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

150 Jahre culture of excellence

How to Measure Network Flexibility?
A Proposal for Evaluating Softwarized Networks

Wolfgang Kellerer

Technical University of Munich

KTH

Stockholm, Sweden

June 18-22, 2018

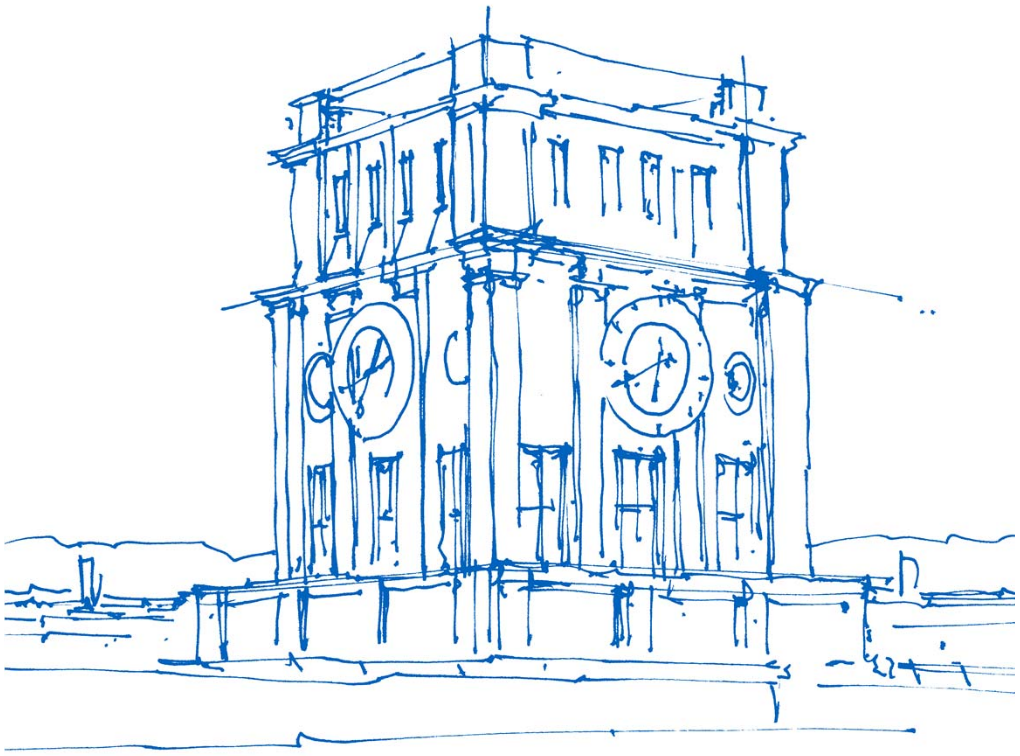

Uhrenturm der TUM

with PhD students and postdocs
A. Basta, A. Blenk, P. Babarczi, M. He
P. Kalmbach, M. Klügel, A. Martinez,
J. Zerwas

FlexNets

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

Research @ Chair of Communication Networks

Flexible, dynamic and resource-efficient internetworking

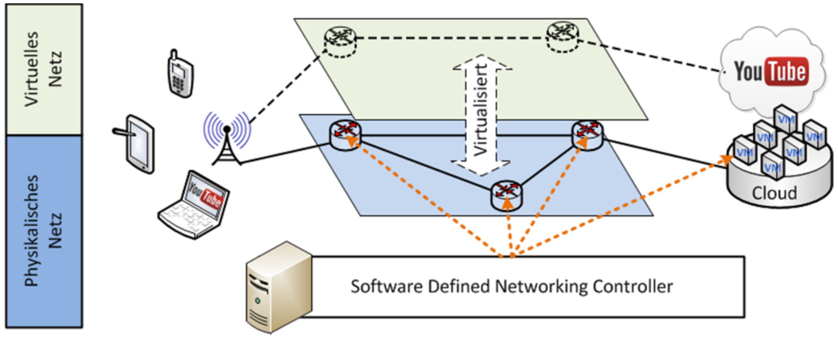

5G mobile networks for reliable, energy efficient and low latency machine-to-machine comunication

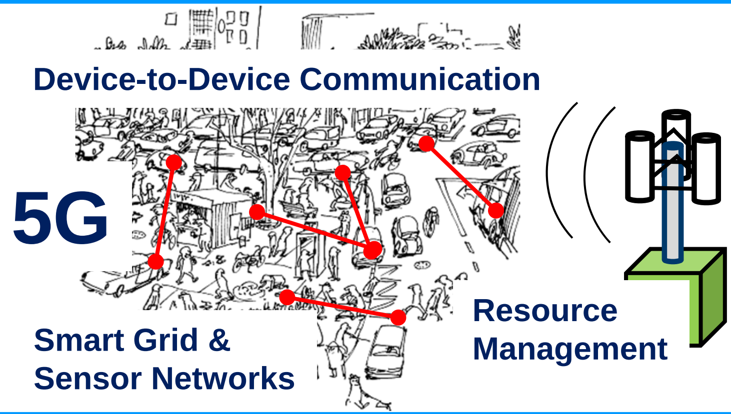

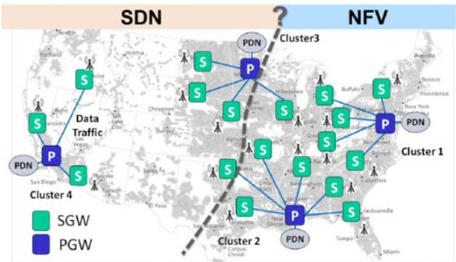

Techno economic analysis & resilient network planning

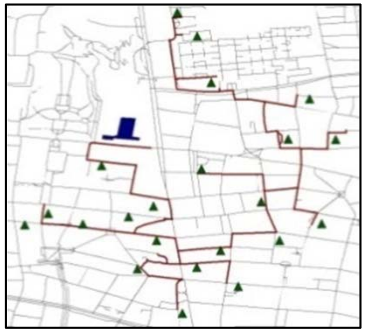

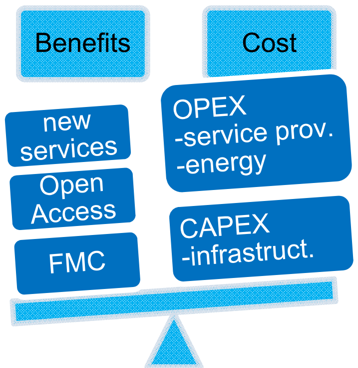

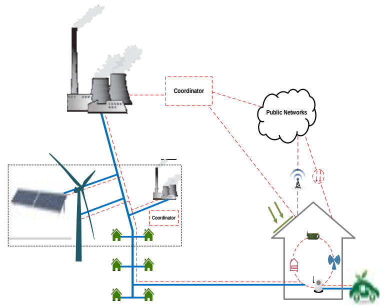

before we start measuring

Why is flexibility important?

- Evolution tells us that the more flexible species can better survive
- What about networks? Will they survive?

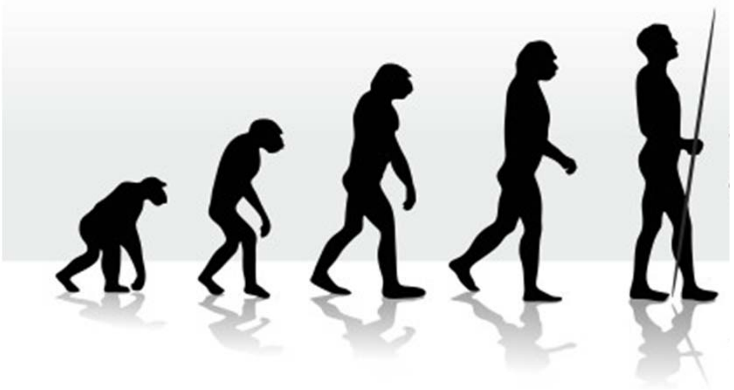

Image source: http://www.paleoplan.com

- So far less explicitly addressed: ***flexibility*** and hence ***adaptation***

- Today, we will …

  … present our definition of network flexibility and a flexibility measure, …

  … give concrete use cases of how to apply …

  … and show ML and *empowerment* methods to speed up adaptations

What does literature say about flexibility?

