TECHNISCHE UNIVERSITÄT DRESDEN

ComNets
Deutsche Telekom Chair of Communication Networks

Robert-Steve Schmoll
Deutsche Telekom Chair of Communication Networks

Mobile Edge Cloud
Motivation, Implementation, Challenges

Oberseminar

Dresden 12.06.2018

Motivation in 5G

Cloud Computing so far
Public Clouds

- Flexibility
  - Varying degree of resoureces
- Scalability
  - Blueprints allow repeated deployment
- Reliability
  - Data-center uptime, fault recognition, backups and snapshots
  - Distributed geographically
- Convenience
  - Administration only on booked VMs, not on underlying infrastructure

ars TECHNICA

BIZ & IT | TECH | SCIENCE | POLICY | CARS | GAMING & CULTURE | FORUMS | SUBSCRIPTIONS | SIGN IN

IRON ROUTER —

In effort to shut down Telegram, Russia blocks Amazon, Google network addresses

Russian censors up ante in effort to block secure chat by political opposition.

SEAN GALLAGHER - 4/17/2018, 8:35 PM

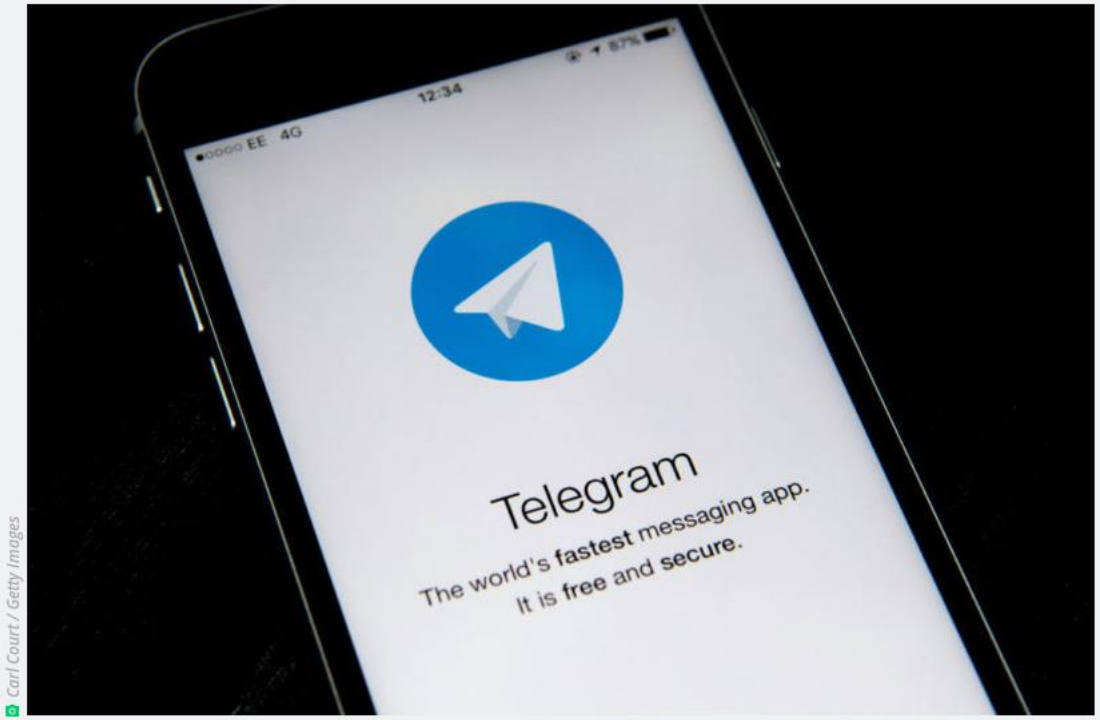

Carl Court / Getty Images

Enlarge

Application

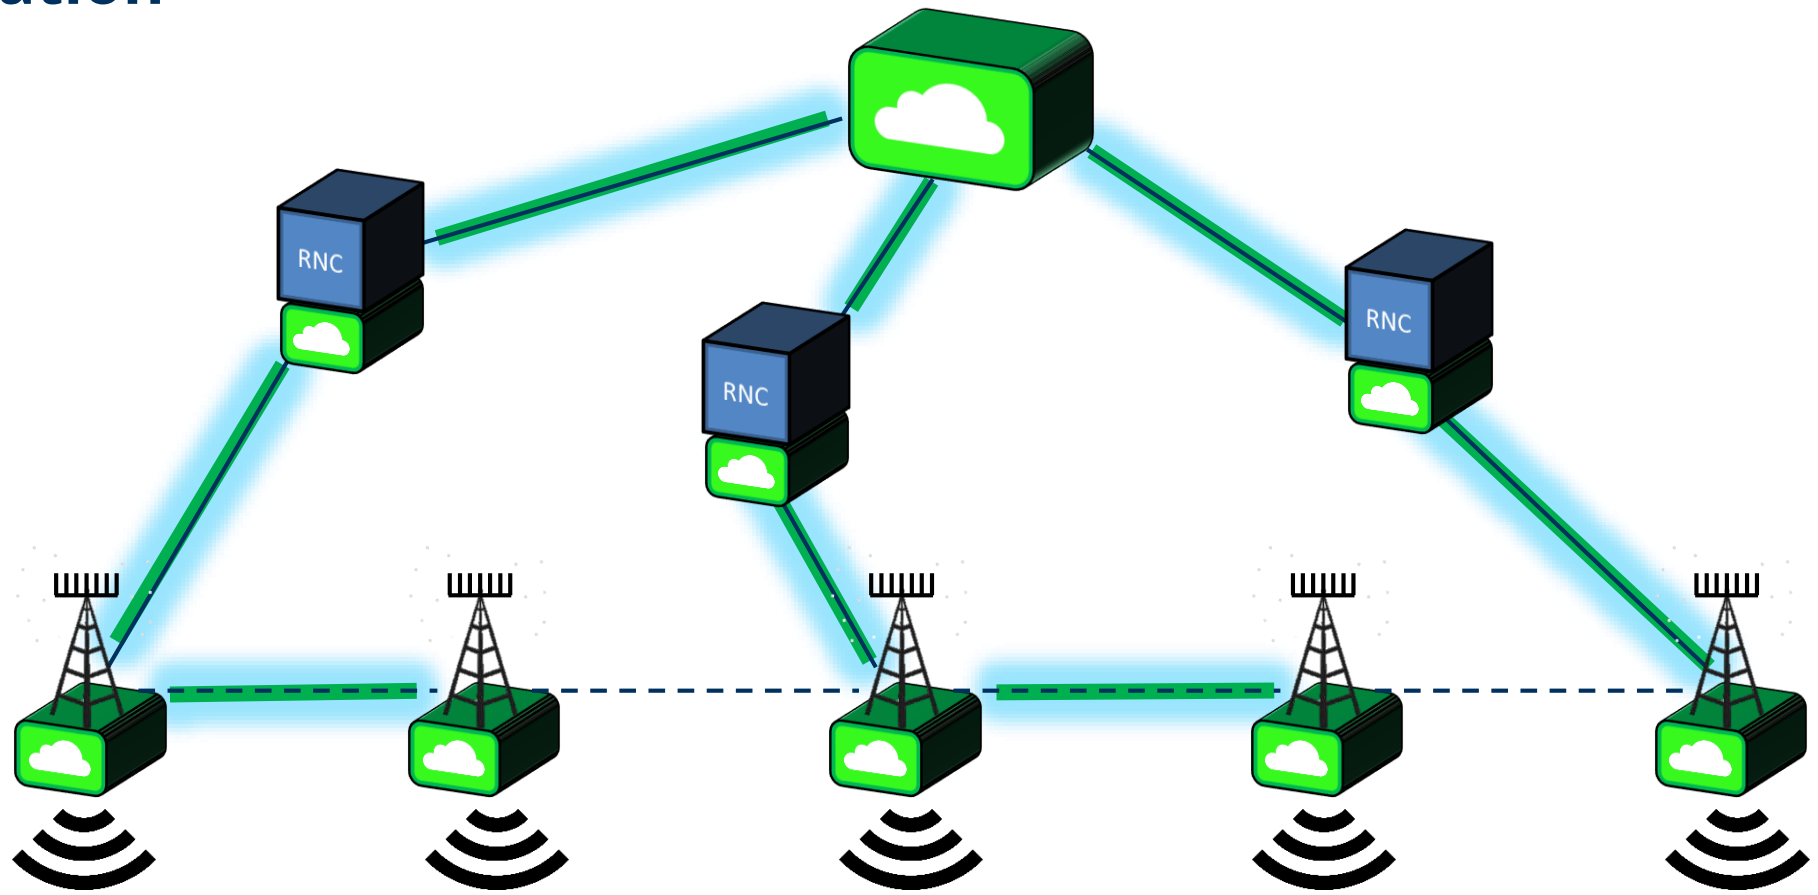

Use Cases - I

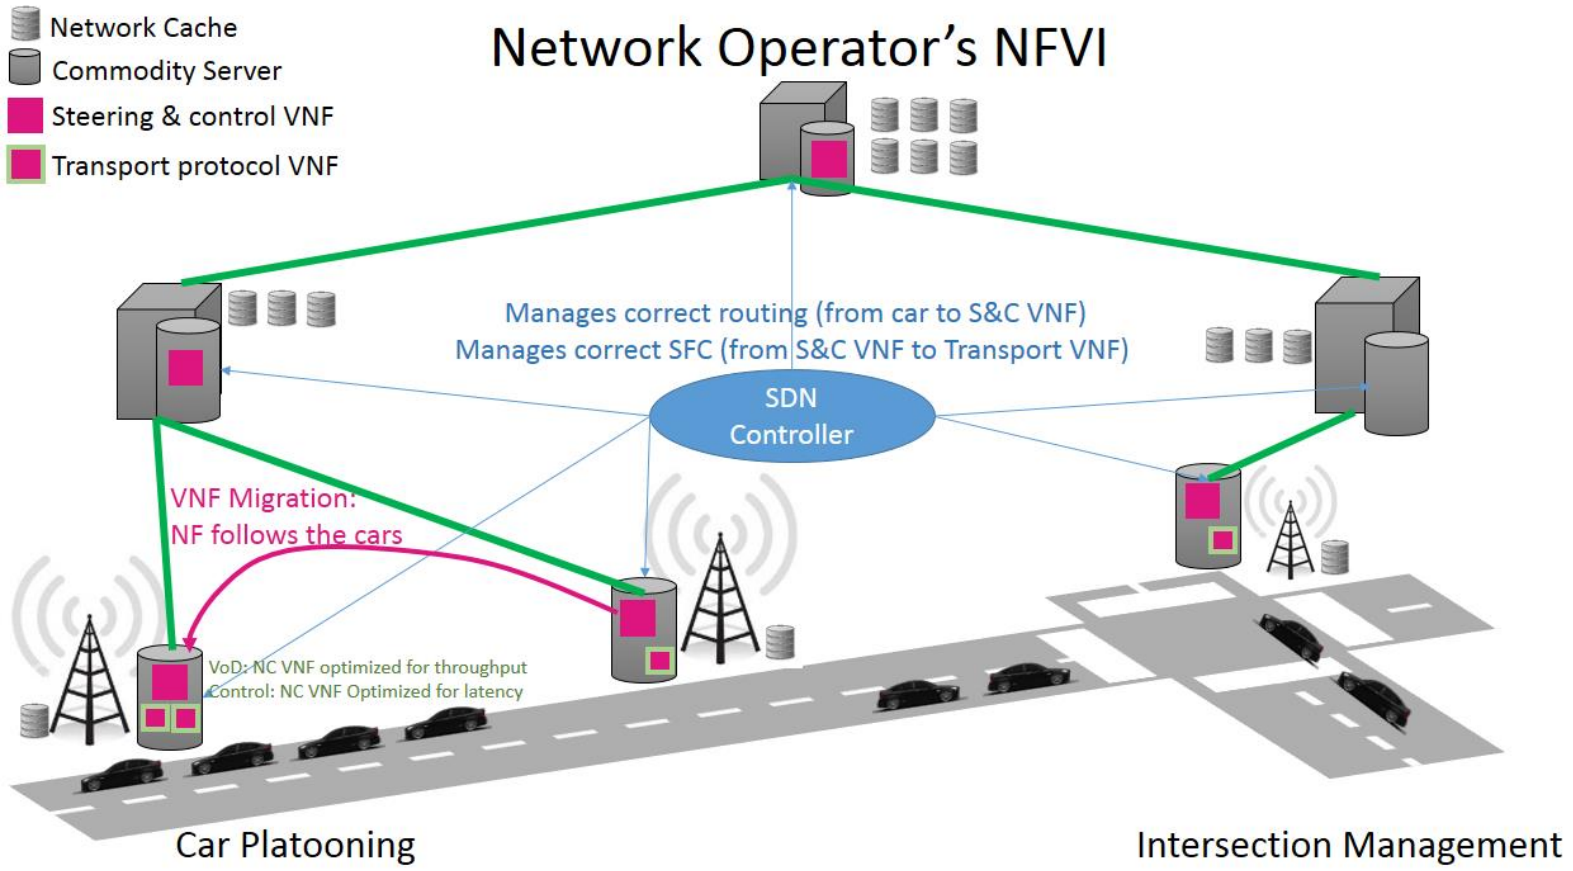
Network Cache
Commodity Server
Steering & control VNF
Transport protocol VNF
Network Operator's NFVI
Manages correct routing (from car to S&C VNF)
Manages correct SFC (from S&C VNF to Transport VNF)
SDN Controller
VNF Migration:
NF follows the cars
VoD: NC VNF optimized for throughput
Control: NC VNF Optimized for latency
Car Platooning
Intersection Management

