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

Towards Flexible Networking in Dynamically Changing Environments

Wolfgang Kellerer

Technical University of Munich

ITC 2016, September 14, 2016

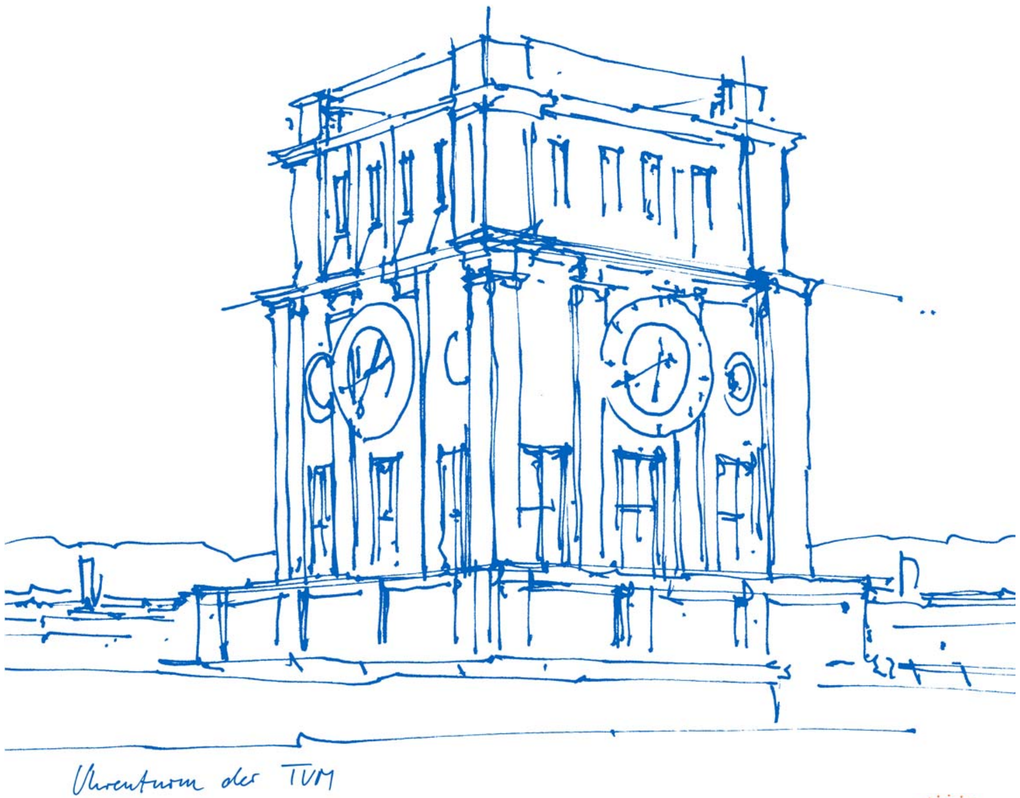

with the support of my PhD students:
A. Blenk, A. Basta, R. Durner, J. Guck,
M. He, A. Van Bemten,…

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 in 2016
  - new requirements from vertical industries
  - new requirements from dynamically changing user behavior
  - new requirements from global digitalization

5G cellular, Industrie 4.0, Smart Grid, Big Data, ITS, Cyber Physical Networking,…

- One challenge that is less (explicitly) addressed is ***flexibility***

- Evolution tells us: be adaptive → network evolution?

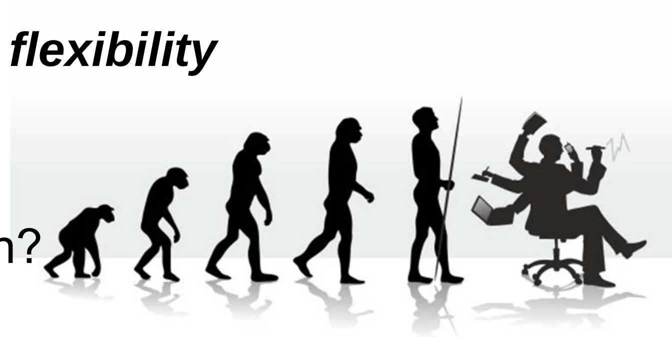

Image source: http://www.paleoplan.com

- In this talk, I try to give answers to …
  … network flexibility discussions and …
  … present some new concepts …
  … and raise possibly more questions

FlexNets

European Research Council
Established by the European Commission

Outline of this Talk

- (more) motivation

- Use Cases (1): *new requirements and the role of SDN, NFV,…*
  - what 5G, windparks and firewalls have in common

- Towards a flexibility measure for network design space analysis focusing on SDN and NFV

- Use Cases (2): *flexibility*
  - 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

very slow adaptation to new requirements

- Industry 4.0, vehicular, tactile → low delay, million sensors,…

adaptation processes limited, costly and inefficient

→ reaction to dynamic changes hardly possible

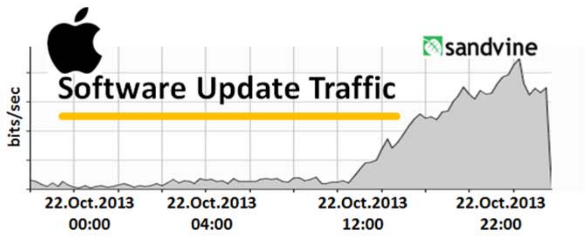

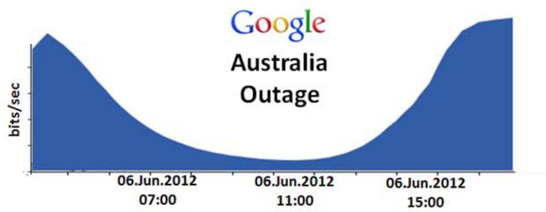

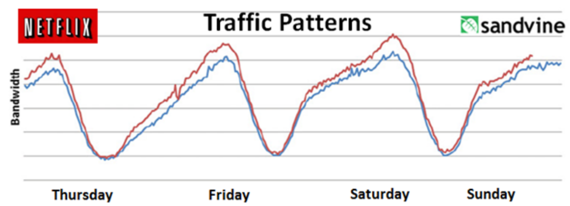

New concepts such as …

Network Virtualization, Software Defined Networking and Network Function Virtualization

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

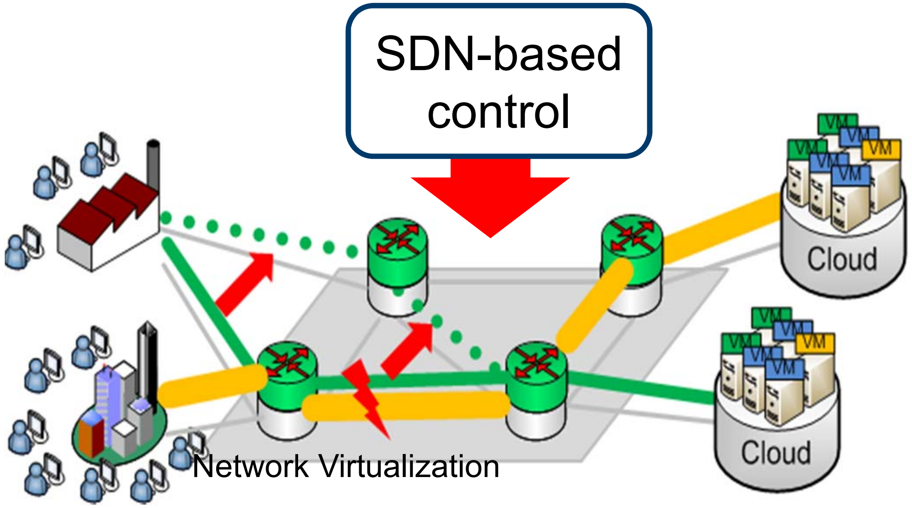

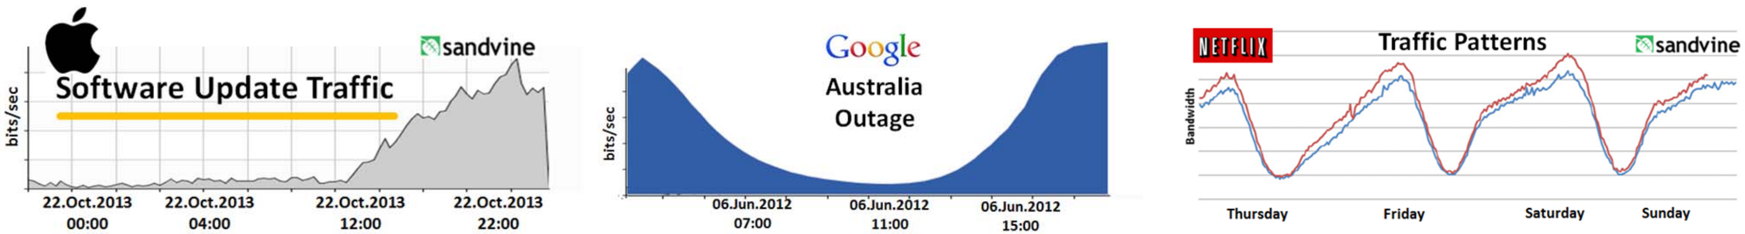

All problems solved?

