13th International Symposium on Communication Systems, Networks and Digital Signal Processing
July 20-22, 2022
Porto, Portugal

CSNDSP 2022
Biennial since 1998

CHALMERS
UNIVERSITY OF TECHNOLOGY

Optical Transport Networks in the 5G Era

Paolo Monti
Electrical Engineering Department
Chalmers University of Technology
Gothenburg, Sweden

Outline

- Network infrastructure evolution
- ML-based network optimization
- Cloud native orchestration
- Open challenges
- Summary

https://cloudify.co/blog/network-automation-at-the-edge/

Telco infrastructure evolution

- Different generation of network infrastructure rollouts:
  - from **one** design/deployment per service
  - to **multipurpose** infrastructure orchestrating diverse resources with different requirements (e.g., latency, capacity, availability)
- Telcos are undergoing a “**digital” transformation** in both how they use their underlying technologies and their interaction with customers

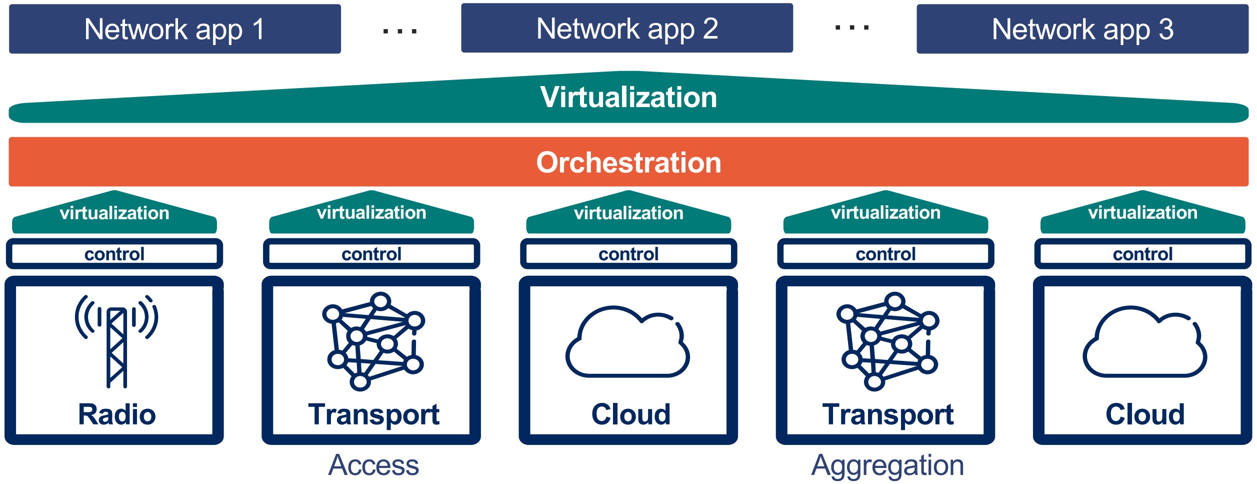

Digital transformation: benefits

Operate networks with *optimized resources*

Better *costumer experience*

Ability to adapt to market changes and lead in *innovation*

Expanding service portfolios addressing new *vertical markets*

Increased return on investment (*RoI*)

Digital transformation: critical aspects

Shift from discrete network elements to an independently managed, ***virtualised communications*** and ***cloud*** *infrastructure*

Data: deployment of ***telemetry frameworks*** for new approach to the collection, analysis/visualization, distribution, and security of data collected from multiple sources

Security: digital services have higher security requirements, need to support full technology stack, the data, the service creation process, and the physical environment

Architecture: **open platforms** and **standardized APIs** to support both internal own-brand and external third-party services

Business models: critical to develop new flexibility in what to sell and how create value for both themselves and their partners