Use Cases - II

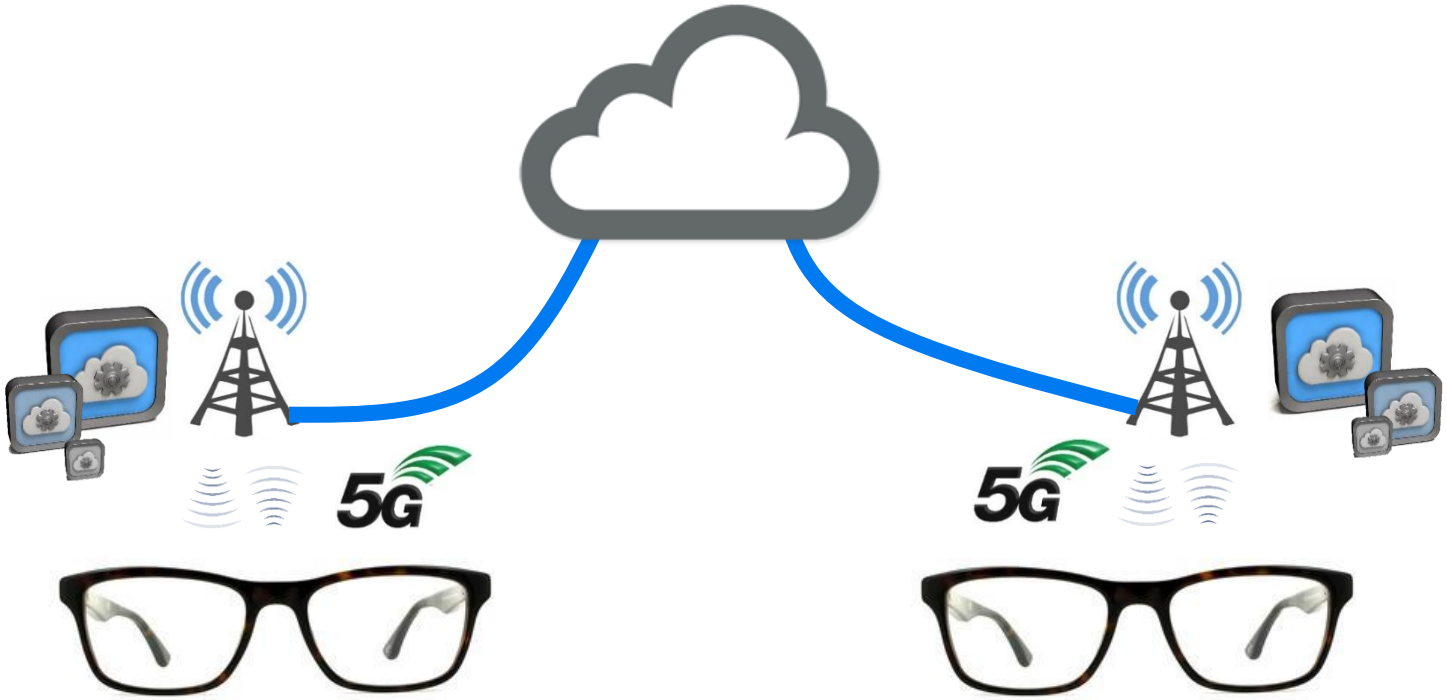

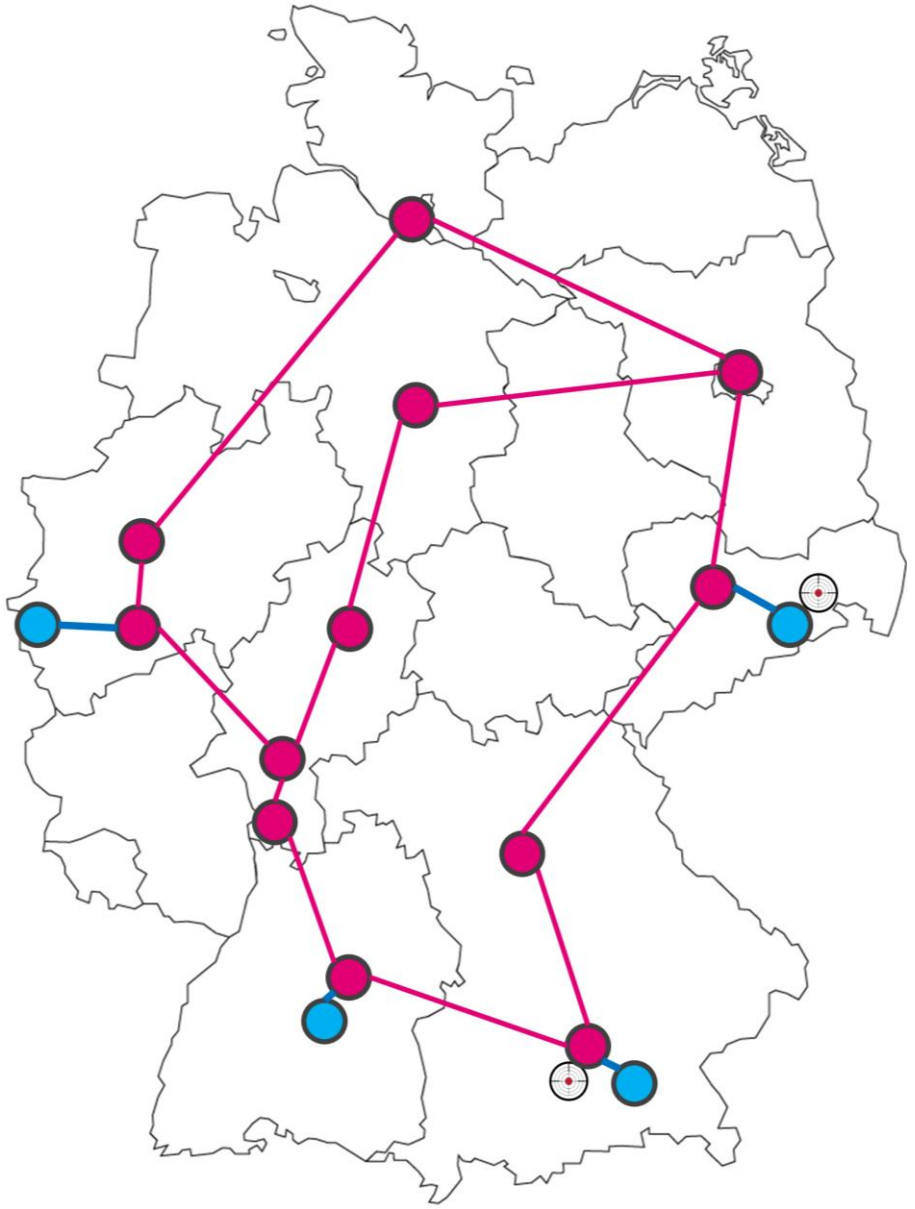

Use Cases - III

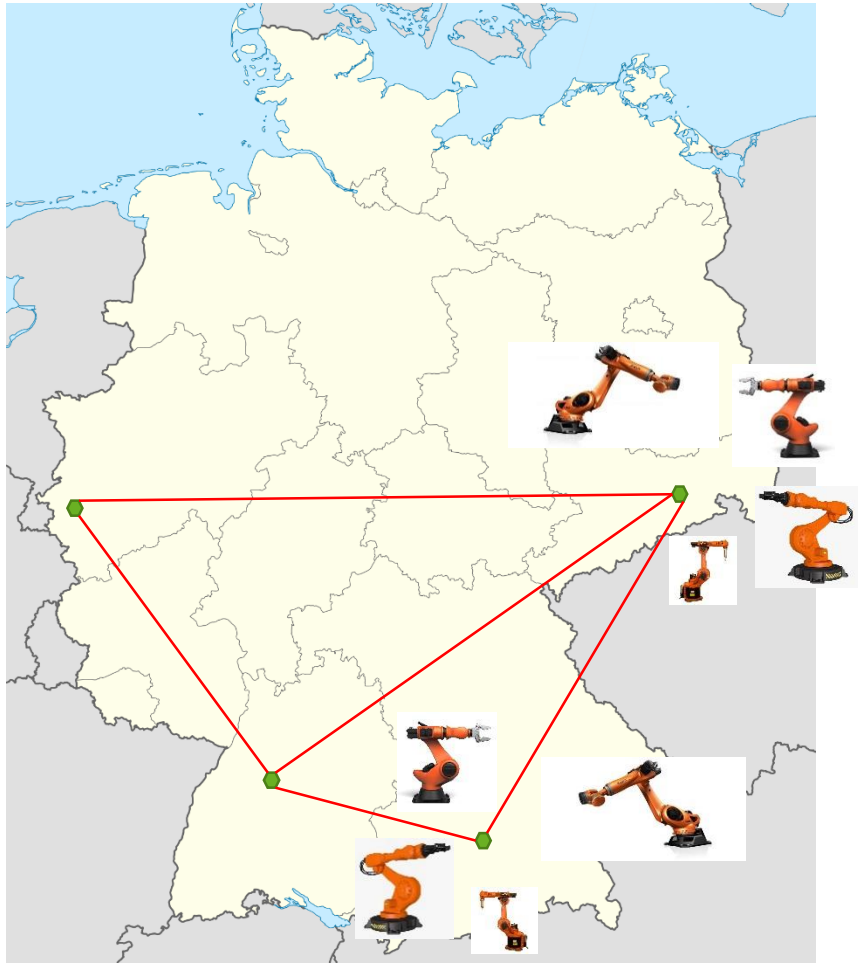

Use Cases - IV

Changes to Cloud Infrastructure?

