Engineering Self-Driving Networks

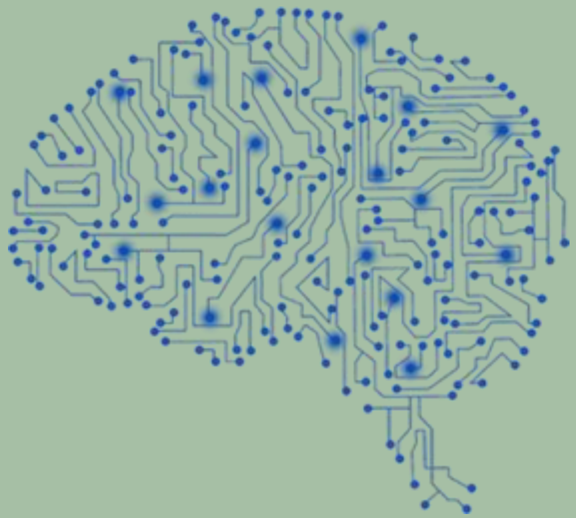

Mariam Kiran
2017 DOE Early Career
2021 N2Women Rising Star in Networking

Group Leader
Quantum Networking and Communications
Oak Ridge National Laboratory
Oakridge, TN

Introductions

- 2011 PhD in Computer Science (University of Sheffield, UK)
  - Optimizing HPC jobs in agent-based simulations
- Postdocs in Cloud Computing (University of Oxford, Leeds)
- Royal Society Scientist at Westminster London
- 2016 Joined ESnet, LBNL
- 2023 Joined ORNL as Group Lead for Quantum networks
- Research Area : AI for Networking - **Self-driving networks**
  - Impacts distributed science workflows
  - Expanding AI from optic networks to wireless and quantum networks
  - Contributions to other science areas:
    - Self-driving **lasers**
    - Self-driving **batteries**
    - Self-driving **quantum transducers**

    and more..

R&E Network for Large-scale Science

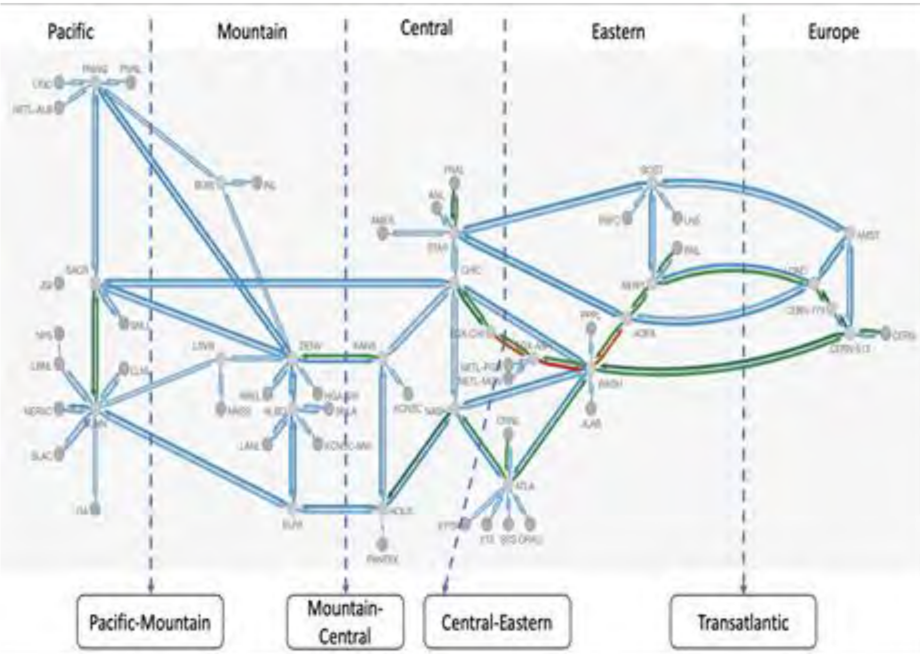

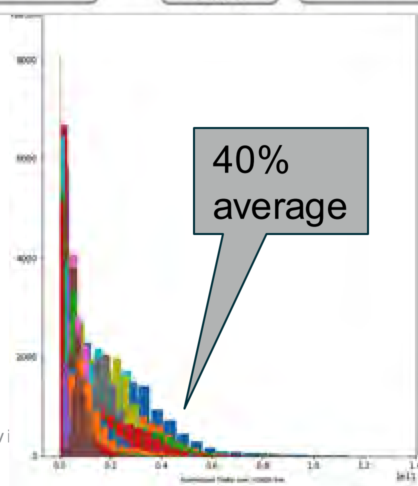

ESnet Network and its utilization in 2019 over 100GB links

ESnet from DOE

- Networks are built for resilience
- Science traffic is highly variable
- Resources are often underutilized and expensive
- Quality of Network Performance is crucial for Science

Need for **predictability** and infrastructure **adaptability**

Challenge: how can we <u>optimize and automate</u> network resources (i.e. links) to adapt to variable workloads?

“Self-Driving” a triggering word!

“Systems manage themselves according to an administrator’s goals. **New components integrate as effortlessly as a new cell establishes itself in the human body**. These ideas are not science fiction, but elements of the grand challenge to create self-managing computing systems.”

- Kephart, Chess (IBM) 2003

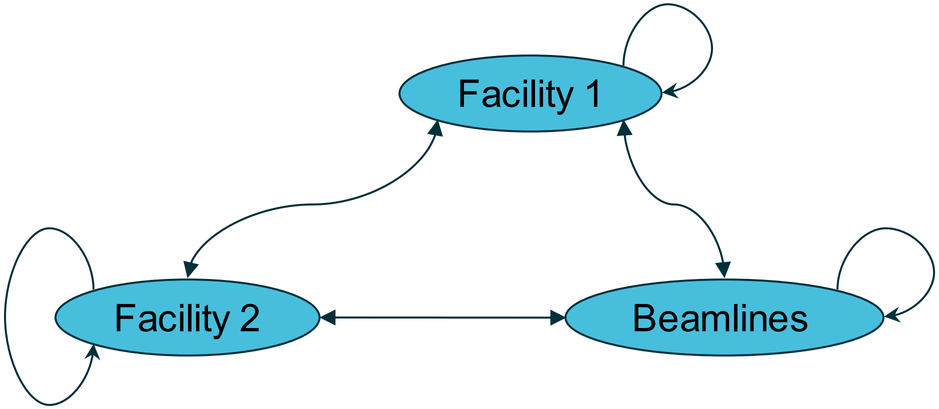

Vision: Future is Hyper Connected! with Self-driving Elements…

Self-Driving Networks for Science

Exponentially growing data rates

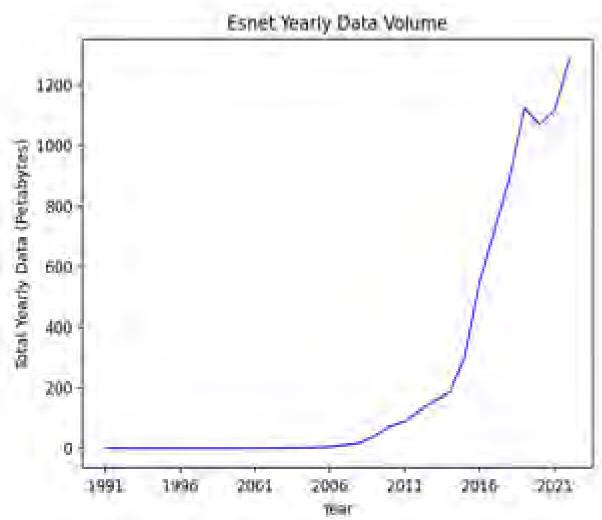

- Upto 58 GBytes/s week avg
- 50% increase/year & 60X/decade

Network beyond "lab borders"

- 10x decrease in latency
- 10x increase throughput
- 100x network performance, traffic capacity
- VR/AR, ability to handle real-time emergency edge intelligence

Evolving Edge and Quantum

- Increased instruction/s
- Machine learning at edge
- Access to more data
- Advent of Quantum Computations and transfers

DOE Early Career Research Project 2017-2023
Self-Driving 5G Network for Science 2021-2024
Poseidon (Intelligent Infrastructure for Science Workflows) 2021-2024

Early Career Research Project (2017-2023)

Large-scale and Deep Learning for Networking

Classification

Compute facility network data 2020

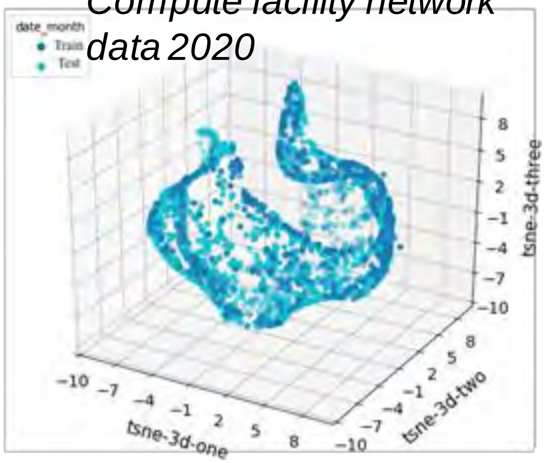

- Big data Challenges
- Data cannot be moved to HPC (due to security)

Prediction

Predicting CHIC-KANS link

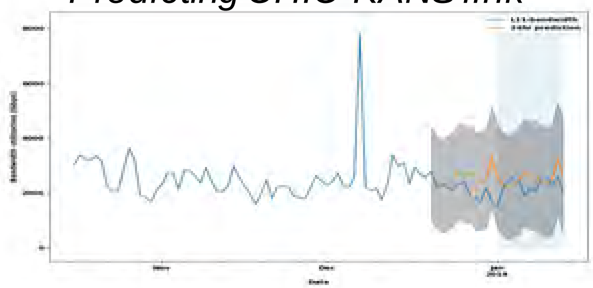

- Random peaks or sudden data transfers
- Capped utilization at 40%

Reaction with Control

Simulation

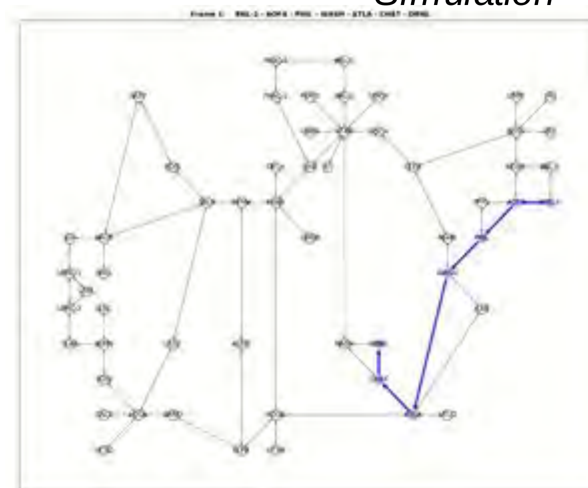

- Random peaks or sudden data transfers
- Capped utilization at 40%
- Integration with Network controllers

“Networks should learn to drive themselves”*

Can do simple actions such as improving availability, attack resilience and dealing with scale. Our argument is AI is needed for mission critical actions.

