Appeared in ICML 2019

A Deep Reinforcement Learning Perspective on Internet Congestion Control

Nathan Jay*, Noga H. Rotman*, P. Brighten Godfrey, Michael Schapira, Aviv Tamar

*Equal contribution

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

TECHNION
Israel Institute
of Technology

AI Week 2019

Internet Congestion Control

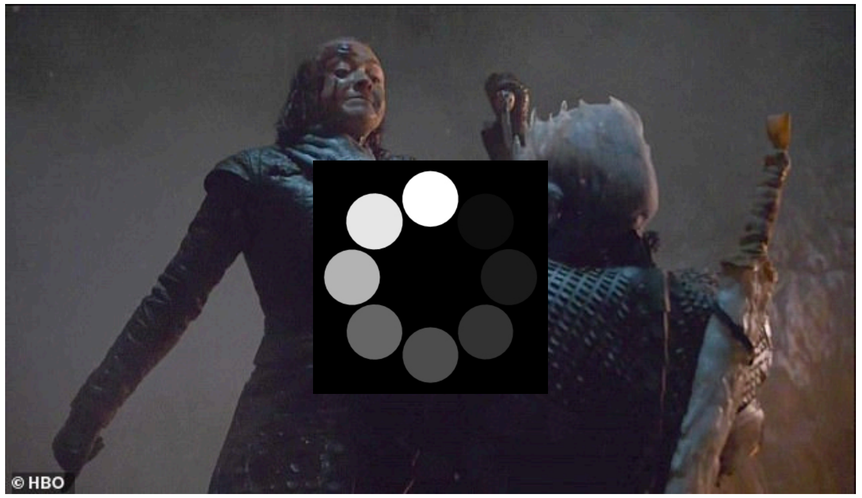

- One of the most **fundamental** and **challenging** problems in communication networks
- Determines **what you get** out of the **internet**
  - At what rate data goes in
- Always running on every connection
  - With no prior knowledge