Open & disaggregated optical networks

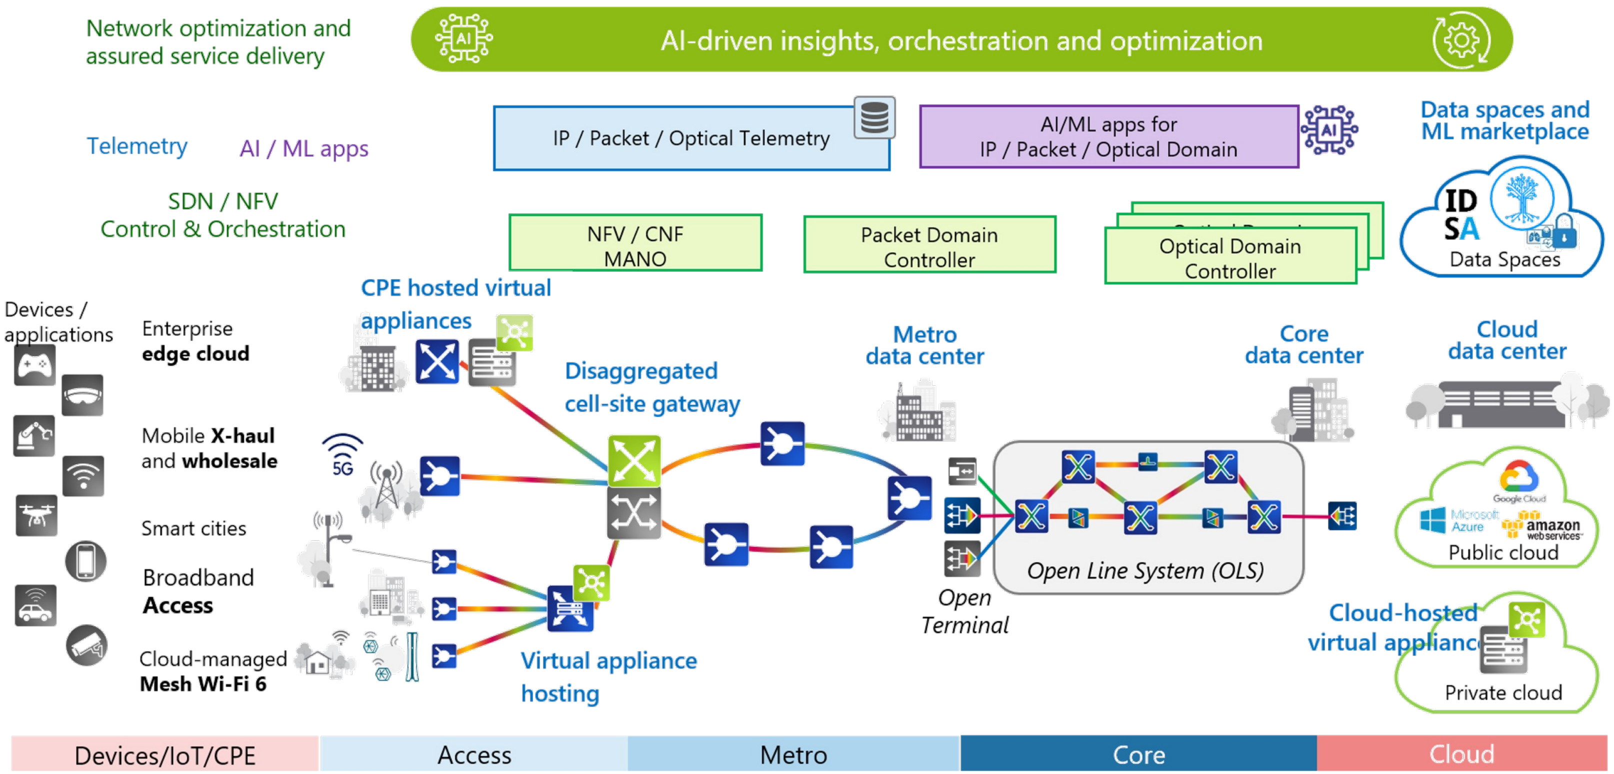

Achim Autenrieth, "Carrier Grade AI/ML for Network Automation", invited talk, OFC 2022, 9 March 2022

End-to-end network optimization

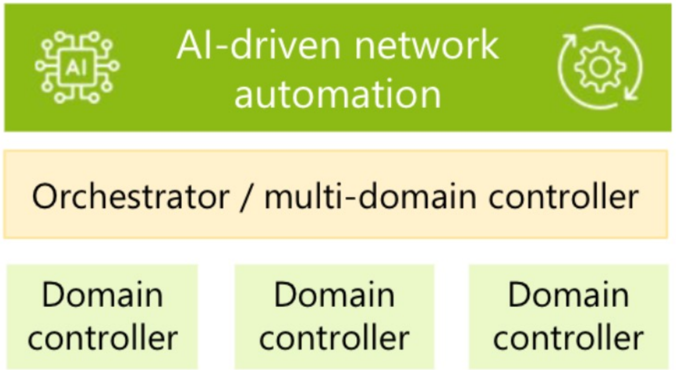

Management and orchestration

- AI driven optimization
- Cloud native
- Open APIs (SDN)

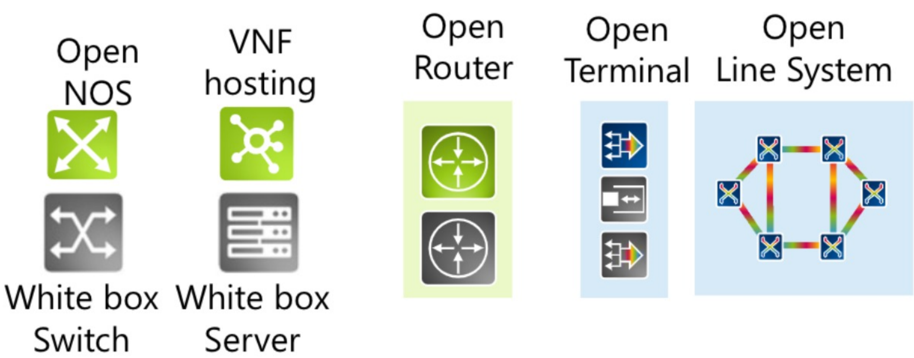

Infrastructure: open and disaggregated

- No vendor lock-in
- Streaming telemetry

Achim Autenrieth, “Carrier Grade AI/ML for Network Automation”, invited talk, OFC 2022, 9 March 2022

AI/ML for automation and optimization

AI/ML

- Data-centric analysis of network conditions
- Streaming telemetry: continuous data collection
- AI/ML use cases:
  - ✓ ML-based transmission performance optimization
  - ✓ Traffic prediction
  - ✓ Failure prediction
  - ✓ Anomaly detection and predictive maintenance
  - ✓ ML-based network optimization

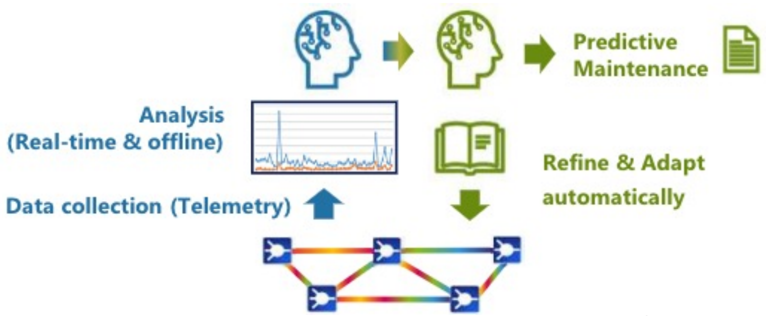

AI-driven network automation and optimization

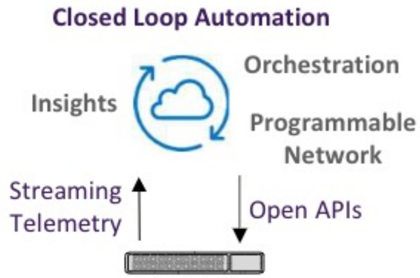

- Data become actionable insights for network and service performance optimization
- Focus on end-to-end optimization
- Combines real-time data collection, AI-driven analysis and orchestration to enable proactive optimization

Achim Autenrieth, "Carrier Grade AI/ML for Network Automation", invited talk, OFC 2022, 9 March 2022

Outline

- Network infrastructure evolution
- ML-based network optimization
- Cloud native orchestration
- Open challenges
- Summary

Network optimization

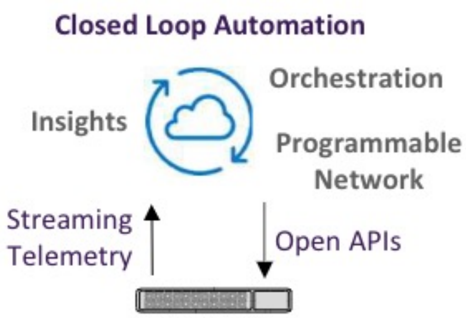

Reinforcement learning

Implementing a policy
(e.g., stochastic policy network)