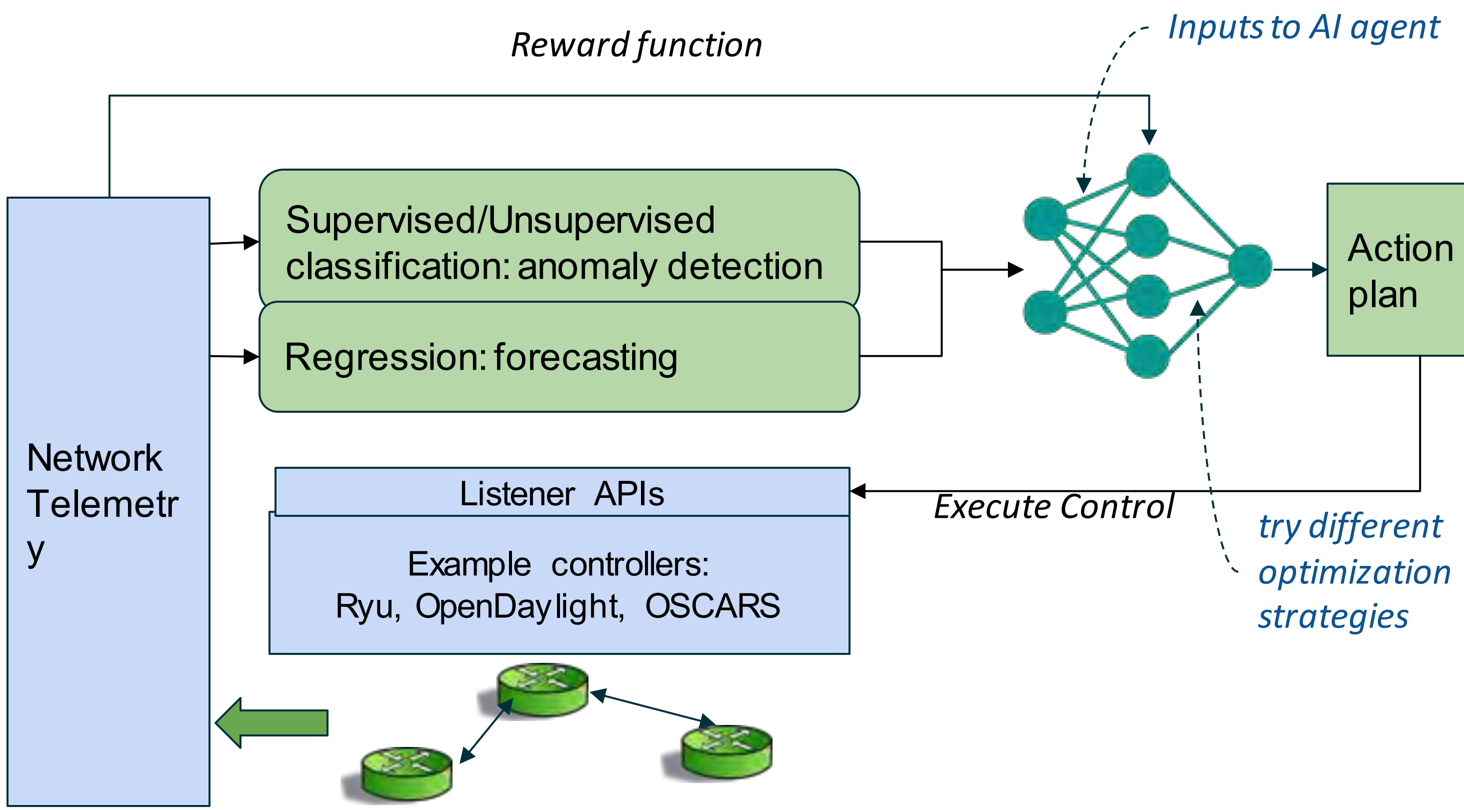

[*] Why (and How) Networks Should Run Themselves, Feamster, Rexford

Problem: Classification in Network Flows

Looking for ‘Cats’ in Flows

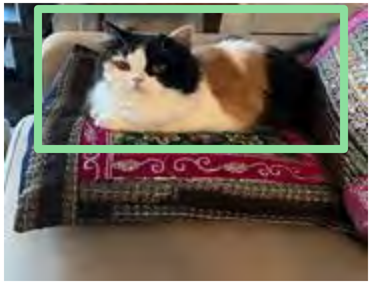

Distinct server connections, changes in traffic profiles

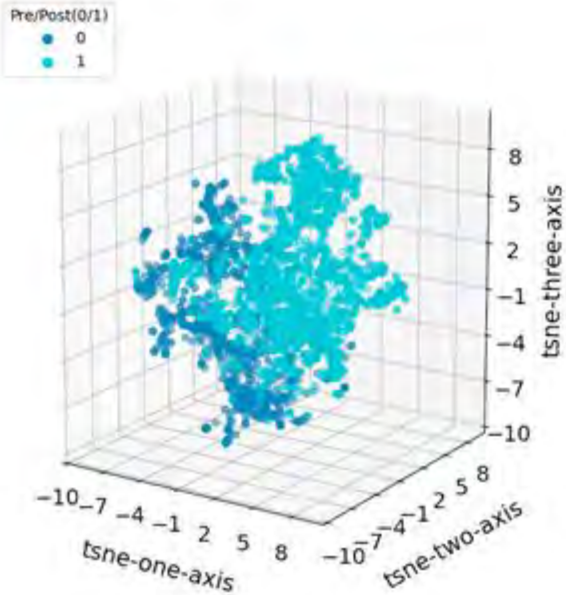

Covid experimentations at DOE Compute

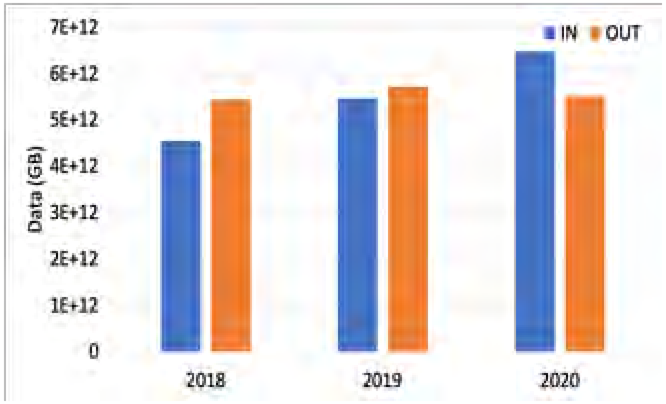

Scientist working patterns have changed

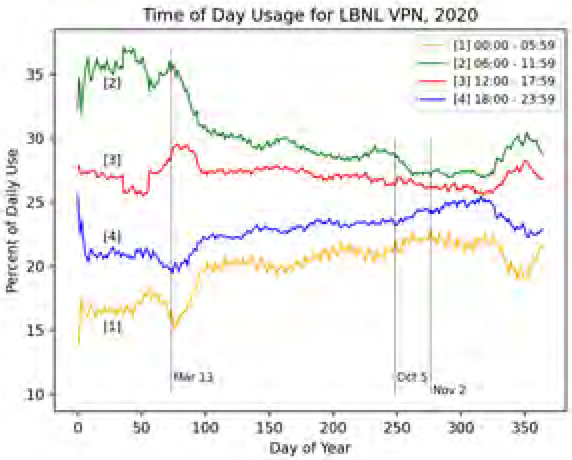

Kiran et al. Machine learning-based Analysis of COVID-19 Pandemic Impact on Research and Education Networks, ACM SIGCOMM Journal 2021

Looking for Patterns in Transfers

- No Periodicity exists
- Produces a smoothed out
- Trend, seasonal or noise can be used to enhance your predictions
- Statistical predictions do not perform well
  - Nature of data

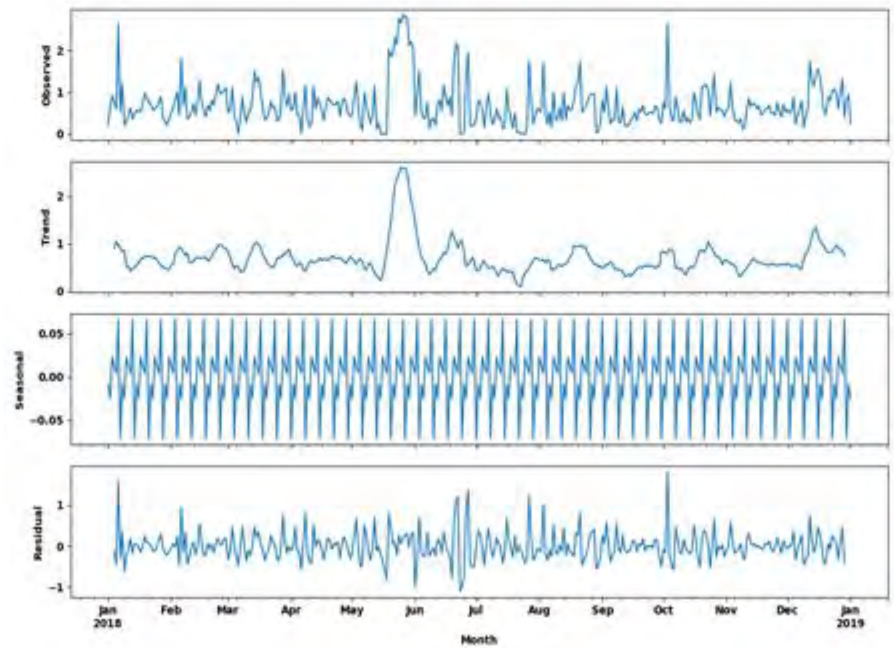

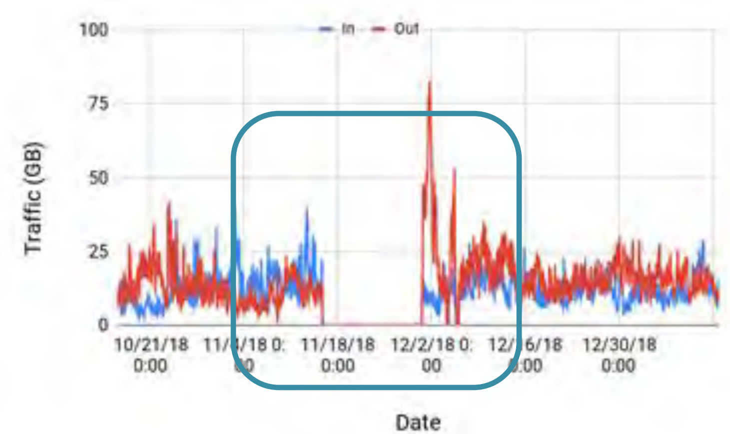

Patterns from Correlating Links can show Common Behaviors among Users

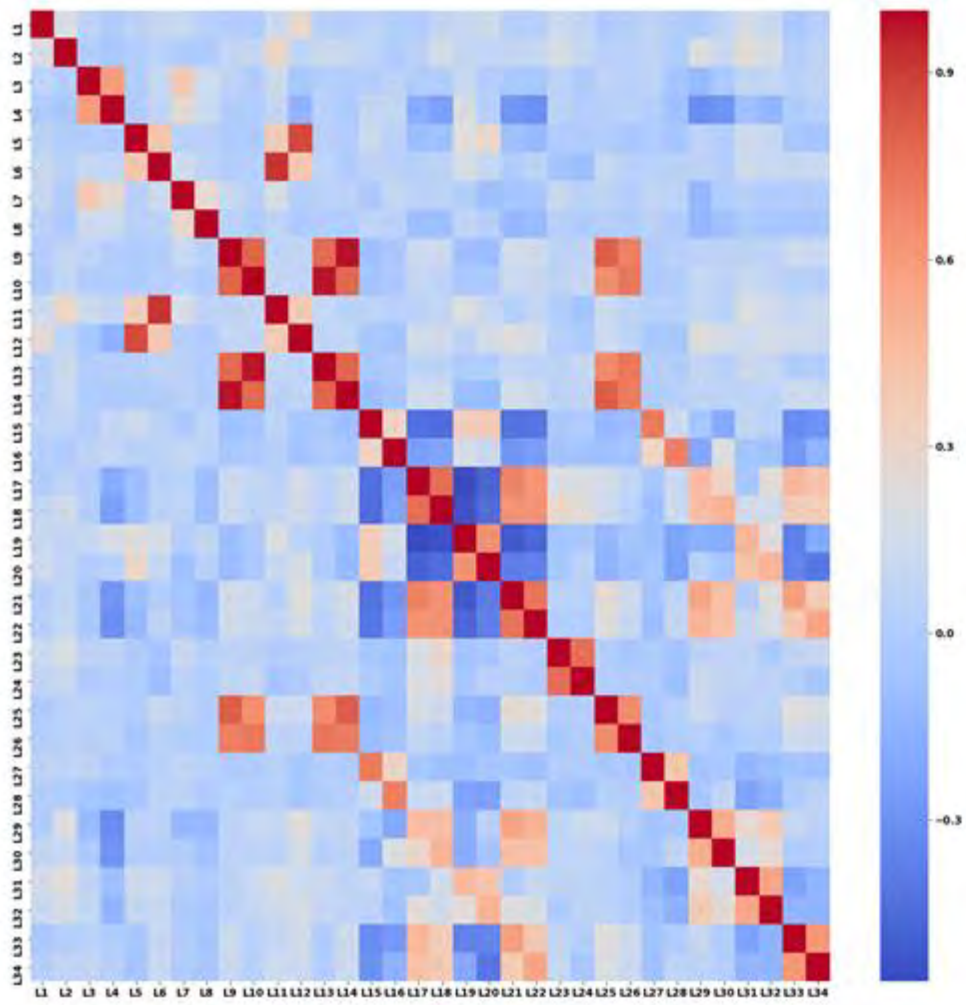

Mapping a Sudden Burst in May 2018:

- Traffic from SACR→DENV flows into DENV→KANS
- Traffic in SUN→ELPA flows into ELPA→HOUS

Problem: Prediction for WAN traffic

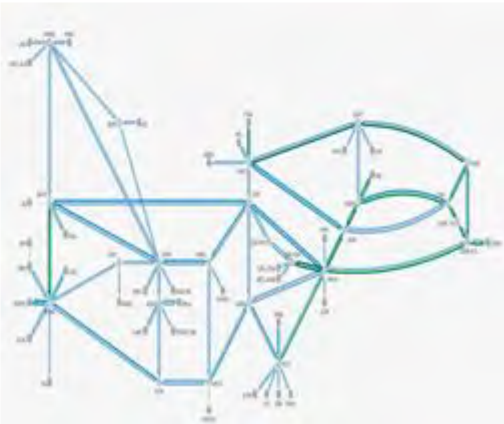

ESnet WAN network

Historic traffic metrics

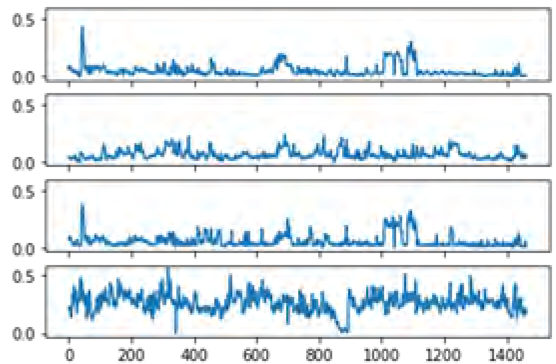

8.00 AM … 10.00AM

Future traffic metrics

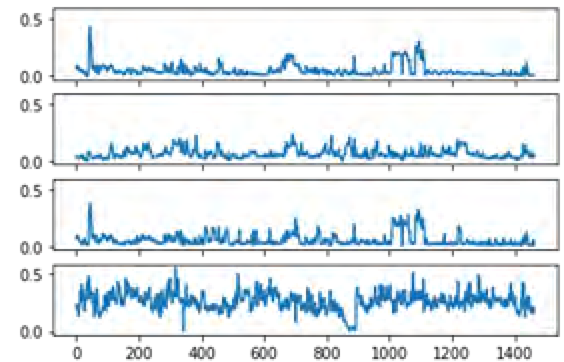

10.00 AM … 12.00PM

Solution: Graph neural networks

- Network as graph
  - V = Vertices (sensors)
  - E = Edges (roads)
  - A = Weighted adjacency matrix