Internet Congestion Control

too slow might harm performance

too fast might overflow the pipe

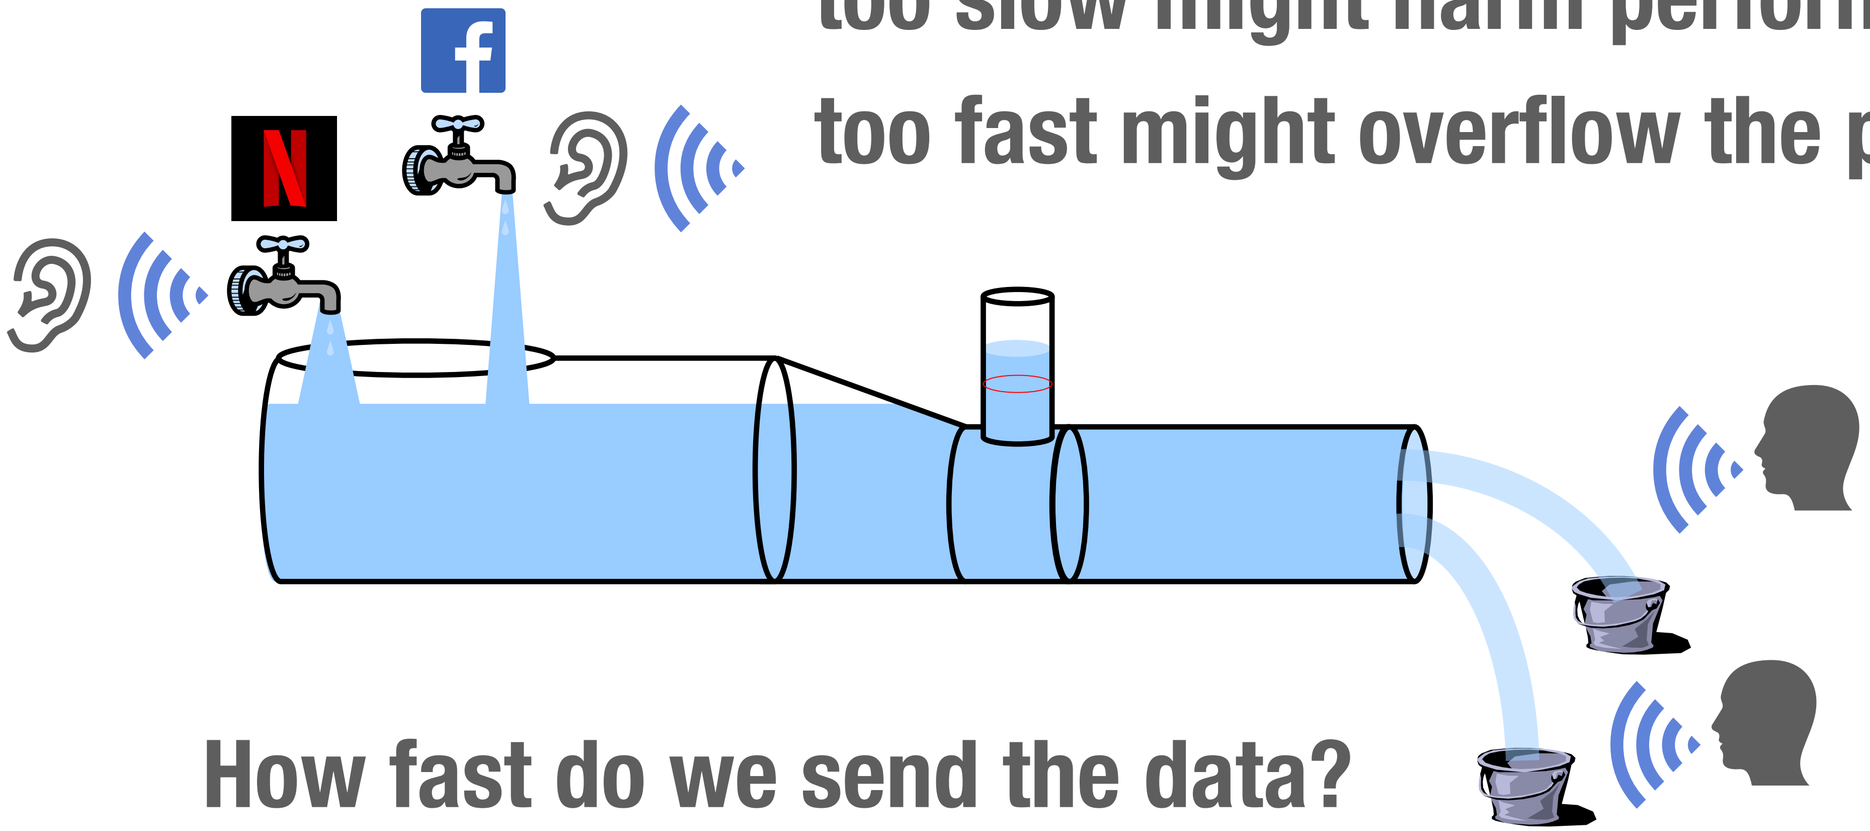

How fast do we send the data?