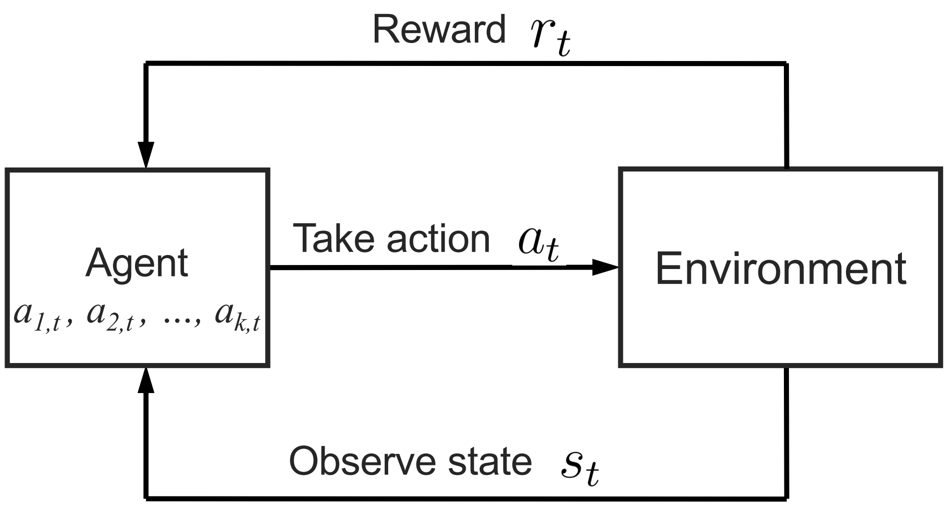

e.g., WDM network

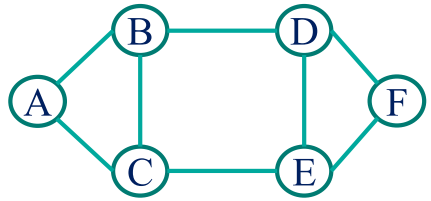

Goal: maximize the cumulative reward $\sum_t r_t$

ML-based slice scheduling

- Slicing: SDN and NFV allow InP to share resources among different tenants
- During provisioning/operation beneficial to adapt resources assigned to a slice to match time varying requirements: **dynamic slicing**
- Slice acceptance ratio can be greatly improved at the cost of small service degradation*
- Crucial to have intelligent policy that accepts only slices not likely to create performance degradation
  - understand when/where resource bottlenecks might appear in the infrastructure
  - deciding which slice to accept in order to maximize the profit of an InP

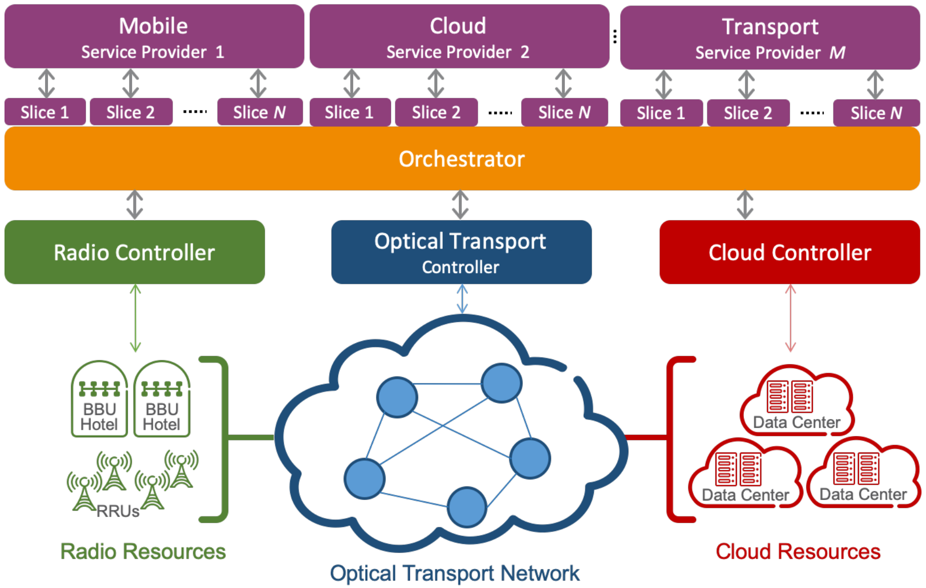

Use case

- Scenario:
  - tenant(s) requests slices with different requirements and priorities
  - different priorities mean different revenue and penalty levels
- Objective:
  - admission policy used by InP to accepts/rejects incoming slice request with aim to maximize Profit = revenue – penalty
- Intuition:
  - beneficial to proactively reject some low priority (low revenue) services to make space for future high priority (high revenue) ones

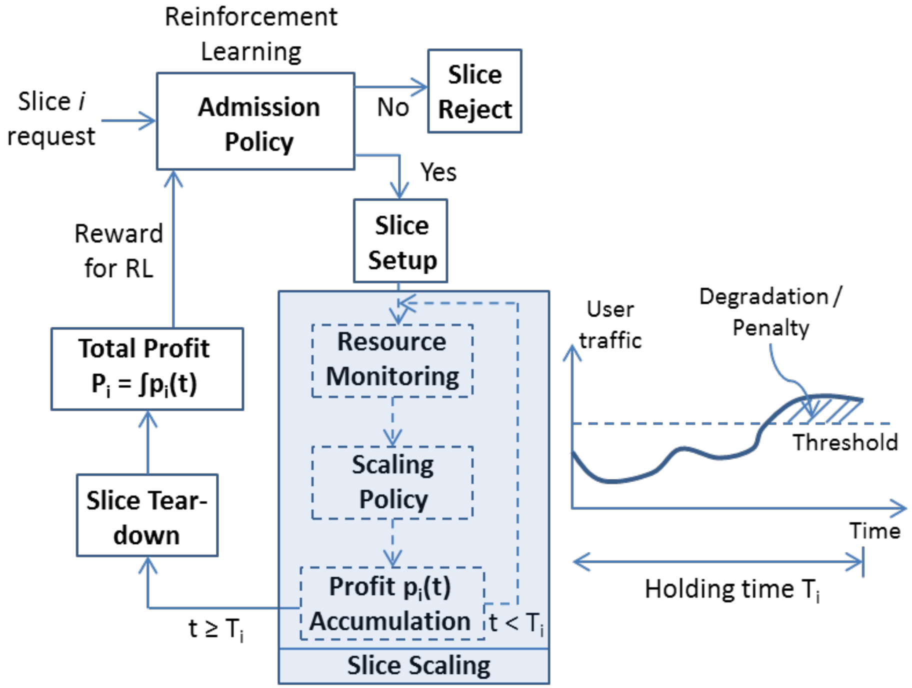

M. R. Raza, et al., "Reinforcement Learning for Slicing in a 5G Flexible RAN," JLT, 2019

Performance evaluation

- Scaling policy = high priority first (HPF), with 50% HP-50% LP services
- L_total= L_rej+ L_deg (sum of the rejection loss and degradation loss), penalty ratio = 1.5
- NN with 4 hidden layer and 40 neurons
- Test results after 2500 training iterations
- RL shows 23% improvement vs. OS-50, 60% vs. RR-50, and 53% vs. Fit