Distributed	Distributed, but bigger spatial diversity
On-demand	On-demand, but with faster deployment
Scaleable	Scaleable, but with even more devices
Reliable	Reliable, but 5G compliant

Implementation

Key Technology

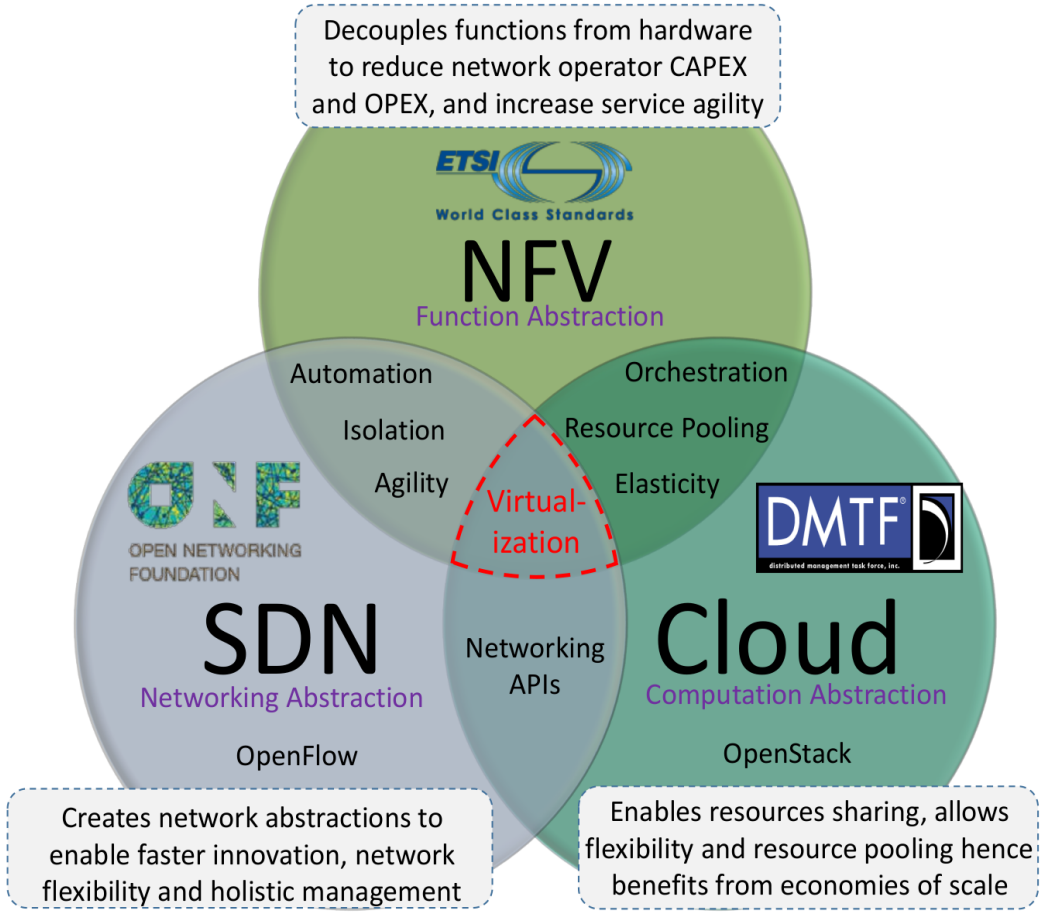

Mijumbi, Rashid, et al. "Network function virtualization: State-of-the-art and research challenges." IEEE Communications Surveys & Tutorials 18.1 (2016): 236-262.

KVM - I

- KVM is a virtualization infrastructure
- It uses the Linux Kernel as a hypervisor
  → Type 2
  - Part of kernel since 2007
- Benefits from hardware extensions: Intel VT-x, AMD-V
- QEMU (Quick EMUlator) provides virtualized hardware

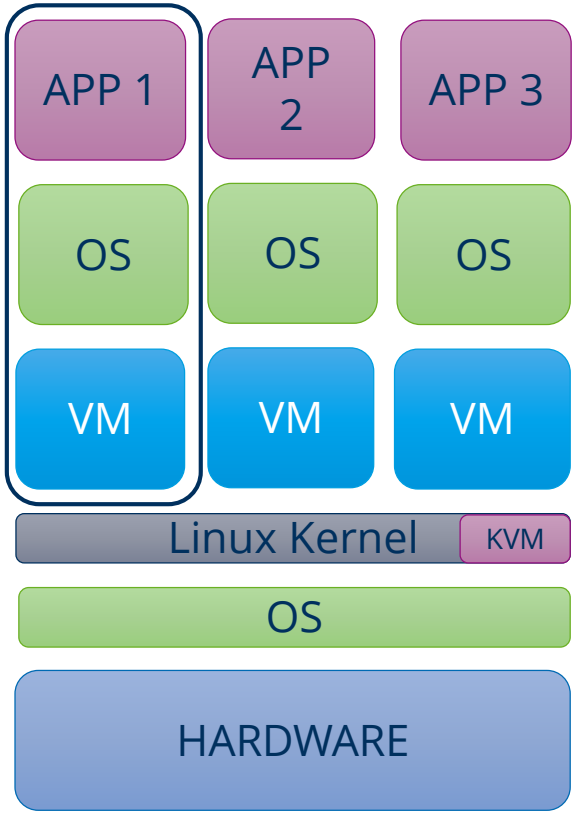

KVM - II

- QEMU (Quick EMUlator) provides virtualized hardware
- Libvirt is the API for lifecycle management
- Possible interfaces for VM management:virsh, virt-manager, Openstack

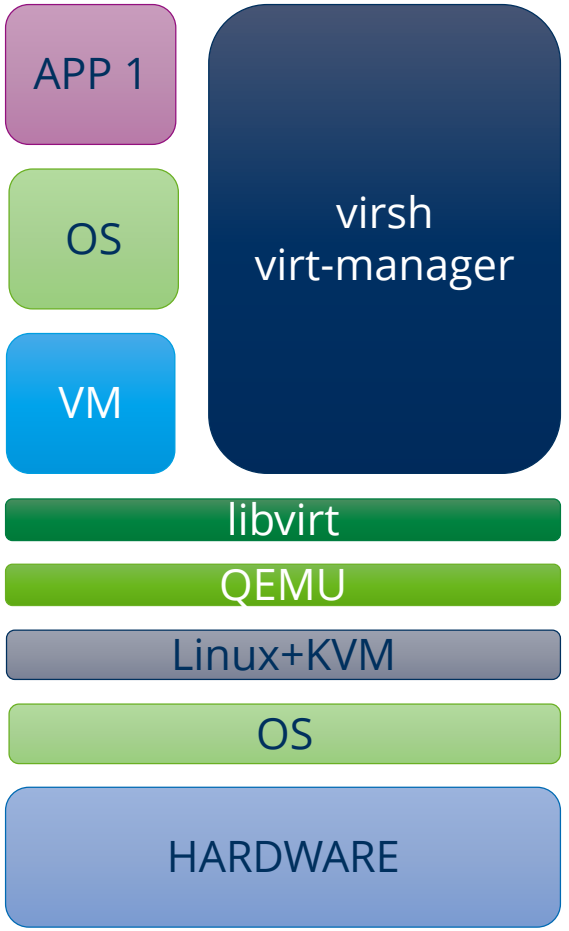

Migration General Structure

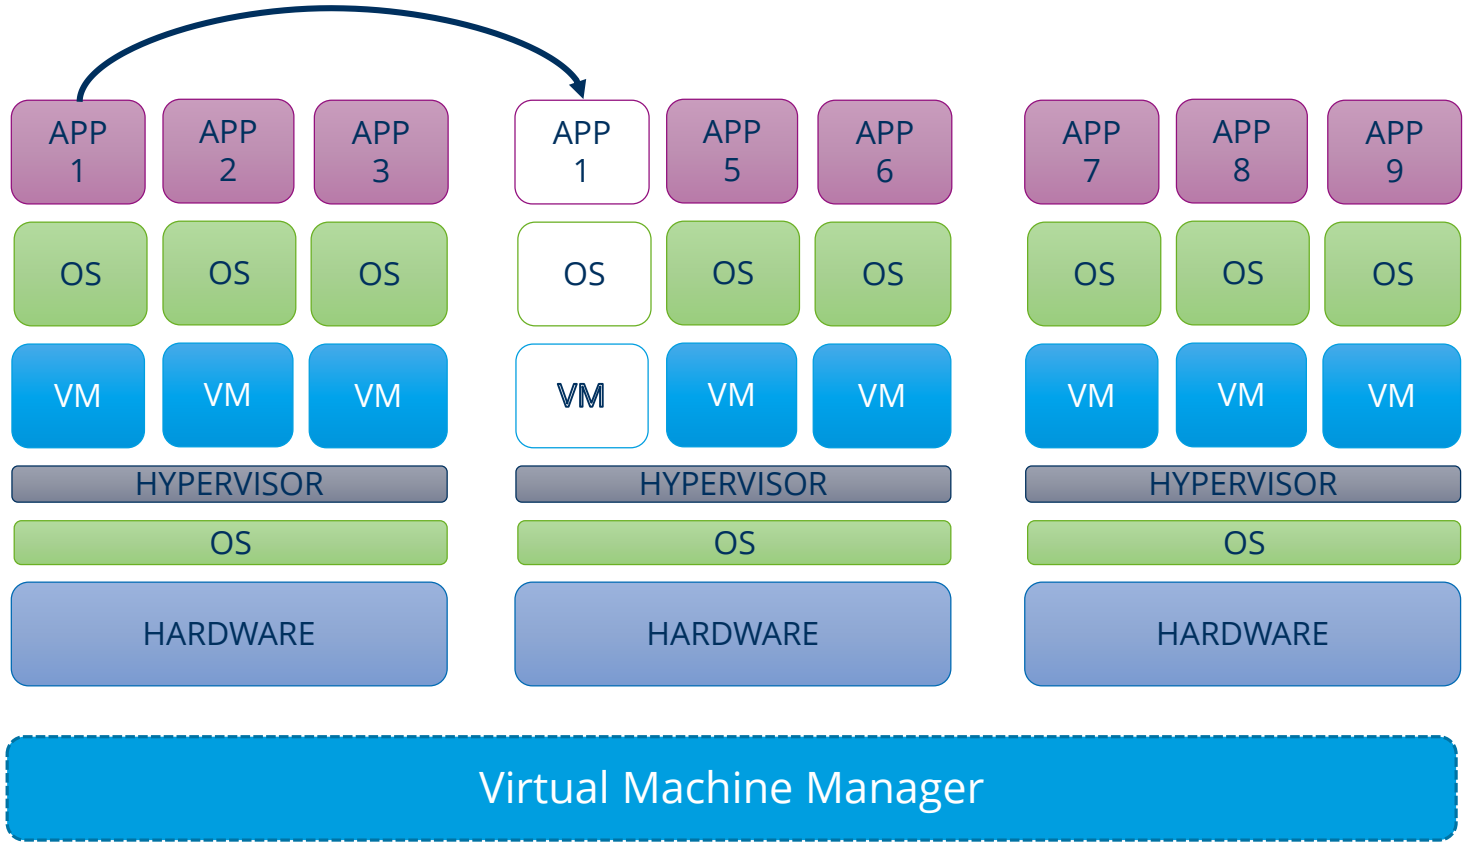

Reasons for Live Migration

Provider Perspective

- Load balancing
  - In case of overloaded network
- Proactive fault tolerance
  - If monitoring indicates imminent failure
- Power management
  - Consolidation of VMs so hosts can be shut down
- Resource sharing
  - In case of overloaded CPU or memory
- Online system maintenance
  - Maintenance of hosts

→ Generally performance-centric

Our Perspective (5G Networks)

- MEC – Continuous Service Delivery
  - The service has to be beneficial to the mobile device → offloading, communication, etc.
  - Stay in (close) proximity to device (for low latency requirements)
  - Always stay in „the middle“ of communicating users/devices (e.g. social VR applications)