Internet Congestion Control

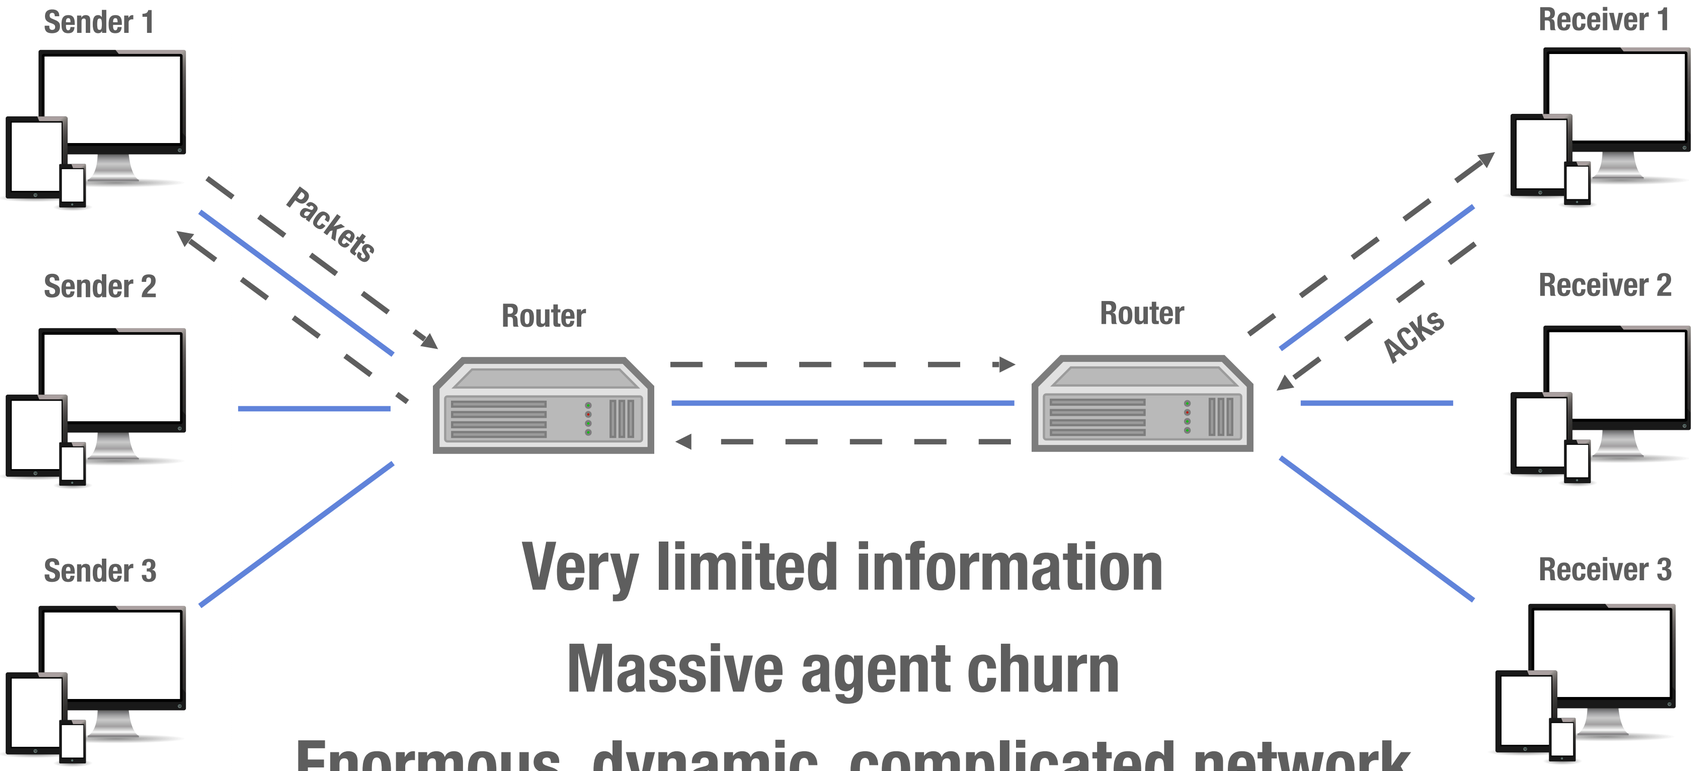

Very limited information

Massive agent churn

Enormous, dynamic, complicated network

Congestion Control Revisited

The Internet's Congestion Control default algorithm: **T**ransmission **C**ontrol **P**rotocol