RR: Resource reservation
OS: Oversubscription

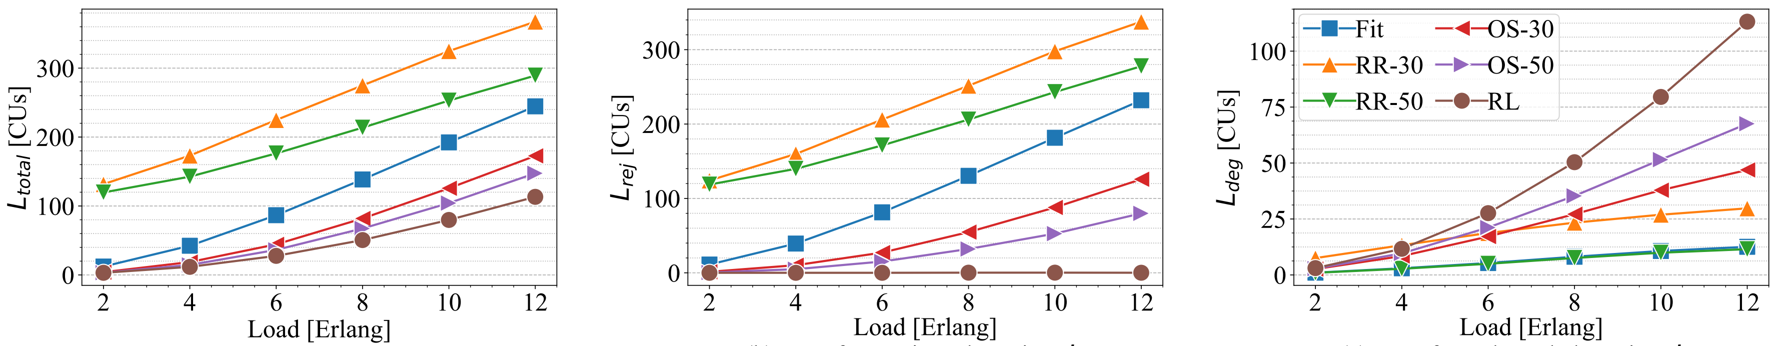

(a) Total loss L_{total} (b) Loss from rejected services L_{rej} (c) Loss from degraded services L_{deg}

M. R. Raza, et al., "Reinforcement Learning for Slicing in a 5G Flexible RAN," JLT, 2019

Intelligent spectrum management

- *Spectrum* is *fragmented* due to services arriving/leaving the network dynamically
- Spectrum Defragmentation (SD): reorganize spectrum allocation of different optical connections
- Main objective: decrease service blocking ratio
- Operators are not happy with frequent SD cycles and large number of reallocations
- Need to decide at run time and in one shot:
  - ✓whether or not to run defragmentation
  - ✓the number, and the order of connections to be reallocated
  - ✓which new spectral resources to use
  - ✓reallocating the connections

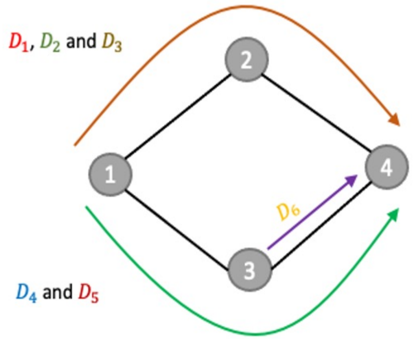

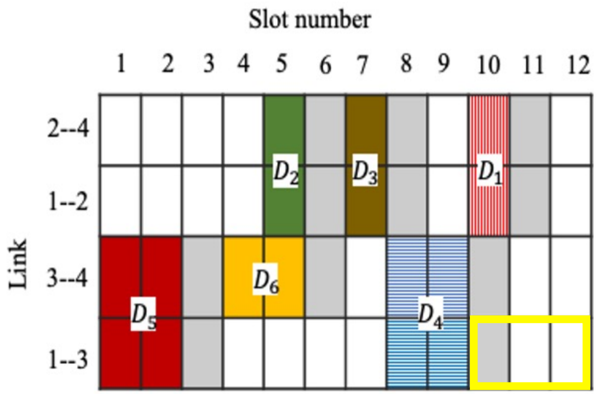

DeepDefrag: spectrum management with RL

- *DeepDefrag* runs at every service departure
- Input (environment):
  - existing connections
  - reallocation options
- Output (action space):
  - SD yes/no
  - Subset of connections to be reallocated
- Reward
  - Cost/benefit of the SD
  - Penalty for connection reallocation
  - Penalty for starting a new SD cycle