- Do they really solve all our problems?

- How far can we go? What is the right network design?

We need
- a **fundamental understanding** of how to provide flexibility
- a set of **quantitative arguments** pro and contra certain design choices
- a set of **guidelines** of how software-based network shall **be designed**

i.e. we need to understand design choices for a network to guarantee a certain degree of flexibility, e.g. graph, controller number and place,…

Some Use Cases to motivate new requirements and new concepts for flexibility

Use Case: 5G

- 5G is one major objective for communication network research today
- Important requirement: support for vertical industries

VirtuWind: Virtual and programmable industrial network prototype deployed in operational windpark

Objectives of VirtuWind [1] include

- TCO reduction through use of SDN and NFV
- Flexible adoption of virtualized network functions
- Industrial-grade QoS for intra- and inter domain multi-tenancy SDN/NFV solution
- Resilience and Security by design

[1] T. Mahmoodi, V. Kulkarni, W. Kellerer, et al, VirtuWind: Virtual and Programmable Industrial Network Prototype Deployed in Operational Wind Park,accepted for Transactions on Emerging Telecommunications Technologies (ETT), Wiley, vol. 27, 2016.

Use Case: VirtuWind

VirtuWind objectives include

- Industrial-grade QoS for intra- and inter domain multi-tenancy SDN/NFV solution

Our approach [2]:

- Central (SDN-based) QoS management based on Network Calculus on standard Ethernet HW/SW to support and guarantee industrial network QoS requirements (delay, resilience)

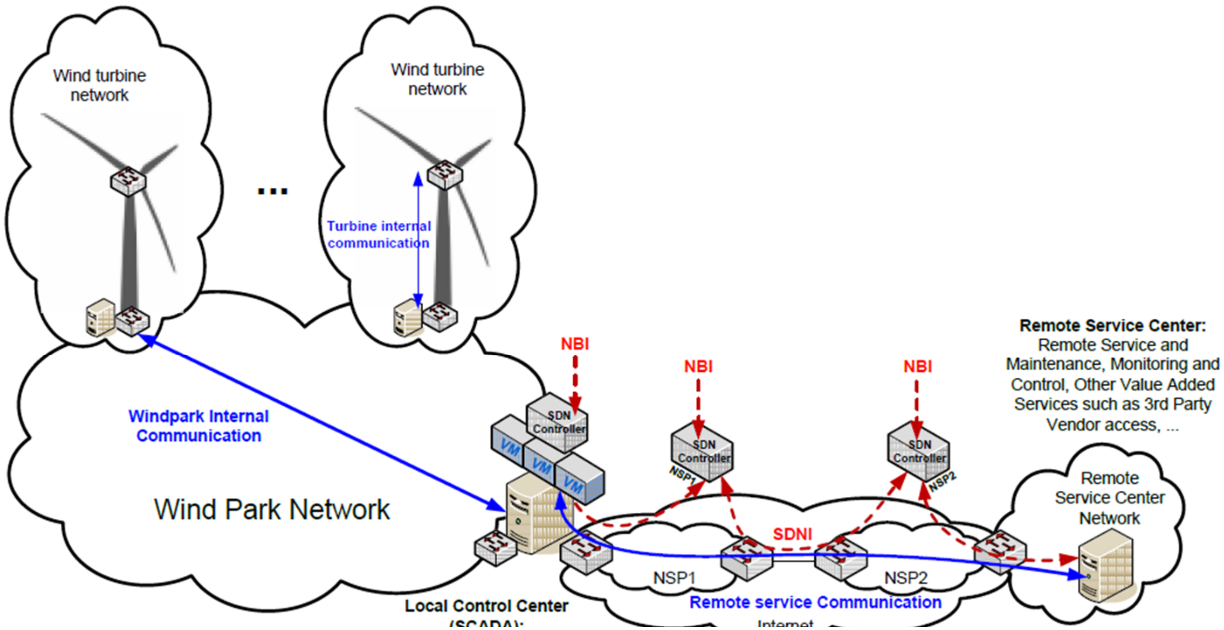

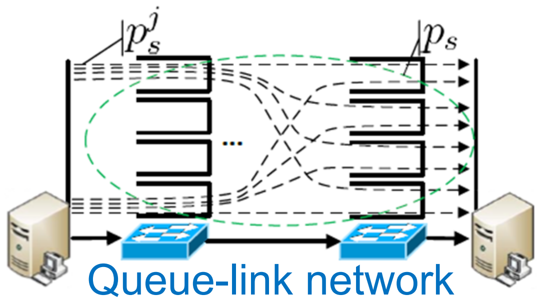

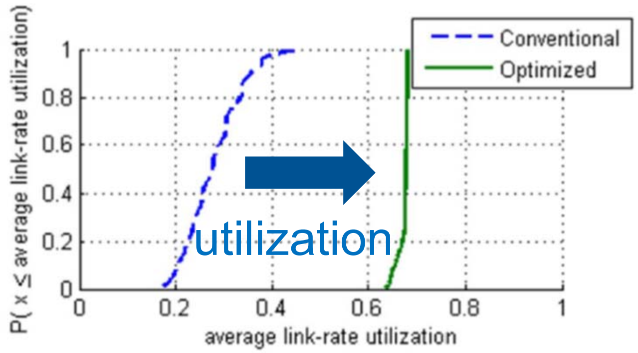

[2] J. Guck, M. Reisslein, W. Kellerer, Function Split between Delay-Constrained Routing and Resource Allocation for Centrally Managed QoS in Industrial Networks. IEEE Transactions on Industrial Informatics, 2016. (open access)

Use Case: Security

Enterprise networks: Need for fine-grained and flexible security solutions

Our approach [3]:

- combine SDN and NFV to adapt to changing demands

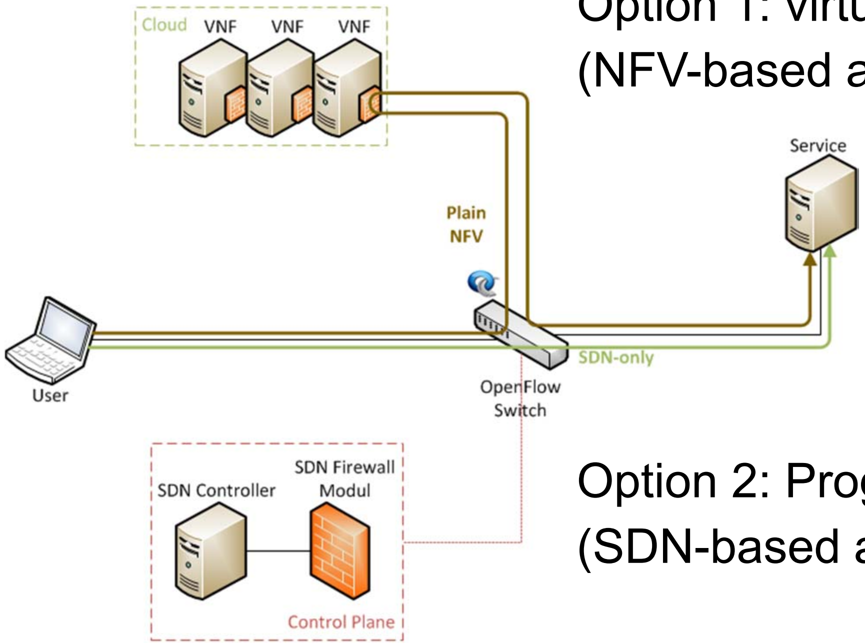

Option 1: virtualised firewall in the cloud (NFV-based approach)

Option 2: Programmable Switch → firewall (SDN-based approach)

[3] C. Lorenz, D. Hock, R. Durner, W. Kellerer, etal.: An SDN/NFV-enabled Enterprise Network Architecture Offering Fine-Grained Security Policy Enforcement. Accepted for IEEE ComMag, 2016.