(A function of the bandwidth, p Pearson correlation coefficient, cov covariance)

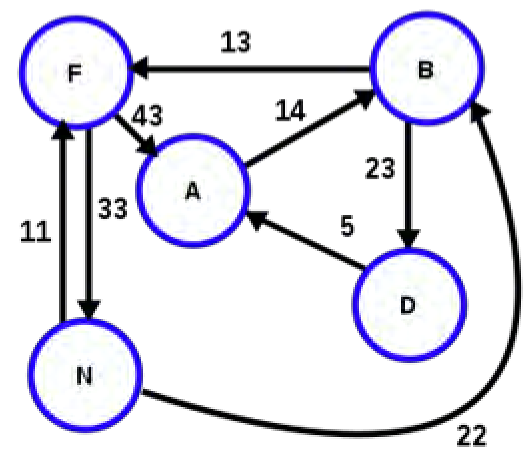

$$A_{ij} = \rho_{X_i,X_j} = \frac{cov(X_i, X_j)}{\sigma_{X_i}, \sigma_{X_j}},$$

Encoder-decoder architecture of D-DCRNN

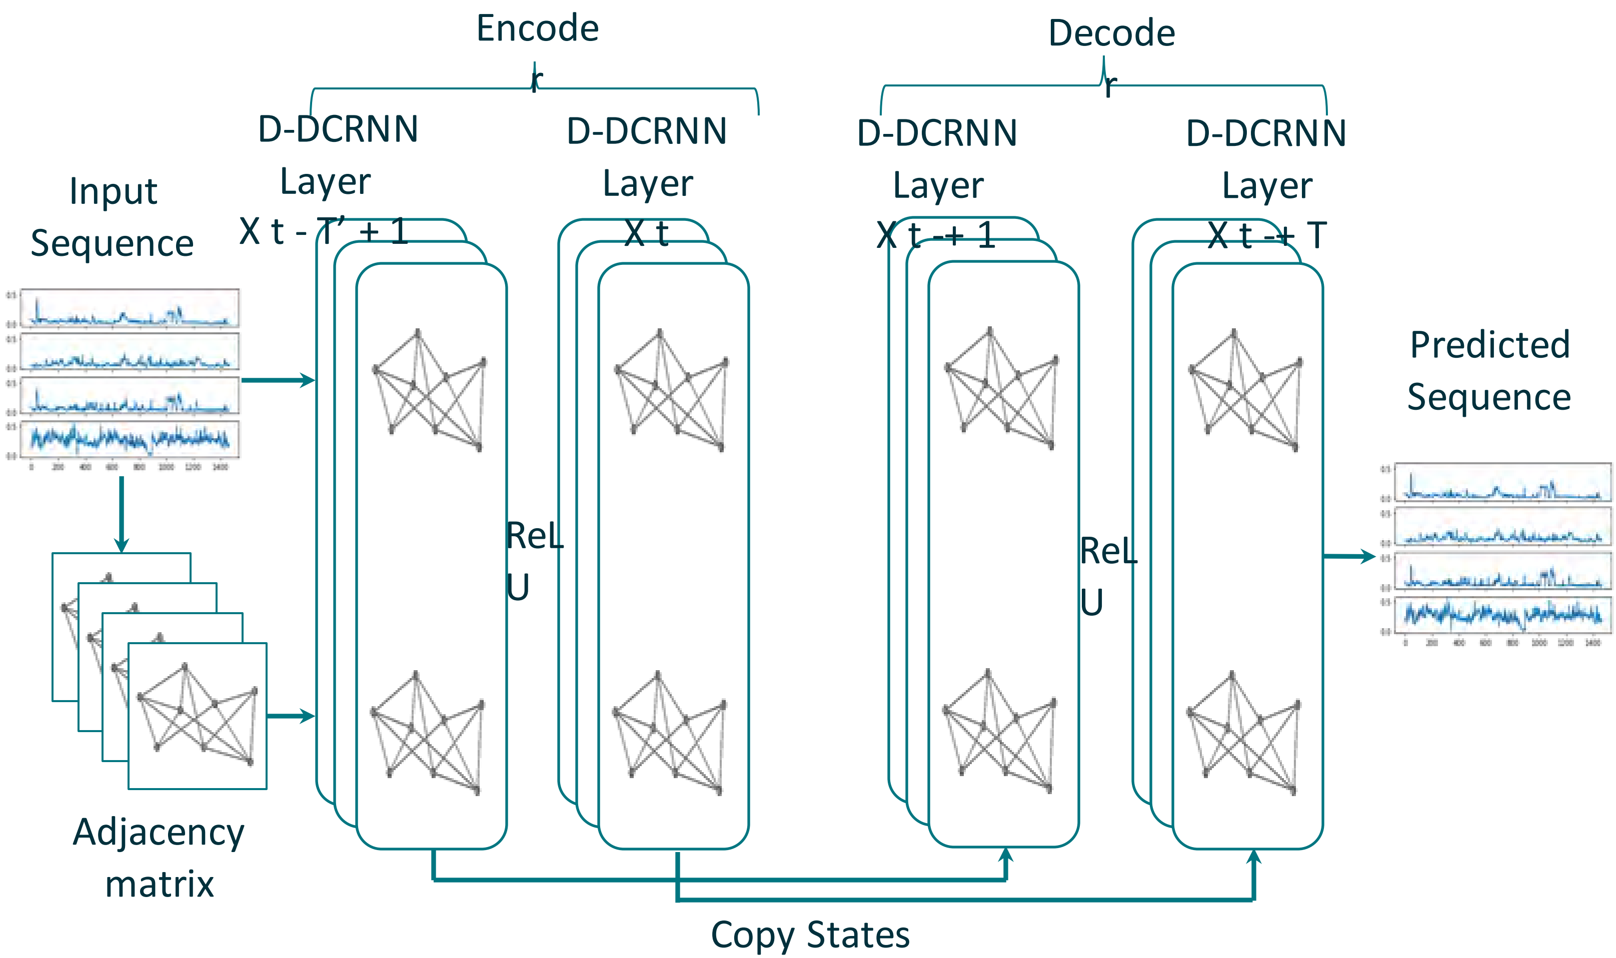

Mallick et al, Dynamic graph neural network for traffic forecasting in wide area networks, IEEE Big Data, 2020

Impact of Input Horizon Duration

- Varied the input horizon duration as 6, 12, 18, 24, 30, 36, 42, and 48 hours to forecast for next 24 hours
- Mean (μ) and standard deviation (σ) of R^2 and MAPE values for varying input horizon durations

Input Horizon Duration	6 hrs	12 hrs	18 hrs	24 hrs	30 hrs	36 hrs	42 hrs	48 hrs
R^2 (μ)	0.58	0.60	0.70	0.72	**0.76**	0.69	0.71	0.77
R^2 (σ)	0.02	0.09	0.10	0.12	**0.07**	0.11	0.05	0.08
MAPE (μ)	22.71	22.60	20.90	19.50	**19.01**	22.38	20.47	20.20
MAPE (σ)	6.16	4.25	6.55	1.74	**4.13**	8.23	3.37	3.77

Best performance achieved with 30 hrs input horizon duration

GNN improves Prediction Accuracy among other models

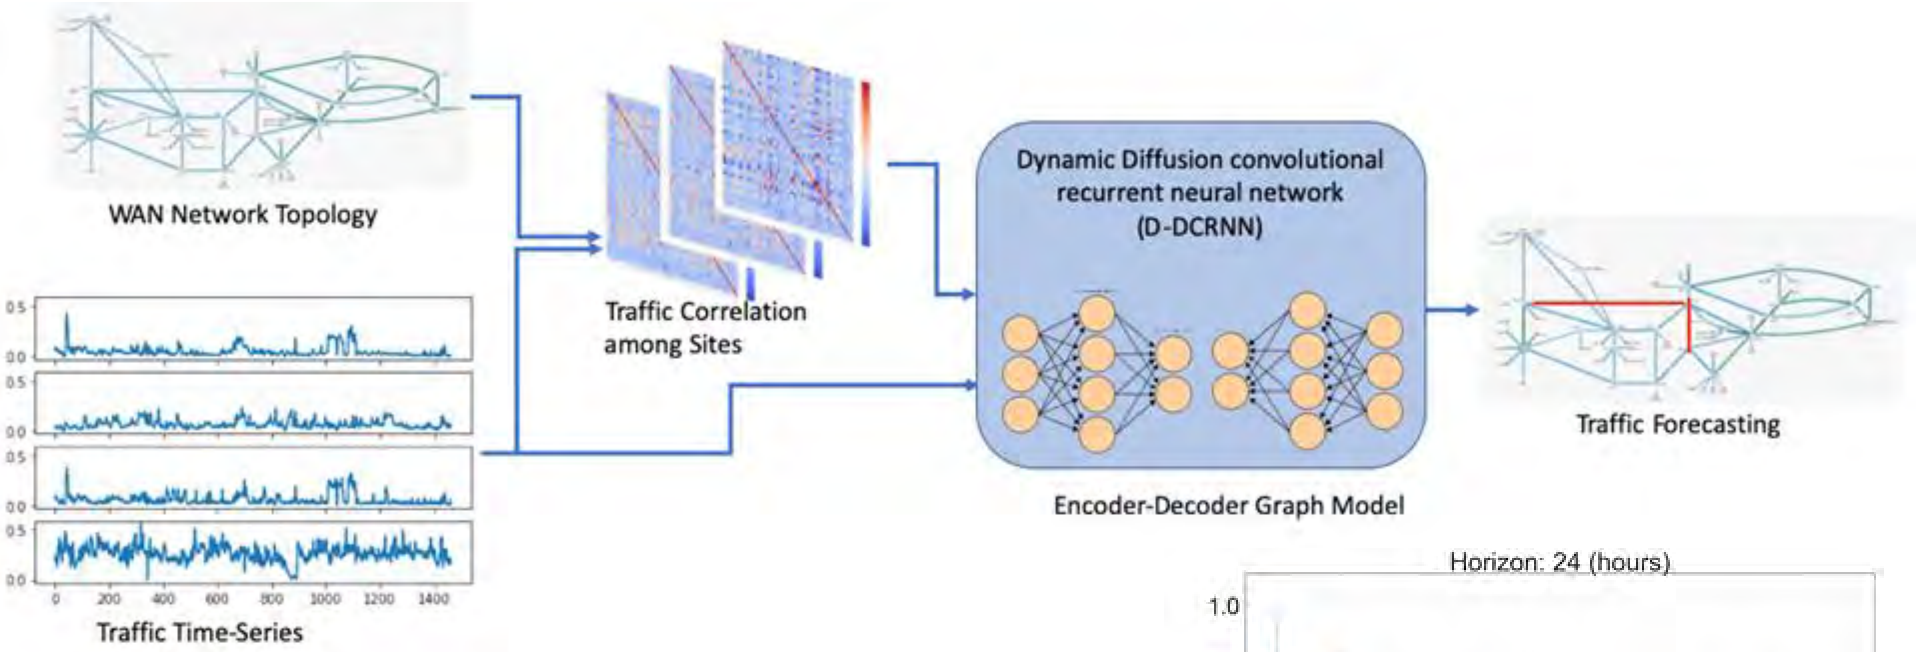

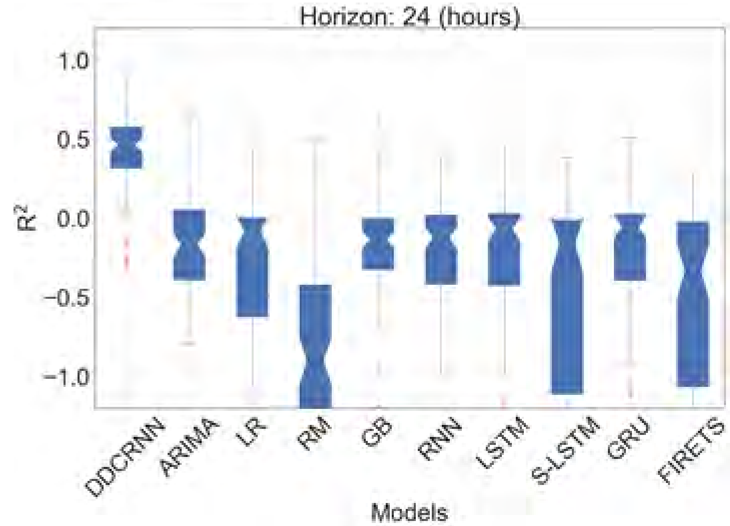

Problem: Reaction with Control

Our requirements:

- High performance throughput with low loss for huge time sensitive data transfers
- Latency sensitive communications: cloud, video, command/control for engineering
- Bandwidth reservation: OSCARS