→ Generally latency-centric

Virtual Machine Memory Layout

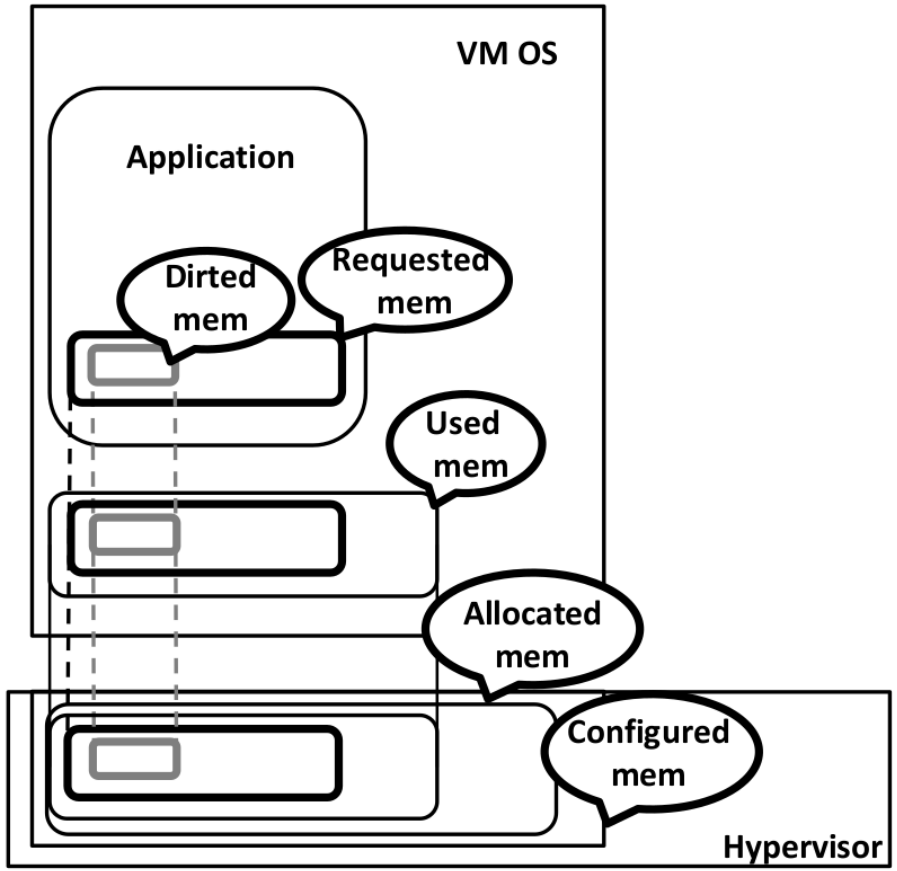

Hu, Wenjin, et al. "A quantitative study of virtual machine live migration." Proceedings of the 2013 ACM cloud and autonomic computing conference. ACM, 2013.

Live Migration – Pre-copy

- Whole VM state is transmitted at the beginning
- Iteration: while running, dirty pages are resend
  - i) until total amount of memory has been sent
  - ii) number of iterations exceed previously set parameters
  - iii) number of dirty pages below threshold

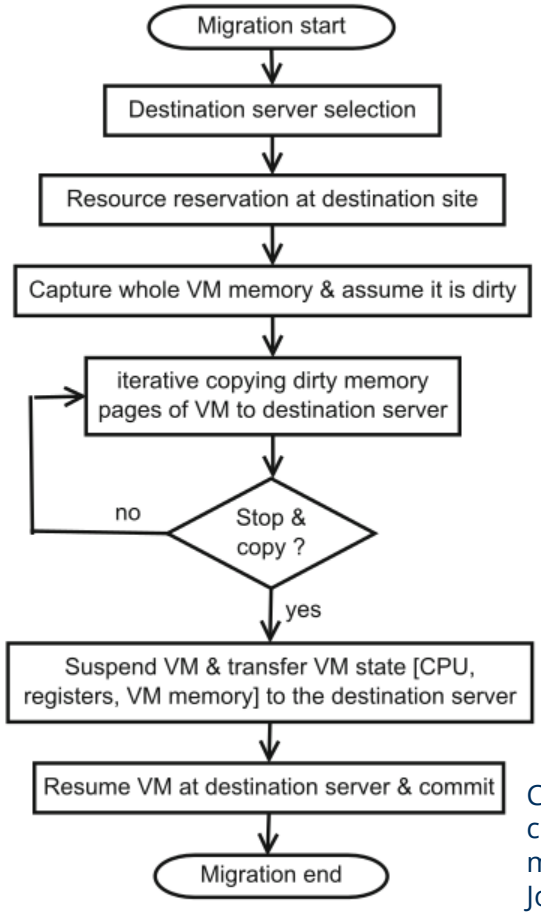

Choudhary, Anita, et al. "A critical survey of live virtual machine migration techniques." Journal of Cloud Computing 6.1 (2017): 23.

Live Migration – Post-copy

- Minimum VM state is transmitted at the beginning
- VM is started at destination
- Resend memory pages until all transferred

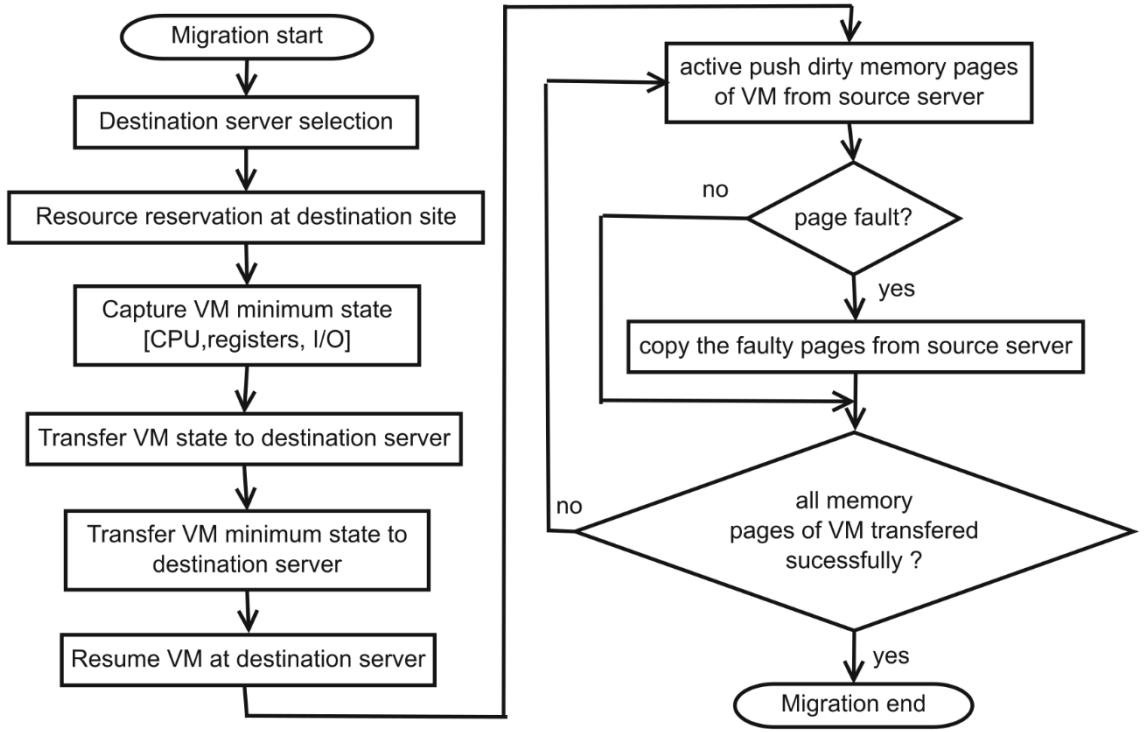

Choudhary, Anita, et al. "A critical survey of live virtual machine migration techniques." Journal of Cloud Computing 6.1 (2017): 23.

Virtual Machine Memory Layout

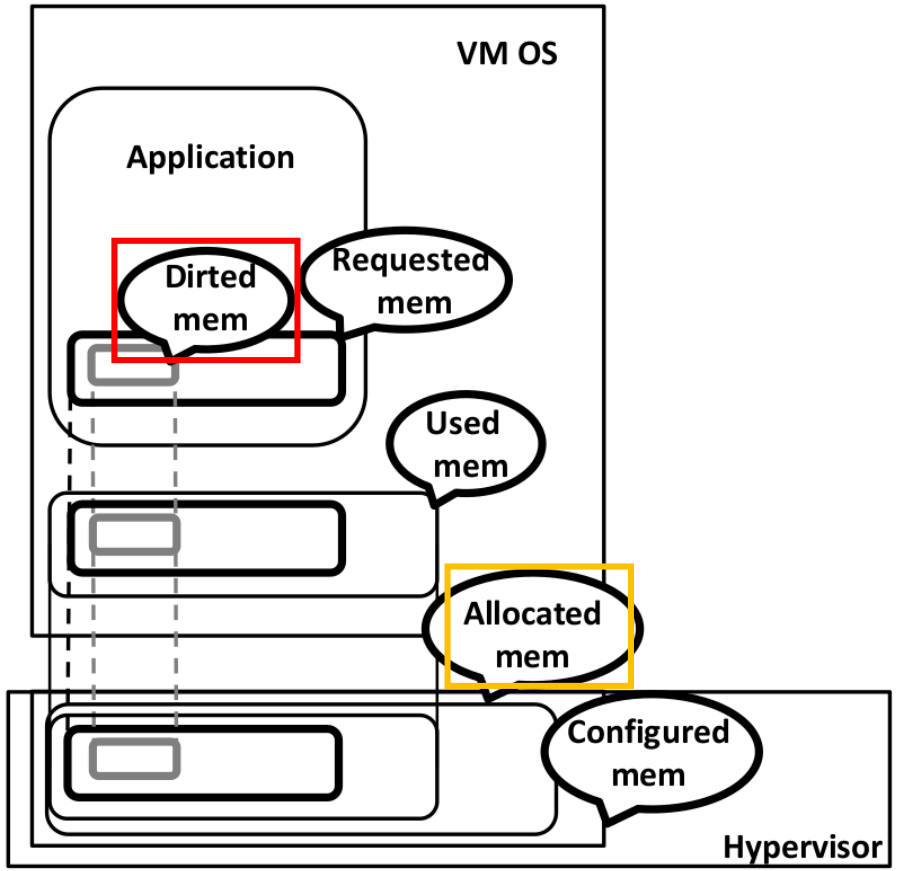

Hu, Wenjin, et al. "A quantitative study of virtual machine live migration." Proceedings of the 2013 ACM cloud and autonomic computing conference. ACM, 2013.