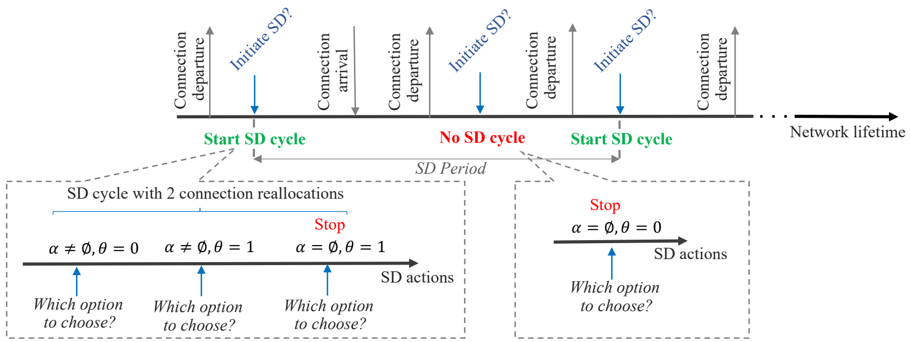

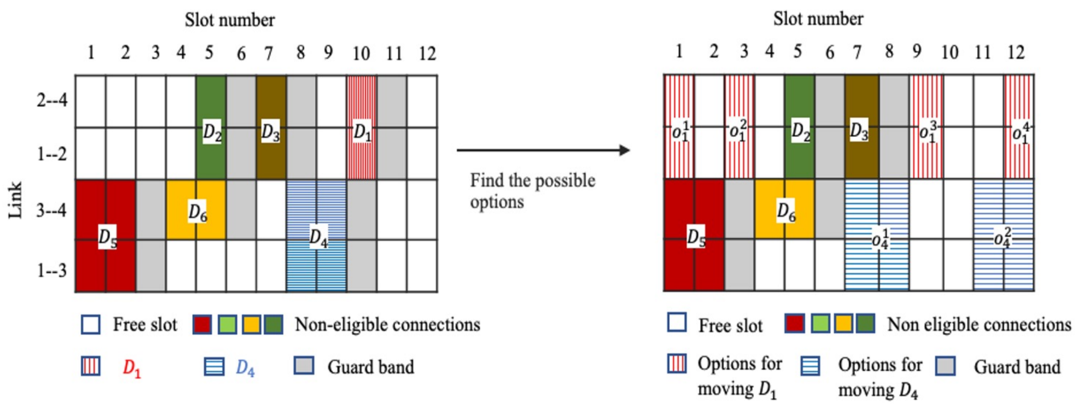

E. Ethezadi, et.al.,, "DeepDefrag: a deep reinforcement learning framework for spectrum defragmentation", online at: https://www.techrxiv.org/articles/preprint/DeepDefrag_a_deep_reinforcement_learning_framework_for_spectrum_defragmentation/20013458

Performance evaluation

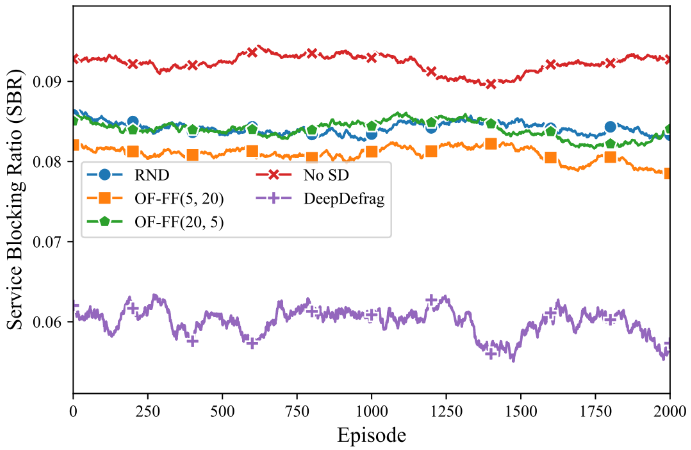

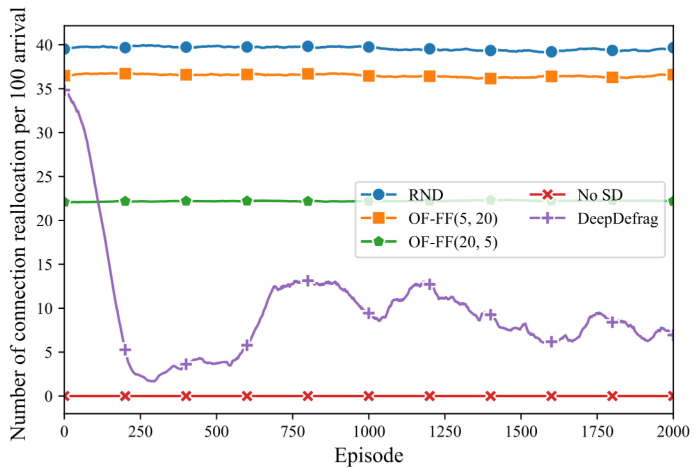

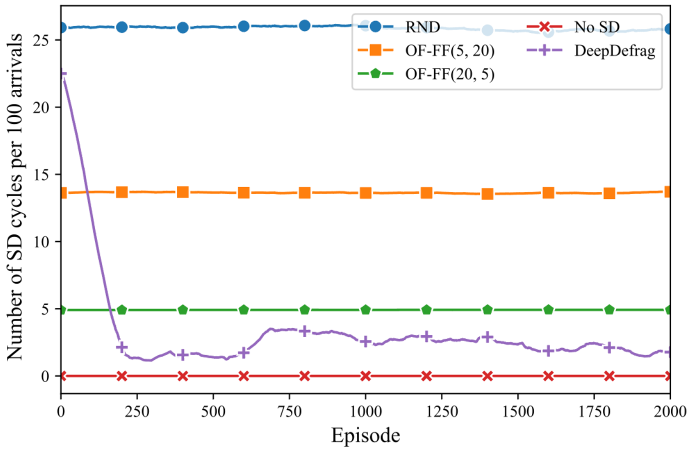

- Use NSFNET topology
- Benchmarks: Random (RND) and Oldest First First-Fit (OF-FF)
- DeepDefrag 25% better SBR performance compared to OF-FF while still having similar or better numner of relocation and SD cycles performance

E. Ethezadi, et.al.,, "DeepDefrag: a deep reinforcement learning framework for spectrum defragmentation", online at: https://www.techrxiv.org/articles/preprint/DeepDefrag_a_deep_reinforcement_learning_framework_for_spectrum_defragmentation/20013458

Outline

- Network infrastructure evolution
- ML-based network optimization
- Cloud native orchestration
- Open challenges
- Summary

Management and orchestration

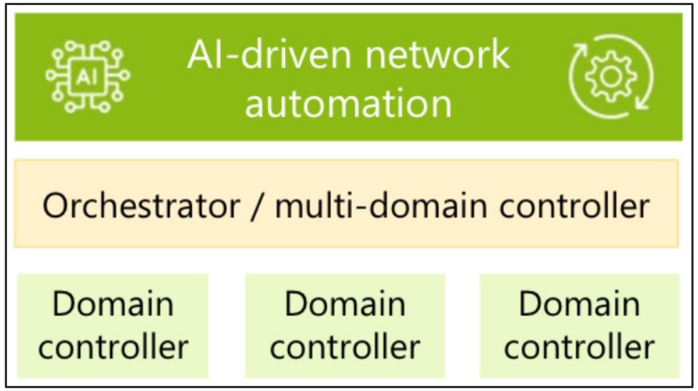

TeraFlow: a cloud native SDN controller

- SDN is a proven architecture
  - logically centralized controller (monolithic software)
  - deployed and scaled/replicated as a whole
  - scalability has always been a concern in SDN [1]
  - synchronization/consistency among replicas is essential [2]

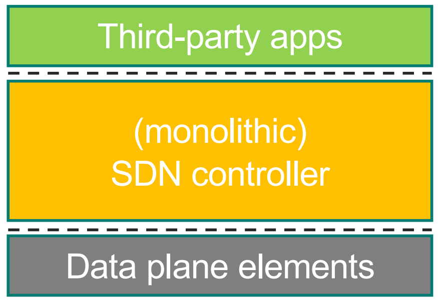

- Microservice architecture is one of the ways to mitigate scalability issues
  - Defining components and the communication among them is crucial
  - TeraFlow[3]: microservice-based SDN controller