This is a Multi-Objective optimization problem

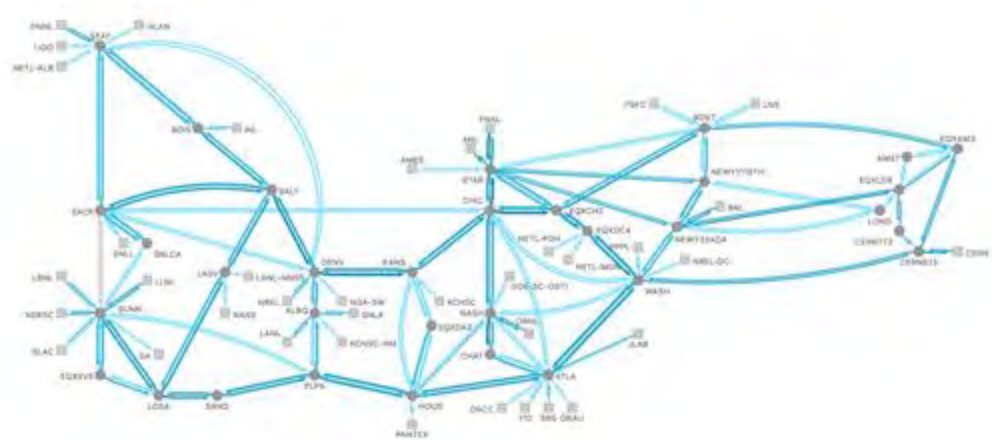

Representing Network as Deep Learning Problem

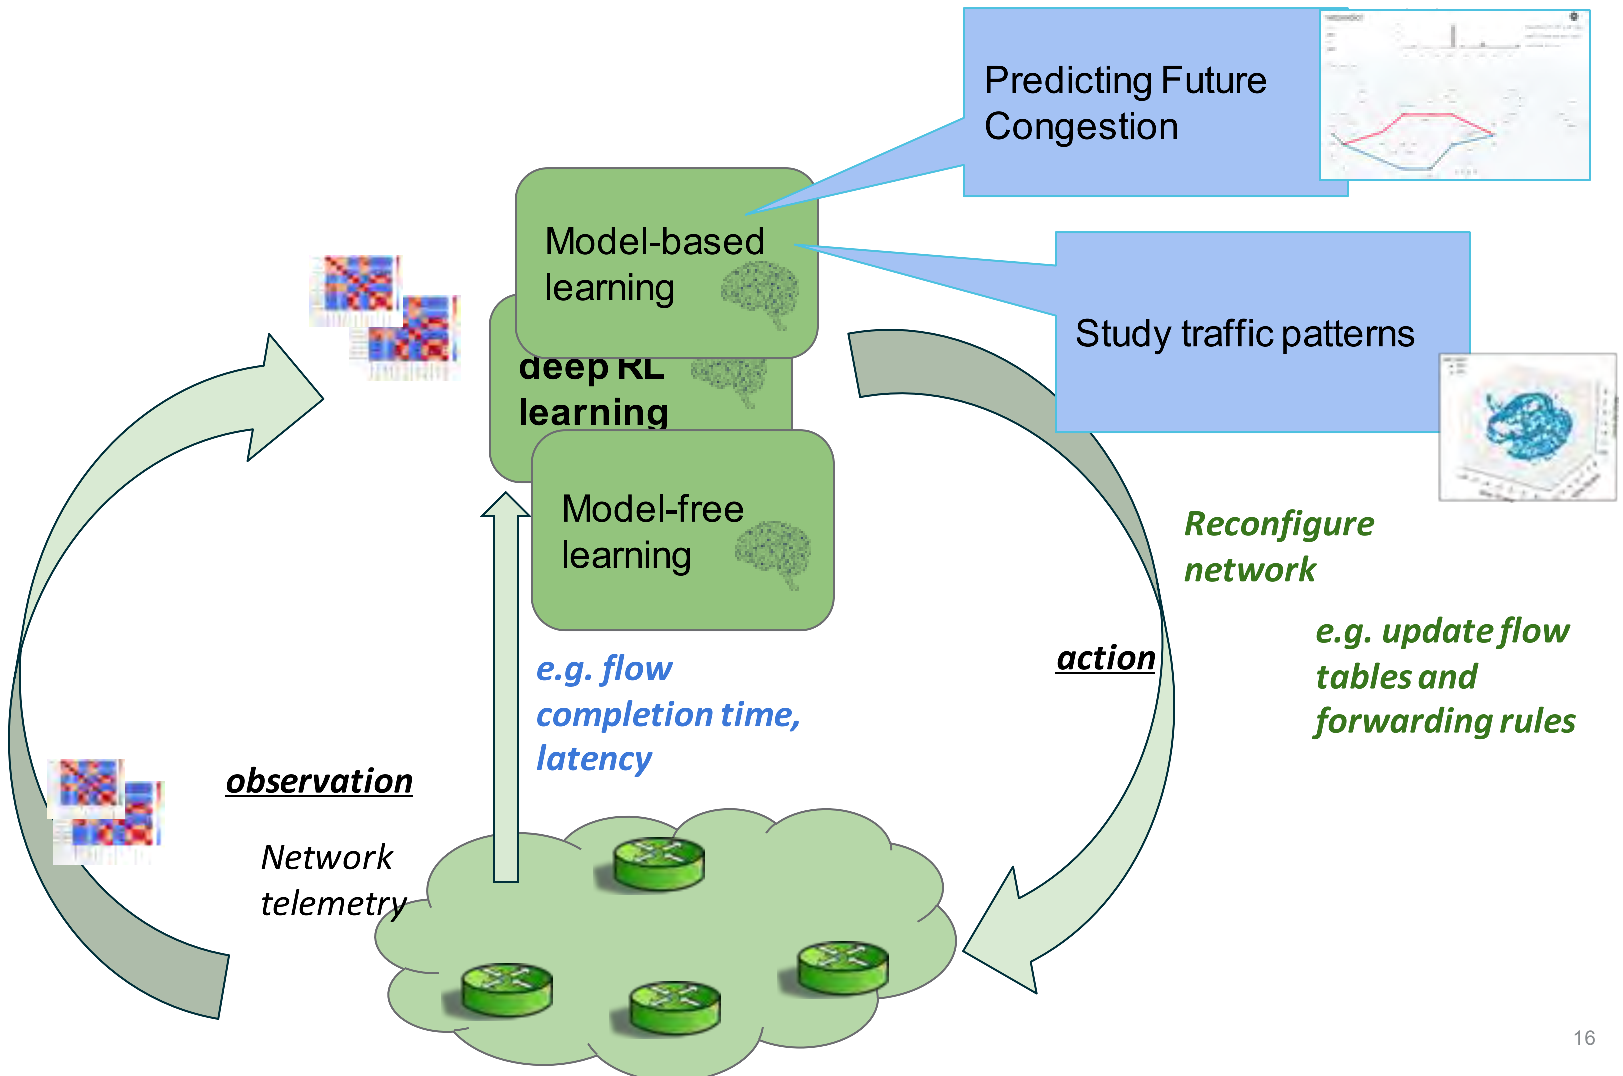

Data-driven (active) learning through experience

Deep Reinforcement Learning

s_t - state
o_t - observation
a_t - action

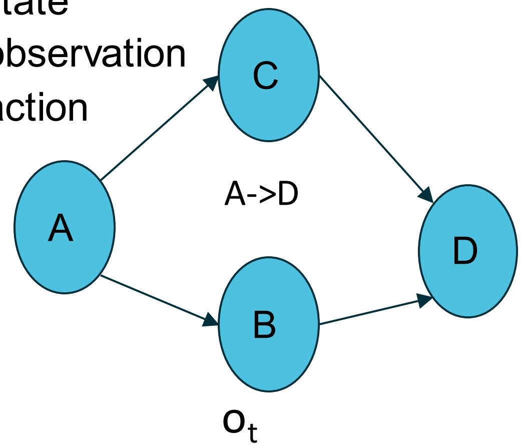

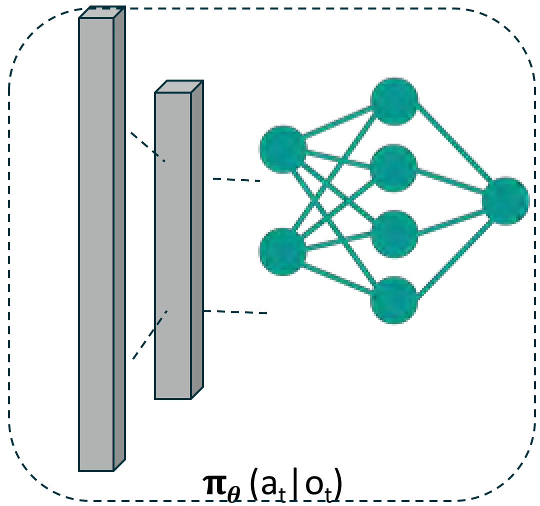

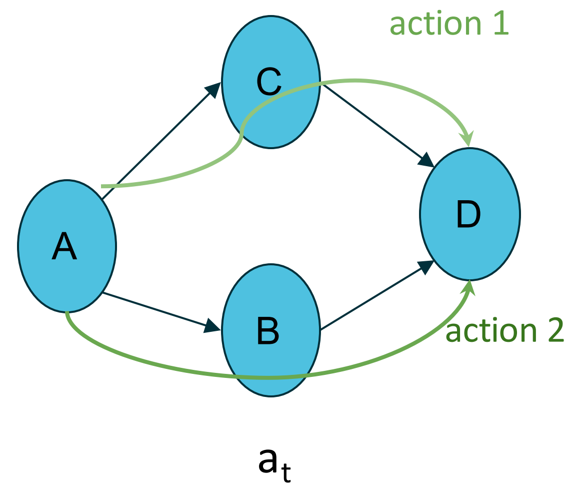

Markov property helps represent system as state-action pairs

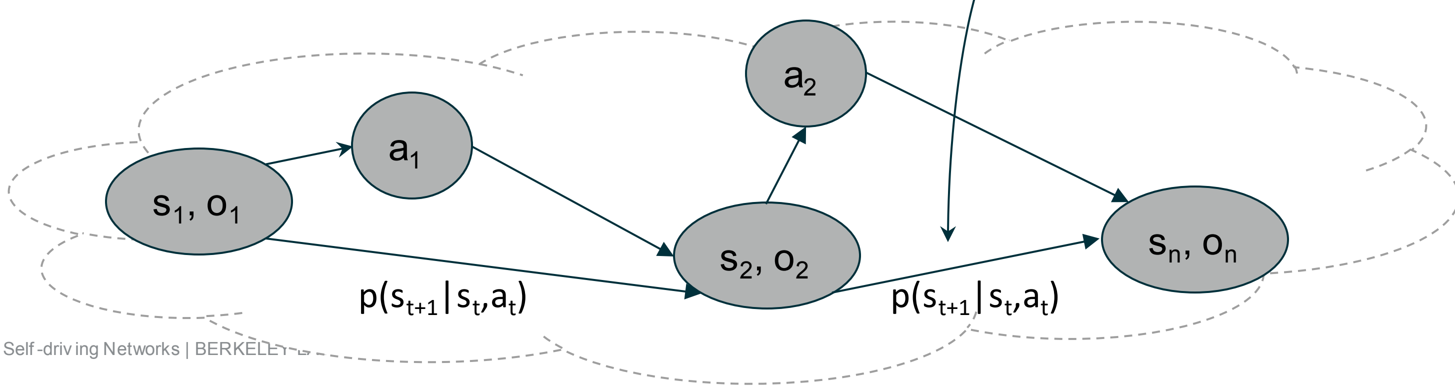

Conductor vs Orchestrator: Data-driven learning

Combines with multi-agent reinforcement learning

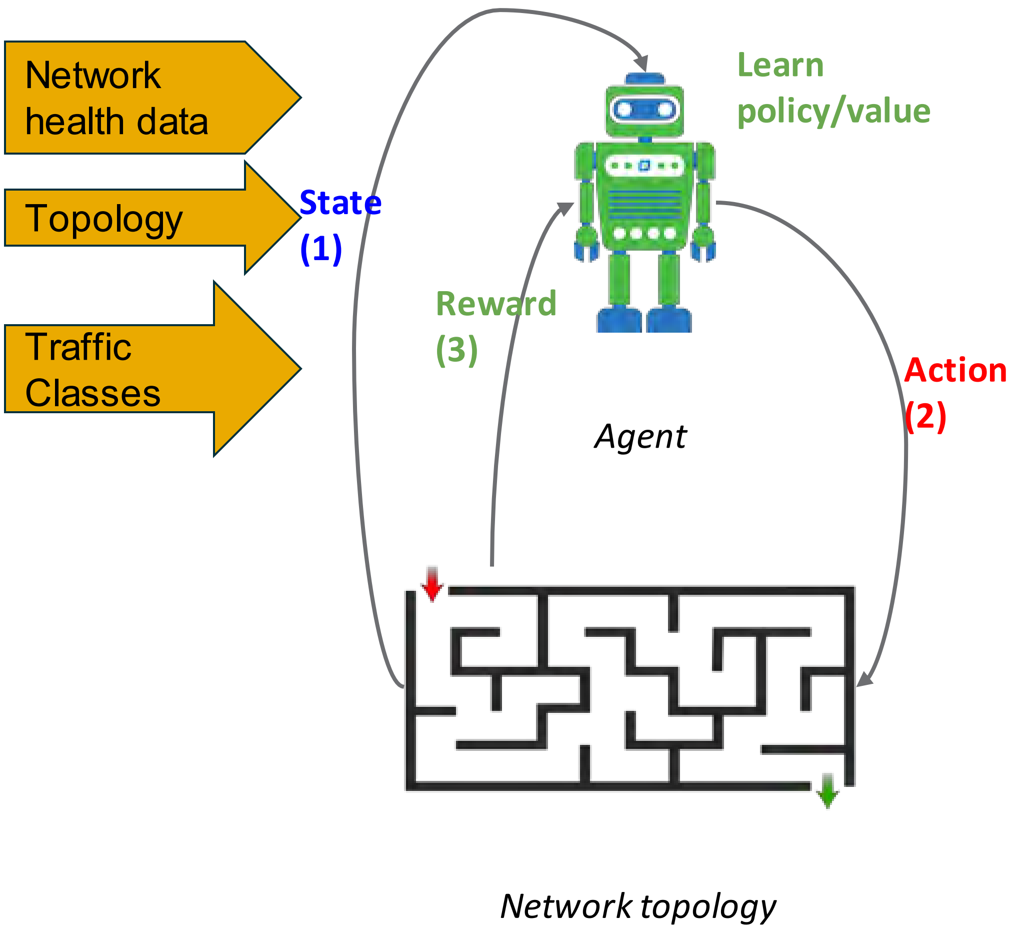

*Reward: Optimize for average utilization, latency, used links

Developing HECATE (learning controller)