Virtual Machine Migration Times - Baseline

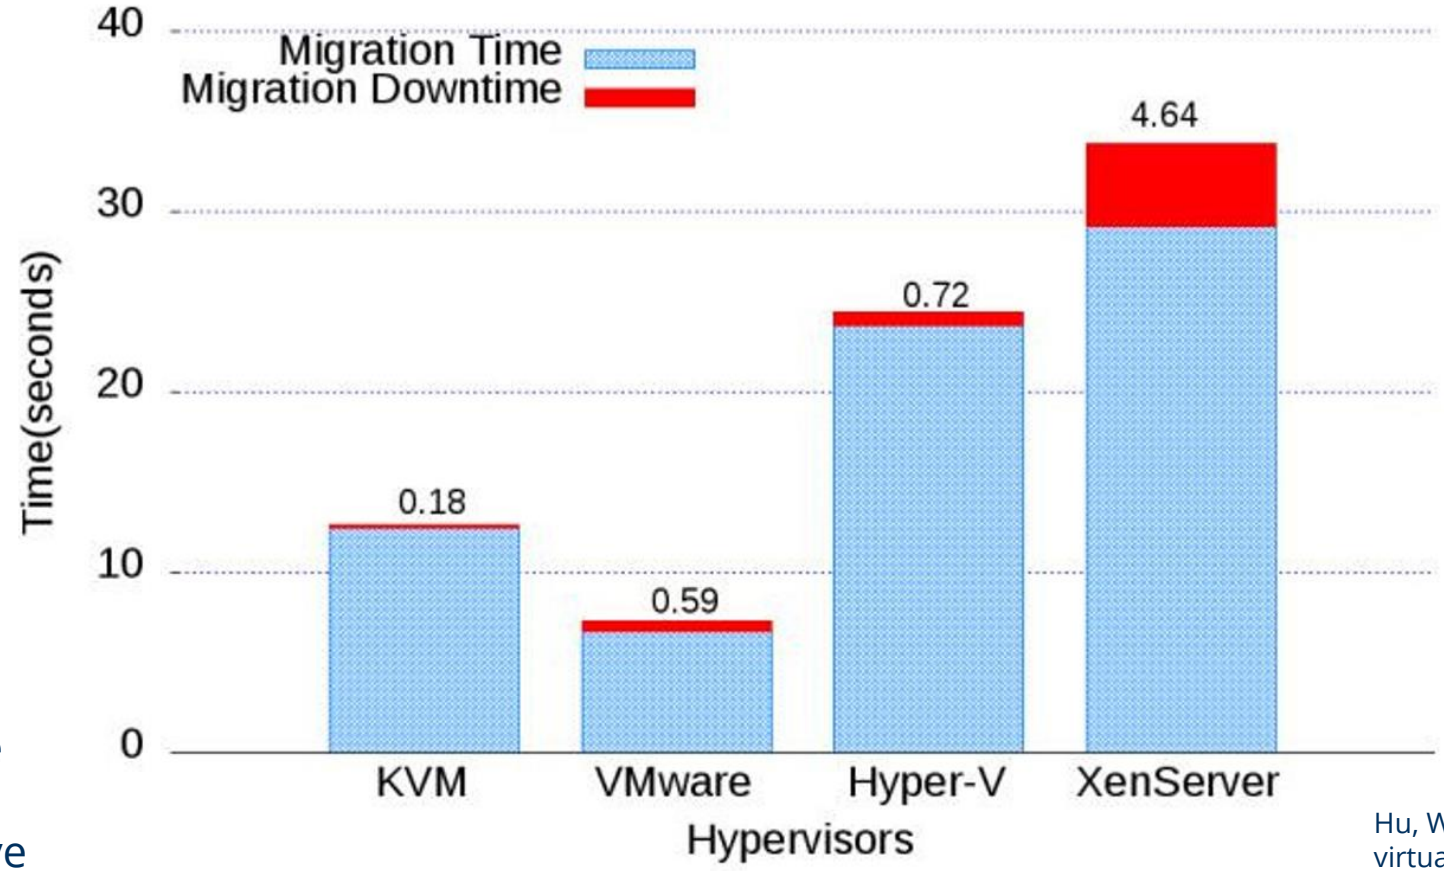

„Downtime represents the time that the VM being migrated was unresponsive to ping requests.“

Hu, Wenjin, et al. "A quantitative study of virtual machine live migration." Proceedings of the 2013 ACM cloud and autonomic computing conference. ACM, 2013.

Migration Scenarios

- Network intensive
  - E.g. web server responding to HTTP requests
- Memory intensive
  - E.g. Database server performing queries on in-memory database, memtest
- Storage intensive
  - E.g. searching inside files (read intensive)
- Compute intensive
  - E.g. calculation of Pi to nth digit, FFT calculations
- Any combination of these
  - E.g. offloading for computer vision (network + computation)


Container
Migration
VMs

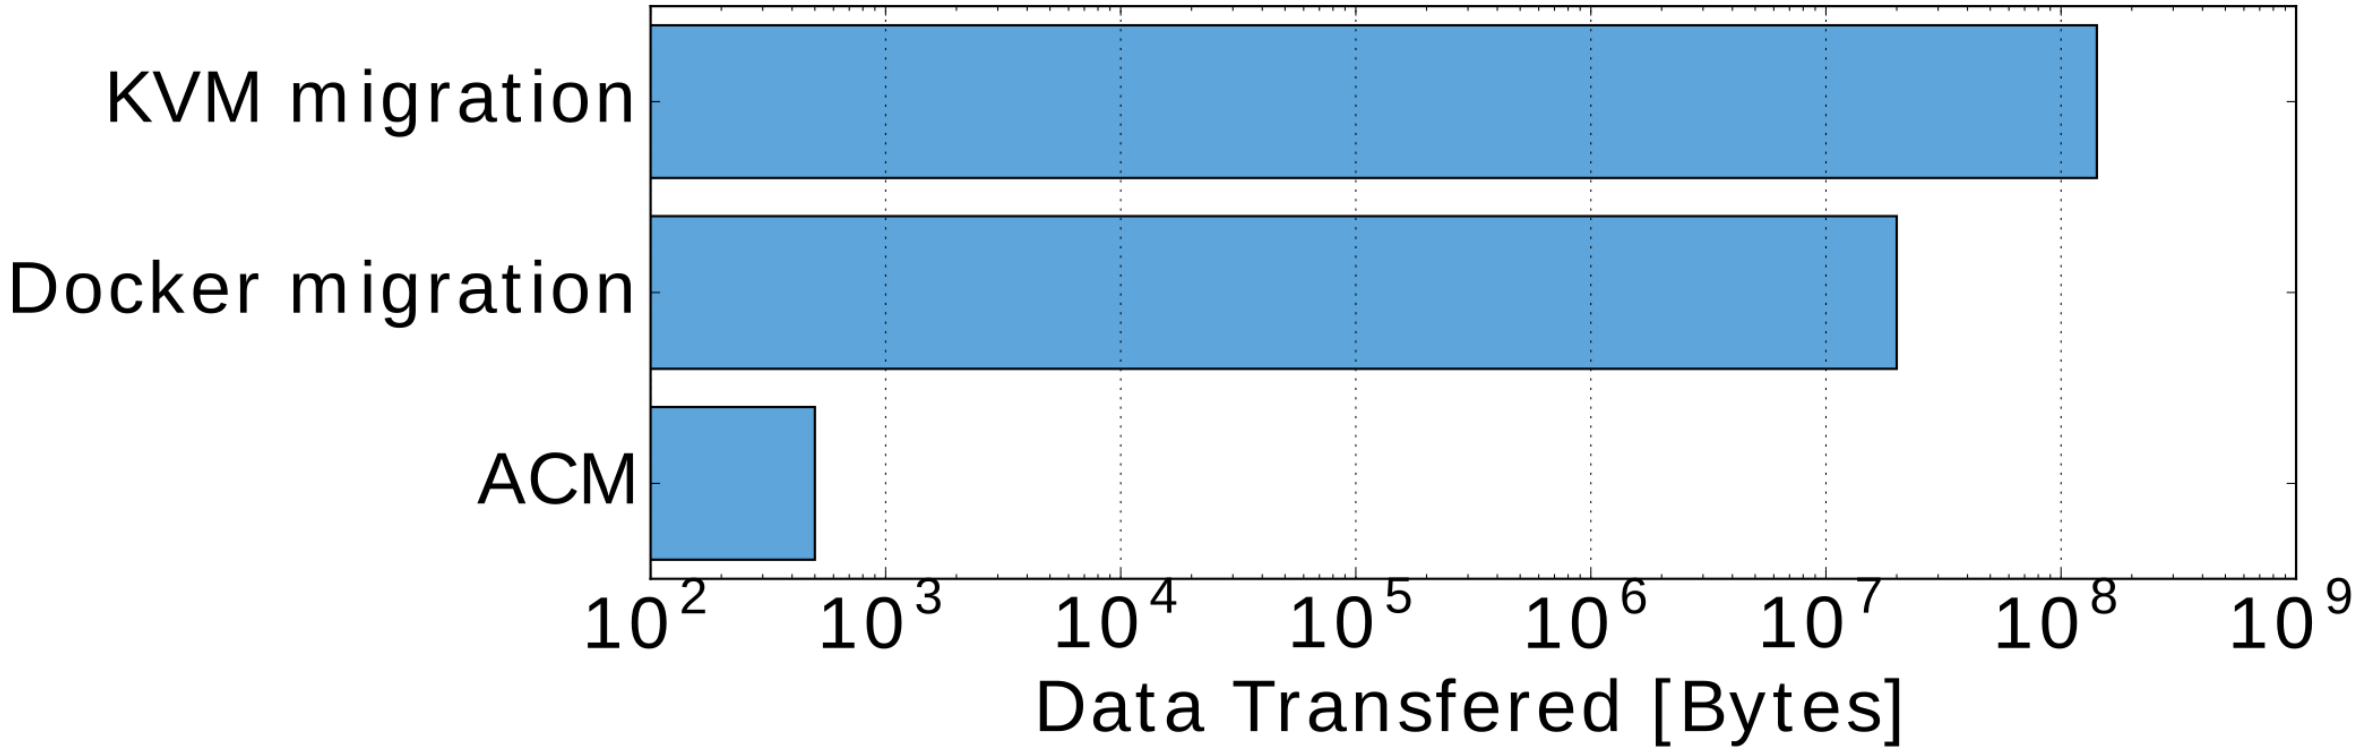
KVM migration
Docker migration
ACM
10^2
10^3
10^4
10^5
10^6
10^7
10^8
10^9
Data Transfered [Bytes]

Possible Use Cases

- Object recognition
  - Computationally intensive
  - For command and control latency-critical
  - High bandwidth necessary when dealing with raw video data
  - Might improve with low-latency video codecs
    - → Alex' research
    - *fast* music