SarDiNe - Netzsicherheit in Unternehmen und Behörden basierend auf Software Defined Networking (funded by the BMBF under grant number 16KIS0260)

Bundesministerium für Bildung und Forschung

Flexibility is a new key term!

to address new use cases

SDN, NV, NFV claim to provide
more flexibility in networks

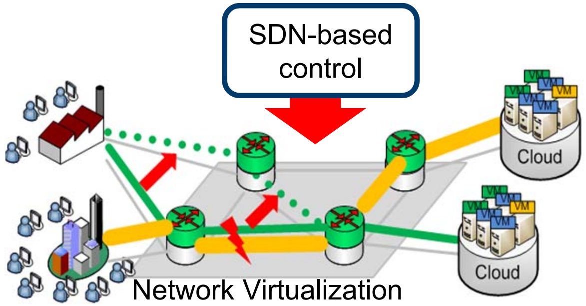

A deeper understanding of what flexibility means and how it could be quantified to compare different network designs remains open

For networks, **flexibility** = ability to *adapt* resources (flows, topology,…) *to changes* of design requirements (dynamic traffic, shorter 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)**.

European Research Council

Flexibility: a new measure?

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

- similar to QoS

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

we propose: ***flexibility aspects*** [4, 5]

- similar as we do with QoS (rate, delay, throughput, jitter,…)

- shall allow us to quantitatively compare two different system designs

[4] W. Kellerer, A. Basta, A. Blenk, Flexibility of Networks: a new measure for network design space analysis?. arXiv preprint arXiv:1512.03770, 2015.

[5] 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.

Flexibility Aspect example 1: Flow steering and reconfiguration

Parameters (for change requests):

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

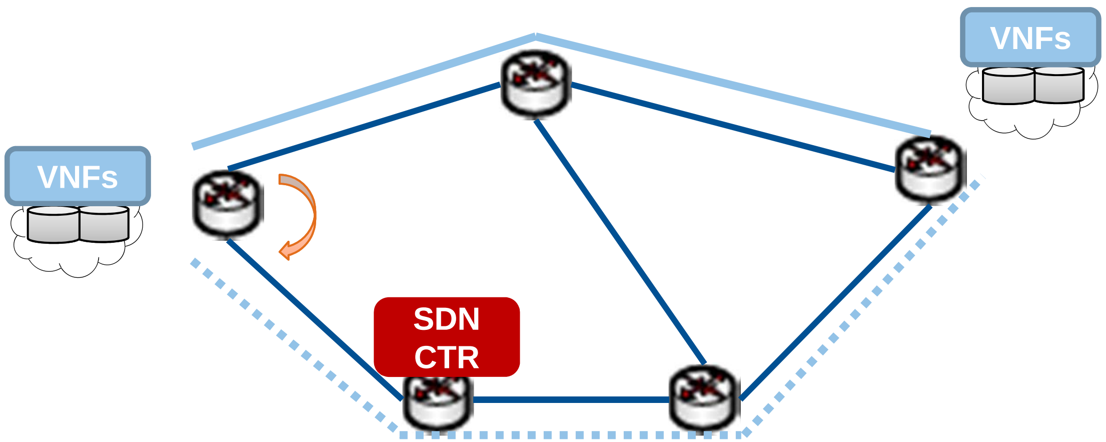

Flexibility Aspect example 2: Function Placement

Parameters:

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

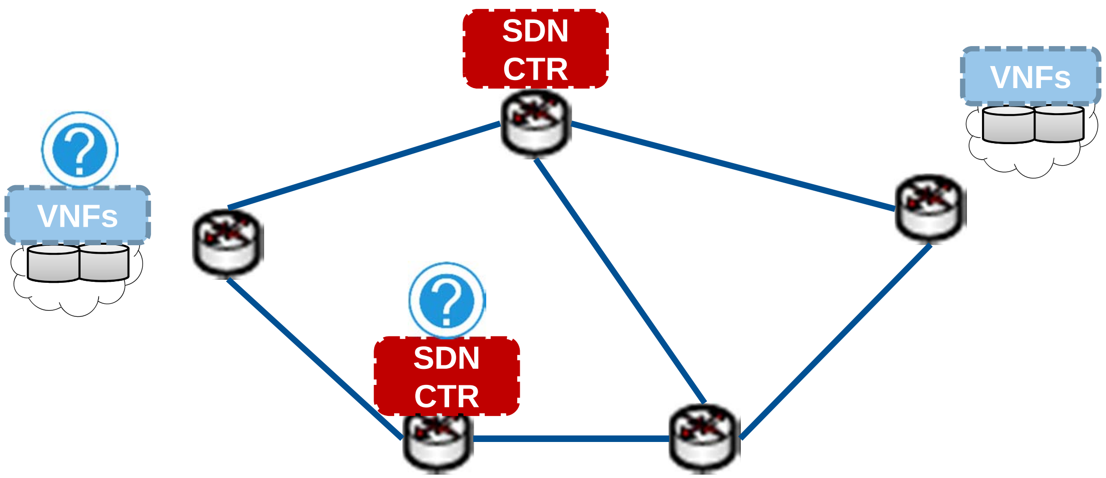

A simple measure

e.g., *placement*

$$\varphi^{aspect}(S) = \frac{|supported\ requests|}{|possible\ requests|}$$

- fraction of the number of **change requests** that can be supported of all possible change requests
- w.r.t. to a certain **flexibility aspect** of a system S
- φ (S) ϵ [0,1] „percentage“

Use Case: **Function Placement** (= aspect)

3 design choices to compare for next generation mobile core network [5]:

(1) SDN design

(2) NFV design

(3) mixed SDN/NFV design

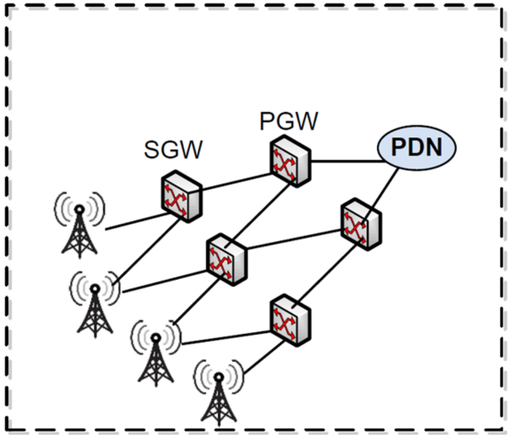

(a) Current LTE Core GW Architecture

Parameter in focus:

- Flexibility to support different latency requirements for
  - control plane latency and
  - data plane latency
- How many latency requirements can be fulfilled by a design choice?