Architectural diagram

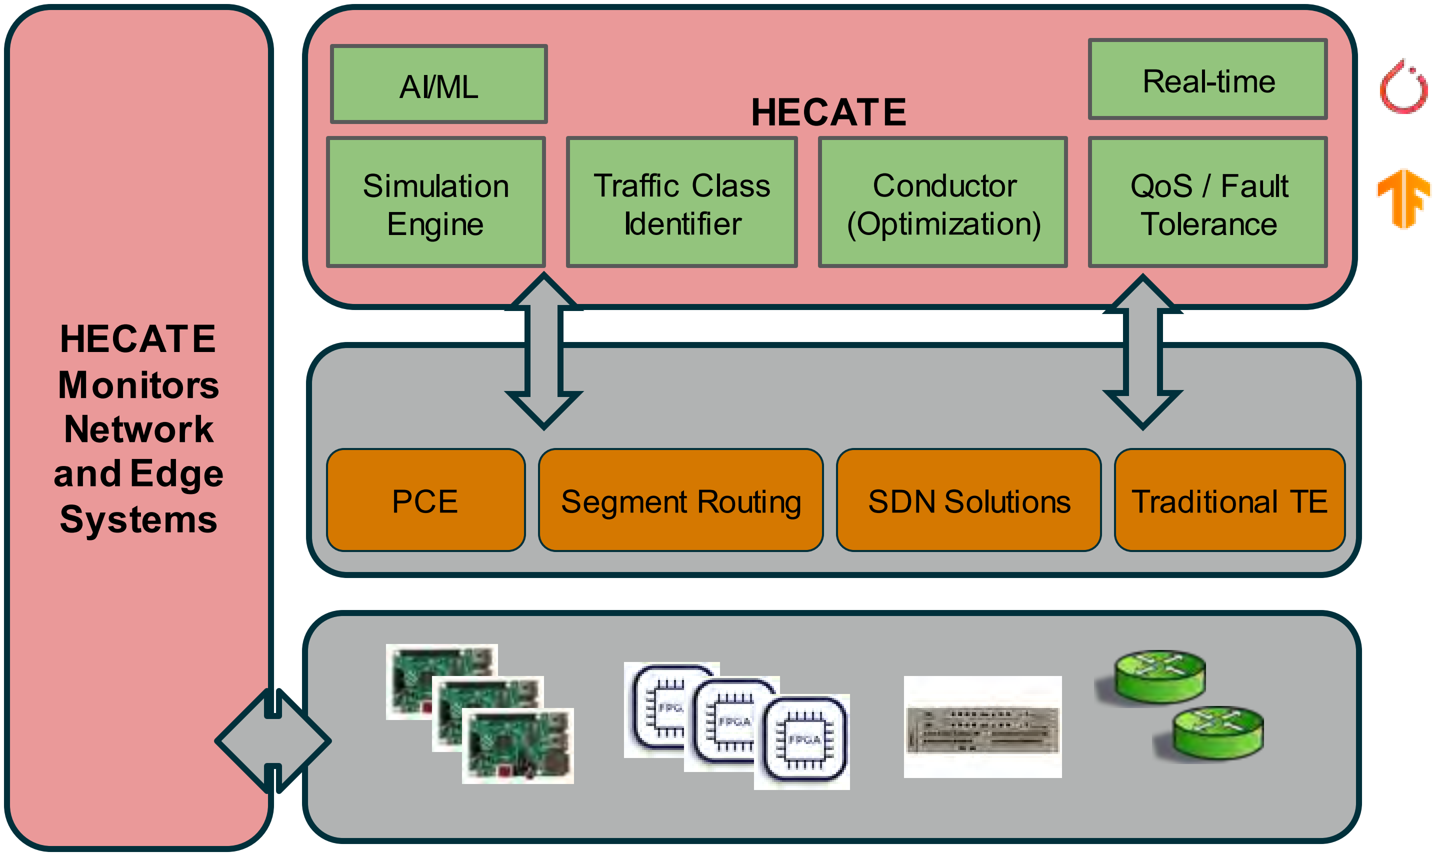

Patent filed: Data driven, machine learning augmented dynamic path optimization, 2022

Hecate Architecture: Overview

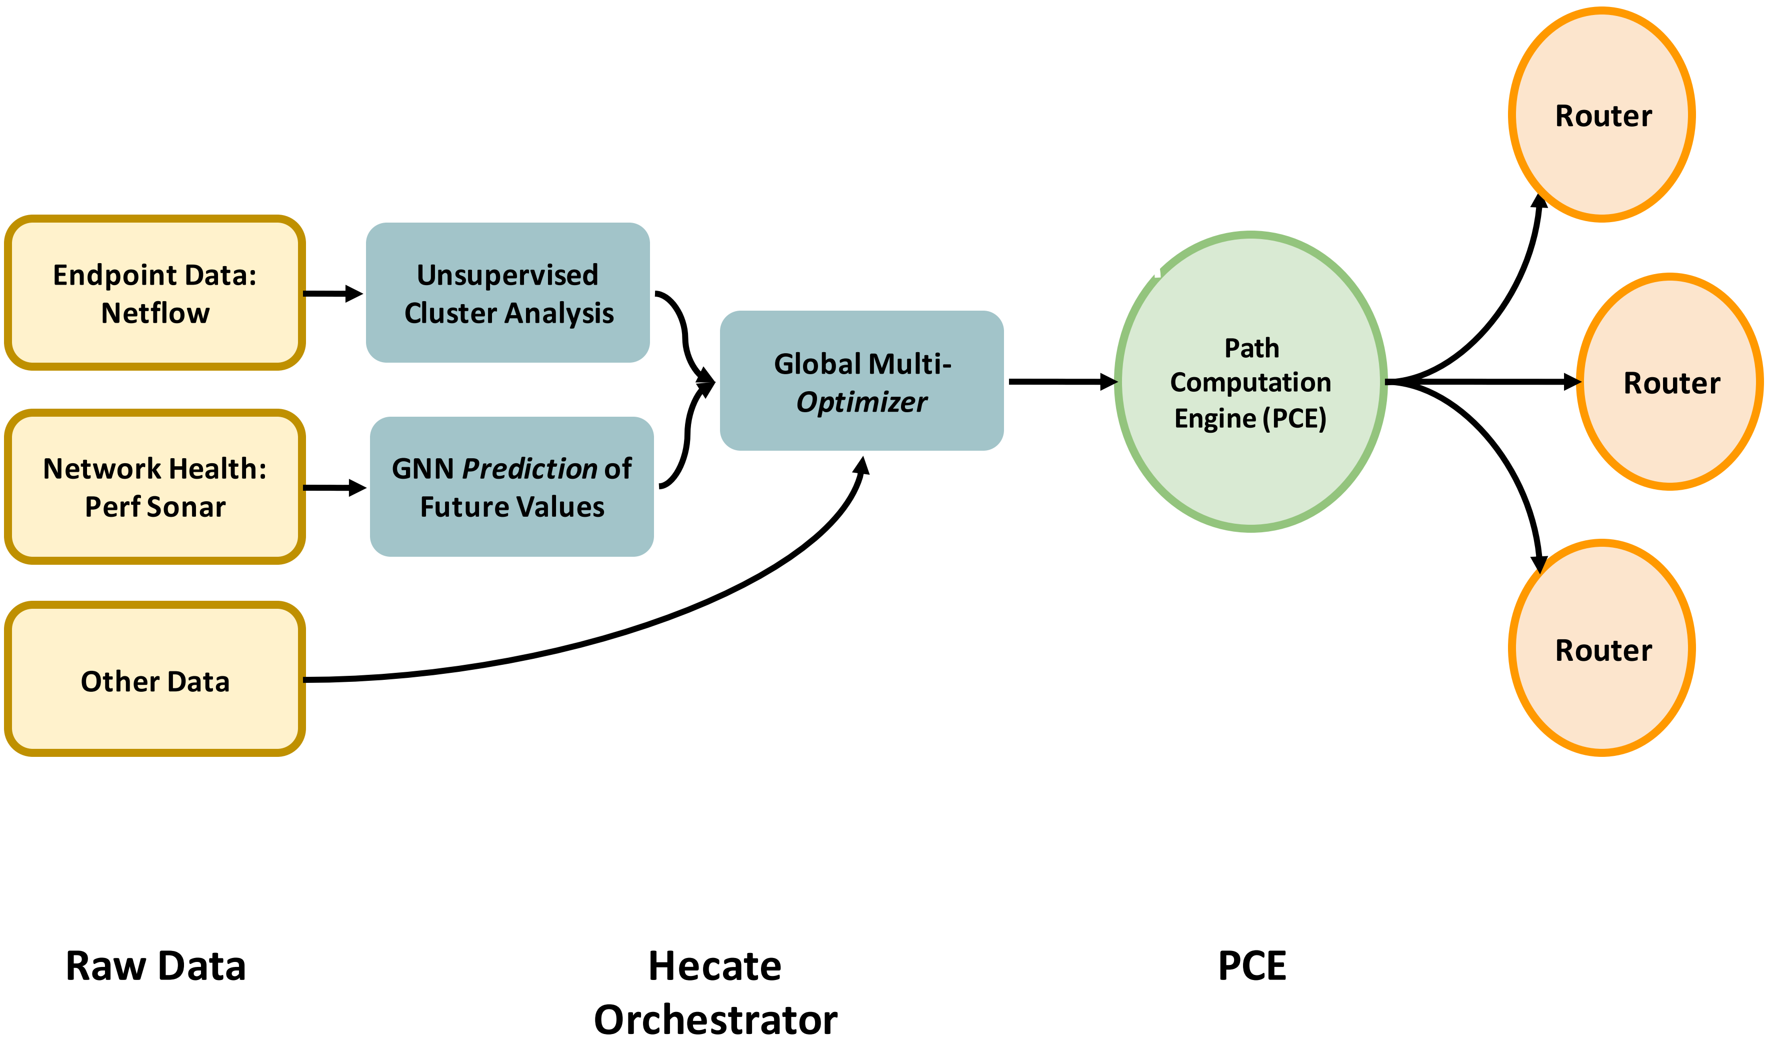

Identifying Traffic Classes

- Unique data analysis using unsupervised learning and clustering algorithms
- Real-time AI learning

Site Characteristics	Link Characteristics
Slow Update cycle	Fast Update cycle
Traffic Classes • Jumbo • Interactive • Default	Link Attributes • Loss • Delay • Jitter
Netflow	perfSONAR
Function of Time and Data Volume	Time Series

Hecate Architecture: Network Health : Predict

Take most current health data and use to predict values for the next several time steps.

Based on work previously done

Model the network as discrete aggregated network traffic at time t, $G_t = (V, E, W)$

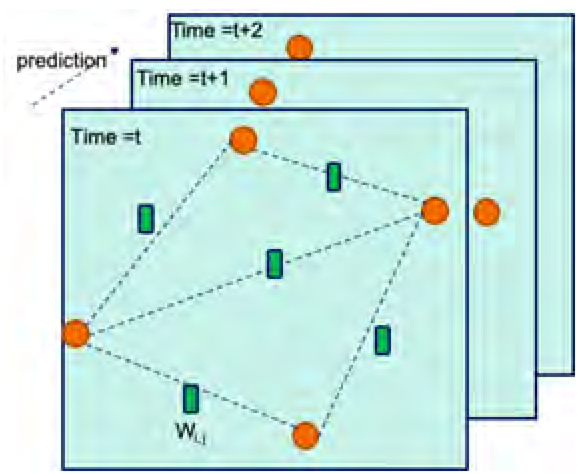

Model:
- Stack of spatio-temporal convolution **blocks**
- **Output** layer.

Each block consists of two temporal gated layer and a spatial graph layer in between.

Output layer consists of convolution, normalized and a fully convolution layer.