Also known as **TCP**

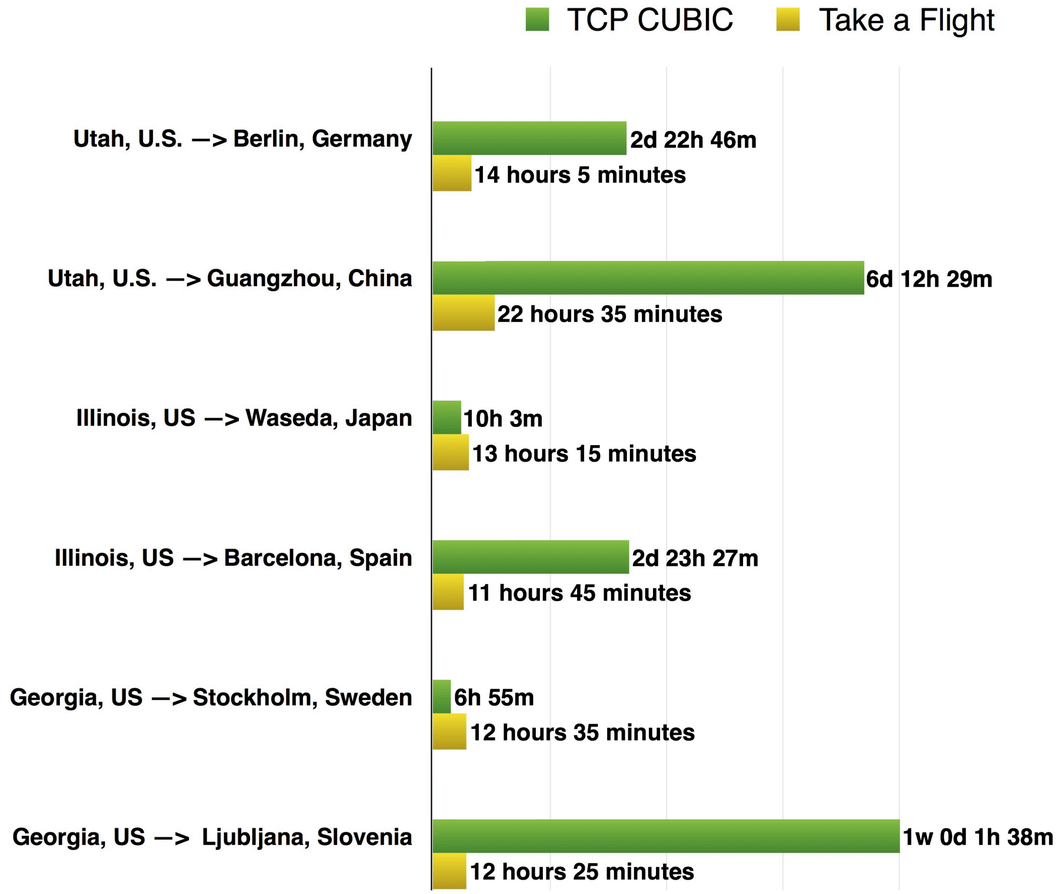

Motivating Deep RL

Locally-perceived history of feedback: Loss rates, latency…

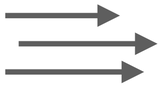

Congestion Control Protocol

Next sending rate

- We hypothesize: this **feedback** contains information about **patterns** that can improve the choice of sending rates
- Maybe deep RL can recognize and use them!
  - Successful in other domains - speech, games,…