[5] 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

Use Case

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

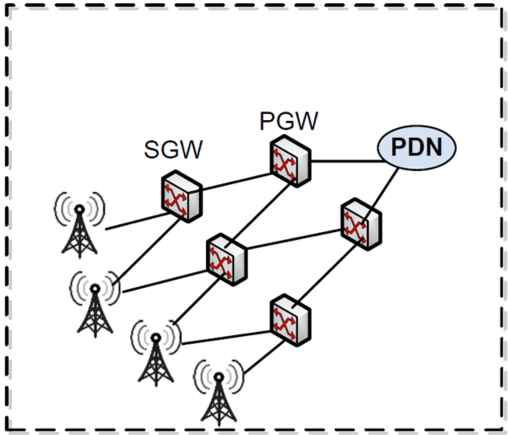

(a) Current LTE Core GW Architecture

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

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

(3)
mixed SDN/NFV design:

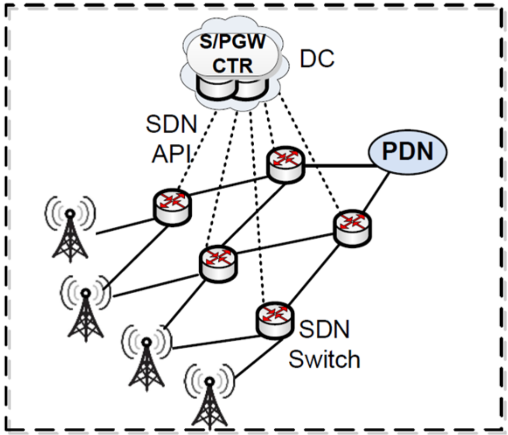

(b) SDN Core GW Architecture

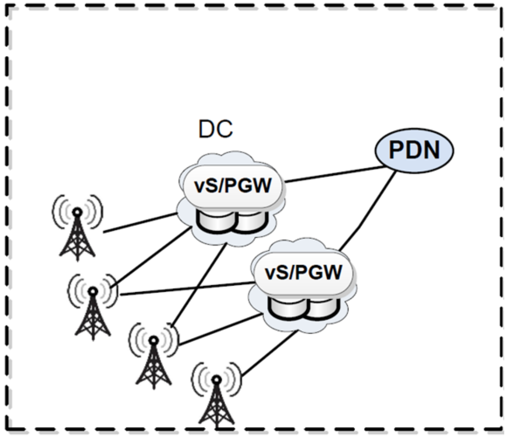

(c) NFV Core GW Architecture

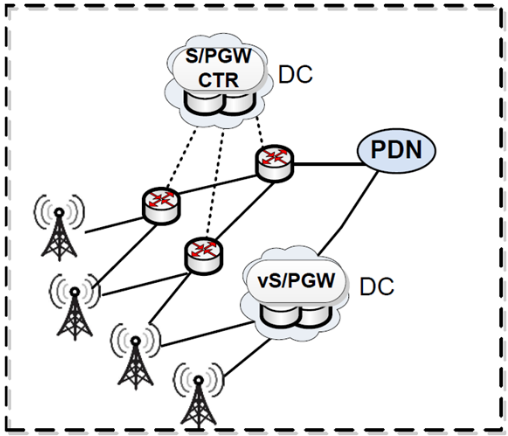

d) Mixed SDN and NFV Core GW Architecture

only control to cloud

control and data to cloud

Flexibility measure and evaluation setup

Use Case

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]

- SDN controllers, mobile VNFs, SDN switches and data centers placement
- constraints on data and control plane latency
- weights

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

[6] 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.

Evaluation parameters

Use Case

Parameters	Values
Data plane latencies to support	{5, 10, 15,…, 45, 50} ms
Control plane latencies to support	{5, 10, 15,…, 45, 50} ms
	*total: 10 * 10 = 100 possible solutions*
Data plane latency weight (α) Control plane latency weight (β)	α = 1 β = 1 α = 10 β = 1 α = 1 β = 10
Design choices	SDN, NFV, SDN/NFV
Data center deployment	Logically centralized (2 DCs) Distributed (8 DCs)
Topology	US

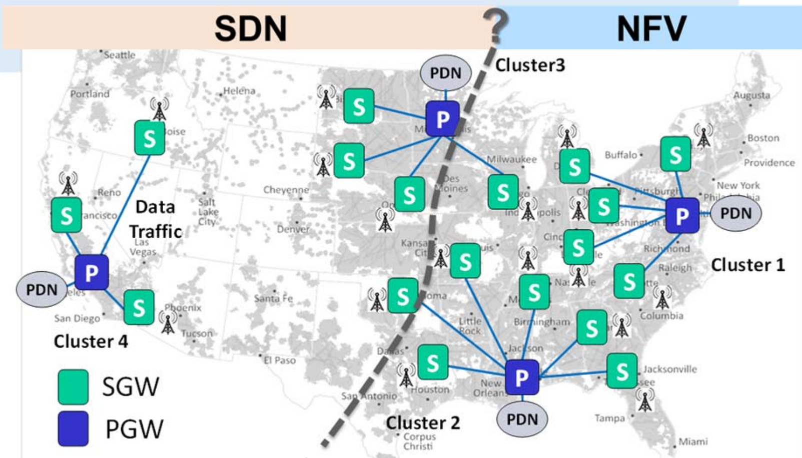

Example placement for mixed SDN/NFV design [6]

[6] 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]

Use Case

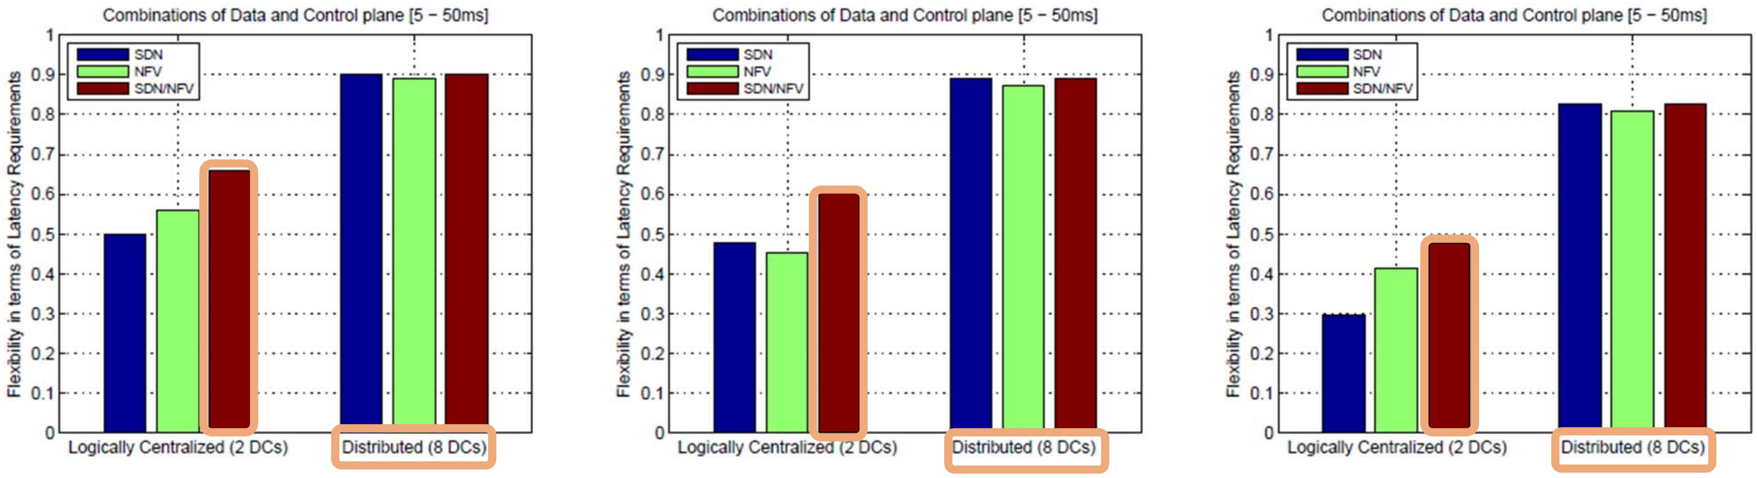

(a) equal weights for data and control latencies (b) weights biased by data latency (c) weights biased by control latency

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

[5] 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.

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***!

Quality of Flexibility – proposed definition

$$\varphi_T^{aspect}(S \mid \text{state i}) = \frac{|supported\ requests\ fulfilled\ in\ T|}{|possible\ requests|}$$

- fraction of the number of change requests that can be supported in a time interval T of all possible change requests
- T is small to capture system and request dynamics (sec to ms)

$$\varphi_{T->\propto}^{aspect}(S) = \frac{|supported\ requests|}{|possible\ requests|}$$

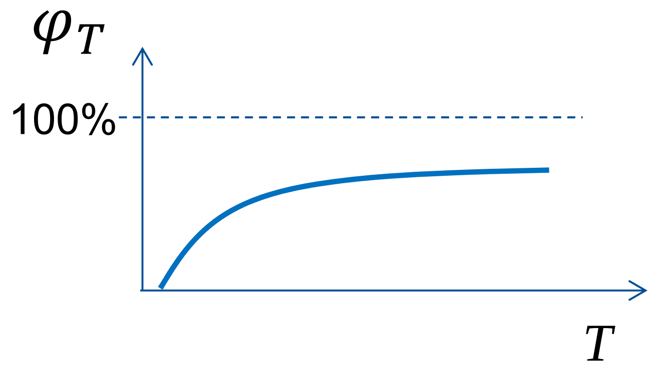

Nothing is for free: Cost of Flexibility

What are the costs of a design for flexibility?