Hecate Architecture: Global Optimizer

Based on work from DeepRoute research project

Use DRL to greedy Q-learning to simulate networks and learn optimal routing strategies for single optimizations

Significant movement in this field - exploring additional options

Hecate uses four types of reward functions in DRL for Graph Optimization: Loss, Latency, Jitter, Utilization

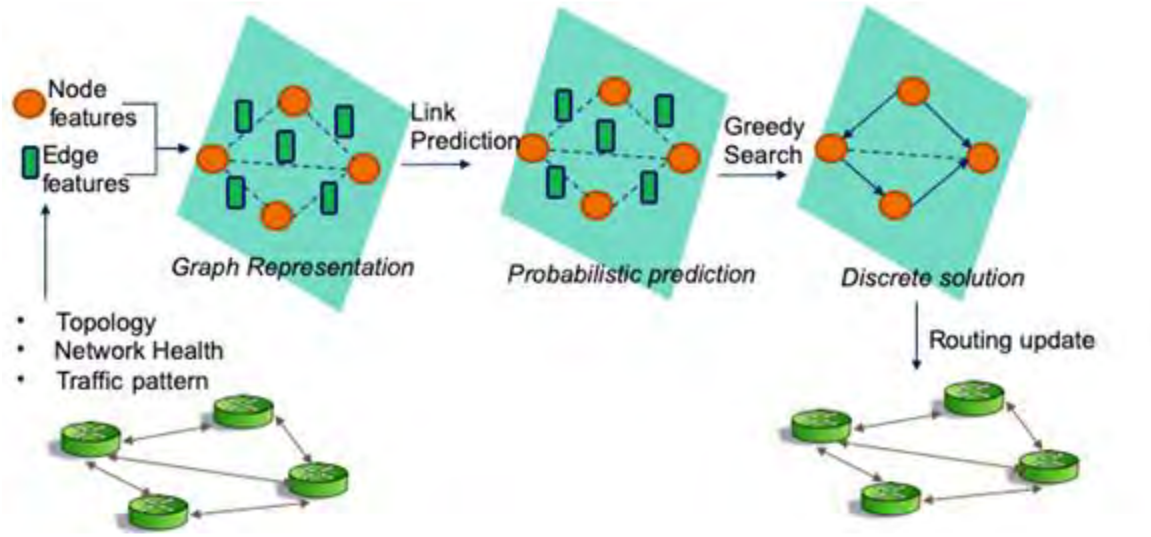

Kiran et al. Hecate: AI-driven WAN Traffic Engineering for Science, INDIS SC 22

Path Computation Engine

"Brains" of the segment counting core infrastructure

Like MPLS, can stack labels which define path through network

Provides programmatic access to network control:

- Read network topology, router details, performance data
- Write to API to provide "suggestions" for path selection

We do not want to replace a routing protocol, just provide good advice

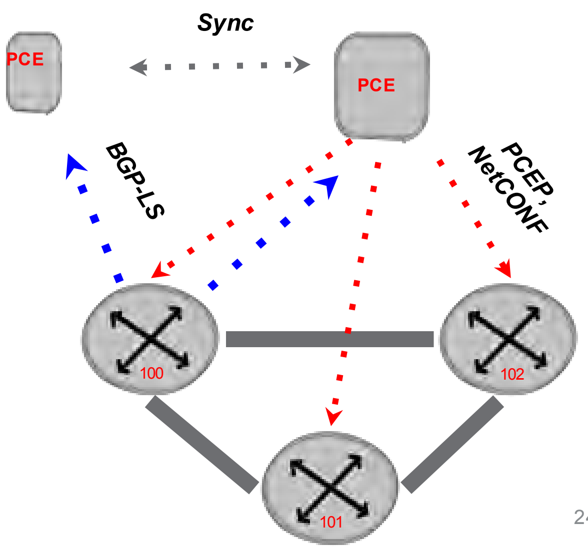

HECATE Simulation

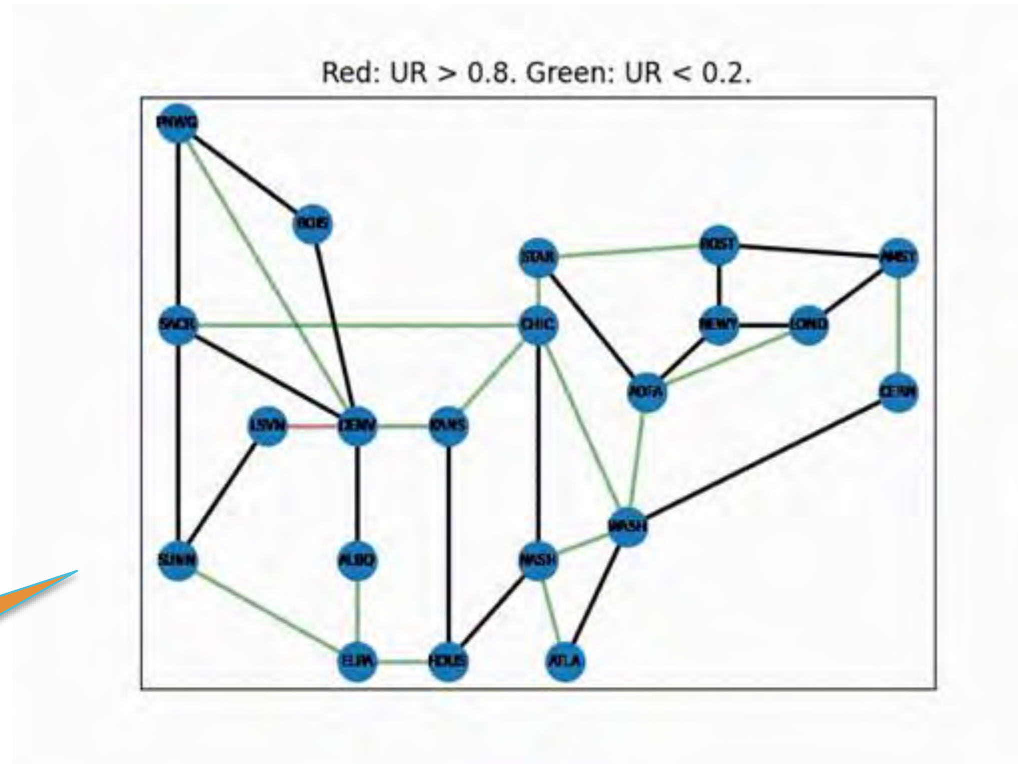

Moves incoming traffic to less used paths

UR: link utilization rate

Improve packet delivery at high loads

- Improves network performance at high loads
- Leverage traffic patterns into learning to cater to different characteristics
- Can be extended to ESnet traffic engineering protocols

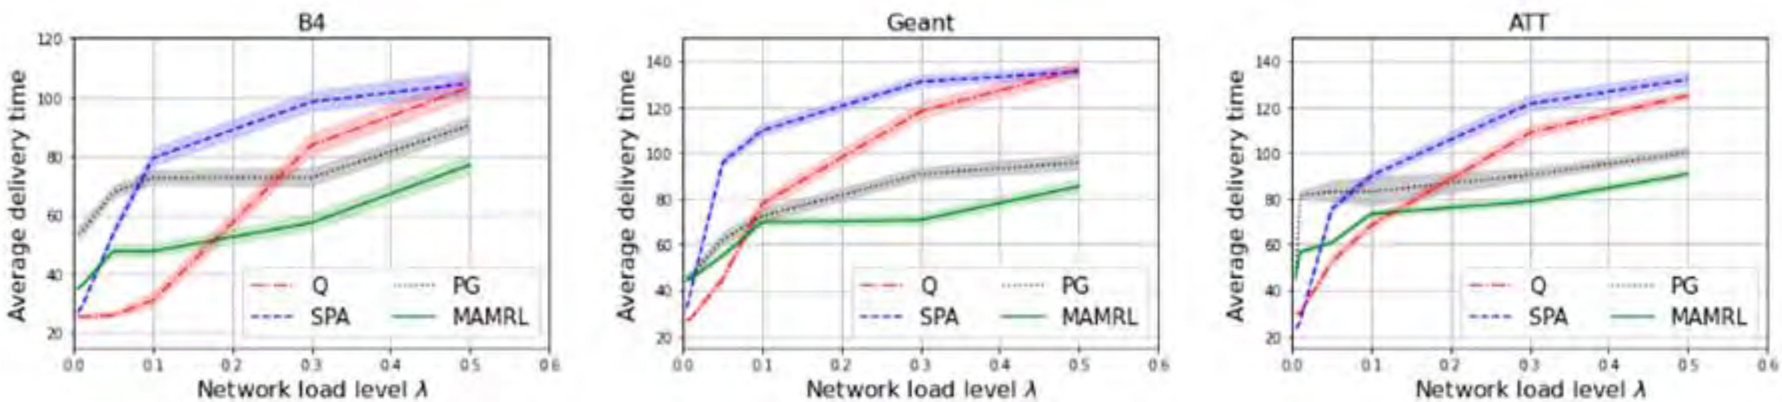

SPA: shortest possible
Q: Q-learning
PG: policy-gradient
MAMRL: multi-agent learning

MAMRL: Exploiting Multi-agent Meta Reinforcement Learning in WAN Traffic Engineering, arxiv, 2021

Average utilization improved

Before

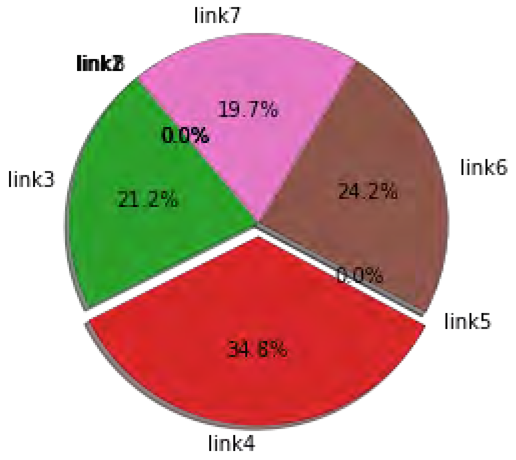

After

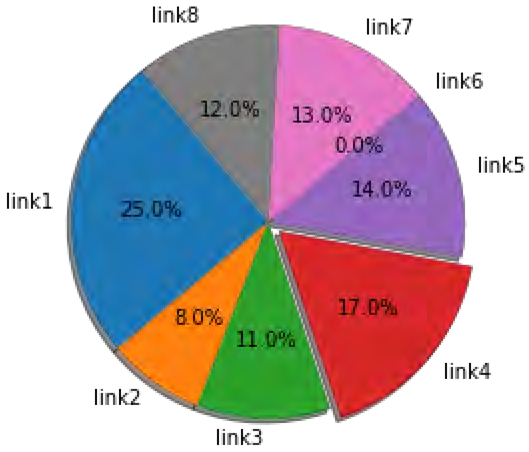

As Science moves to Wireless, Edge and Quantum…

We bring our expertise into these networks as well

Problem: Networks Beyond “Lab Borders”

- Optimizing Radiation science in the field
- Multiple data sources bringing fast data to HPC
- Real-time Edge Control
- Emergency corrections e.g. drones
- AI rewritten for Edge

5G connections for DOE NNSA

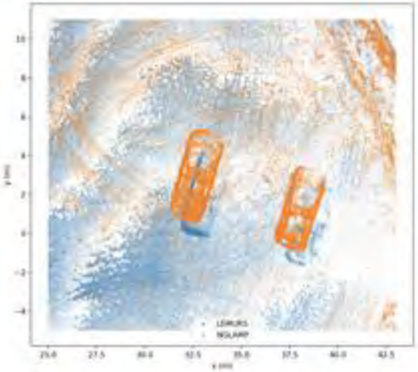

Digital Twins in Biology….

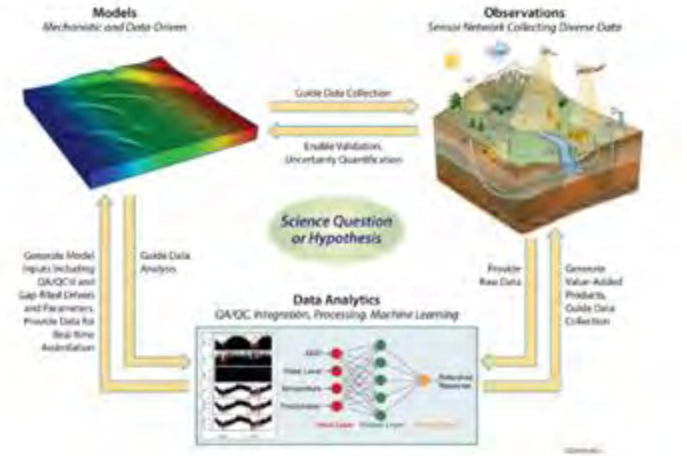

Novel "TempoSpatial" Control Plane for DOE Network

With 5G requirements and capabilities, we see the need for a tempostatial control plane that works with heterogeneous hardwares to tackle edge-to-core challenge