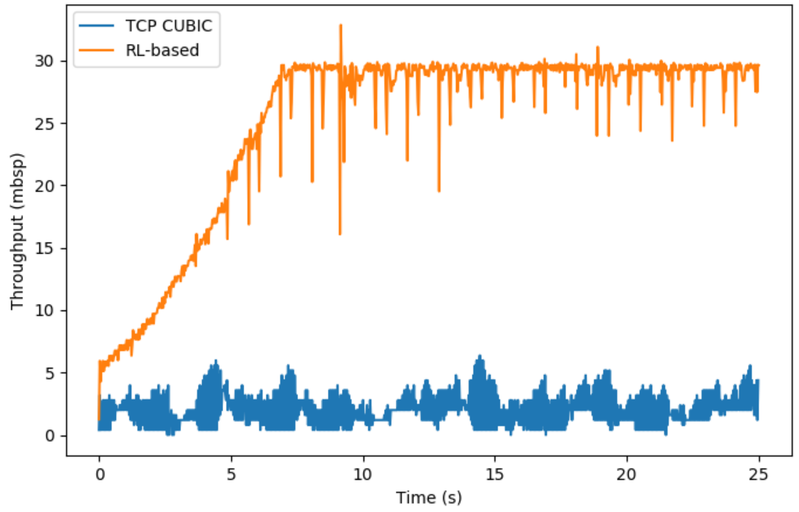

Throughput of a link with 1% random loss

RL Formulation

Goal: maximize the cumulative reward $\sum_t r_t$

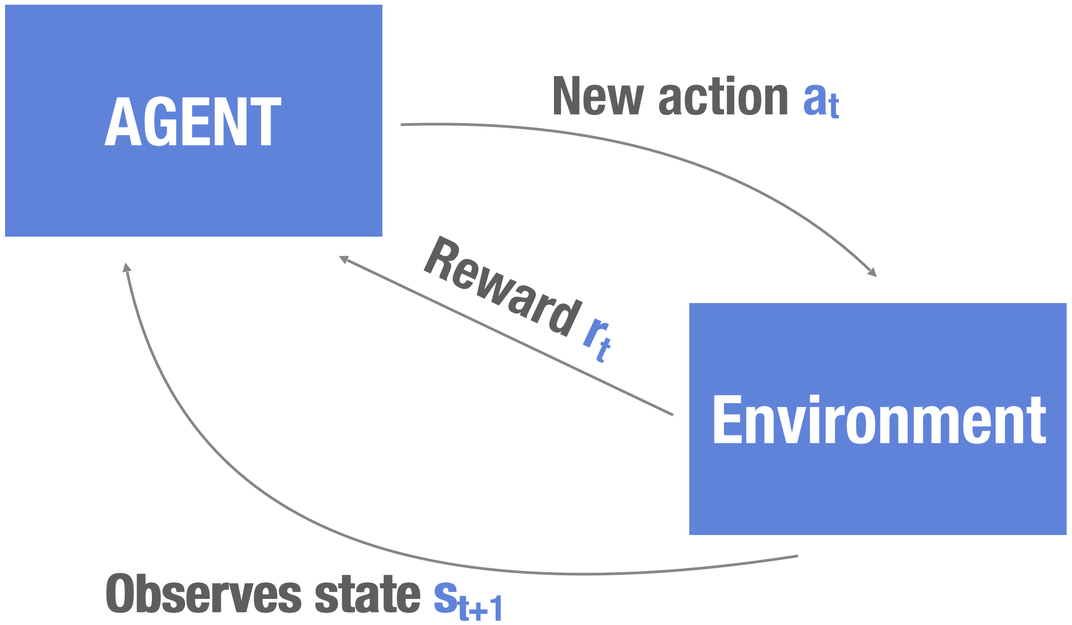

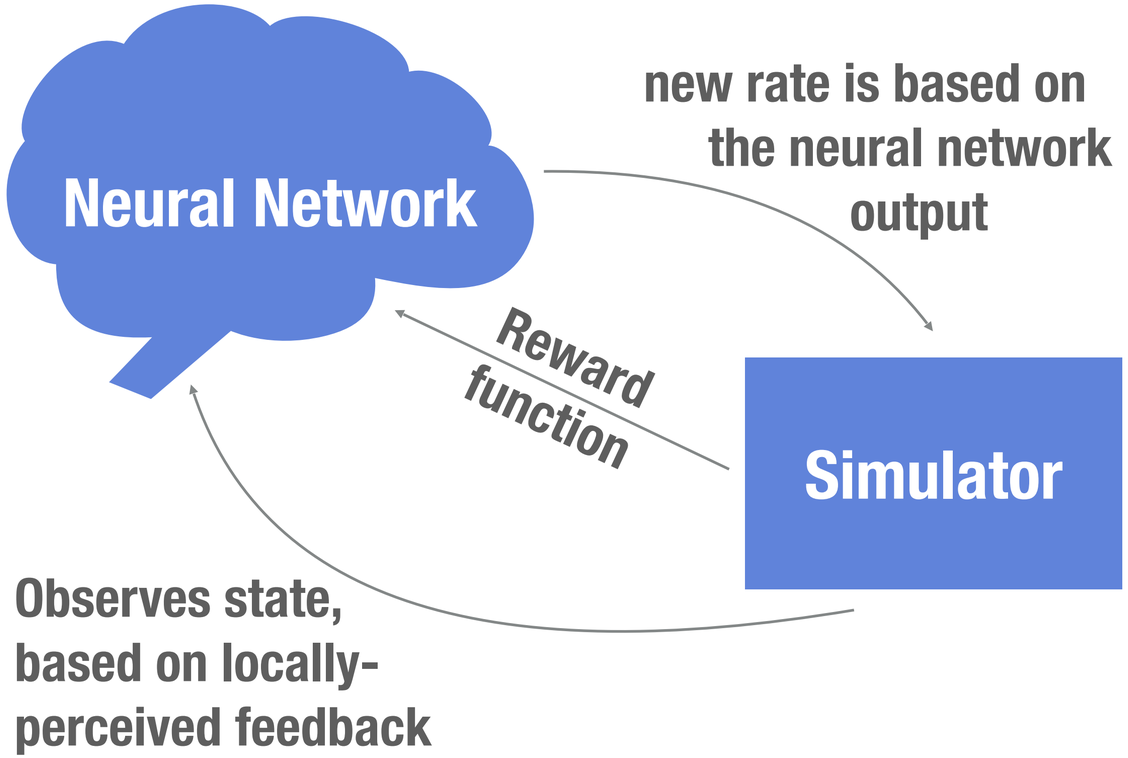

Introducing Aurora

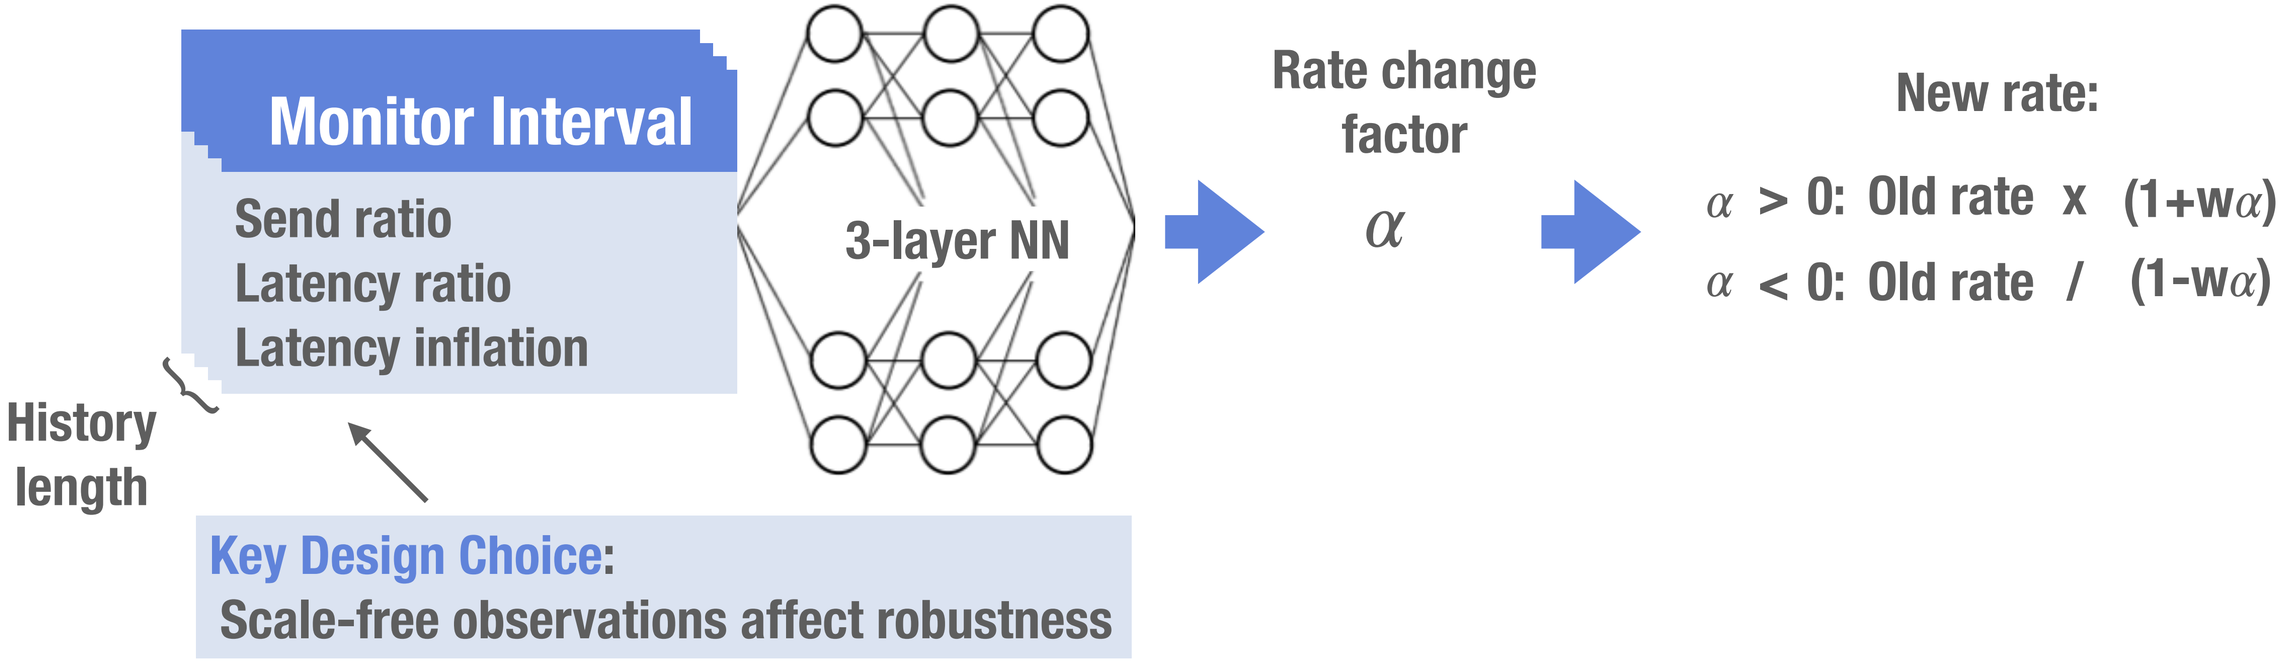

Aurora: Training / Testing

Training