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

Possible relationship (to be confirmed):

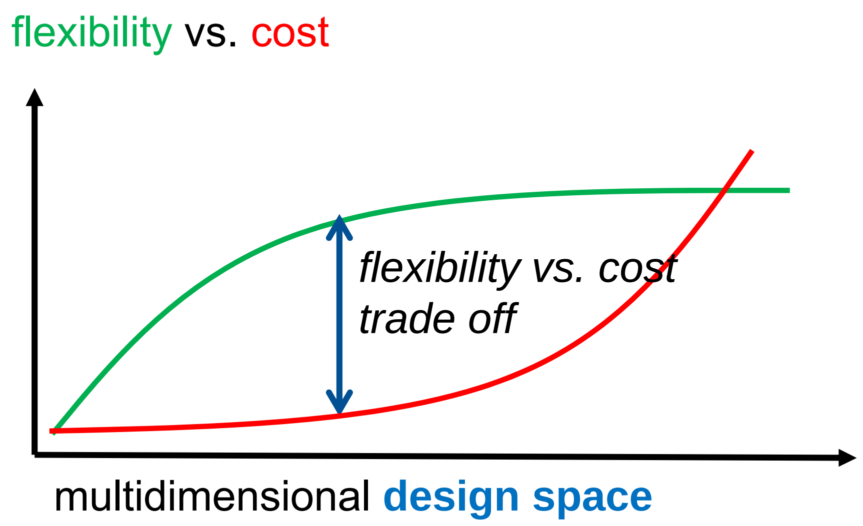

Use Case: Dynamic Controller Placement Problem

- Controller Placement Problem [7]:
  find optimal position for 1,…,n controllers given flow input
- Dynamic Controller Placement Problem:
  do the above for time varying input → controller migration/reconfiguration

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

[7] B. Heller, R. Sherwood, N. McKeown. The controller placement problem. HOTSDN 2012, Helsinki, Finland, Aug. 2012.

Simulation Results

Use Case

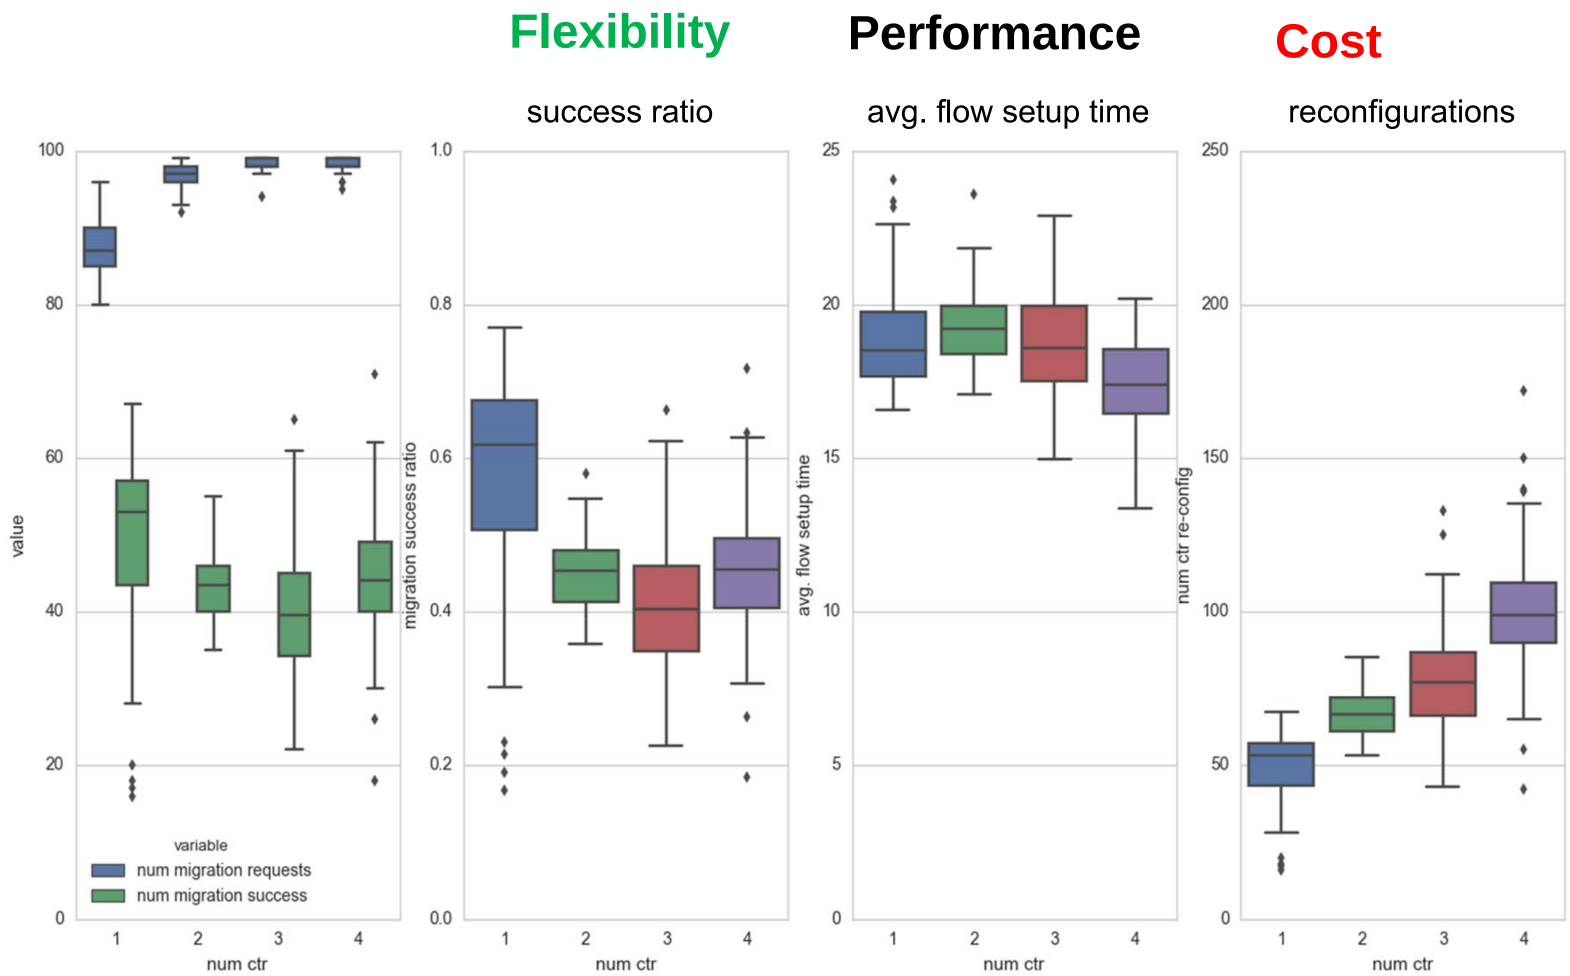

migration time threshold = 806 ms

1 controller has highest flexibility at low cost

But: performance is not good (flow setup time)