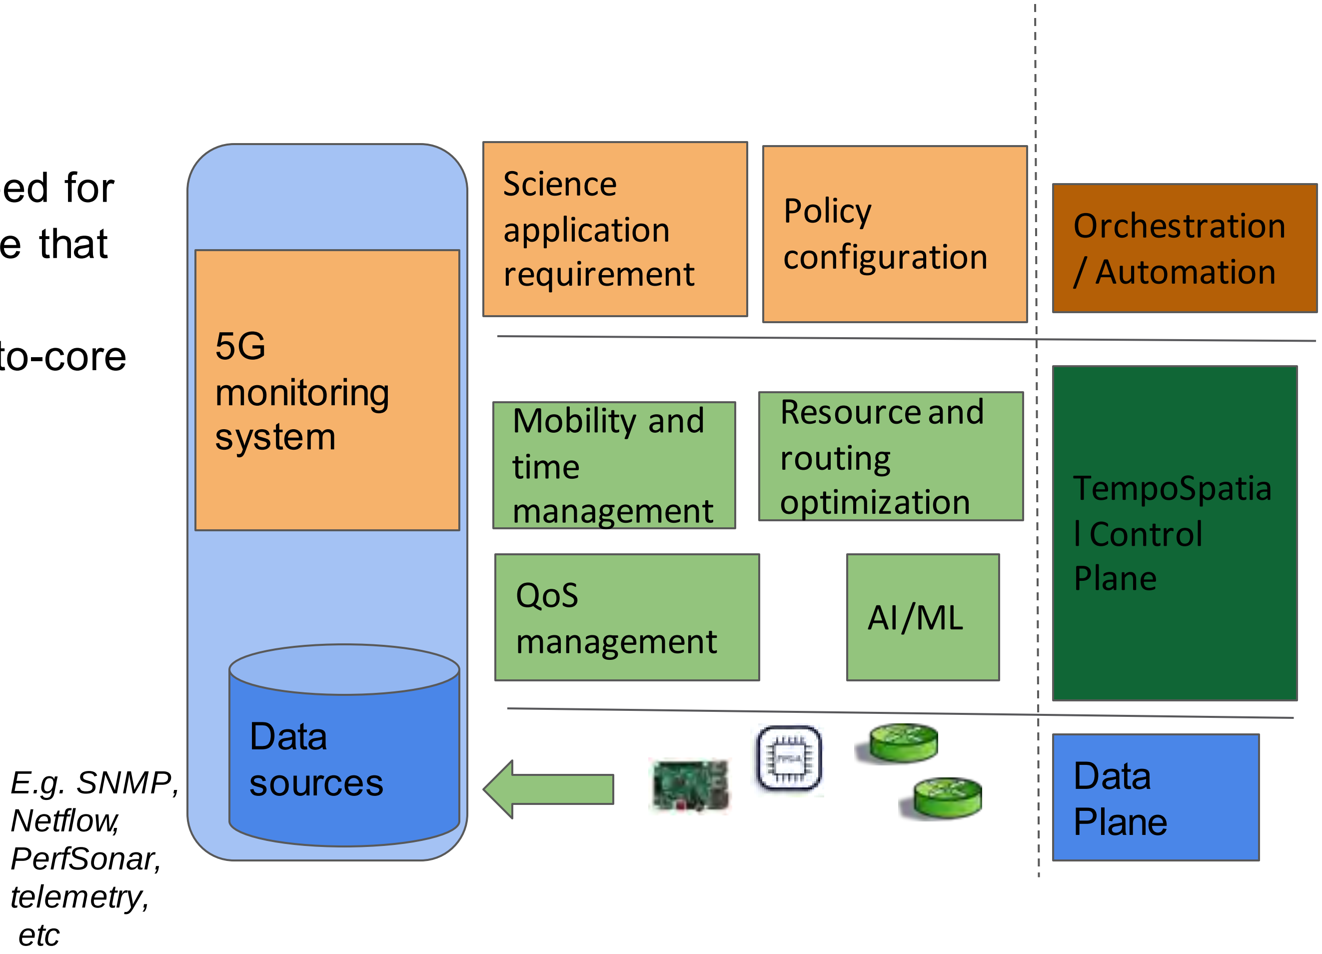

Introducing ϴ-AWR (*tower*)

Θ-AWR in B59 LBNL

Raspberry Pis with 5G hats

Programmed Sim cards

With Firecell kit

Why:

- Open Source
- Our own controllers to build on
- Horizontal expansion
- Research innovation and universities

Packet loss and throughput to find optimum deploy position

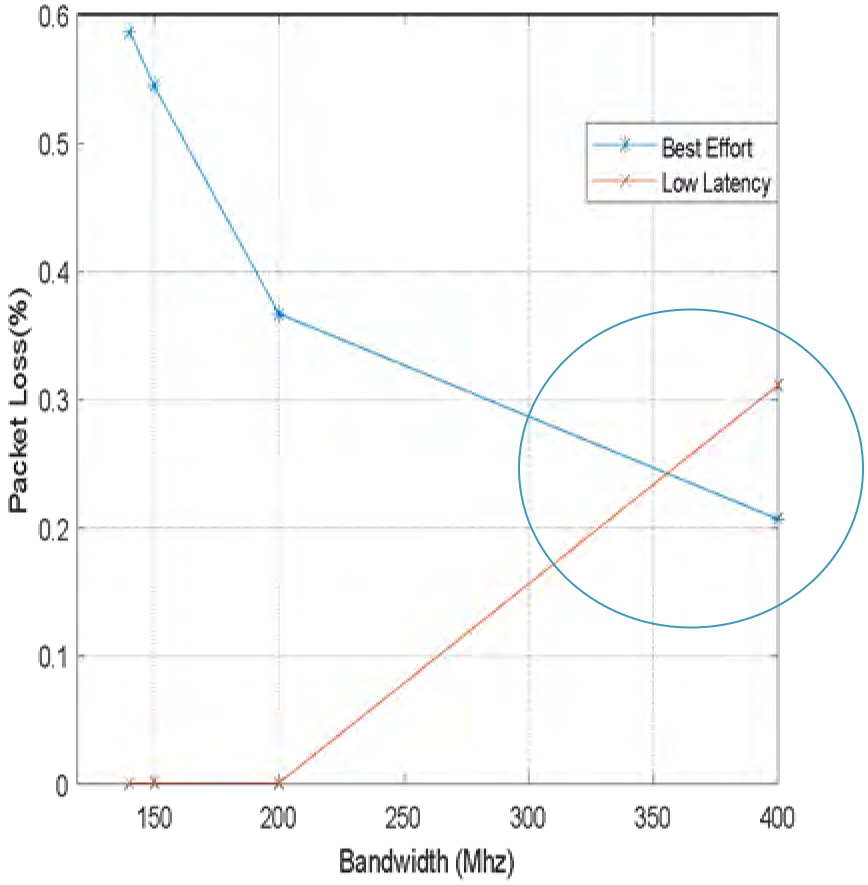

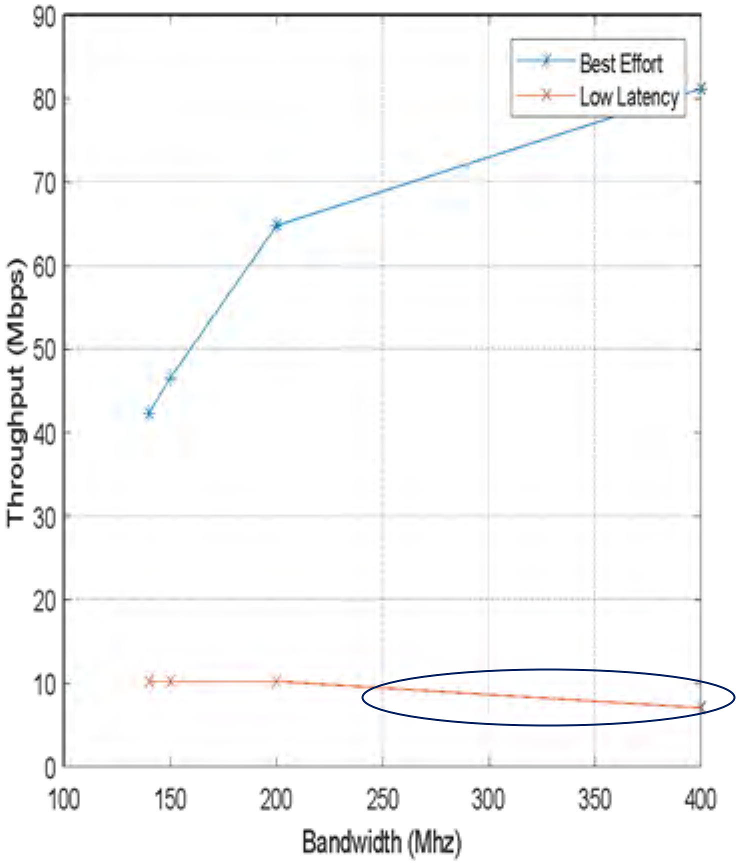

Challenges in simulation, not real world

Problem: Network of Quantum Computers

- Facilitate communication using distributed quantum states
- Reimagined science applications
- New networking challenges being introduced with quantum entanglement, repeaters, infrastructure and more…

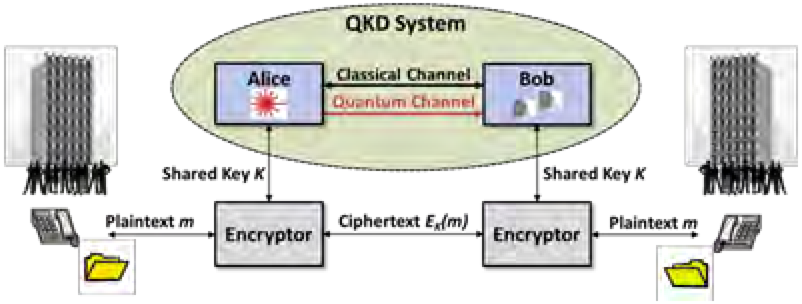

One way use Quantum Transducers

Reinforcement Learning for Quantum Transducers

Connecting quantum computers using fiber optic cables necessitates the conversion of quantum information from electrical to optical frequency regimes.

How can we use control protocols to enhance the efficiencies of transducers?

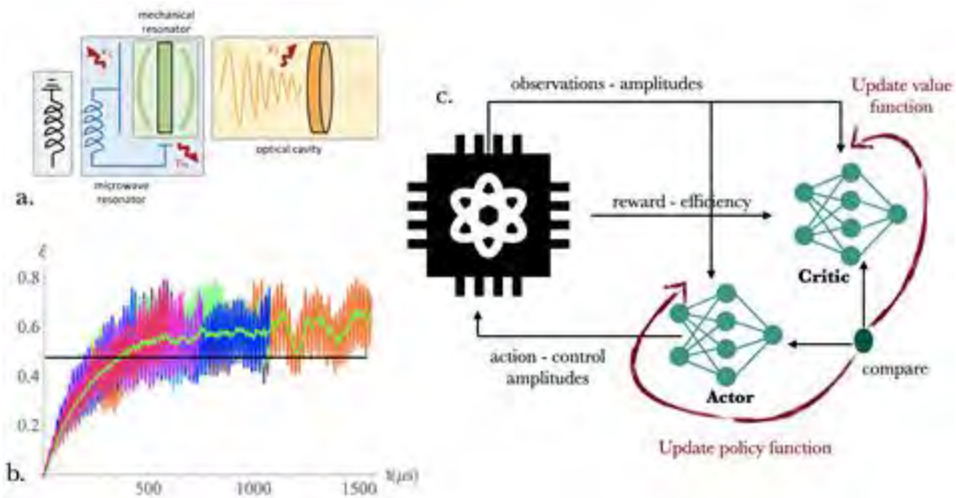

***Figure:** a) Optomechanical quantum transducer; b) Transduction efficiency using trained Reinforcement Learning Agent; c) Reinforcement Learning protocol.*

Solution

- We use the DDPG actor-critic algorithm with a continuous action and state space. Our trained neural networks are robust to deviations in experimental parameters in inference.

Simulation shows upto 60% efficiency (better than Google's experiment), but we now know this is not a practical approach!

Funded LBNL LDRD (2020-2022)

Another way explore Clock Synchronization for Quantum Network

Build simulation of the quantum network

- Transportation of quantum bits, consider attenuation

Simulation of the classical network

- Nodes make a request for sync through classical network
- Generate network packets to simulate traffic
- Based upon mathematical probability theory

But very preliminary work,

- no verification with real data
- many simulation platforms exist
- Can be relevant for distributed quantum sensing experiments

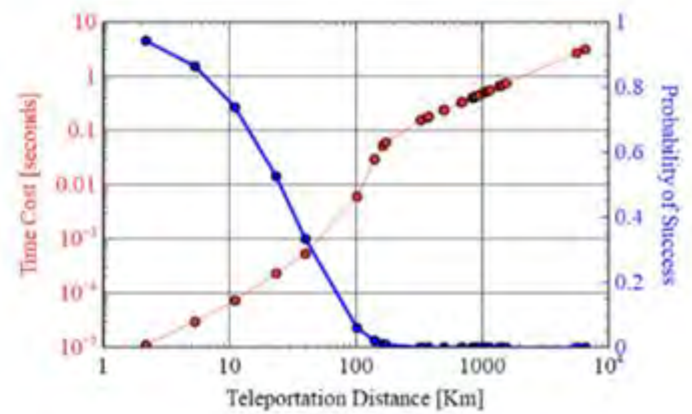

Fig. 2. The average time to perform teleportation as a function of distance with the loss rate set at δ.

Eric Yu et al. A Simulation Study of Quantum Clock Synchronization Using Teleportation, IEEE International Conference on Communications 2023

Another way explore Dumbbell Topology for Quantum/optic information exchange

- Explored Dumbbell topology for measuring how TCP interacts with other TCP flows
- 2 senders, 2 receivers and 1 router (bottleneck at the router)
- All the nodes use dedicated nics at 25Gbps
- Avg. RTT between senders and receivers is 16ms
- Tuning Params:
  - 32MB max buffer size
  - Jumbo frames
  - MTU probing
  - Fair queuing
- Used multiple iperf3 processes and aggregated the avg. goodput over 5 min. tests

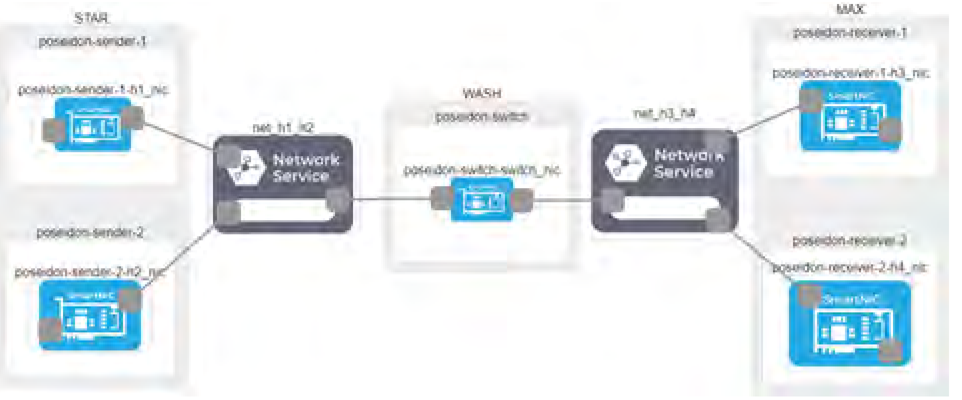

FABRIC Testbed

Quantum Experiment Reimagination:
- Explore quantum and optic information in the same link?
- Can we tune the optic to improve quantum signal quality ?
- or explore noise among the two?