- Very simple simulated network
- Each episode chooses different link parameters
- Entire training platform is **available** as a **standard Gym** environment at:
  Github.com/PCCProject/PCC-RL

Real-World Testing Setup

- Real packets in Linux kernel network emulation
- Using **inference only**
  - Allowing faster adaptation
  - Still producing state-of-the-art results!

Aurora: Robustness

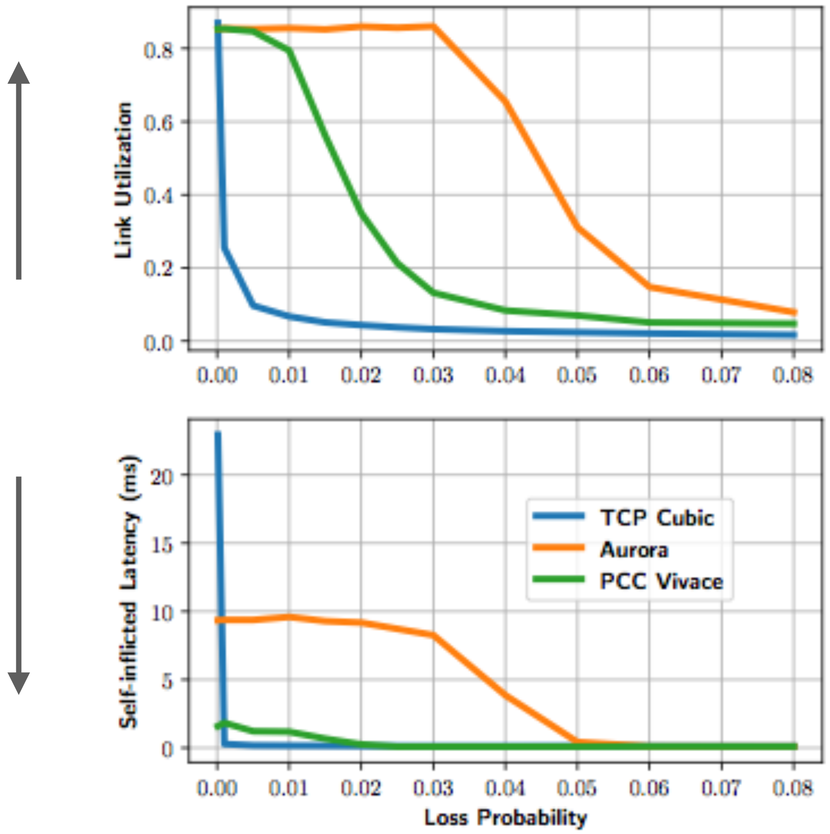

- Compared with TCP Cubic (prevalent) and state-of-the-art PCC Vivace
- Tested over links with changing parameters, some up to **50x lower** and **20x higher** than training range
- Aurora was found to be with comparable or better than state-of-the-art algorithms