Simulation Results

Use Case

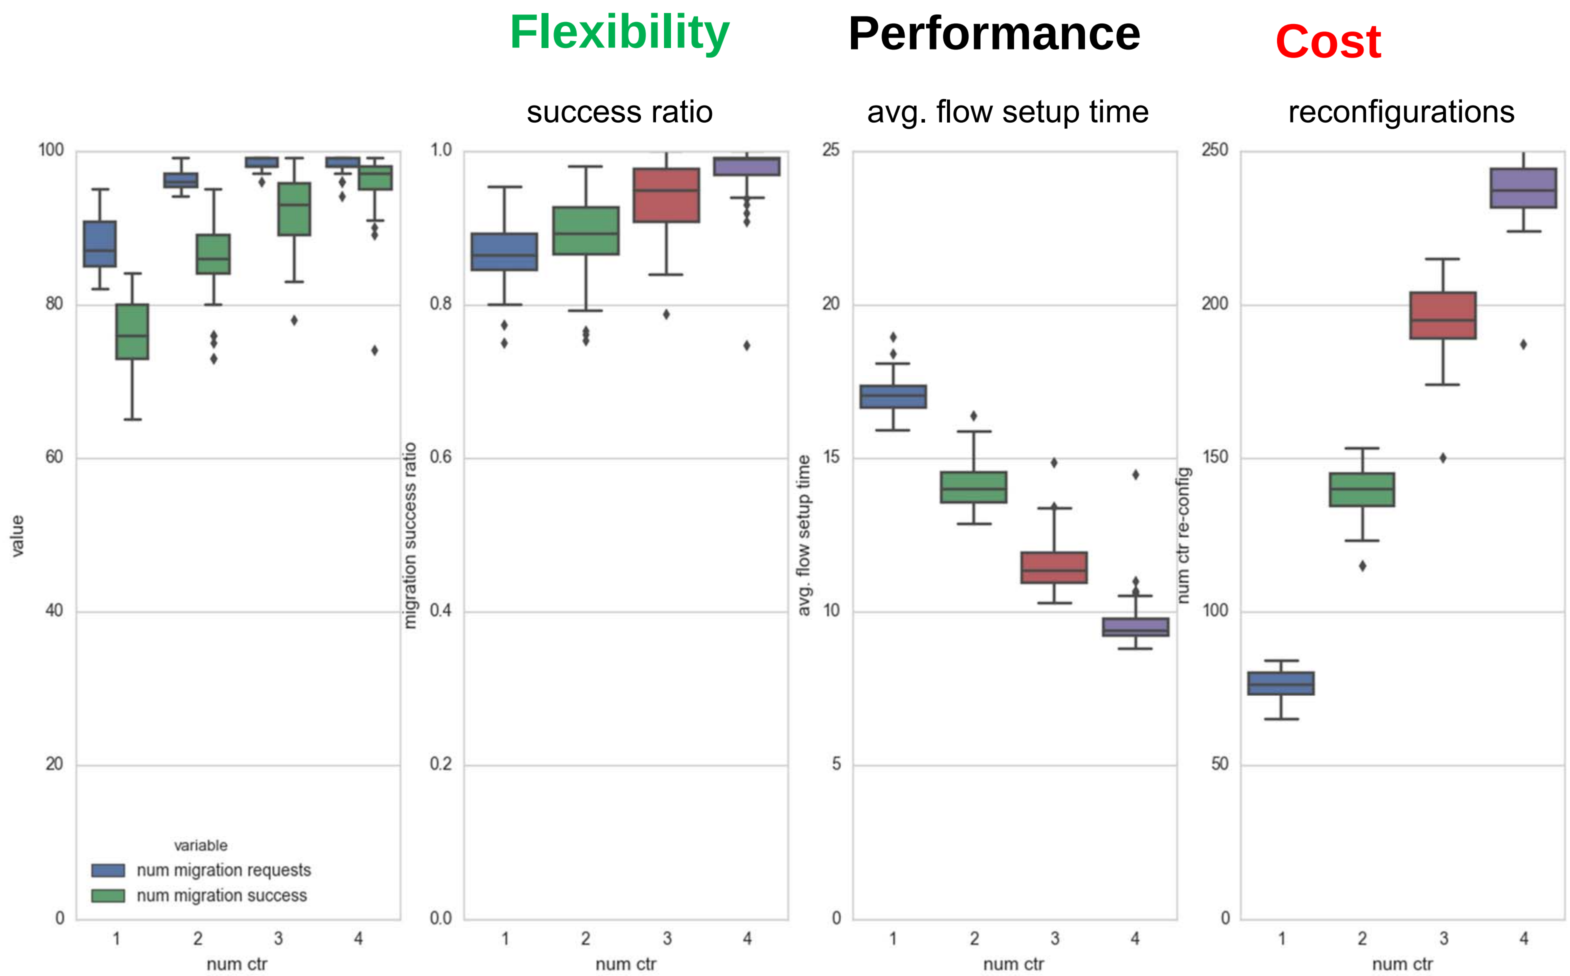

migration time threshold = 811 ms

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

from fundamental research to practice:

an implementation solution for flexibility

Designing for Flexibility: Network Slicing

- Why do we need network virtualization “slicing“?

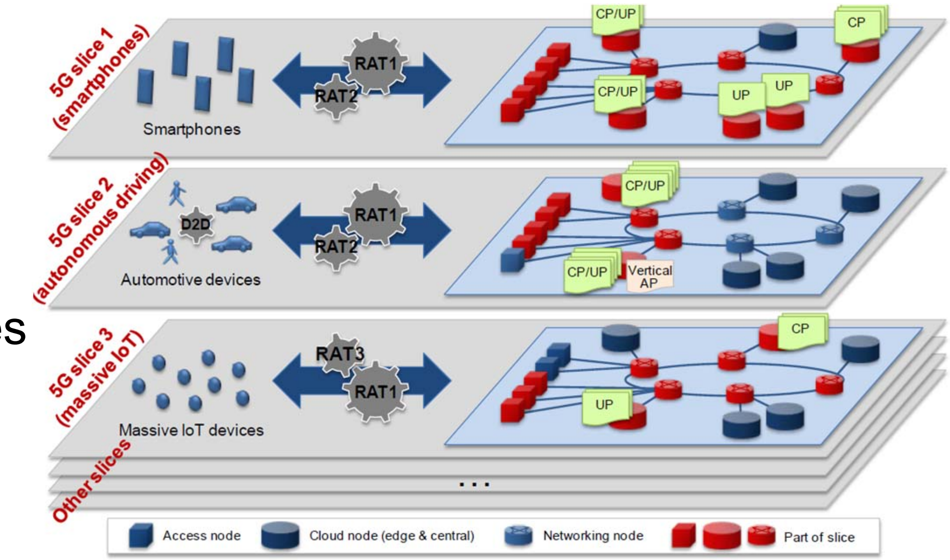

Source: NGMN 5G white paper

- NGMN 5G white paper [8]
  - logical virtual mobile network slices
  - reliable and on-demand slices

- METIS 5G system concept and technology roadmap [9]
  - application and service differentiation
  - logical virtual mobile network slices
  - heterogenous and dynamic slices

[8] 5G Initiative Team, NGMN 5G White Paper, 2015, https://www.ngmn.org/uploads/media/NGMN-5G-White-Paper-V1-0.pdf
[9] 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