Conclusions

- DOE science is moving towards two main areas in networks - Wireless and Quantum
- We can study these using simulations, but lack of real data is currently hindering our impact in these two fields. Can we build data sets for the community?
- As networking researchers, we need to lead with networking protocols and understand novel challenges
  - 5G/6G latency and packet drops matter
  - Quantum latency, quality and hardware matter
- Utilizing the optic infrastructure is key, and then building new quantum control planes to manage optic/quantum information in channels
- Using AI/ML for quantum, e.g. state estimation, quantum distribution protocols, control, and more, could be game changing

<u>Changing landscape: Need to understand how our scientists will be using networks with quantum processing for a complete end-to-end solution</u>

Acknowledgements

Science is Collaborative and Fostering New Talent Along the Way!









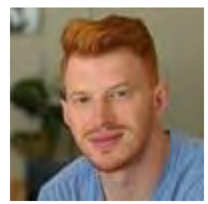

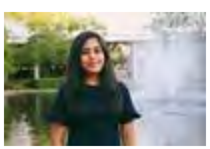


including colleagues and mentors at USC, RENCI, ANL, ORNL, NIST, UCB, UC Davis and more..

BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY
Office of Science

Thankyou for listening!

mkiran@es.net
kiranm@ornl.gov

Clock synchronization on a quantum network

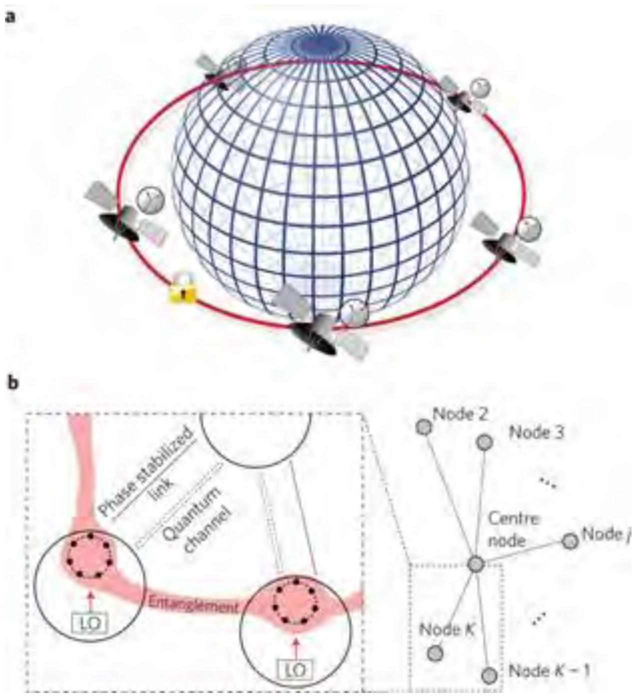

A quantum network of clocks (2014)

1. Assume that in a network with N nodes, a clock is located on each node (N clocks). We will label them C = {c1, c2, c3, ..., cN}, each with n clock qubits, Q={q1, q2, q3, ..., qn}.
2. Consider c1 as the central node, which will have an additional 2(N-1) ancilla qubits. For ease of notation, we will split the ancilla qubits evenly into two sets, A = {a2, a3, a4, ..., aN}$ and B = {b2, b3, b4, ..., bN}.
3. Perform local entanglement with q1 and one set of the ancilla qubits B={b2, b3, b4, ..., bN} (apply a CNOT gate with c1 and each ancilla qubit of B).
4. With the remaining clocks not including the central node c2, c3, c4, ..., cN, form EPR pairs with the other half of the ancilla qubits A. That is, (c2, a2), (c3, a3), (c4, a4), ..., (cN, aN).
5. On the central node c1, perform Bell-state measurements on the ancilla qubits from each set on the corresponding index. That is, (b2, a2), (b3, a3), (b4, a4), ..., (bN, aN).

Simulation parameters

Symbol	Definition	Value
c^*	Speed of light in the fiber	0.7 * c
δ	Loss rate (in dB/Km)	0.12
M	Max teleportation attempts	100
T	Simulation time	100000
N	Number of nodes in the network	31

$$L = 1 - 10^{(-loss/10)}$$

$$loss = \delta * \text{distance}$$

Results

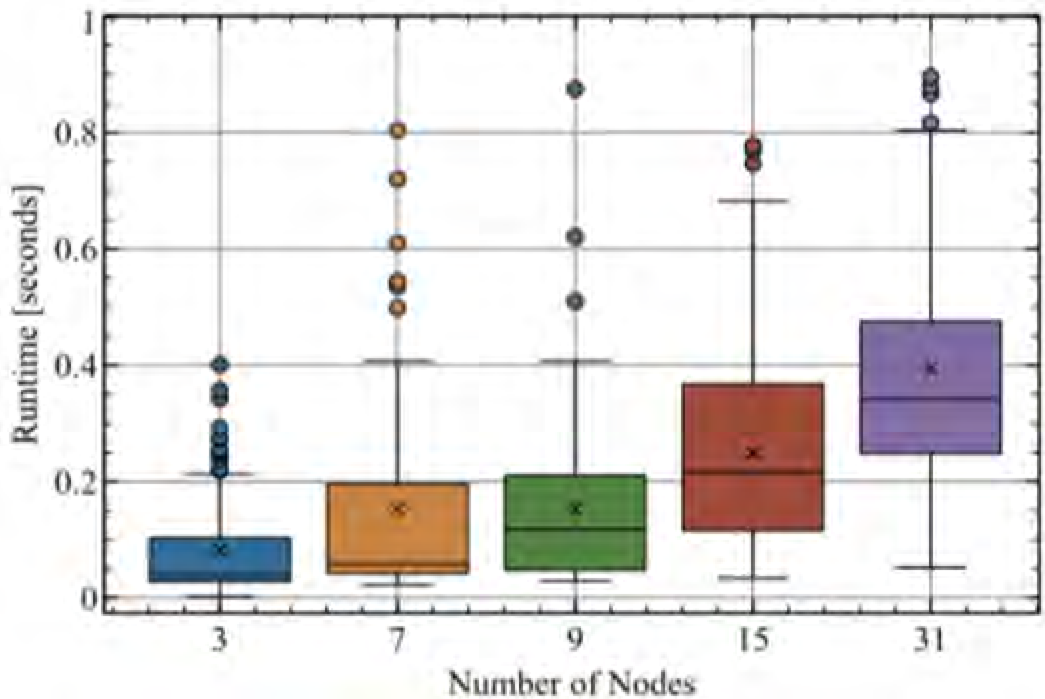

Section head white background

Section Subtitle — can also be Division and/or author name

Can be date or other less important info

Section head color backgrounds

Can be date or other less important info

5G brings Digital Twins to Life

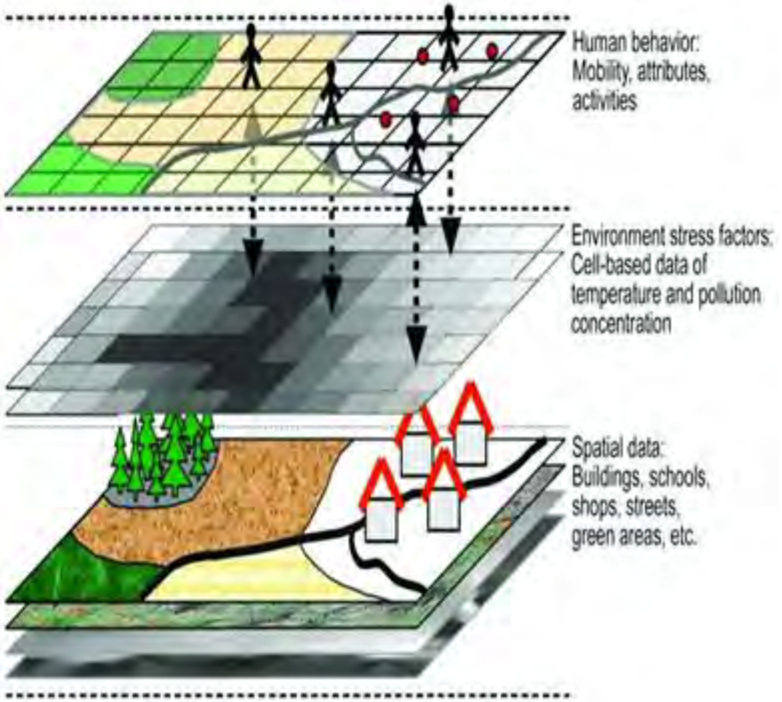

- Agent-based models connected to real-time sensor data
- Simulations for what-if analysis
- Analyzing massive amounts of data on DOE HPC facilities

Simulating agent-based model with real data

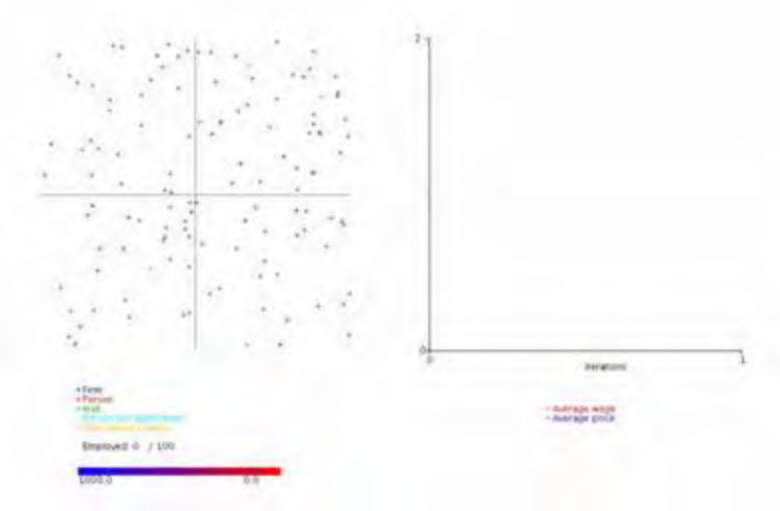

Self-driving ESnet optimized for 5G

- Optimized for **fast multiple data collection to HPC** and delivery for **real-time emergency control**
- Optimal Edge server: allocate content just 1 hop away for scientists
- **Caching at the edge** for optimal network performance
- **Optimal Routing** gives fasting path between source and destinations
- TCP optimization for 5G
- **AI determines prefetching** to move needed data closer
- Deliver "**bespoke network**" for science

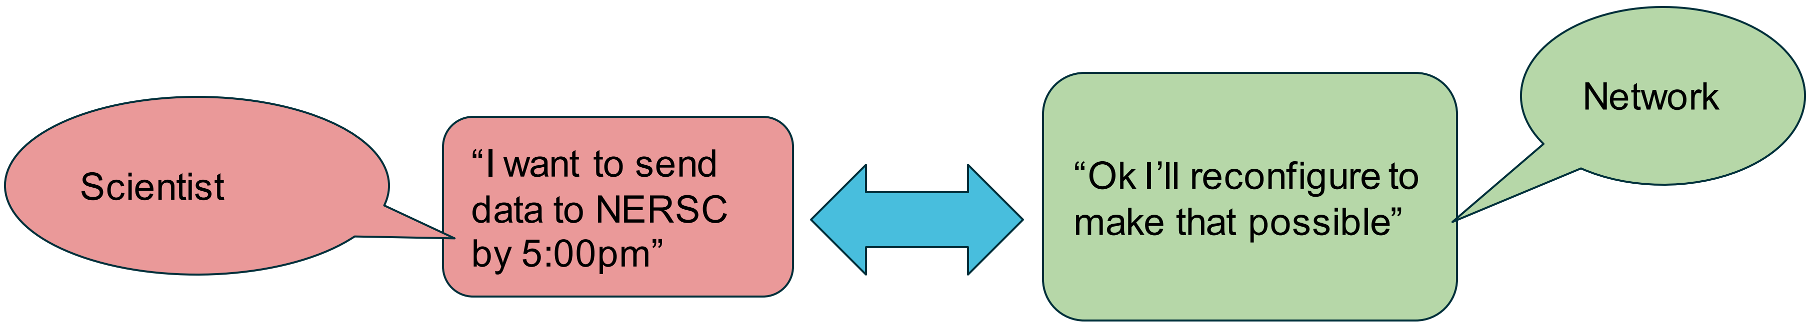

Conclusions: AI and 5G are essential to Self-driving Labs

- New networking challenges in 5G and Beyond will **truly transform DOE science**
- Quantum will work orthogonal to 5G advances, but merge
- Rise to experimentation in virtual worlds - DOE *Metaverse*
- Intelligent control with optimal data movement
- New research needed in
  - AI data validation - uncertainty, adversarial attacks and faulty data
  - Correctness and testing AI
  - encourage wireless adoption across Science areas
  - new Science with new data
- New moving data sources seamlessly integrate into network fabric like human nerves

Our work with the new testbed, self-driving systems, agent-based models and a real radiological case study are all the needed elements to push this new frontier

Thank You