[1] A. Voellmy and J. Wang, “Scalable software defined network controllers,” ACM SIGCOMM 2012
[2] K. Foerster, et al., "Survey of Consistent Software-Defined Network Updates," in IEEE Communications Surveys & Tutorials, 2019
[3] R. Vilalta, et. al. “TeraFlow: Secured Autonomic Traffic Management for a Tera of SDN Flows,” in Proc. EUCNC, 2021

TeraFlow architecture and scenarios

CHALMERS
UNIVERSITY OF TECHNOLOGY

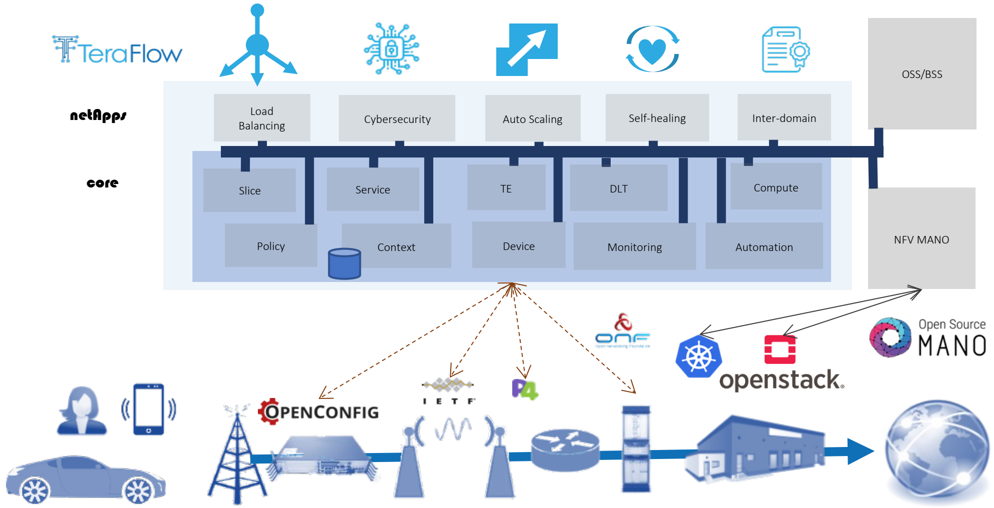

Autonomous network beyond 5G

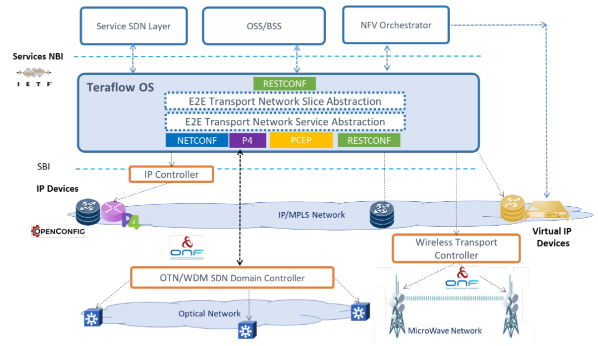

Inter domain

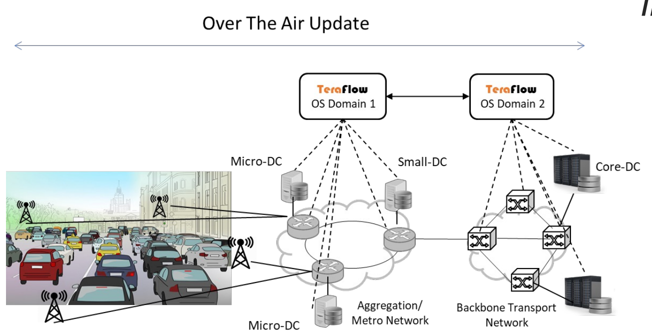

Cyber-security

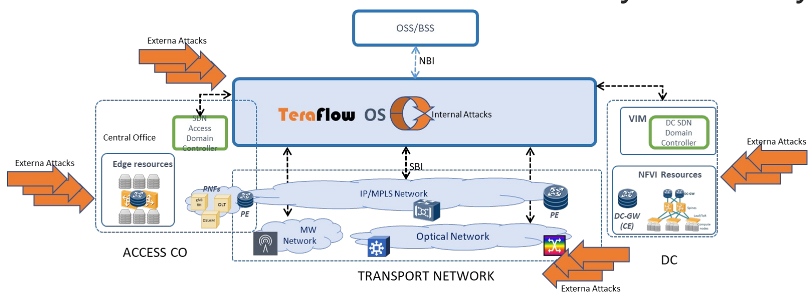

Scalable security component

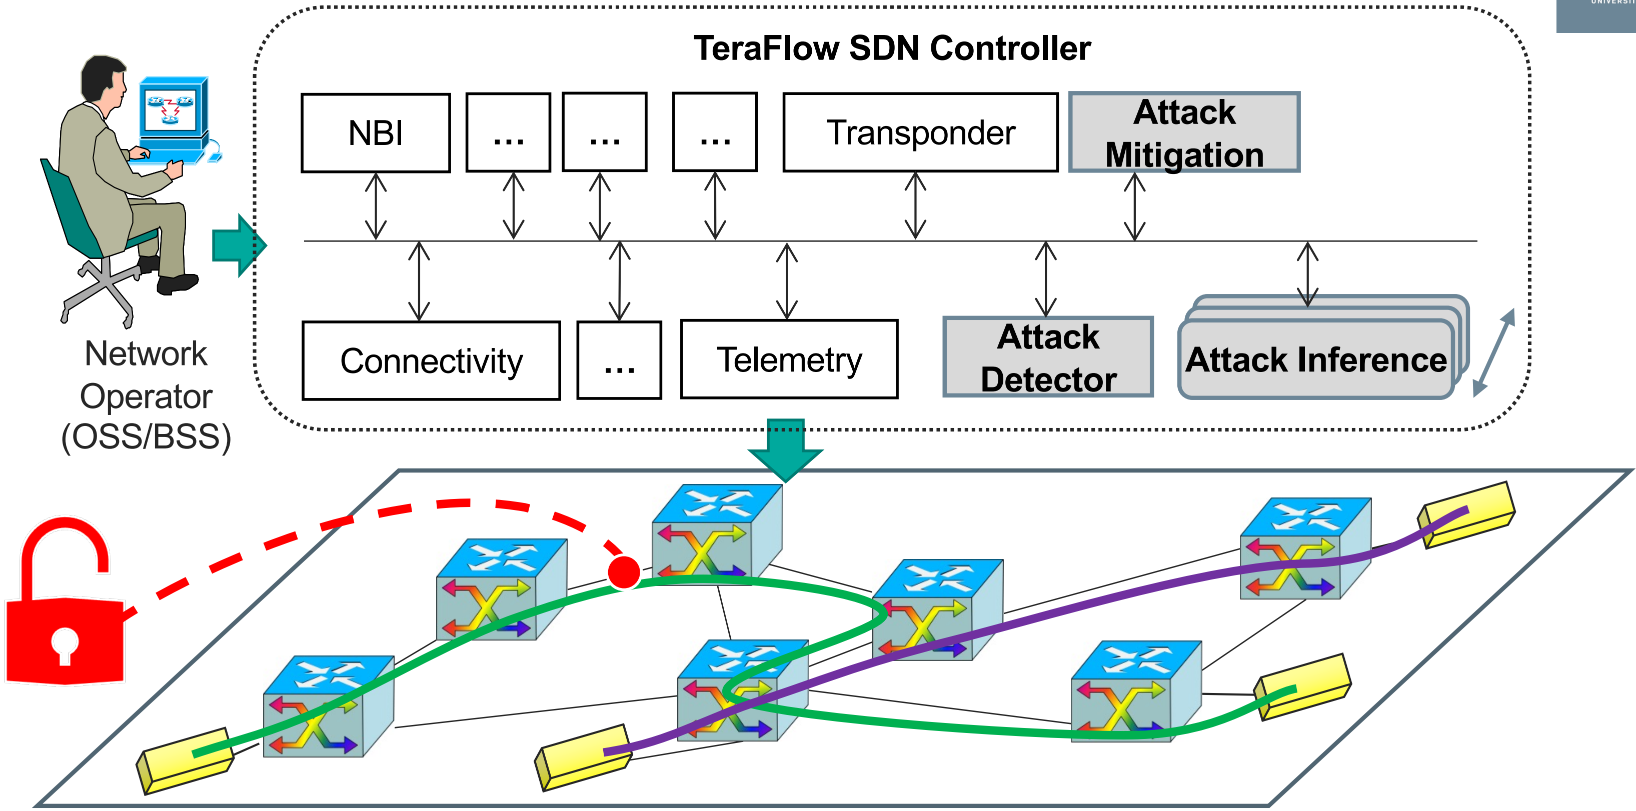

Experimental setup