- Flexibility is gaining increasing **attention** and **importance**

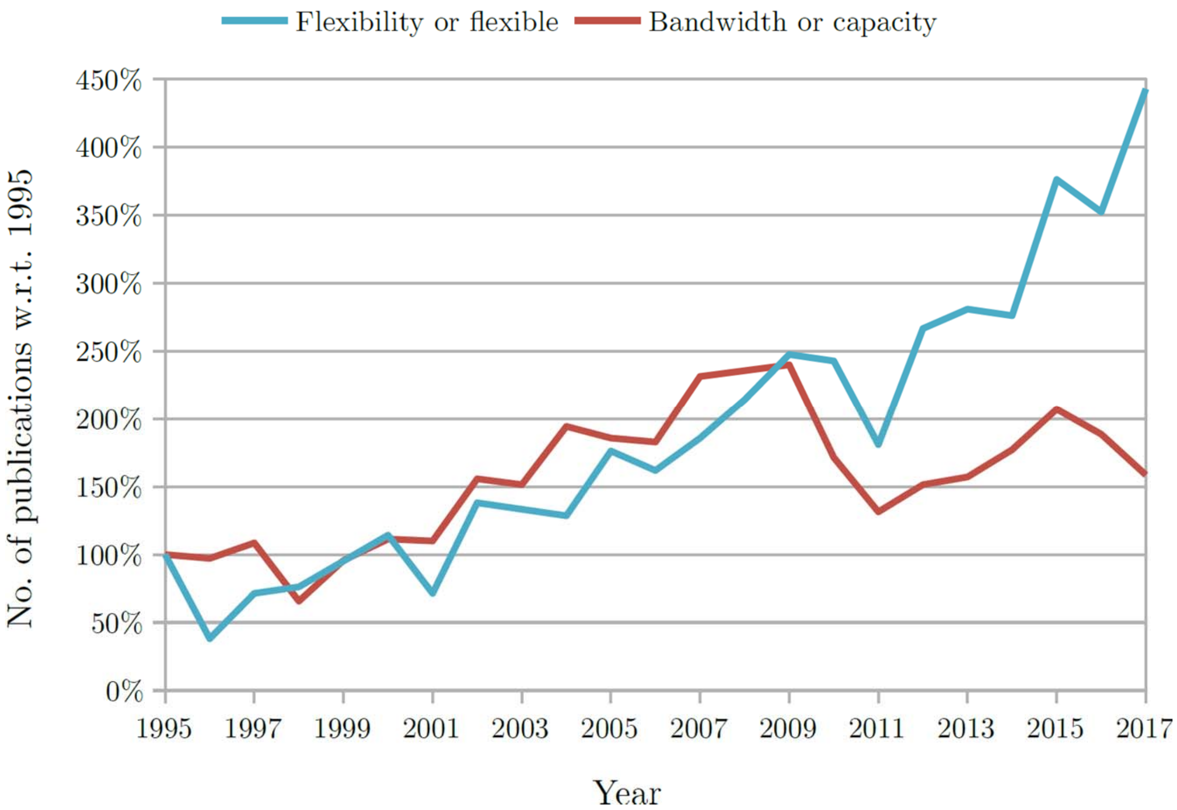

Evolution of the number of publications containing the words ”flexible” or ”flexibility” in contrast with those containing ”bandwidth” or ”capacity” in four major IEEE journals and magazines on communication, with respect to the number of publications in 1995.

Towards softwarized networks

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
→ reaction to dynamic changes hardly possible

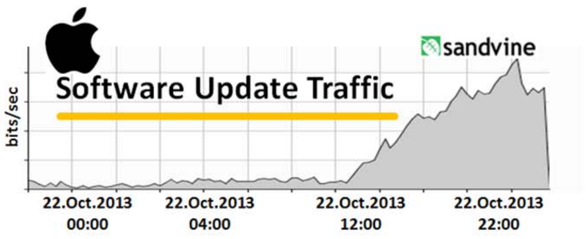

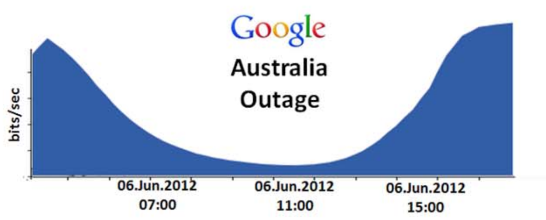

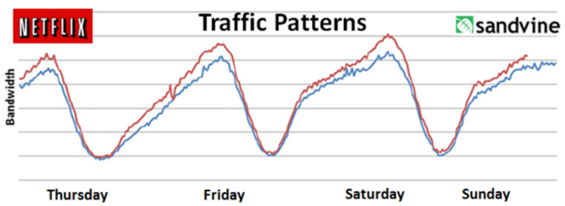

New concepts such as …

Network Virtualization (NV),

Network Function Virtualization (NFV), and

Software Defined Networking (SDN)

Softwarized Networks

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

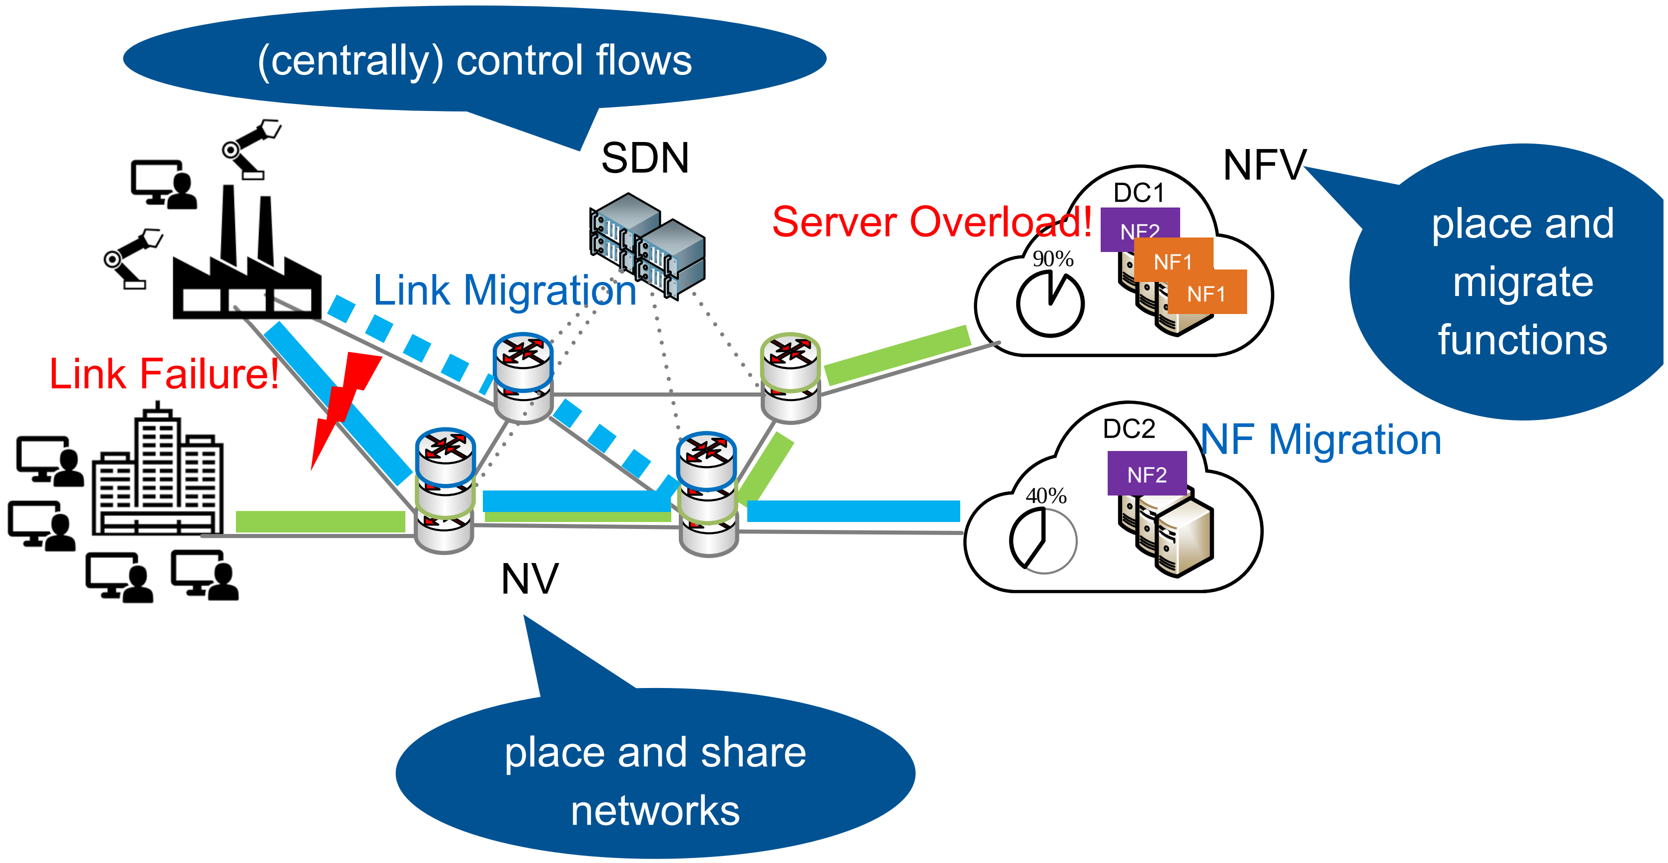

All problems solved?

- Are we fully flexible already?

- How far can we go? What is the optimal 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,…) in a *timely* manner via adaptation of resources (topology, capacity, ...) if needed

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

erc

European Research Council

Why do we think flexibility analyis is important?

- Enables operators to **cover the future**!
  - react to regulatory changes and fast arrival of new technologies

- A key **decision factor** between network designs
  - can be a tie-breaking decisive advantage for a certain network design

- For research and development
  - which technical concepts lead to more flexibility in network design ?
    → **optimize** networks **for flexibility**
    → **design guidelines** for more flexible networks

- SoA: lack of a concrete definition and a quantitative analysis!

How to define network flexibility? [3]

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

System?

- communication network (topology, flows, node functions, resources) serving a certain objective (e.g. highly reliable communication)

Note: in most cases, flexibility is not the objective

Request?

- “new challenges”, e.g., new flows, new (virtual) topology or new latency requirements

Note: explicit list or via a distribution (e.g. flow arrivals)

So: *the more requests are supported the more flexible a system is?*

Time?

[3] W. Kellerer, A. Basta *et al.*, “How to measure network flexibility? A proposal for evaluating softwarized networks,” *IEEE ComMag*, 2018.

How to define network flexibility?

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

System?