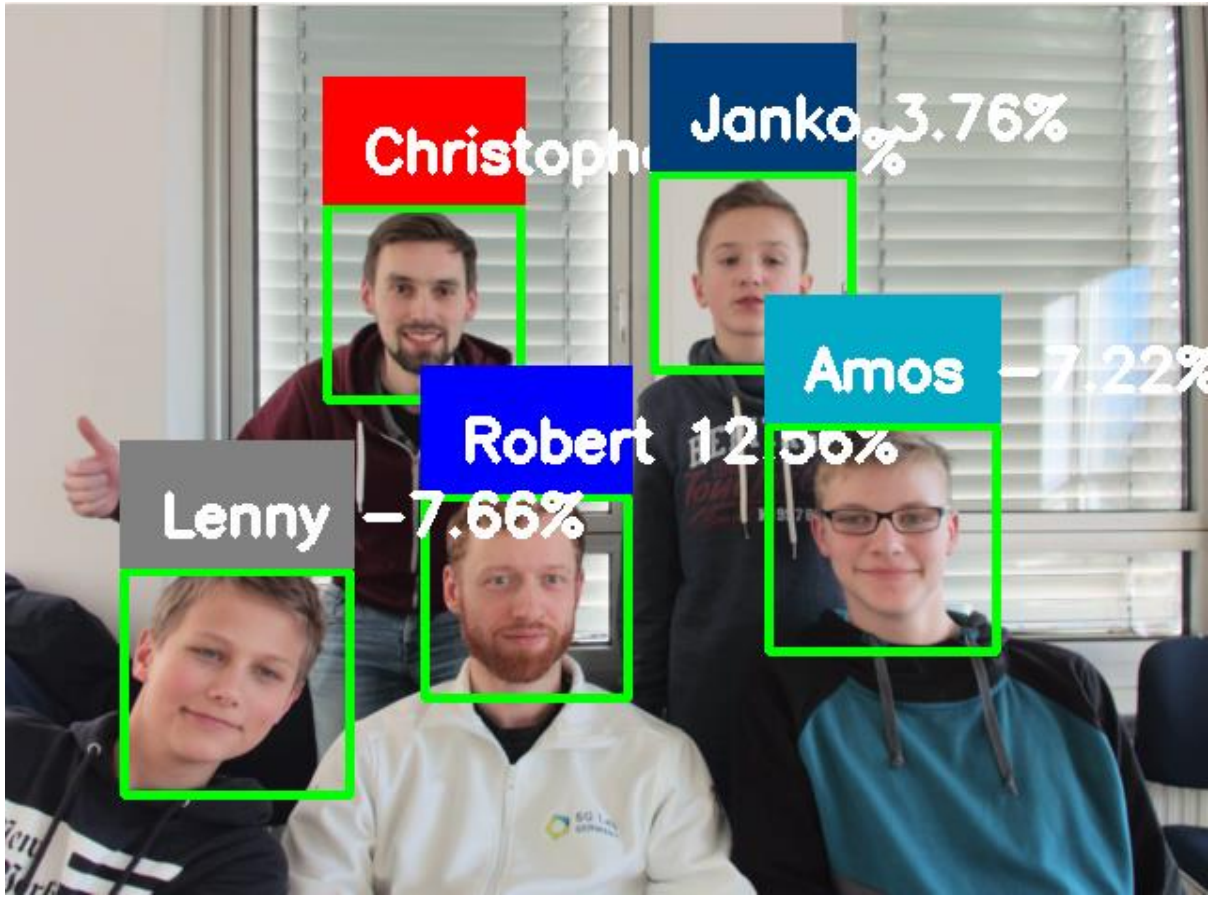

Computer Vision Testbed

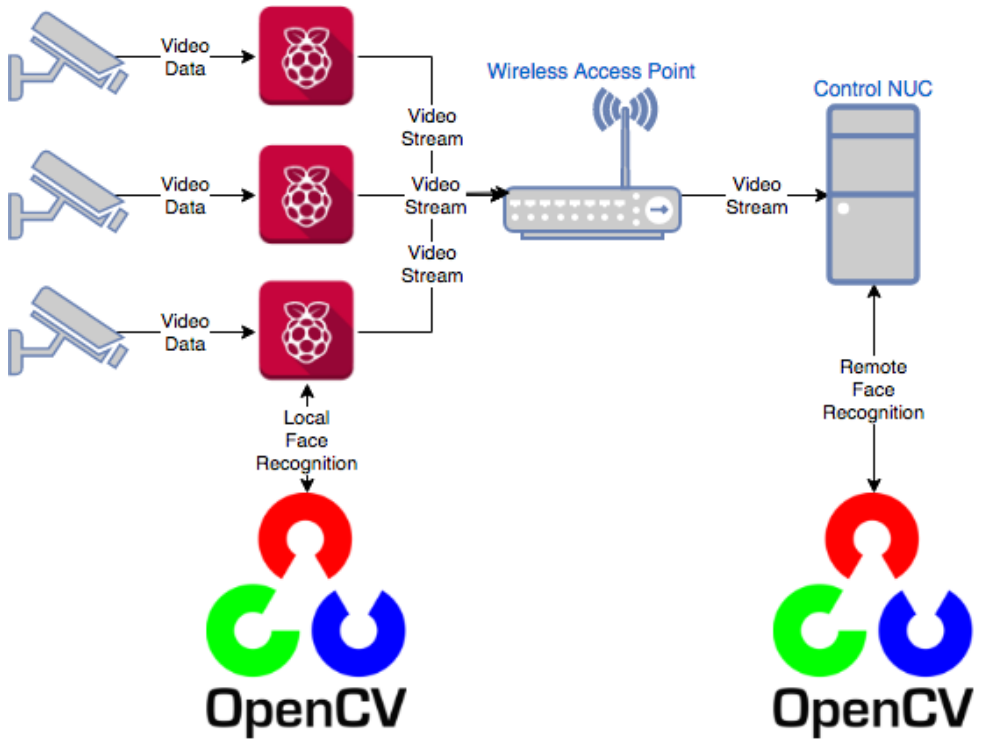

Performance Comparison

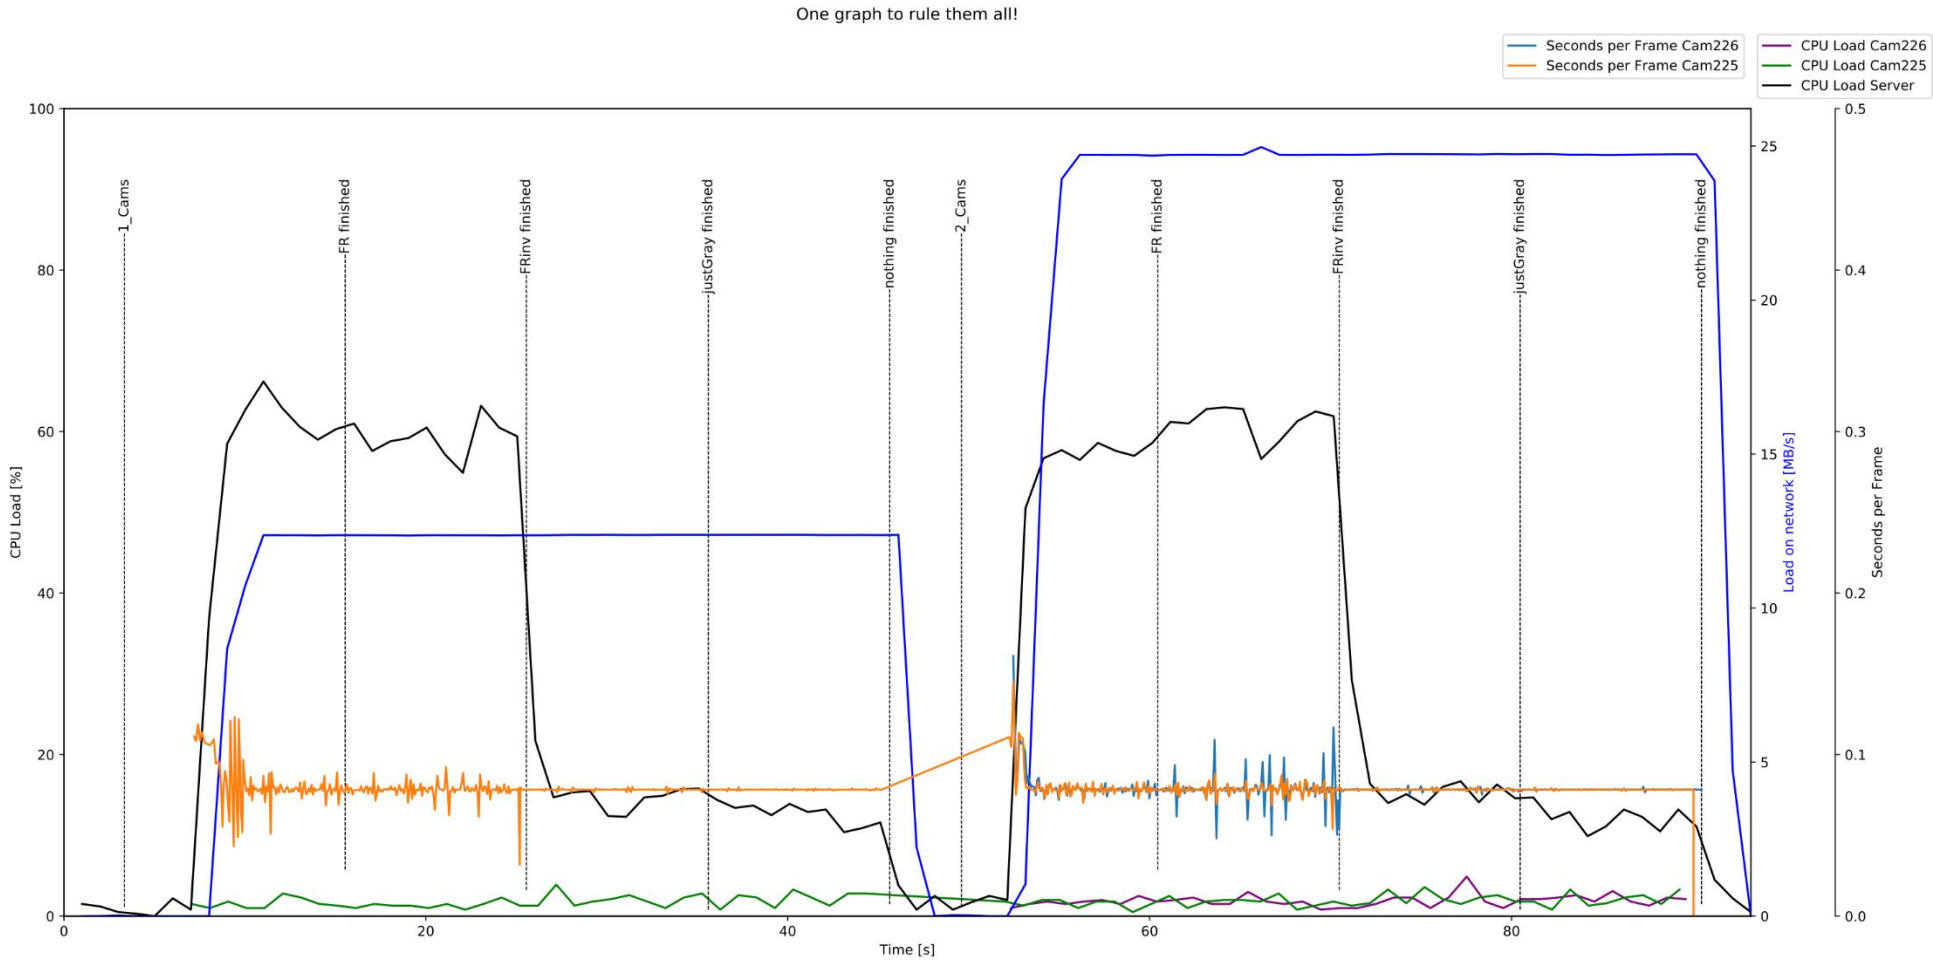

Face Recognition through streaming to Intel NUC

Conclusions

- Previous research focusses on data centers with live migration
- Baseline migration time is already very fast
  - Depending on the technology that is used
  - Might speed up with Containers <-> VMs
- Further experiments must include apprioriate scenario
  - Modify scenario to fit possible application

