Hypervisors for Software Defined Networking,
  IEEE Communications Surveys & Tutorials, vol. 18, no. 1, pp. 655-685, January 2016.

- R. Sherwood et al., Carving research slices out of your production networks with OpenFlow, ACM CCR, 2010

- A. Al-Shabibi et al, OpenVirteX: A network hypervisor, Open Networking Summit, 2014

- 5G Initiative Team, NGMN 5G White Paper, 2015, https://www.ngmn.org/uploads/media/NGMN-5G-White-Paper-V1-0.pdf

- Mobile and wireless communications Enablers for the Twenty twenty Information Society (METIS), Final report on architecture
  (Deliverable D6.4), 2015, https://www.metis2020.com/wpcontent/uploads/deliverables/METIS-D6.4-v2.pdf