- communication network (topology, flows, node functions, resources) serving a certain objective (e.g. highly reliable communication)

Request?

- "new challenges", e.g., new flows, new (virtual) topology or new latency requirements

Time?

- the network may need to adapt the state of the topology, flows, functions, or resources → it should meet a time constraint

[3] W. Kellerer, A. Basta *et al.*, "How to measure network flexibility? A proposal for evaluating softwarized networks," *IEEE ComMag*, 2018.

Flexibility is important

- network softwarization (SDN, NFV, NV) provides flexibility

Flexibility definition is important

- for a meaningful system analysis and comparison
- to design for flexibility

- *how to quantify?*

We need a measure

more ingredients needed

before we can come up with a measure, more context needs to be considered:

- is flexibility a simple singular measure?
- what trade-offs that come with flexibilty need to be considered?

Flexibility aspects

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

- similar to QoS: to be regarded by case

we propose: ***flexibility aspects*** [1, 3]

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

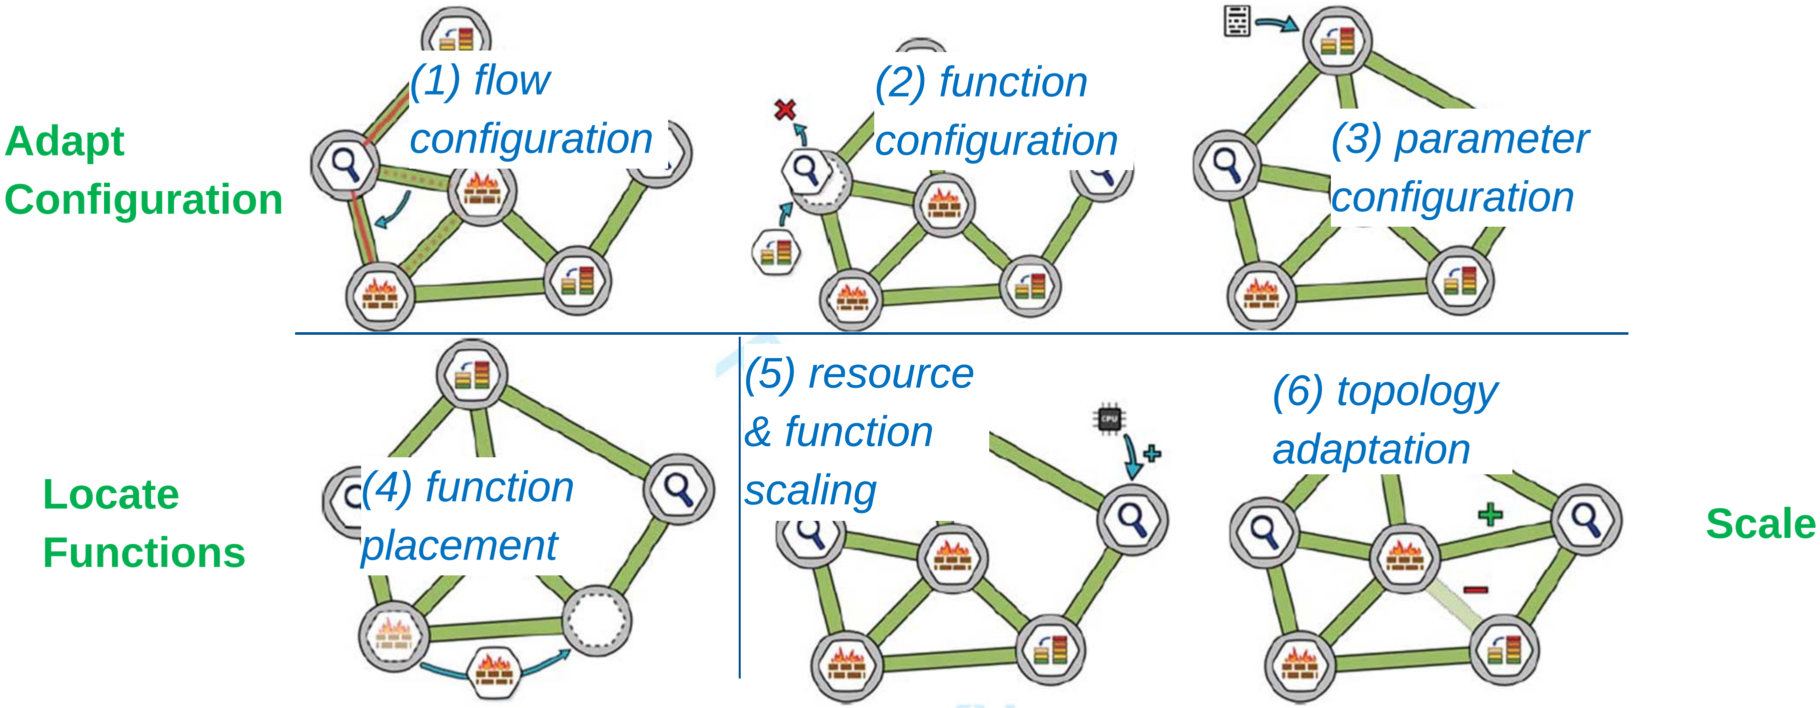

[3] W. Kellerer, A. Basta *et al.*, "How to measure network flexibility? A proposal for evaluating softwarized networks," *IEEE ComMag*, 2018.
[2] W. Kellerer, A. Basta, A. Blenk, Using a Flexibility Measure for Network Design Space Analysis of SDN and NFV, SWFAN'16, IEEE INFOCOM Workshop, April 2016.

Flexibility aspects to technolgies mapping

- SDN: is about **flow control**, also supports **network resources scaling**
- NFV: targets flexible **placement**, degrees of freedom in **configuration** and **function scaling**
- NV: targets flexible (virtual) **topologies**, also provides degrees of freedom for **configuration** and **scaling** of these (virtual) networks

Aspect (see Sec. III-B)	SDN	NFV	NV
Flow Configuration: flow steering	•		
Function Configuration: function programming		•	
Parameter Configuration: change function parameters		•	•
Function Placement: distribution, placement, chaining		•	•
Resource and Function Scaling: processing and storage capacity, number of fuctions	•	•	•
Topology Adaptation: (virtual) network adaptation			•

Cost vs. Flexibility

- Flexibility has to be evaluated against cost

- It is not clear if flexibility adds more **cost overhead**
- A flexible system can also achieve **cost savings** on the longer run

→ **trade-off** needs to be studied and evaluated

- We need to consider all different cost factors

Resources (CAPEX)	Operation (OPEX)	Adaptation/Migration	SLA
resource overhead	control, data plane throughput	synchronization overhead	fines
network complexity	control, data plane latency	configuration latency	flow interruption
software complexity	energy consumption	topology adaptation latency	network interruption

[3] W. Kellerer, A. Basta *et al.*, "How to measure network flexibility? A proposal for evaluating softwarized networks," *IEEE ComMag*, 2018.

- *how to quantify?*

We need a measure!

- Let‘s start with a qualitative measurement exercise first

Flexibility qualitative measure exercise

vs.

Source: Magazin.com

Fixed-set tool **Re-configurable tool box**

- Which tool is more flexible?
  - re-configuration shows more potential to be **more flexible**

- When can both exihbit the same flexibility?
  - maybe there is **no need to change** → probability of requests make a difference
  - maybe both cannot satsify my requests → **infeasible**

- When can the re-configurable tool be less flexible?
  - **adaptation time** → might make the re-configurable object not very useful

Flexibility Measure – Proposal

constraint on max.
adaptation Time

System

Request: “challenge“

$$\varphi_T{}^{aspect}(S) = \frac{|supported\ new\ requests\ within\ time\ constraint\ T|}{|given\ new\ requests|}$$

- fraction of the number of **new requests** that can be supported in a **time interval T** of all given new requests [3]

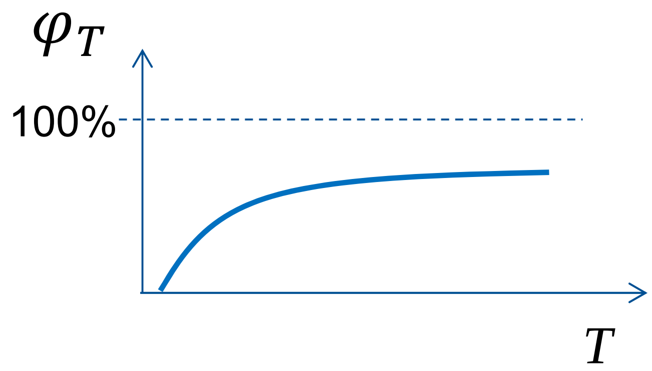

$$\varphi_{T->\infty}(S) = \frac{|supported\ new\ requests|}{|given\ new\ requests|}$$

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 Node Capacity (nc)
- BUT: migration time **cannot exceed** “1 hop“ (T)