Challenges

Network

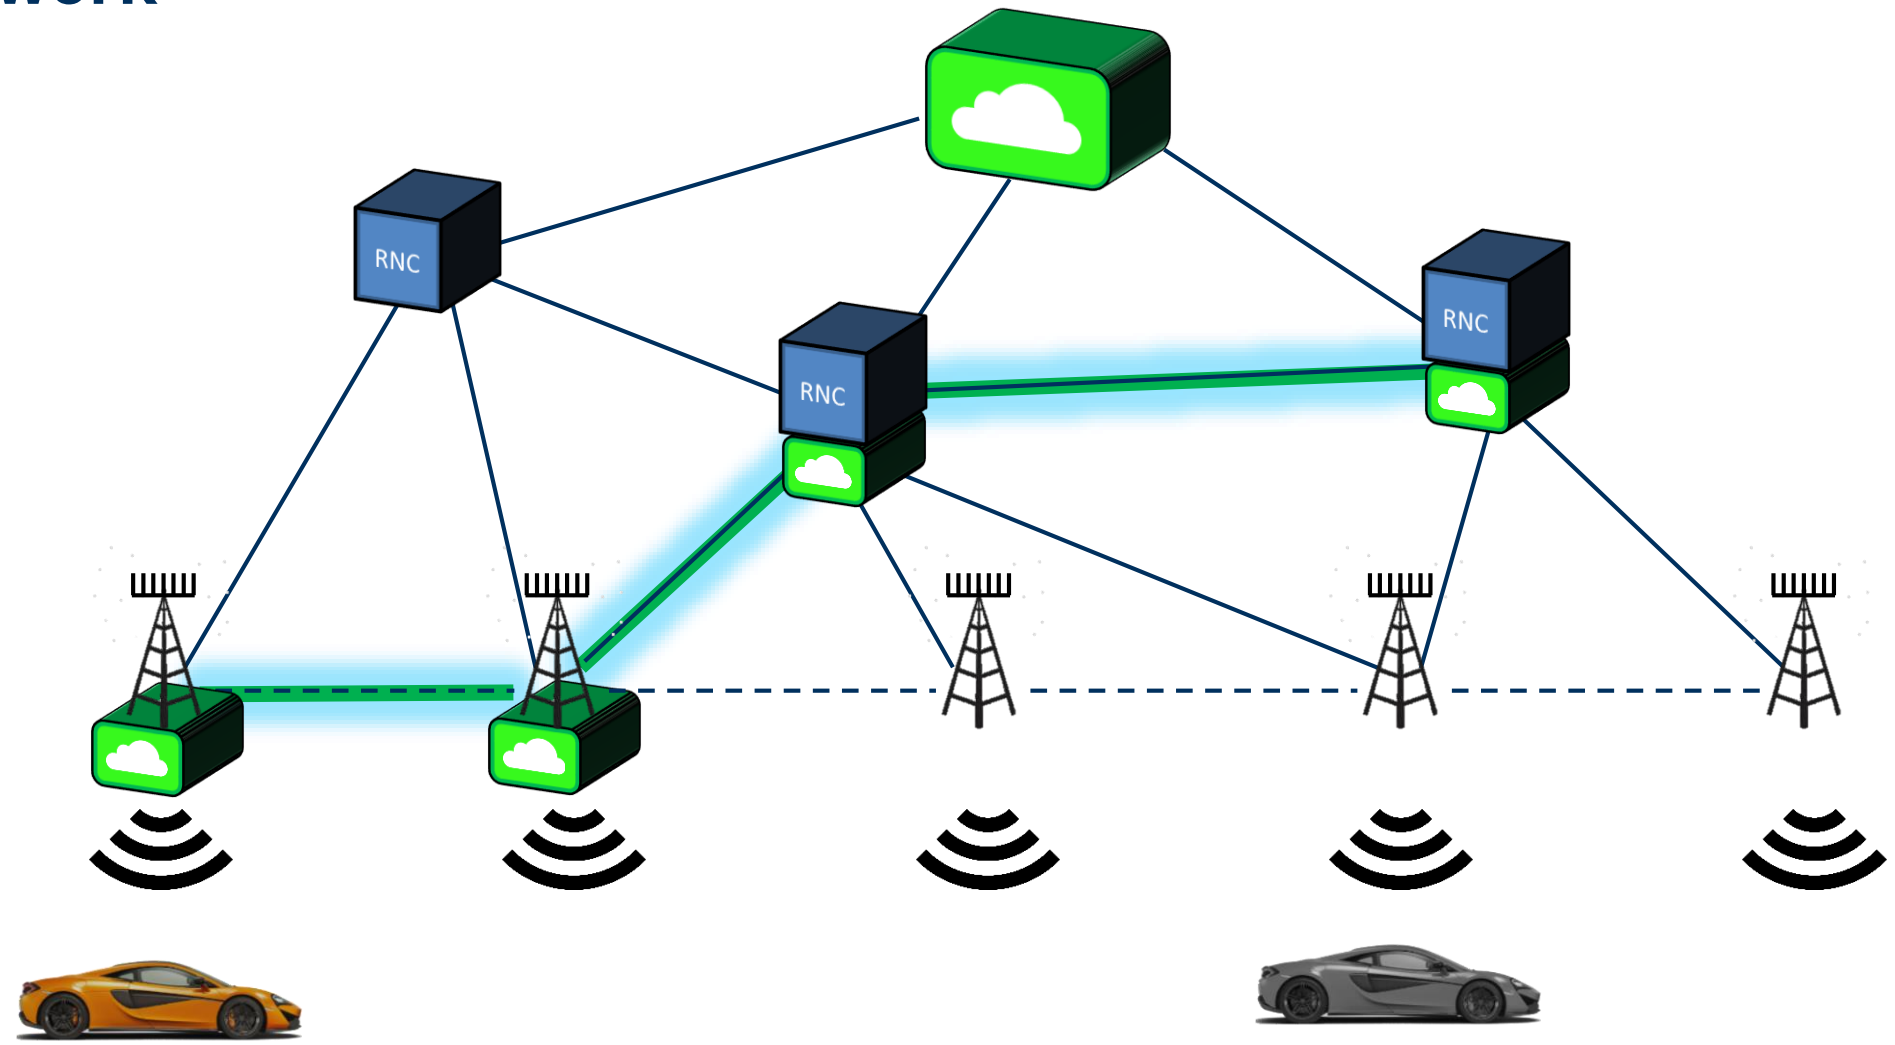

Data - I

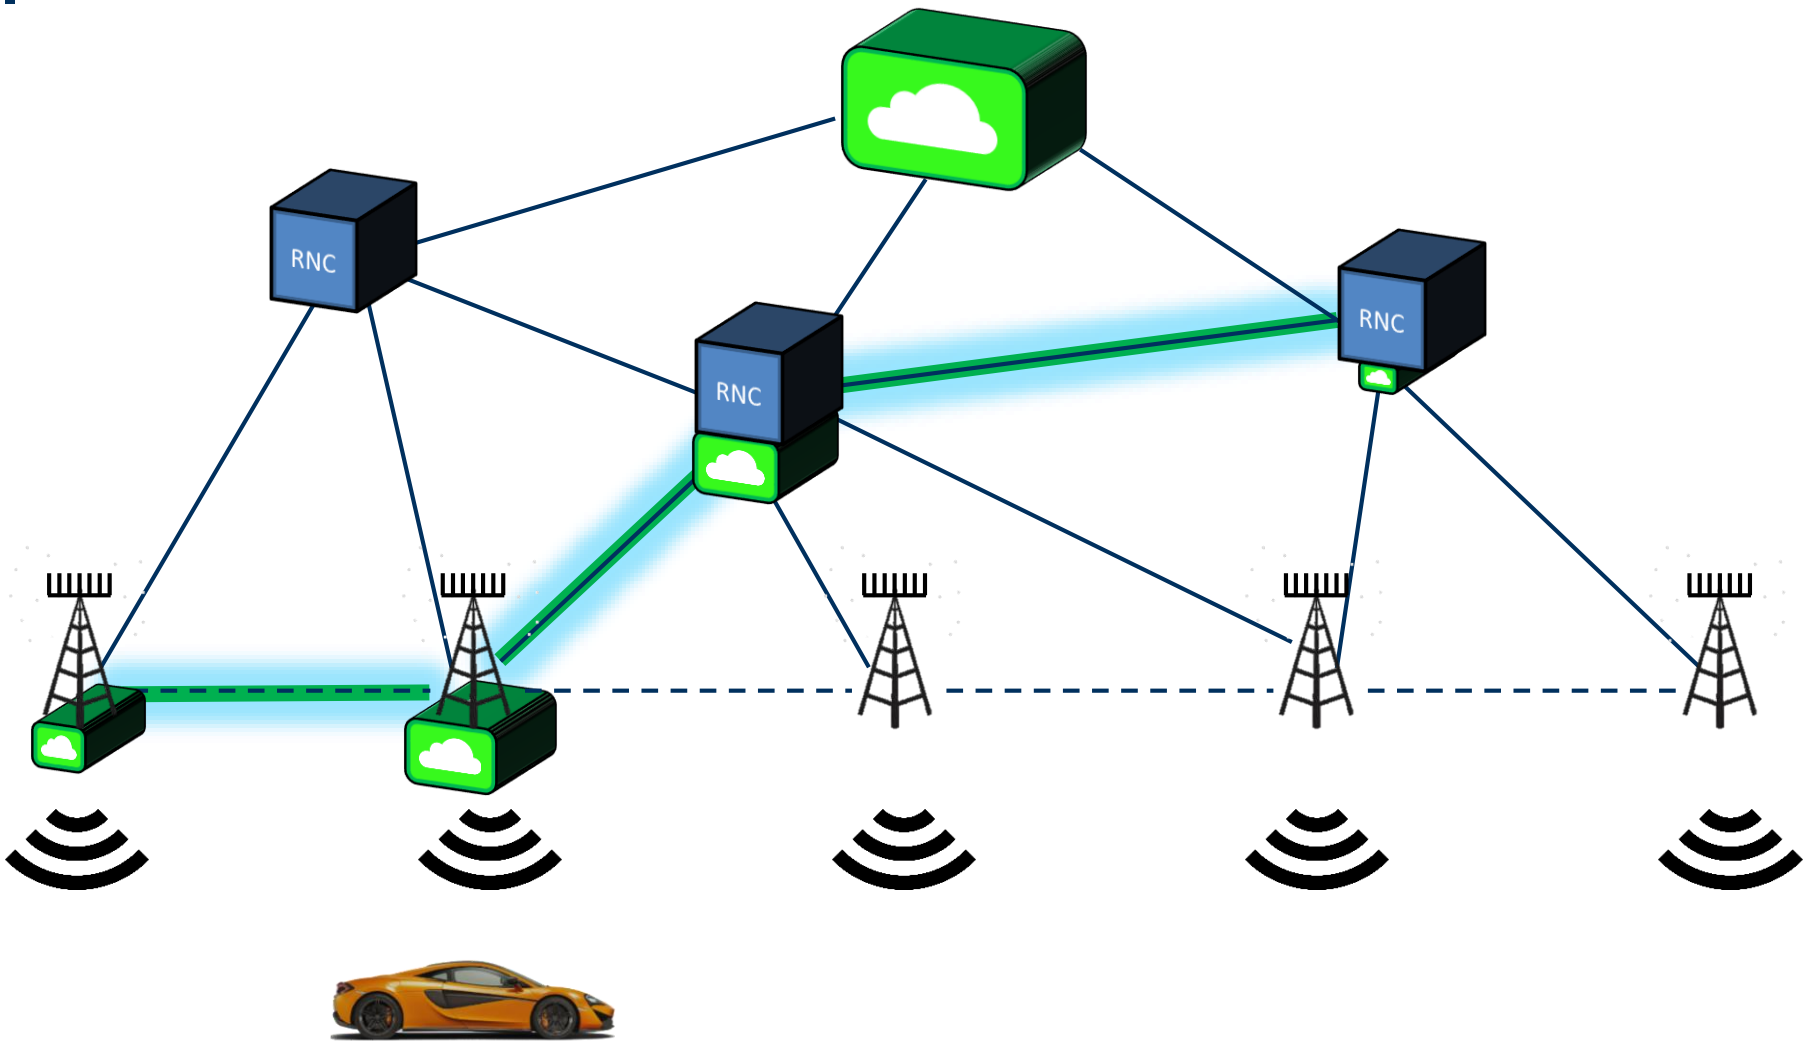

So…separating Framework and State?