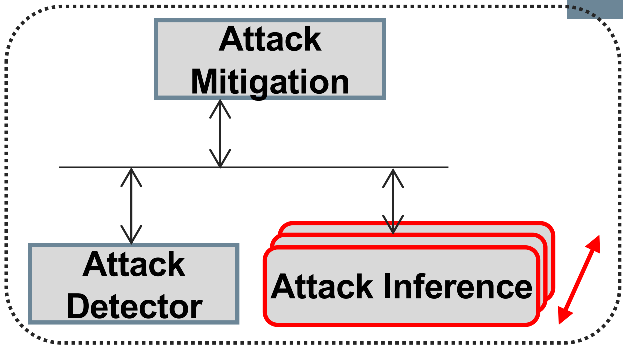

- TeraFlow deployment over Kubernetes
- Dataset obtained from testbed [1]
  - Two scenarios:
    - Normal working conditions
    - Two types of attacks: in-band jamming and out-of-band jamming
- *Attack inference* module
  - Feed-forward neural network[2] receives OPM data and performs attack detection and identification
  - Minimum of 1 replica* (for the sake of this work, 2 is the minimum for production environments)
  - Kubernetes scales when CPU utilization exceeds a threshold (80% CPU usage)
- Monitoring loop is executed every 30 seconds

[1] M. Furdek et al., “Machine learning for optical network security monitoring: A practical perspective”, JLT, vol. 38, no. 11, pp. 2860–2871, 2020.