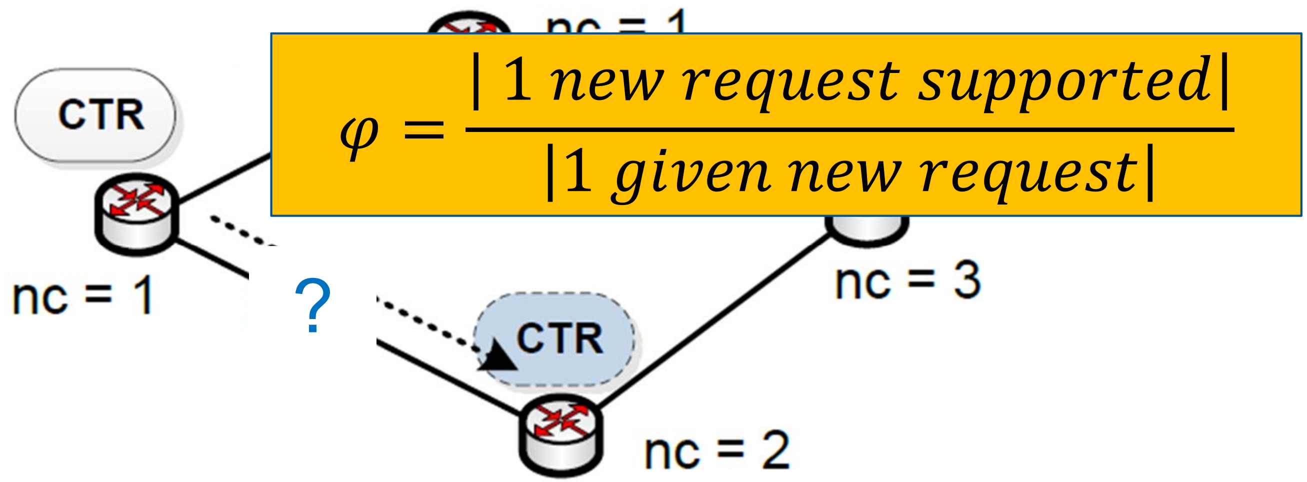

$$\varphi = \frac{|\,1\ new\ request\ supported\,|}{|\,1\ given\ new\ request\,|}$$

A simple illustration (2): more requests

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

We have a measure

- *How to validate this measure?*

*Flex*Nets Project Goals

Arsany Basta
Peter Babarczi
Andreas Blenk
Mu He
Patrick Kalmbach
Johannes Zerwas
Markus Klügel
Alberto Martinez Alba
Wolfgang Kellerer

- EU ERC Consolidator Grant (5 years)
- Project runs from September 1, 2015 – August 30, 2020
  - March 2018: mid project

- What should we have at the end?
  - a clear **definition of flexibility** in communication networks with a focus on softwarized networks
  - a **measure** for flexibility (and a procedure for how to use it)
  - a set of **design guidelines** for flexible network system designs