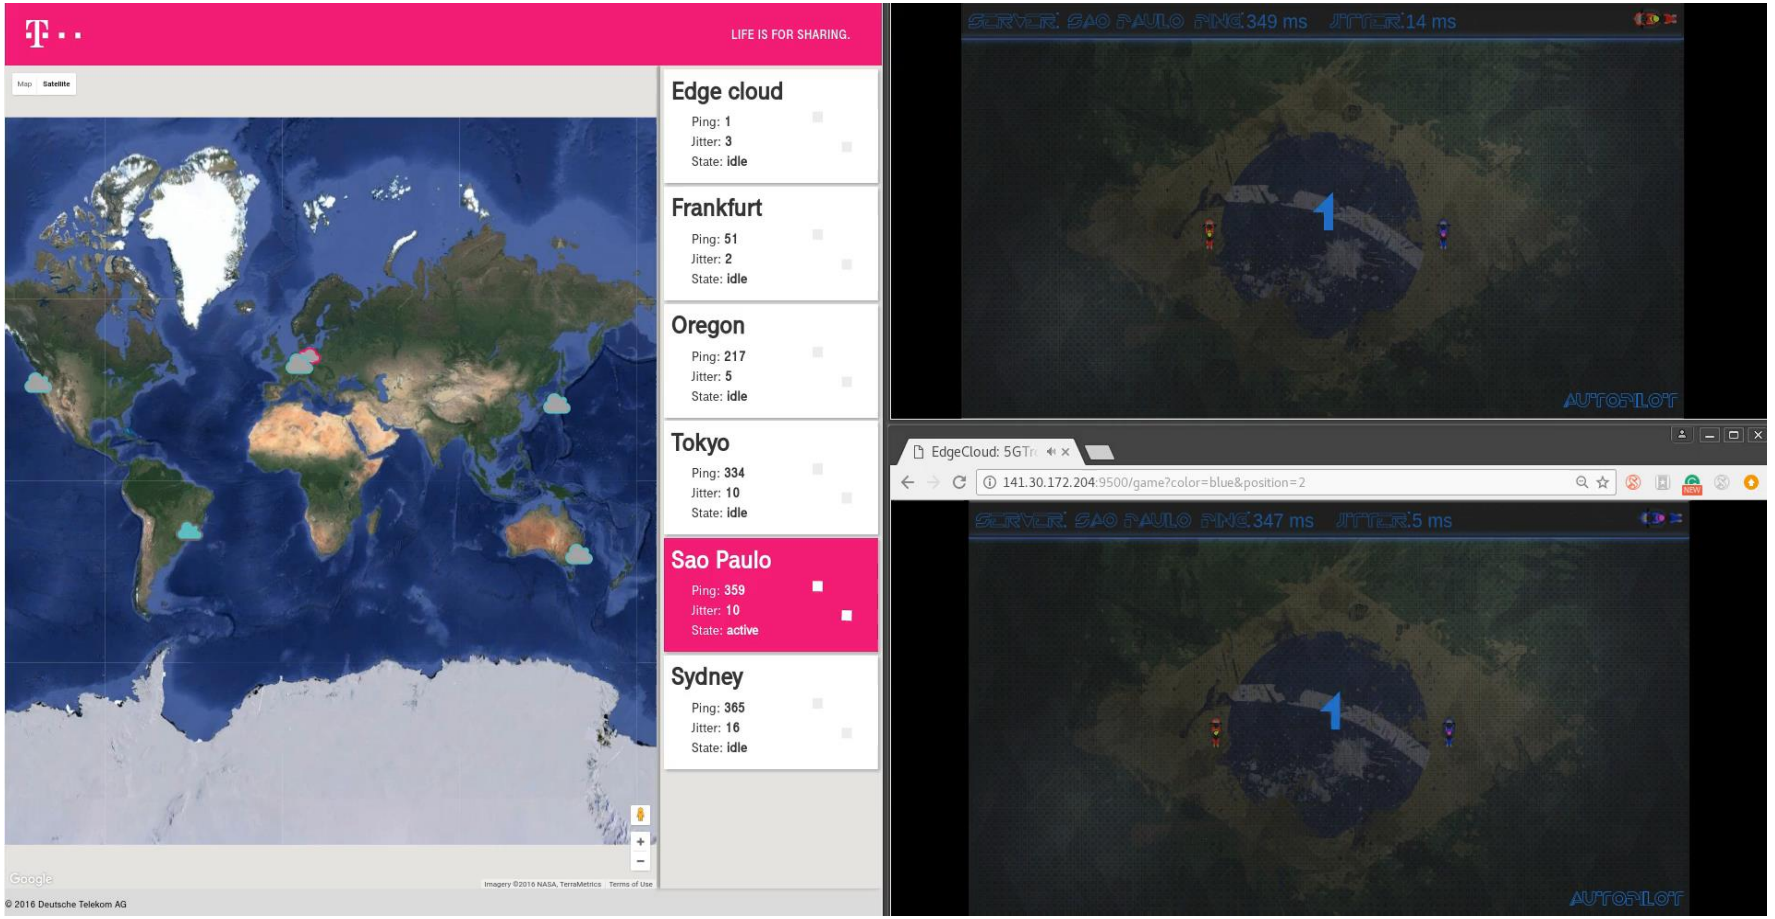

Data - II

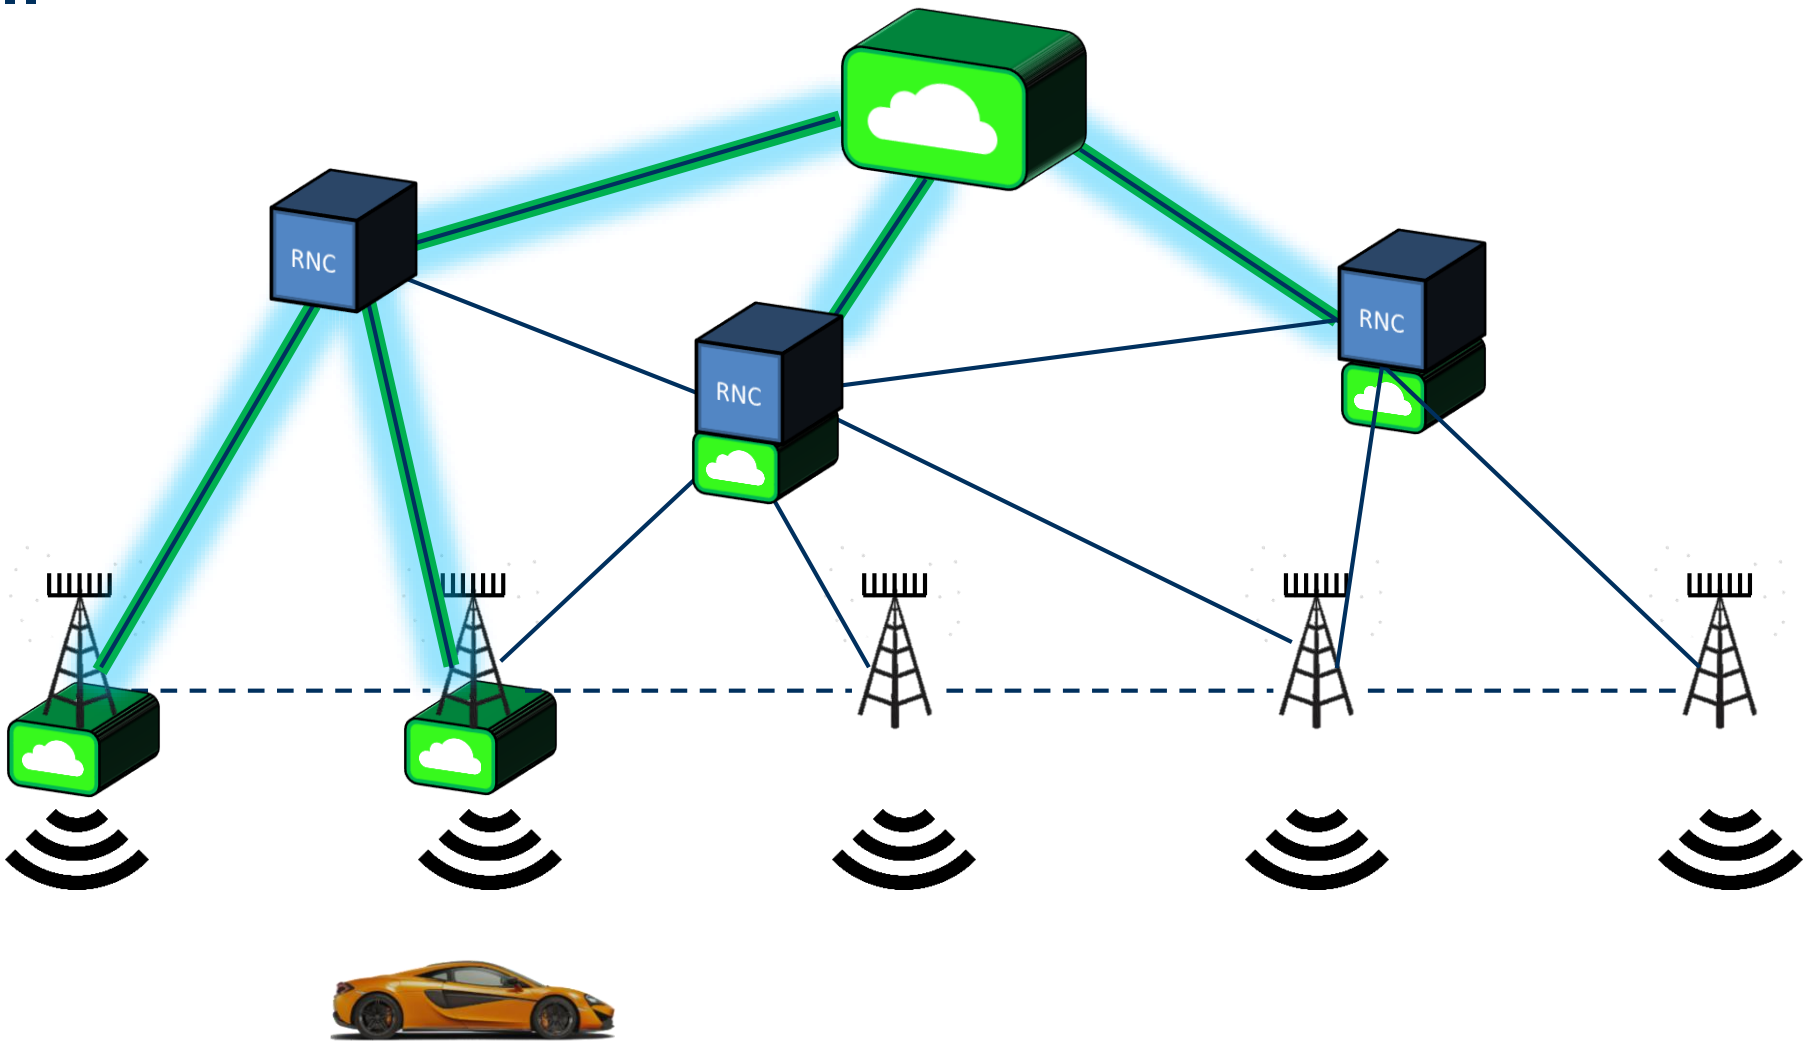

So…distributed data?

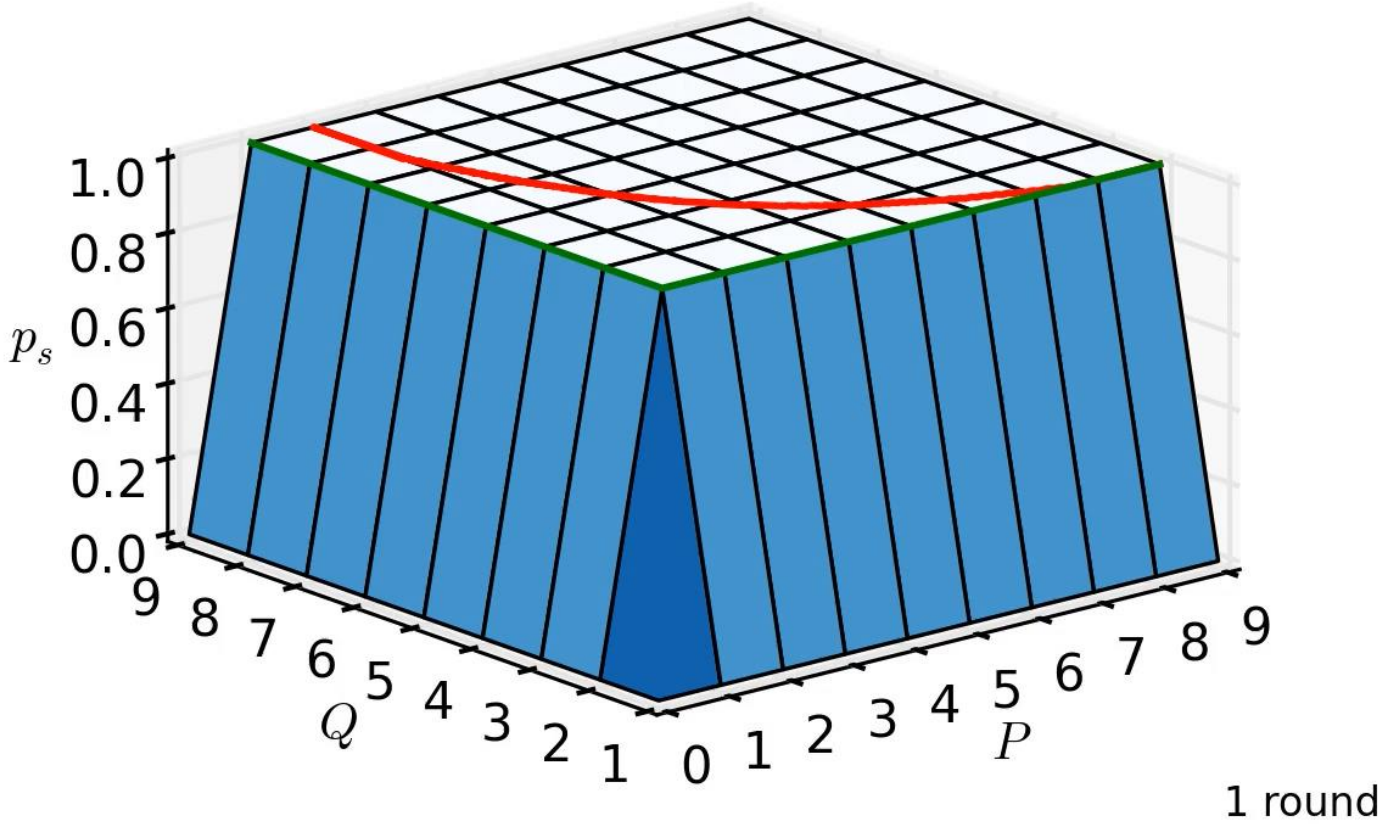

Thank you for your attention