Attack inference performance

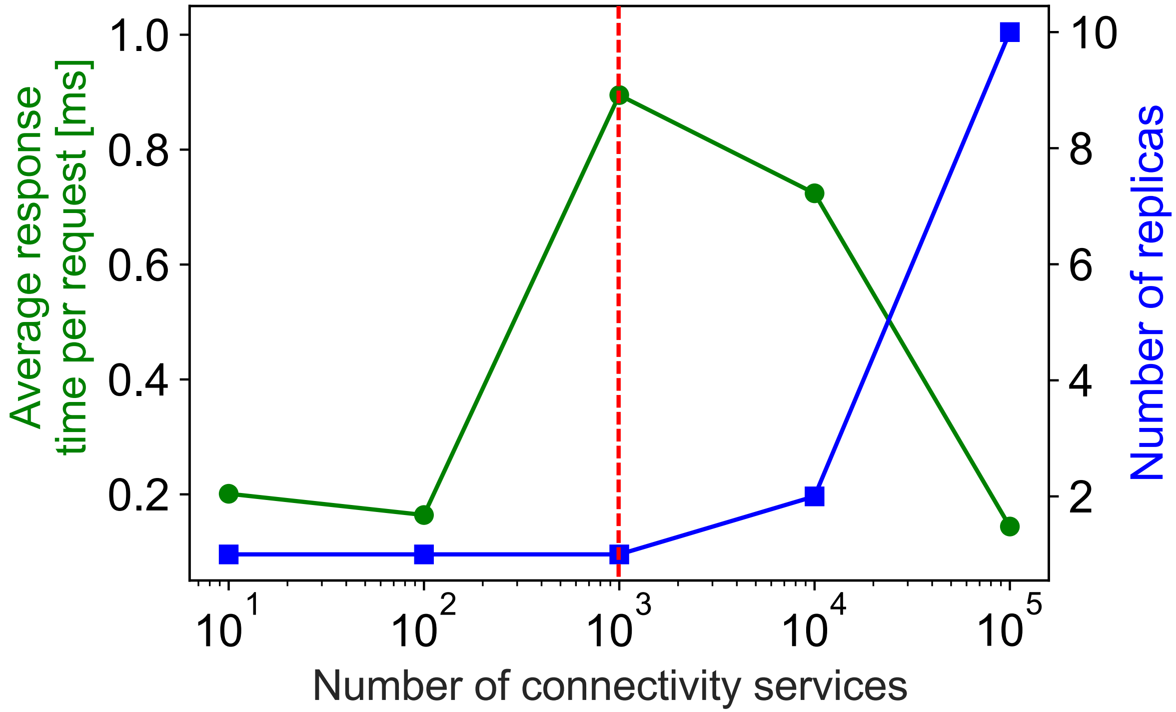

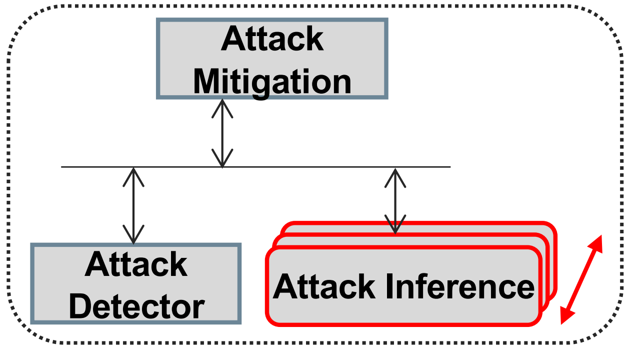

- TeraFlow deployment over Kubernetes
- Up to 1,000 connectivity services served by a single replica
- Response time is contained within 1 ms

C. Natalino, et. al., "Scalable Physical Layer Security Components for Microservice-Based Optical SDN Controllers," in Proc. 47th European Conference on Optical Communication (ECOC), Sept. 13-16, Bordeaux, France, 2021

TeraFlow demo @ OFC 2022

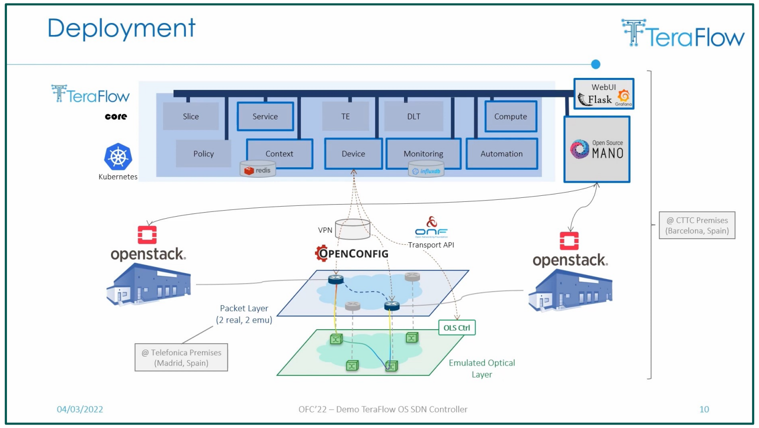

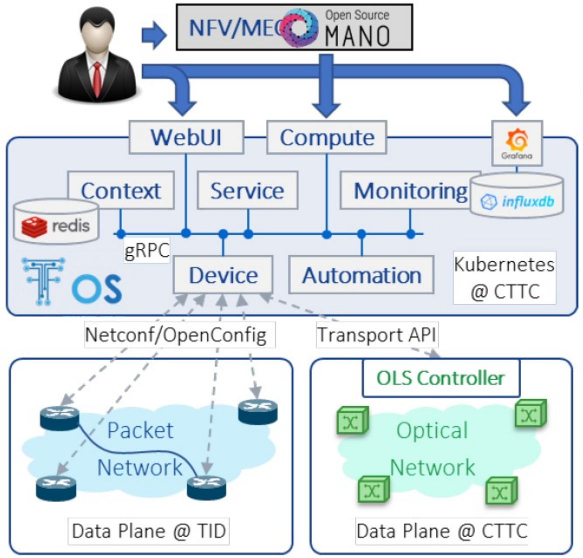

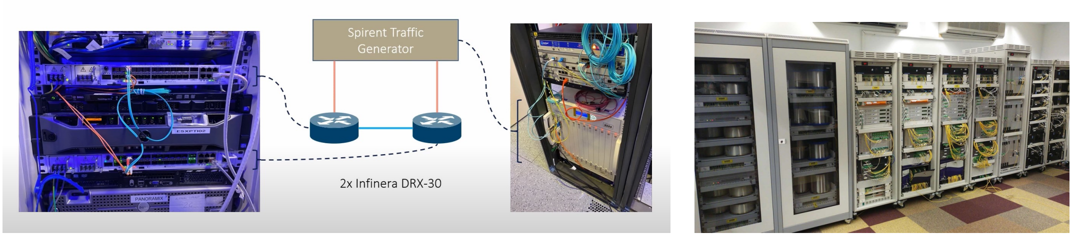

Some open challenges

Where to start and with which tasks? Decide use case of interest, and then the elements and functionalities that should be automated

Trust AI to act on the network autonomously? What's the level of human intervention? What skill will be required for network engineers?

Cost of retrieving data? Data availability, storage in the network elements for online access, cost of H/W and S/W are key elements

Confidentiality requirements. Which data to share among vendors and operators? How can this data sharing/data exchange be standardized?

How to perform E2E planning/operation of multi-domain networks in an automated context

Scalability with the number of services of the "knowledge loop"

Impact of network automation on other network KPIs (e.g., energy consumption)

Accountability and traceability of action taken

In summary

Virtualization, telemetry frameworks, AI, open architectures enable a paradigm shifts towards self driving networks

Potential benefits are evident:

- optimization of network resources
- better costumer experience
- flexible service portfolio and short time to market
- increase RoI

… but substantial work still needed:

- data availability
- standardized APIs
- security/privacy/accountability/traceability
- scalability

Success dictated not only by ability to overcome the technical challenges but also on how to best leverage the new businnes ecosystem that will be created

Acknowledgments

Twitter: @CELTIC_AI_NET
@AI_NET_PROTECT
Web: https://ai-net.tech/protect

Twitter: @TeraFlow_h2020
Web: www.teraflow-h2020.eu
TeraFlow GitLab: https://gitlab.com/teraflow-h2020/controller
TeraFlowSDN ETSI group: https://tfs.etsi.org/about/

The CELTIC-NEXT project AI-NET-PROTECT (Project ID C2019/3-4) is partly funded by the German Federal Ministry of Education and Research BMBF, the Swedish Governmental Agency for Innovation Systems Vinnova, the Netherlands Ministry of Economic Affairs and Climate EZK, and the Polish National Center for Research and Development NCBR

This project has received funding from the European Union’s H2020 research and innovation programme under the grant agreement No. 101015857

13th International Symposium on Communication Systems, Networks and Digital Signal Processing
July 20-22, 2022
Porto, Portugal

Optical Transport Networks in the 5G Era

mpaolo@chalmers.se

AVANCEZ
1829
CHALMERS
UNIVERSITY OF TECHNOLOGY