FlexNets Project Overall Approach

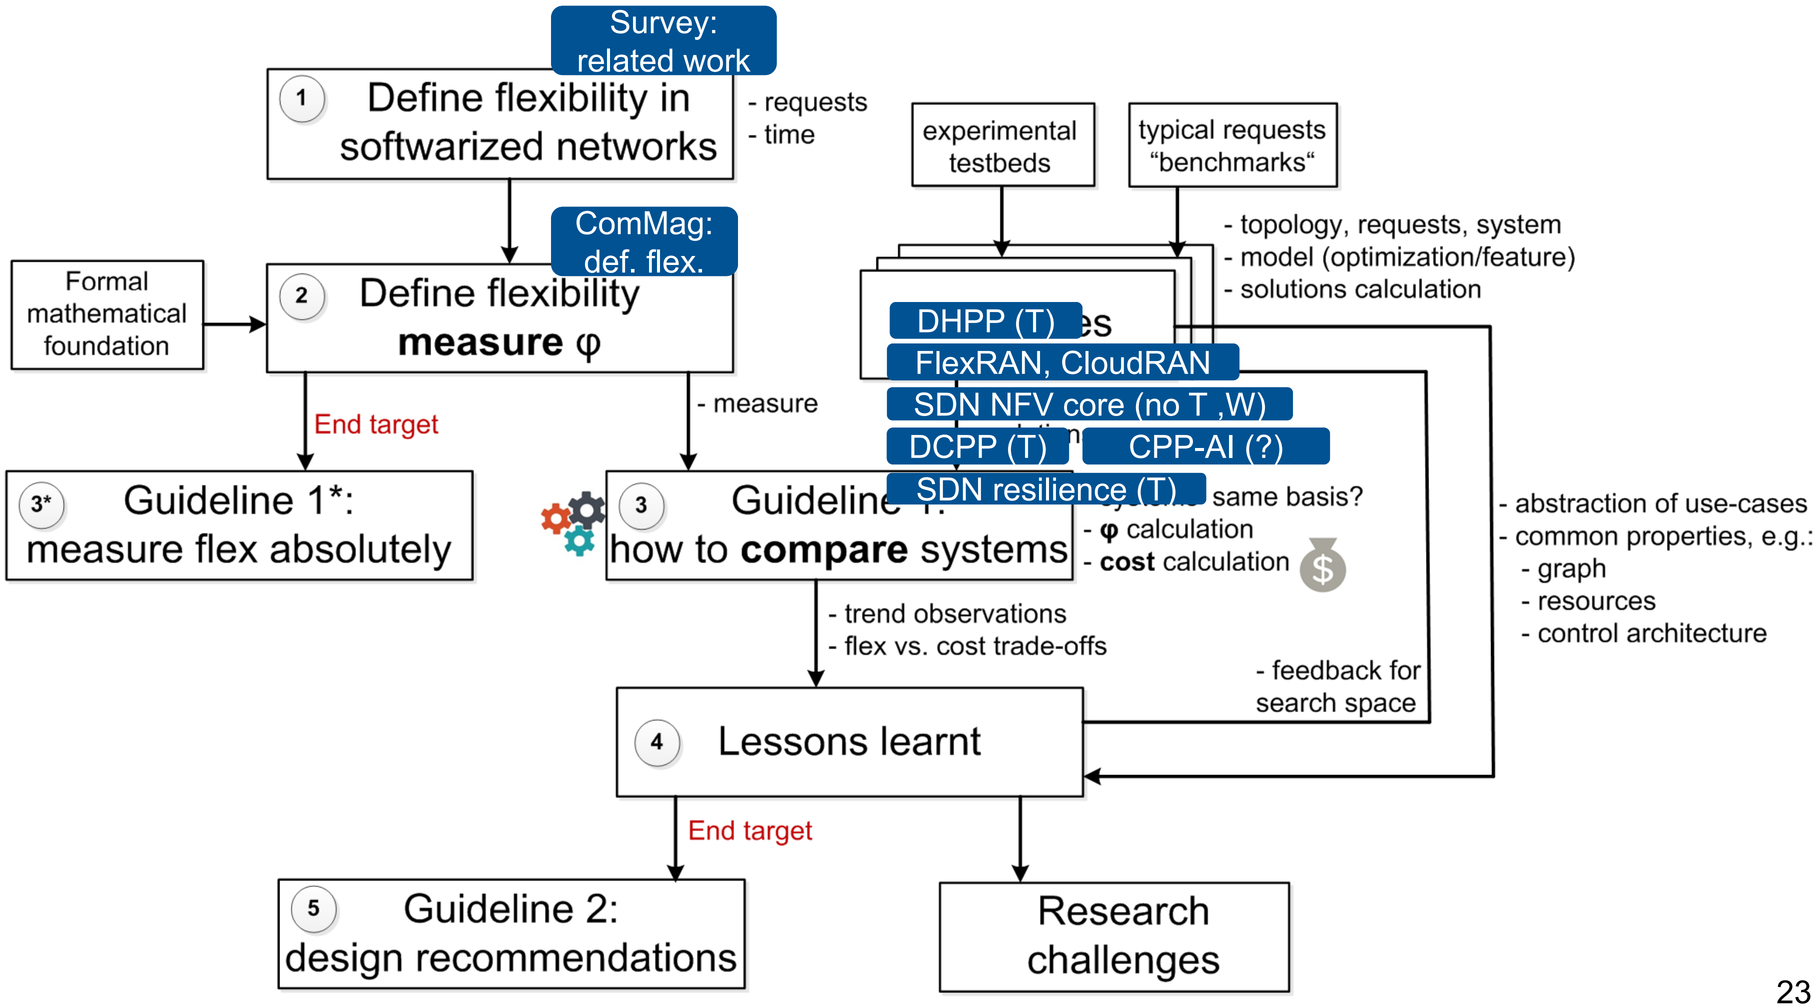

Case study 1: Dynamic Controller Placement

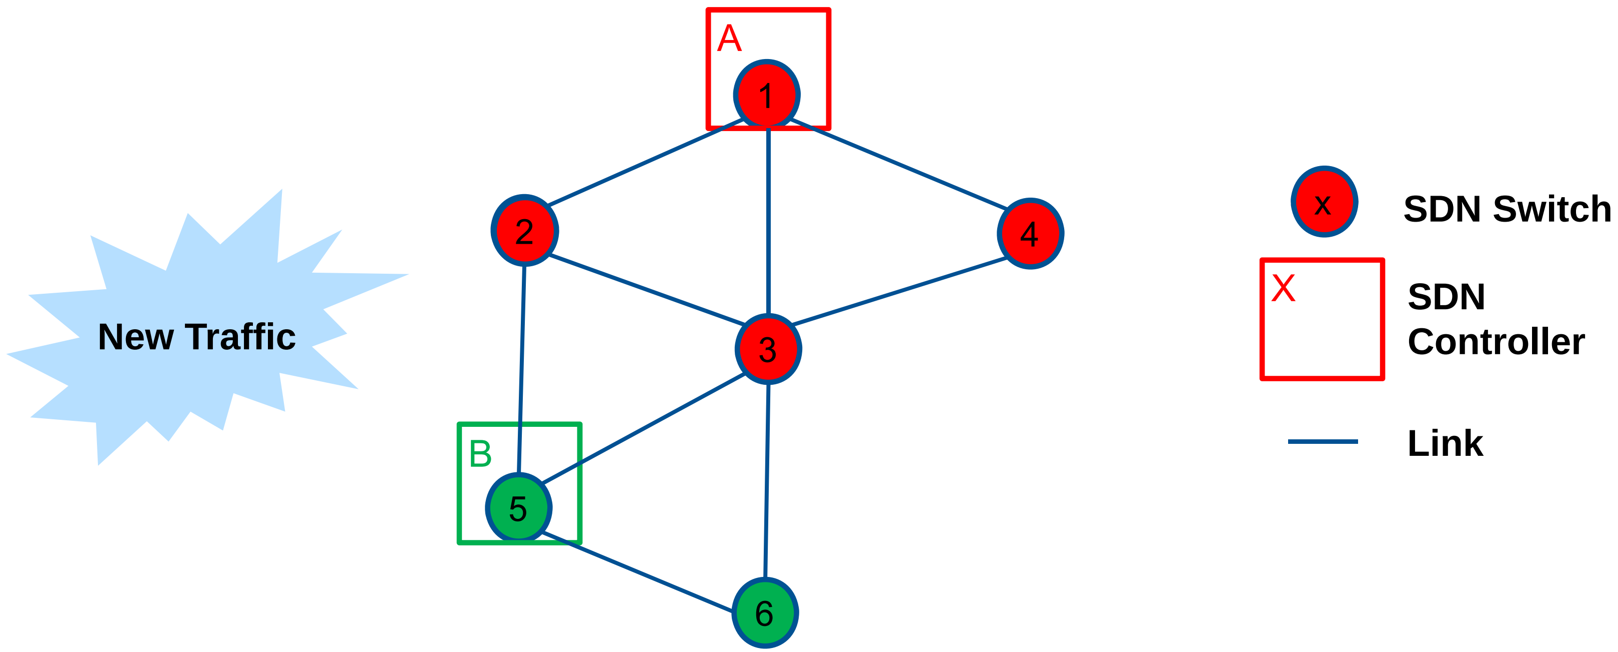

- Traffic fluctuations require control plane to adapt in order to achieve better control performance → *Dynamic Control Plane [4]*
  - SDN controller migration
  - SDN switch reassignment

Case Study	Flexibility Aspect	New Request	Flexibility Measure	System Objective	Cost in focus
dynamic SDN controller placement	function placement	new flow arrival (from distribution)	fraction of successful controller placements	control performance: (min. avg. flow setup time)	operation latency (OPEX): avg. flow setup time

Case study 1: Dynamic Controller Placement

- Application of the flexibility measure

Varying traffic flow profiles

max. adaptation time threshold
(will be varied)

$$\varphi_T(S) = \frac{|supported\ requests\ within\ T|}{|given\ new\ requests|}$$

SDN controller migration and switch reassignment can be done within T

- Flexibility → Migration Success Ratio
  - Calculate controller migration and switch reassignment time *T_migration*
  - If *T_migration* smaller than T → count as a supported request

Case study 1: Dynamic Controller Placement

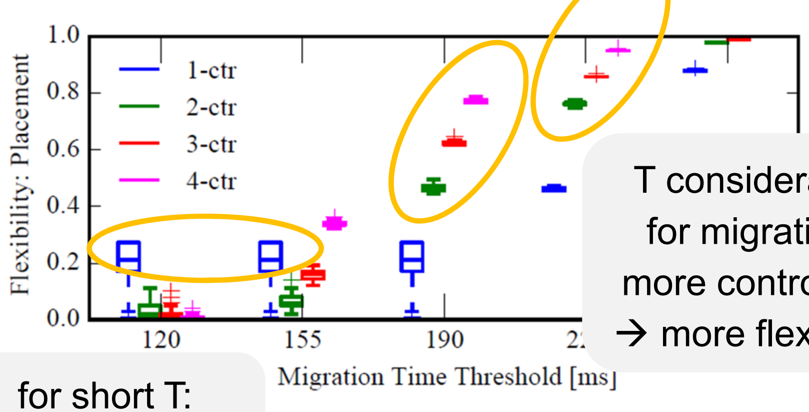

… in terms of successful control plane migration.

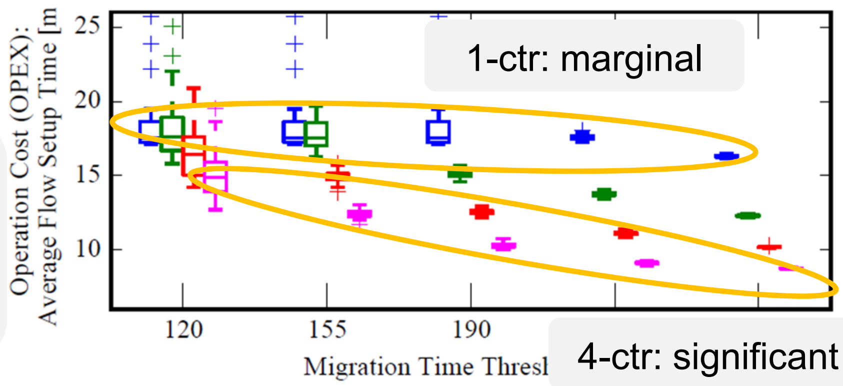

(b) Operation cost (OPEX) in terms of the average flow setup time.

intuitive

- **More controllers** (larger migration time threshold) → higher flexibility
- **Single controller** case: more flexible for **tight time threshold** as probability that single controller stays in optimal location is high

unexpected!

- 1 controller → **marginal** performance improvement vs. adaptation T
- 4 controllers → **significant** performance improvement vs. adaptation T
- However, if we consider **all cost factors**, we can reach a trade-off!

Case study 2: SDN Resilience

- Flexibility aspect of **flow configuration** for a **resilience** scenario in an SDN network under a given **recovery time** threshold T [3].
- Compare 3 systems: 1:1 protection vs 1+1 protection vs restoration
- New requests: all possible **single and dual link failures**
- Objective: system recovery
- Flexibility measure: fraction of recoverable failures

Case Study	Flexibility Aspect	New Request	Flexibility Measure	System Objective	Cost in focus
failure recovery in SDN (protection vs. restoration)	flow configuration	all possible single and dual failures	fraction of recoverable failures	system recovery: (single and dual failures)	resources overhead (CAPEX): node and link reservation

Case study 2: SDN Resilience

- **1:1 protection**
  - primary and backup paths pre-calculated
  - backup path is inactive
  - need switching time between primary and backup in case of a failure

- **1+1 protection**
  - primary and backup paths pre-calculated
  - primary and backup paths are both active
  - recovery time is almost instantaneous!

- **Restoration**
  - no backup path in advance
  - switch detect failure → controller informed → re-routes affected flows
  - recovery time is very critical

Case study 2: SDN Resilience

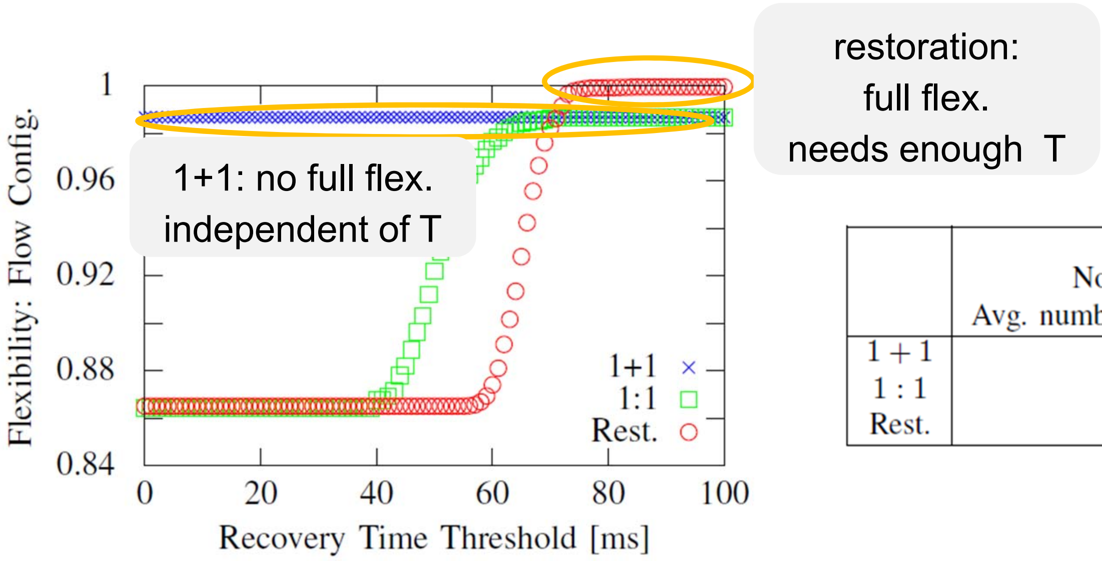

(a) Flexibility in terms of covered single and dual link failures.

	Resources Cost (CAPEX)	
	Node reservation: Avg. number of flow table entries	Link reservation: Number of required links
1 + 1	11.78	13038
1 : 1	11.78	13038
Rest.	5.05	5400

(b) System resources cost (CAPEX) in terms of nodes and links used for reservation.

intuitive

- 1+1 **can not** reach **full flexibility**
- However, 1+1 is obviously **independent** of **recovery time**

intuitive

- Restoration can cover **all failures** if given enough recovery time

- Protection imposes more than **2x capex overhead** than restoration
- Again, if we consider **all cost factors**, we can reach a trade-off!