Aurora: State-of-the-Art Results

Test settings:

- Emulated network, with real Linux kernel noise
- Time-varying link

Aurora is on the Pareto front of state-of-the-art algorithms

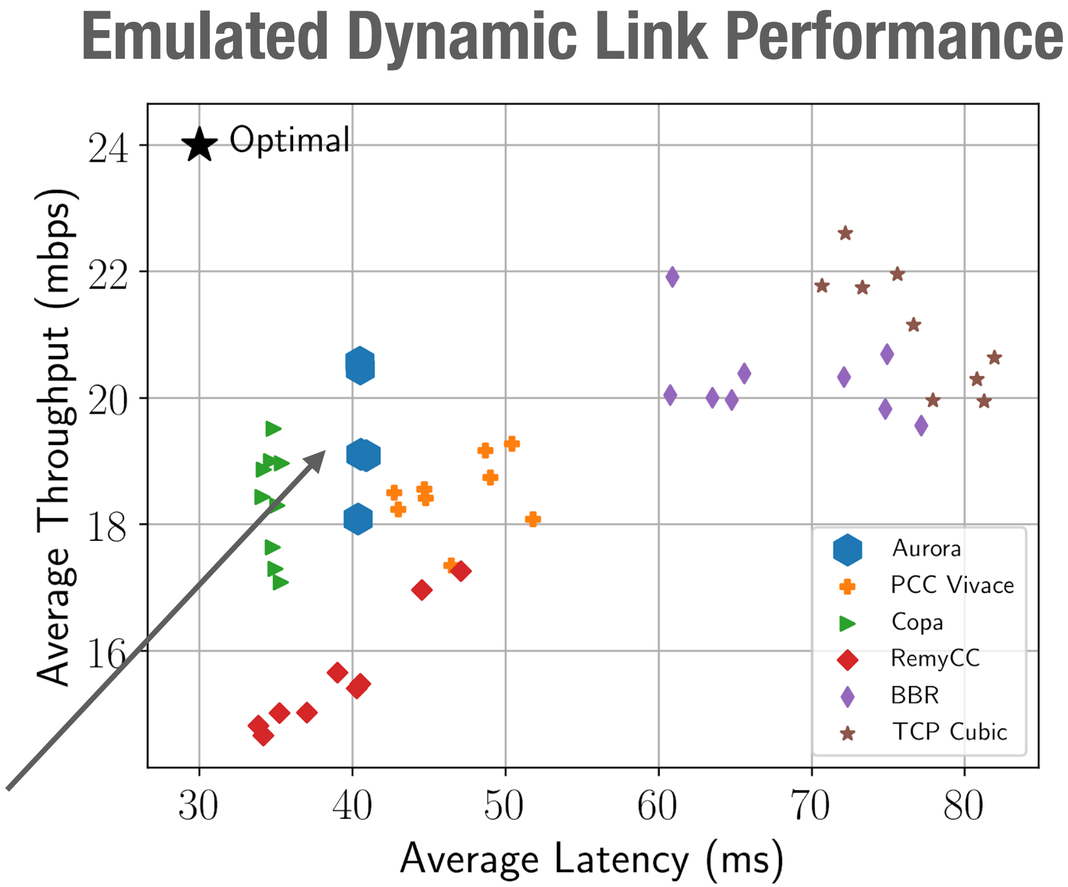

Exciting Directions

- Multi-agent scenarios:
  - Cooperative
  - Selfish
- Online training:
  - Few-shot training
  - Meta-learning
- Multi-objective Learning:
  - File transfer
  - Live video

Github.com/PCCProject/PCC-RL