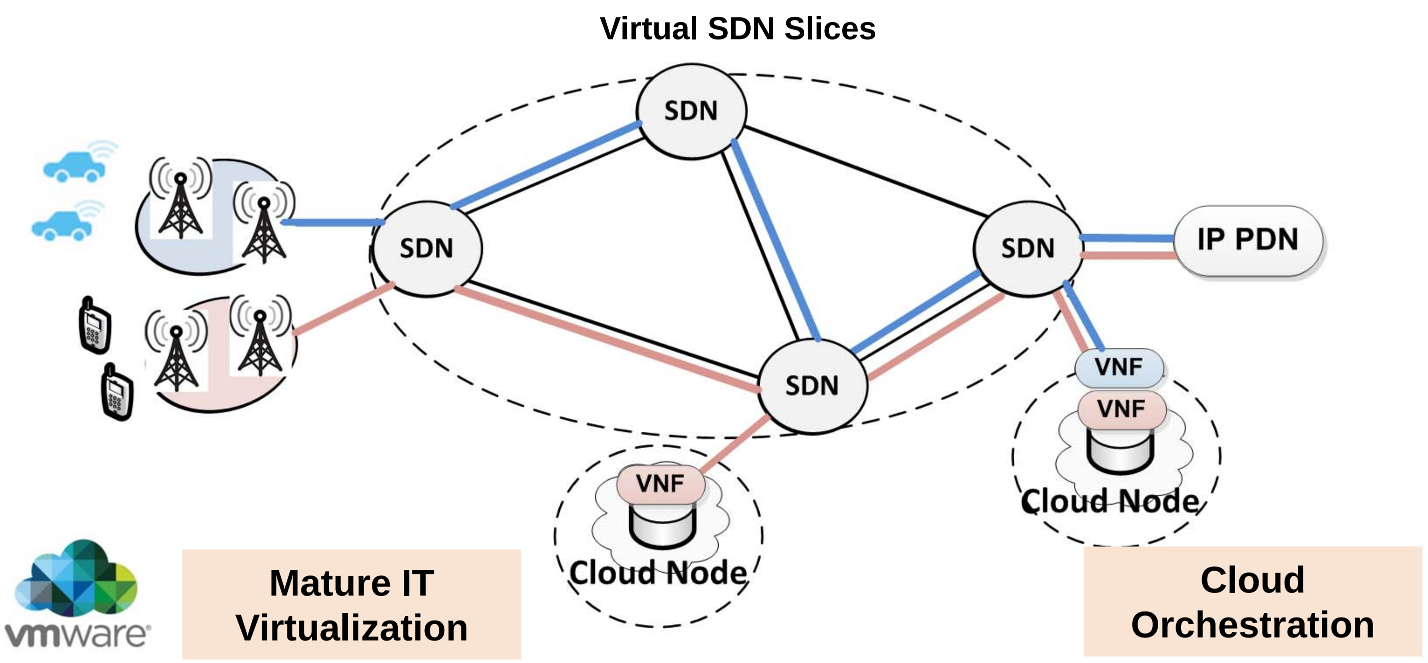

SDN Virtualization Overview

- How to achieve slicing for SDN networks?
  - SDN virtualization layer, i.e., SDN hypervisors
  - e.g. FlowVisor [10], OpenVirteX [11]

- What should an SDN hypervisor do?

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

- … in a most flexible way

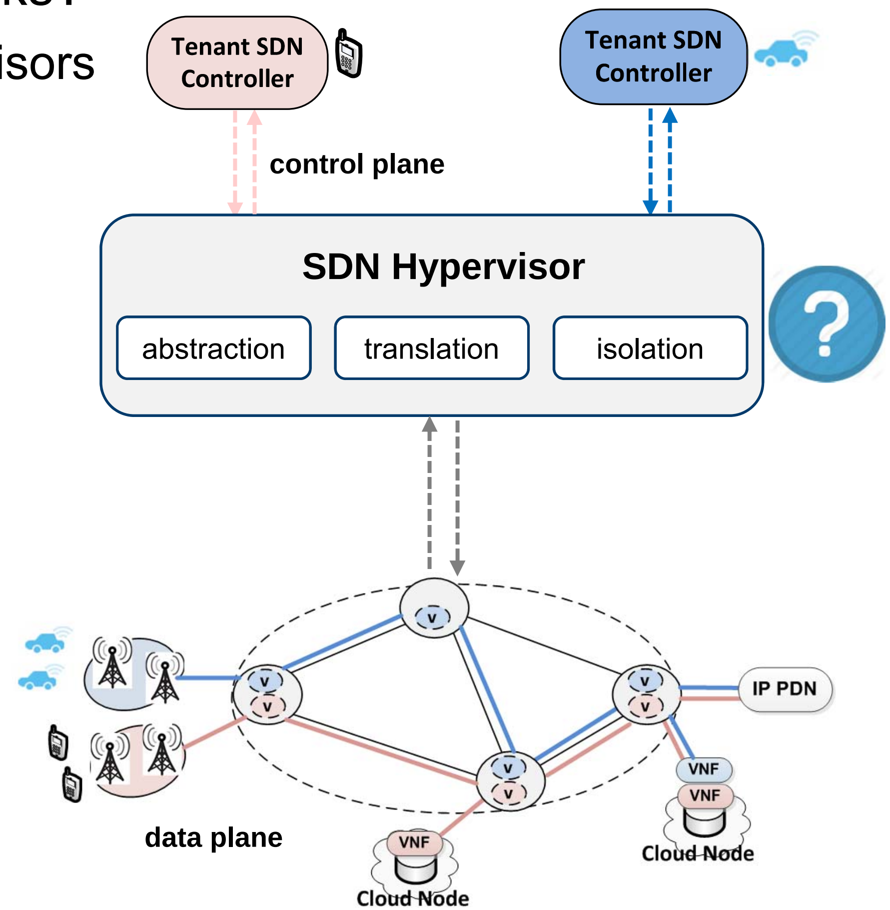

[10] R. Sherwood et al., Carving research slices out of your production networks with OpenFlow, ACM CCR, 2010

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

State-of-the-art Limitations [12]

- 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

[12] 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 [13,14]
  - 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

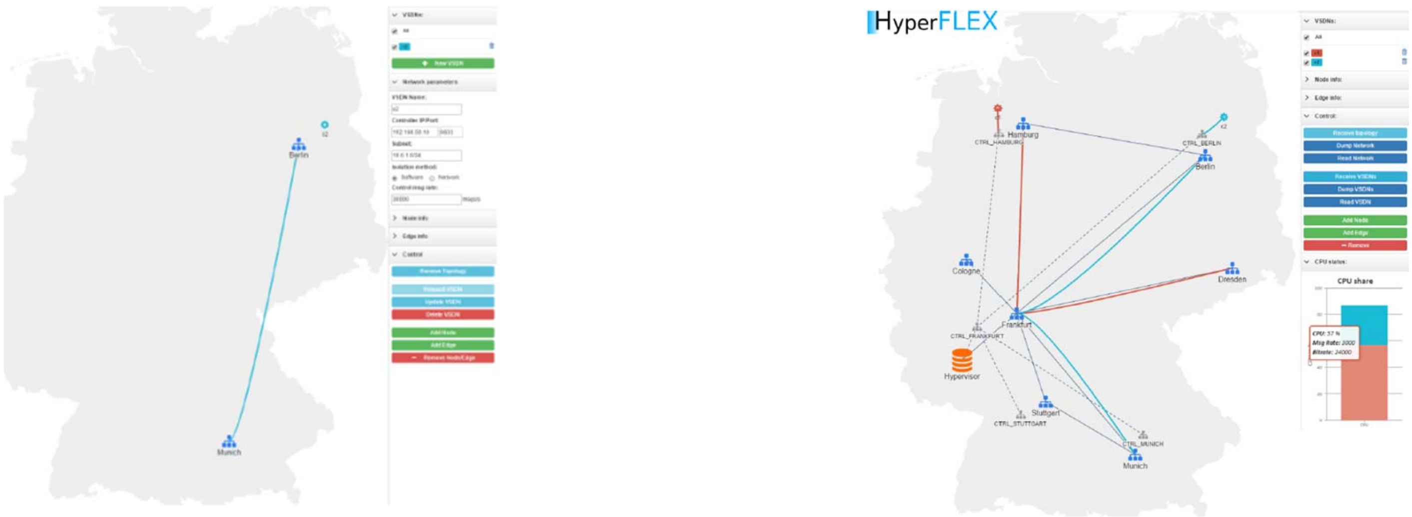

(a) Tenant View **(b) HyperFlex View**

[13] A. Blenk, A. Basta, W. Kellerer, HyperFlex: An SDN Virtualization Architecture with Flexible Hypervisor Function Allocation. IFIP/IEEE IM, pp. 397-405, 2015
[14] A. Basta, A. Blenk, Y. Lai, W. Kellerer, HyperFlex: Demonstrating Control-plane Isolation for Virtual Software-Defined Networks. IFIP/IEEE IM, pp. 1163-1164, 2015
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 [13,14,15]
  - monitor the performance of the running hypervisors, e.g., CPU
  - monitor the performance of the SDN slices
    - control plane latency
    - control plane loss rate