Case study 3: FlexRAN (ongoing work)

- Radio Access Network plus SDN/NFV
  → unexplored flexibility
- our use case:
  coordinated scheduling
- initial results: PoC

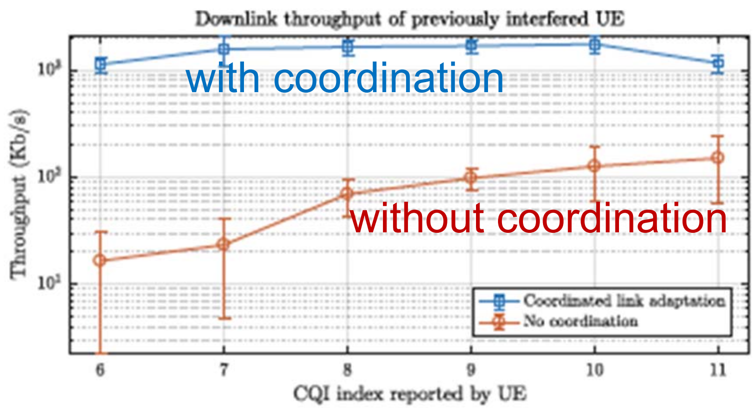

- next: quantify flexibility
  flexibility: ratio of successful handling of request

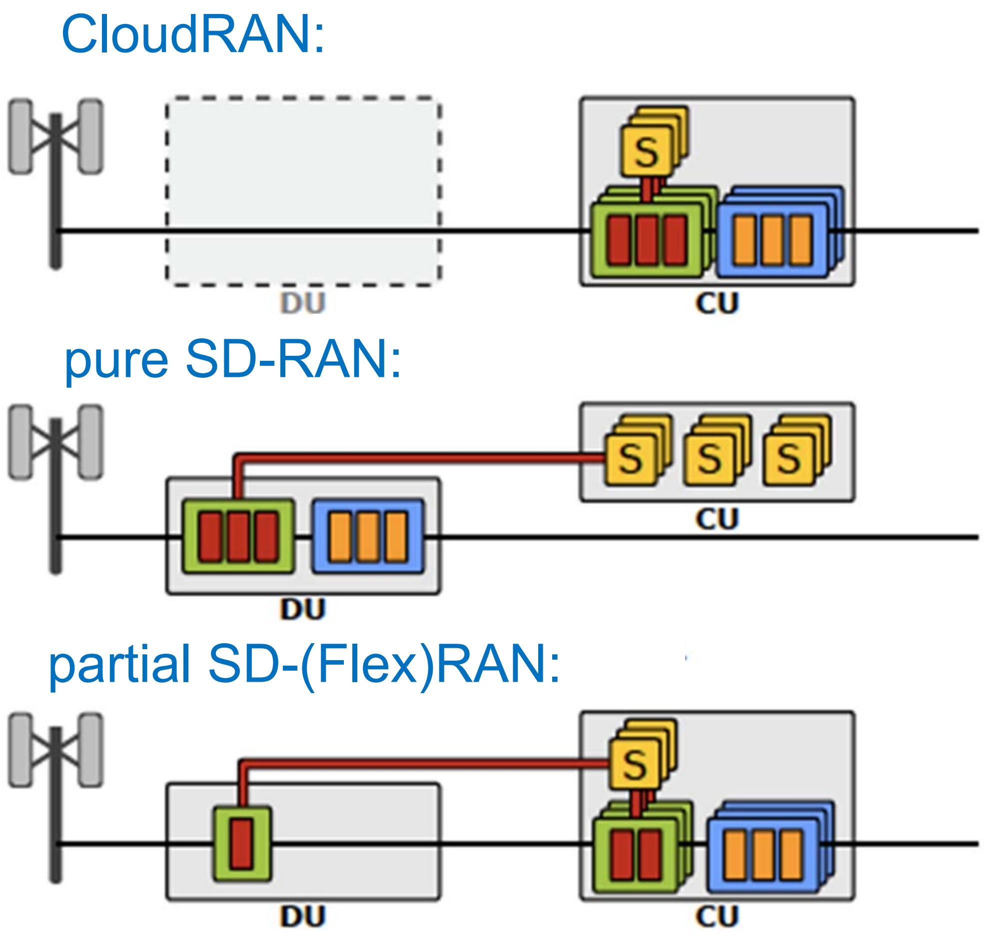

(c) Example of a partially centralized architecture (also SD-RAN).

Lessons learnt → design guidelines
(so far intuitive)

- A graph with a **high node degree** is highly flexible?!
- **Distributed cloud** data centers offer more flexibility to VNF chaining?!
- **Coordination** between cell schedulers provides more flexible resource allocation using information from multiple cells?!
- **Decoupling** data and control planes is highly flexible?!

… however, to what extent?

… at which point would these guidelines are not valid anymore?

step back and reconsider

- One way to **measure flexibility**
  so far only relatively between multiple systems
- Results can be **less intuitive** than one might think
- Flexibility tends to decrease **cost** but also comes at a cost
- Measure can be used to **design for flexibility**

Ongoing work: Optimize for Flexibility φ

Measure Phase

Optimize for performance metric (e.g. latency and throughtput), quantify flexibility value

Design Phase

Optimize for flexibility metric, decide system design parameters (e.g., bandwidth, # base stations, etc.)

Use Case: **Dynamic Controller Placement Problem**

- Requests: traffic profiles with target average flow setup time
- Objective: max. flexibility (success: # accomodated traffic profiles)
- Design parameters: # data centers and their locations

step back and reconsider

- One way to **measure flexibility**
  so far only relatively between multiple systems

- Results can be **less intuitive** than one might think

- Flexibility tends to decrease **cost** but also comes at a cost

- Measure can be used to **design for flexibility**

- **Design methods to improve flexibility** (based on AI)
  - adaptation time speedup through machine learning
  - *empower* a network

Speedup adaptation time

- **Adaptation time** is very important for a flexibility measure
- Adaptation examples:
  - Function placement, e.g., SDN controller
  - (re-)embedding of virtual networks/flows, e.g. for resilience

- How can we speedup?
- Yet another heuristic for a specific case study?

We propose:

- Keep your favourite optimization algorithms and
- ***Boost your network algorithm with ML preprocessing***

How can we boost the solving of the related optimization problems?

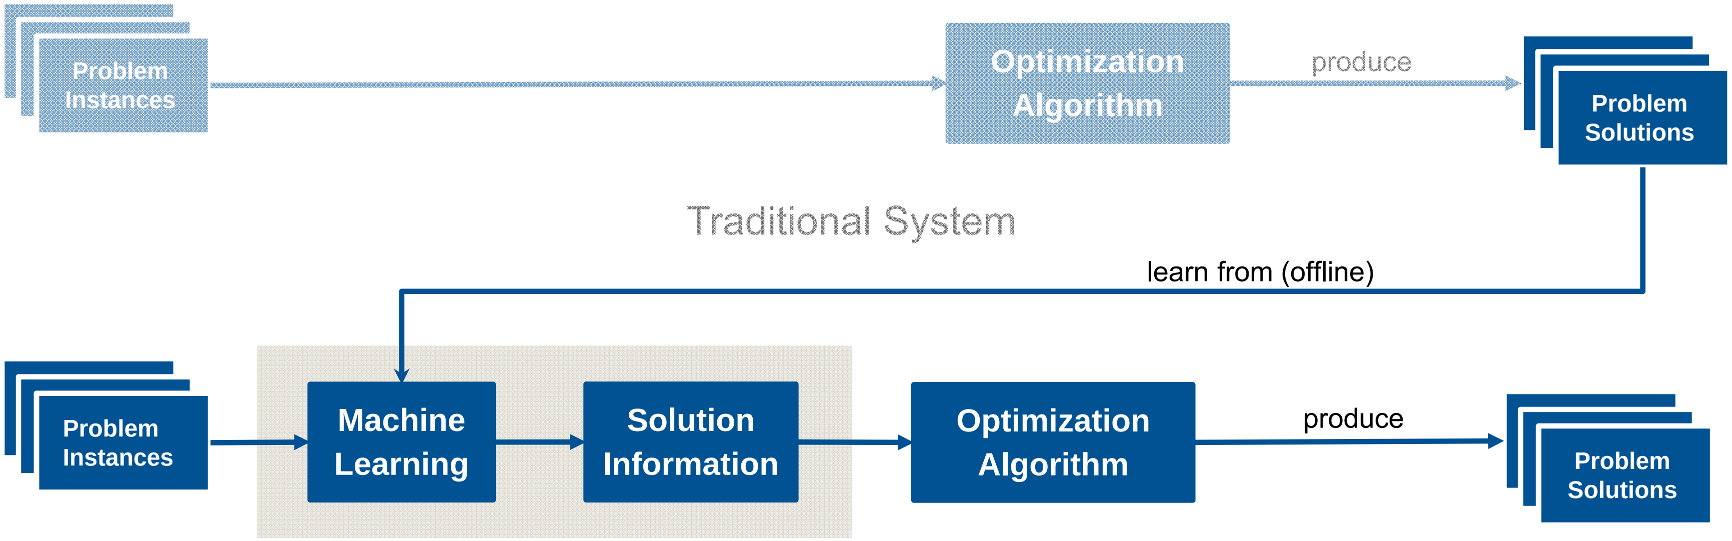

o'zapft is framework [5]

State-of-the-art: Neglects produced data!

Idea: Use problem/solution data generated by algorithms regularly solving problems

[5] A. Blenk, P. Kalmbach, S. Schmid, W. Kellerer: ***o'zapft is: Tap Your Network Algorithm's Big Data!*** ACM SIGCOMM 2017 Wrksp. on Big Data Analytics and Machine Learning for Data Communication Networks (Big-DAMA), 2017.

Data Available: P. Kalmbach, J. Zerwas, M. Manhart, A. Blenk, S. Schmid, W. Kellerer. Data on "o'zapft is Tap Your Network Algorithm's Big Data!",2017 https://doi.org/10.14459/2017md1361589

Case Study: Predicting Acceptance Probabilities of VNE Requests

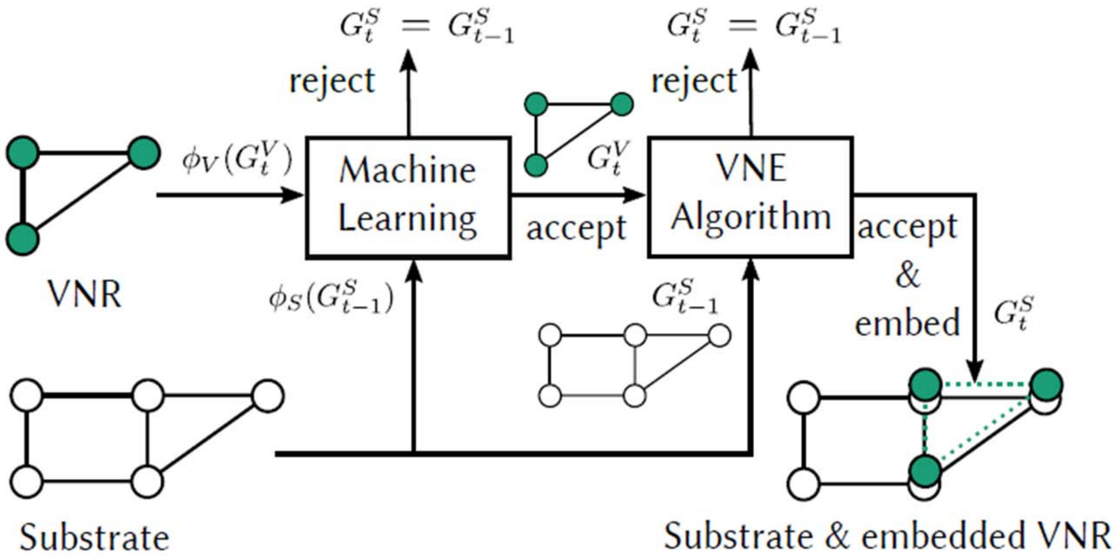

- Supervised learning: use data with accepted and rejected requests! Offline training!
- Recurrent neural network (RNN) for classification
- Filter infeasible and requests with unacceptable algorithm runtime ("no solution")

Can we speed-up optimal algorithms using admission control?

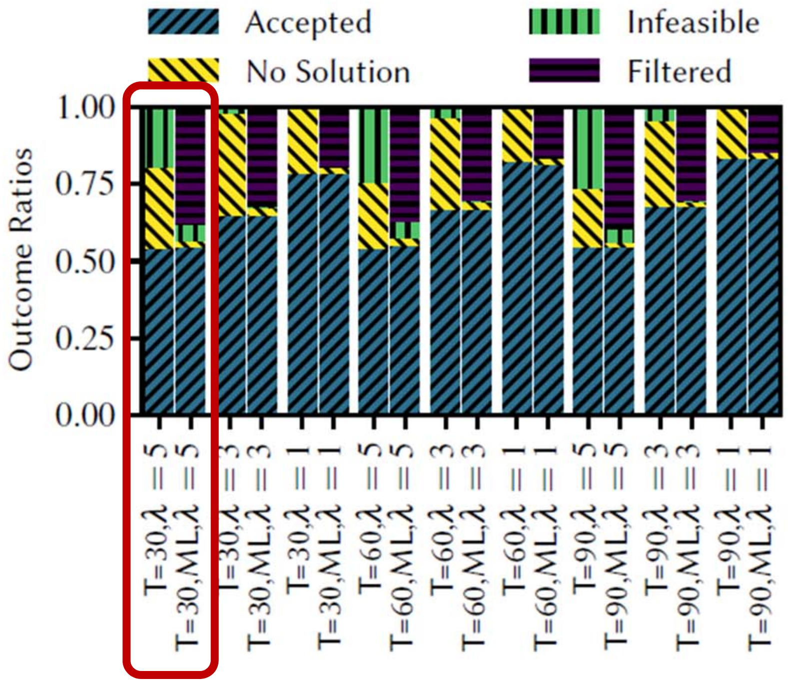

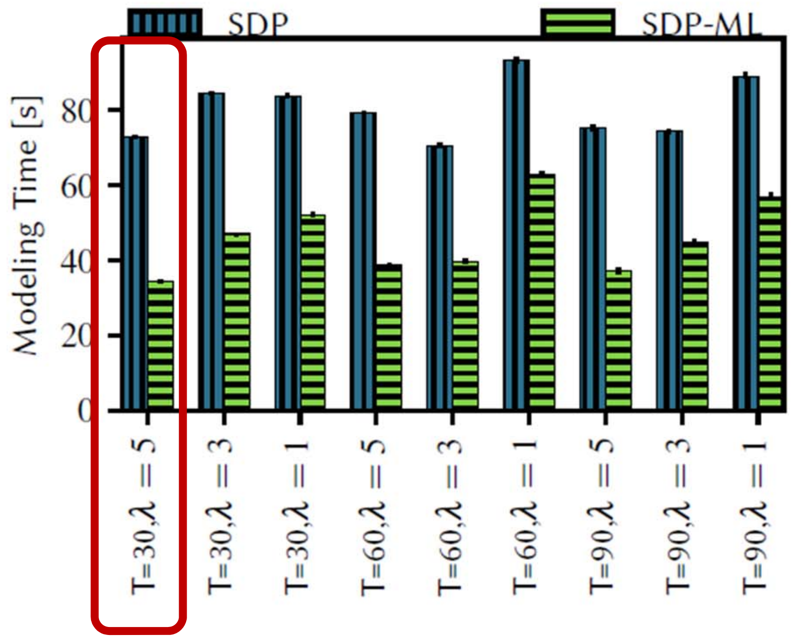

Efficient Filtering of infeasible and unacceptable requests

Efficient saving of model creation time

Latest Results: Neurovine [6]

Hopfield neural network to preprocess (subgraph extraction) VNE algorithms
– *tailored filtering*

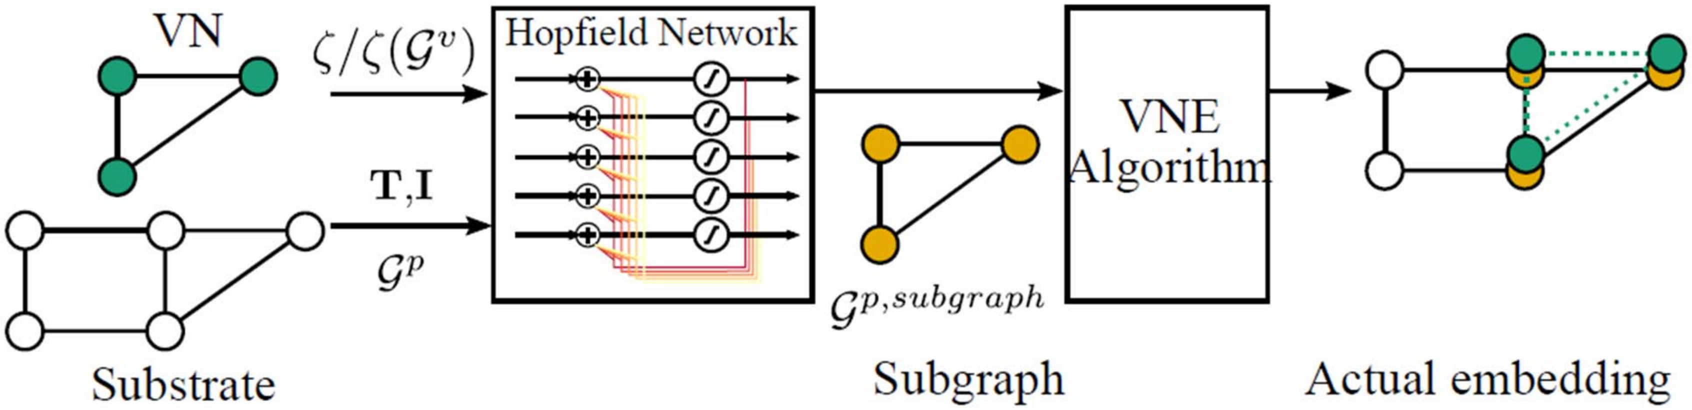

- Idea: Extract subgraph with physical nodes close to each other and high available capacities

[6] A. Blenk, P. Kalmbach, J. Zerwas, M. Jarschel, S. Schmid, W. Kellerer: ***NeuroViNE: A Neural Preprocessor for Your Virtual Network Embedding Algorithm*** IEEE INFOCOM 2018 (main conference), Honolulu, HI, USA, April 15-19, 2018.