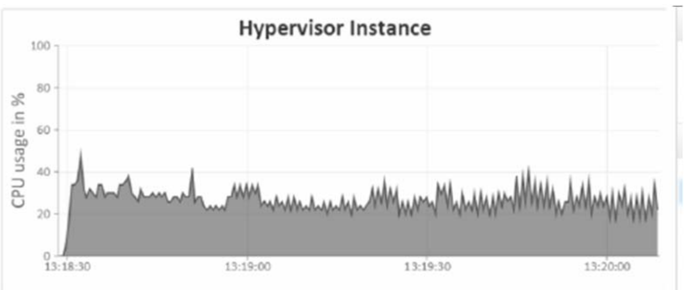

(a) Hypervisor performance

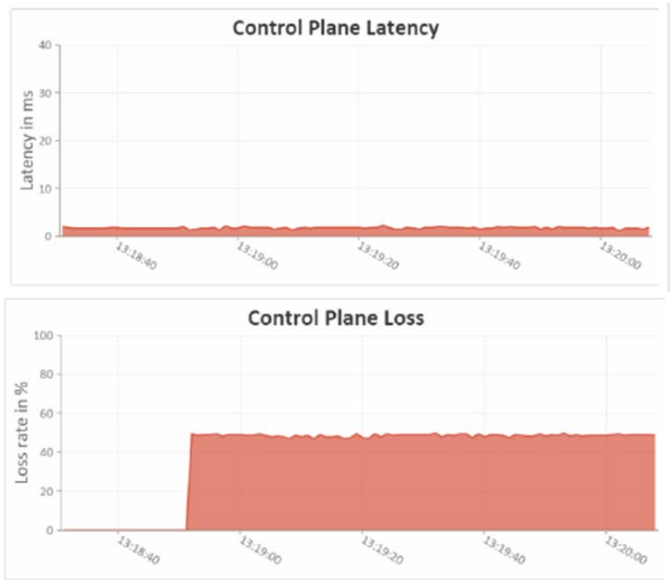

(b) Tenant control performance

[15] C. Sieber, A. Basta, A. Blenk, W. Kellerer, Online Resource Mapping for SDN Network Hypervisors using Machine Learning, IEEE NetSoft, Seoul, Juni 2016.

HyperFlex Features

- Dynamic Deployment “Orchestration” [16]
  - 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

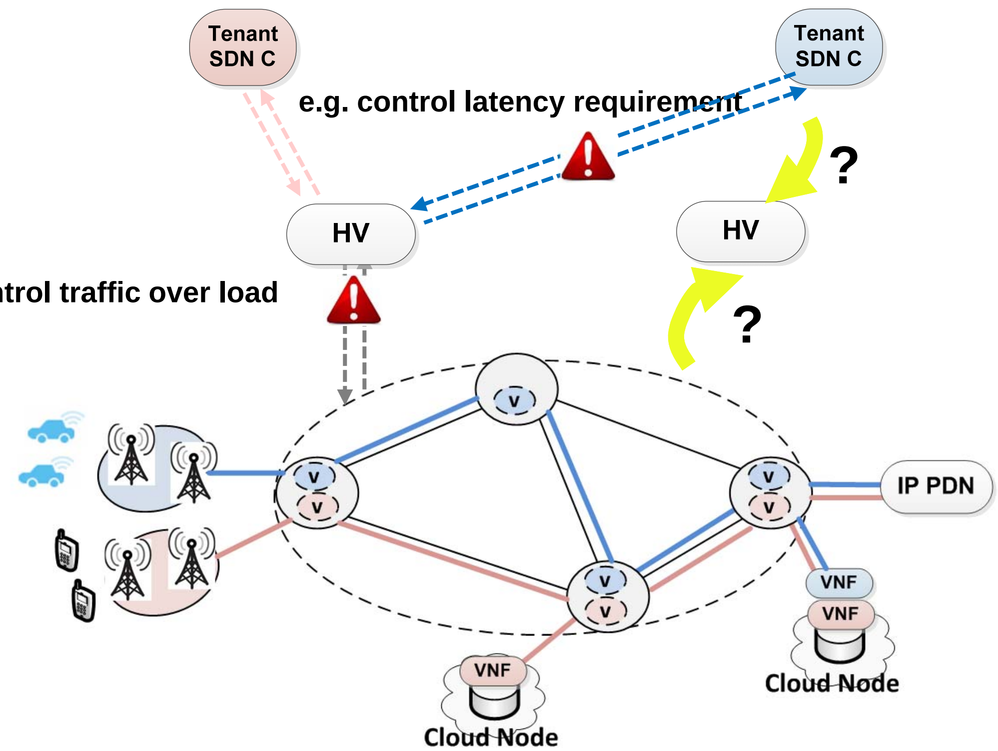

[16] A. Basta, A. Blenk, H. Belhaj, W. Kellerer, Towards a Dynamic SDN Virtualization Layer: A Control Plane Migration Protocol, 11th International Conference on Network and Service Management (CNSM), pp. 354-359, 2015

HyperFlex Features

- Algorithms/Models for SDN Network Hypervisor Placement Problem [17,18]
  - Control plane performance, e.g., latency, affects data plane performance
  - SDN hypervisors can add to the control plane latency
  - There is a need to investigate the optimal design and placement of hypervisors

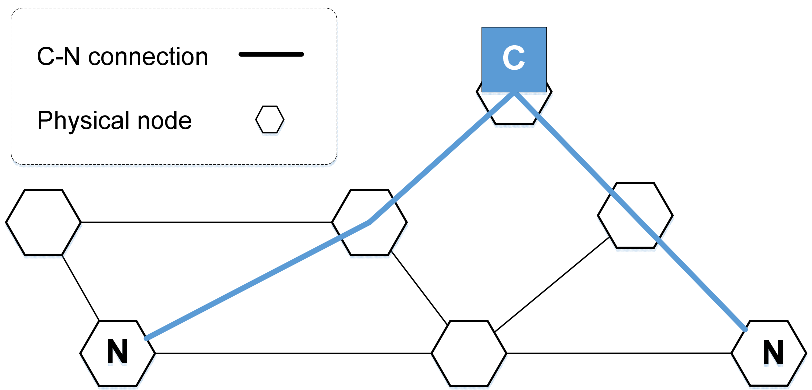

control connection for non-virtual SDN network

control latency requirement?

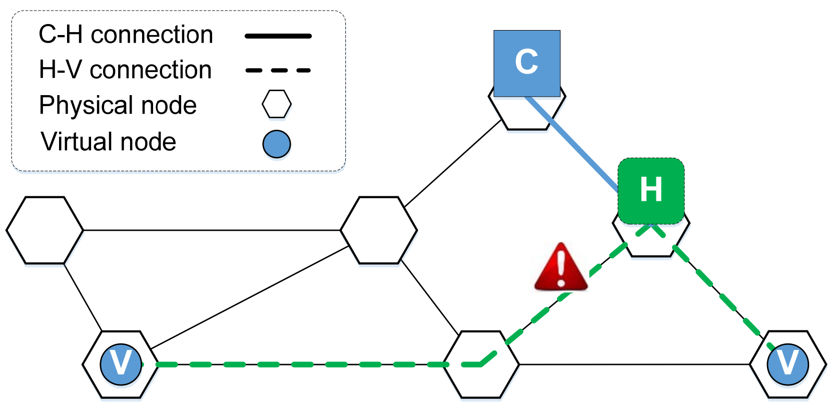

control connection for virtual SDN network

[17] A. Blenk, A. Basta, J. Zerwas, W. Kellerer, Pairing SDN with Network Virtualization; The Hypervisor Placement Problem, IEEE NFV-SDN Conference, pp. 198-204, 2015

[18] 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

Our flexibility testbed (SDN switches)
www.lkn.ei.tum.de

Rack 1

10x DELL switches

(Bare metal switches)
8 DELL S3048 - 48x1G,4*10G
2 DELL S4048 - 48x10G
- FTOS (OF 1.3)
- Cumulus Linux,
- Switch Light OS
  (big switch)

2x HP switches

Rack 2

2x Pronto switches

2x Pica8 switches

2x NEC switches

2x Net Optics traffic monitors

2x Sun Fire Servers

Spirent TestCenter C1
Provides layer 2-7 router, switch, application and security test solutions.
Supports line-rate 1GE or 10GE test ports.

Conclusion

Key Takeaways

- Network research is faced with **new requirements** from emerging networked industries
- These include **flexibility**
- Need for: new **flexible concepts** (e.g. SDNQoS, HyperFlex)
- Need for: a **measure** to compare flexibility among designs