Neurovine:
Efficiency on Real Network Topologies

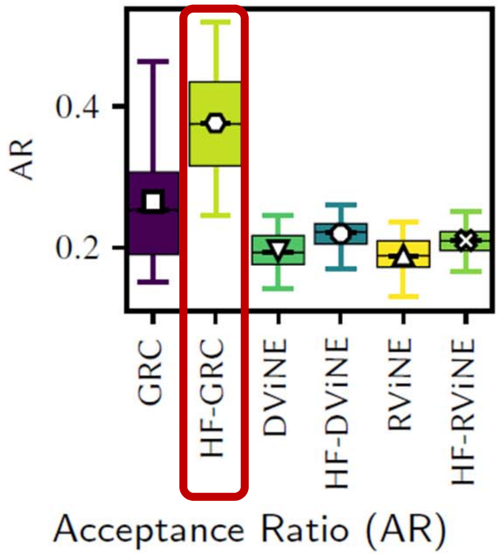

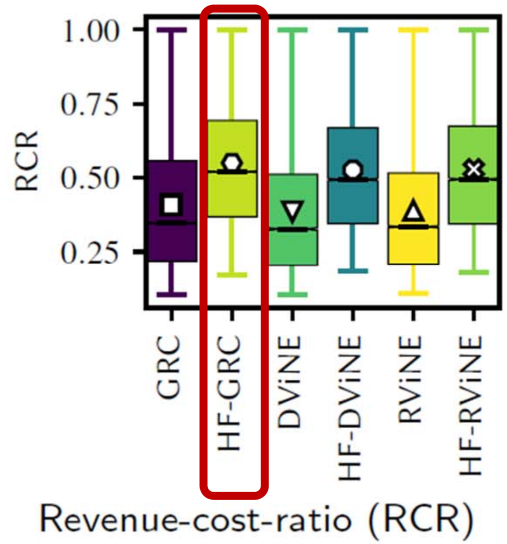

- VNE algorithms (GRC, DViNE, RViNE) vs. Hopfield variants (HF-GRC, HF-DViNE, HF-RViNE)

- NeuroViNE accepts more networks with less costs

Empower your network

recent emergence of *self-driving* networks (Rexford, Feamster):
networks which measure, analyze and control themselves in an automated manner, reacting to changes in the environment (e.g., demand), while exploiting existing flexibilities to optimize themselves

(self-)optimizations shall also prepare for possibly unexpected events → **preparedness** → **flexibility**

We propose:

- use ***empowerment*** for preparedness
  (information-theoretic measure to quantify the influence of an agent on its environment, e.g. used in robotics)

[7] P. Kalmbach, J. Zerwas, P. Babarczi, A. Blenk, W. Kellerer, S. Schmid: ***Empowering Self Driving Networks***, accepted for ACM SIGCOMM 2018 workshop on self-driving networks August 2018.

Empowering Self Driving Networks

empowerment: quantify the influence of an agent on its environment: agent (several actuators, 1 sensor) restructures networks to maximize options (c) - not an objective as in optimization (a) and (b)

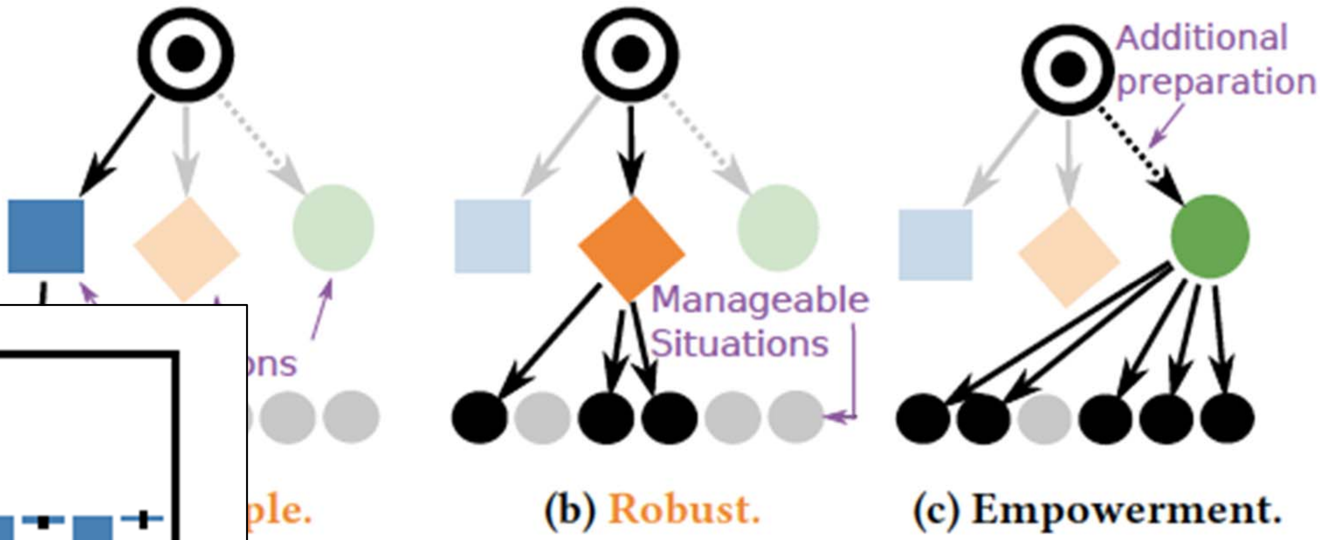

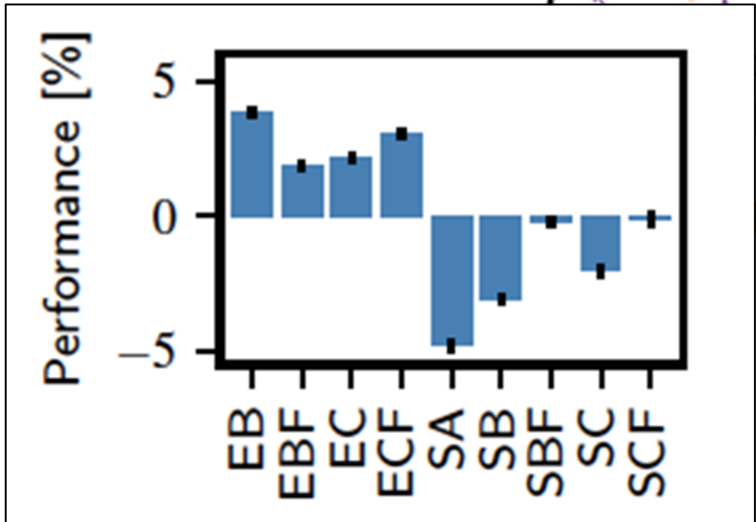

initial result:
empowerment-based approaches (EB, EC) can outperform ILP („0“-line) and heuristics (SA) [7]

[7] P. Kalmbach, J. Zerwas, P. Babarczi, A. Blenk, W. Kellerer, S. Schmid: ***Empowering Self Driving Networks***, accepted for ACM SIGCOMM 2018 workshop on self-driving networks August 2018.

Key Takeaways & outlook

Flexibility matters!

- We propose a

definition and measure for flexibility

- to compare flexible systems
- to explicitly **design for flexibility**

- **Adaptation/optimization time** is important for flexible systems

Speedup optimization algorithms through
Machine Learning-based preprocessing

- Recent work: **Empowerment** concept to design for flexibility

References

[1] Sdxcentral. Carriers 5G Plans are Rooted in SDN/NFV, says Ixia Survey. [Online]. Available:https://www.sdxcentral.com/articles/news/carriers-5g-plans-rooted-sdnnfv-says-ixia-survey/2017/09/?c action=related articles

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

[3] W. Kellerer, A. Basta *et al.*, "How to measure network flexibility? A proposal for evaluating softwarized networks," *IEEE Communications Magazine*, 2018.

[4] M. He, A. Basta, A. Blenk, W. Kellerer, *How Flexible is Dynamic SDN Control Plane?,* IEEE INFOCOM Workshop, SWFAN'17, Atlanta, USA, May 2017.

[5] A. Blenk, P. Kalmbach, S. Schmid, W. Kellerer: o'zapft is: Tap Your Network Algorithm's Big Data! ACM SIGCOMM 2017 Wrks. on Big Data Analytics and Machine Learning for Data Communication Networks (Big-DAMA), 2017.

[6] Andreas Blenk, Patrick Kalmbach, Johannes Zerwas, Michael Jarschel, Stefan Schmid, Wolfgang Kellerer: NeuroViNE: A Neural Preprocessor for Your Virtual Network Embedding Algorithm IEEE INFOCOM 2018 (main conference), Honolulu, HI, USA, April 15-19, 2018.

Abstract

In order to address network dynamics and highly varying requirements, flexibility has emerged as a key property for networks to cope with increasing dynamics and to be prepared for future demands. Softwarized networks including concepts such as Network Virtualization, Software Defined Networking and Network Function Virtualization promise flexibility. However, so far flexibility is mainly used as a qualitative advantage for a certain design choice where the meaning of flexibility is varying a lot in literature. To provide a better understanding of how to design flexible networks, we propose a definition for flexibility and present an approach for a quantitative measure of flexibility in softwarized networks. In our proposal, we refer to flexibility as the ability to support new requests, e.g., changes in the requirements or new traffic distributions, in a timely manner. We illustrate with use case studies for function placement and SDN resilience, how this measure can be used to evaluate and compare different network designs quantitatively. To address adaptation time in flexible networks, we further present approaches to speed up the execution of algorithms based on machine learning. Examples include virtual network embedding and function placement. With our proposed approach for the definition and evaluation of flexibility, we intend to stimulate the discussion towards a more quantitative analysis of softwarized networks and beyond.