


# A Deep Reinforcement Learning Perspective on Internet Congestion Control

by **Nathan Jay\***, Noga H. Rotman\*, Brighten  
Godfrey, Michael Schapira, and Aviv Tamar


\*Equal contribution

# Internet Congestion Control


End Host




The Internet  
(maybe?)



Server




# Internet Congestion Control

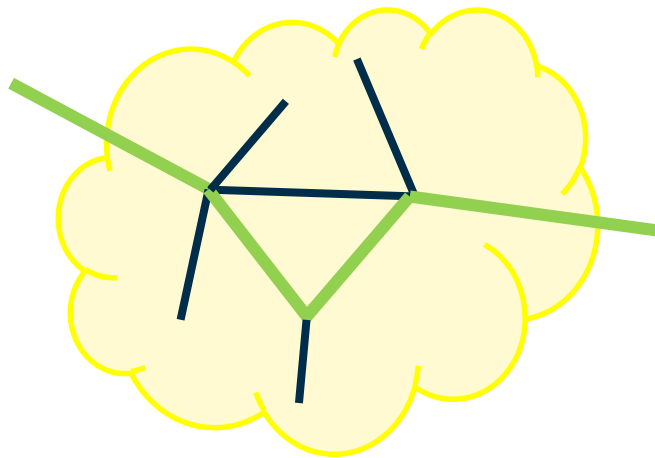


# Internet Congestion Control


End Host



The Internet  
(maybe?)




Server




# Internet Congestion Control


End Host




The Internet  
(maybe?)



Data




Server




# Internet Congestion Control


End Host



The Internet  
(maybe?)




Server

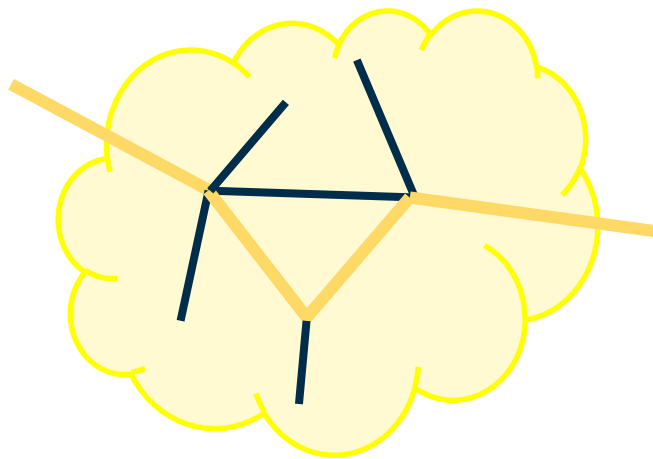


# Internet Congestion Control


End Host



The Internet  
(maybe?)





Server



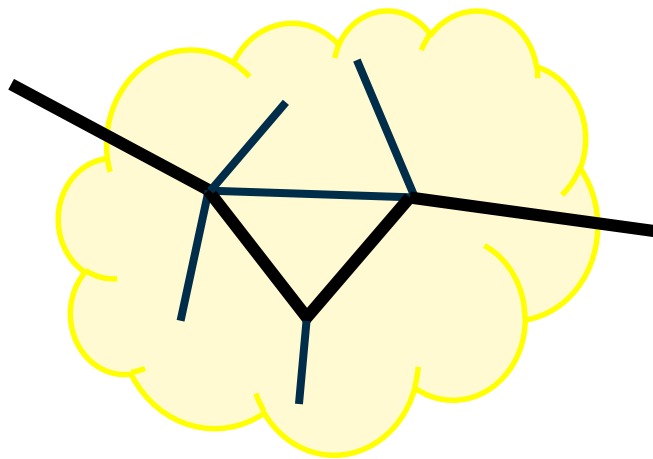
# Internet Congestion Control

End Host




Ack

The Internet  
(maybe?)

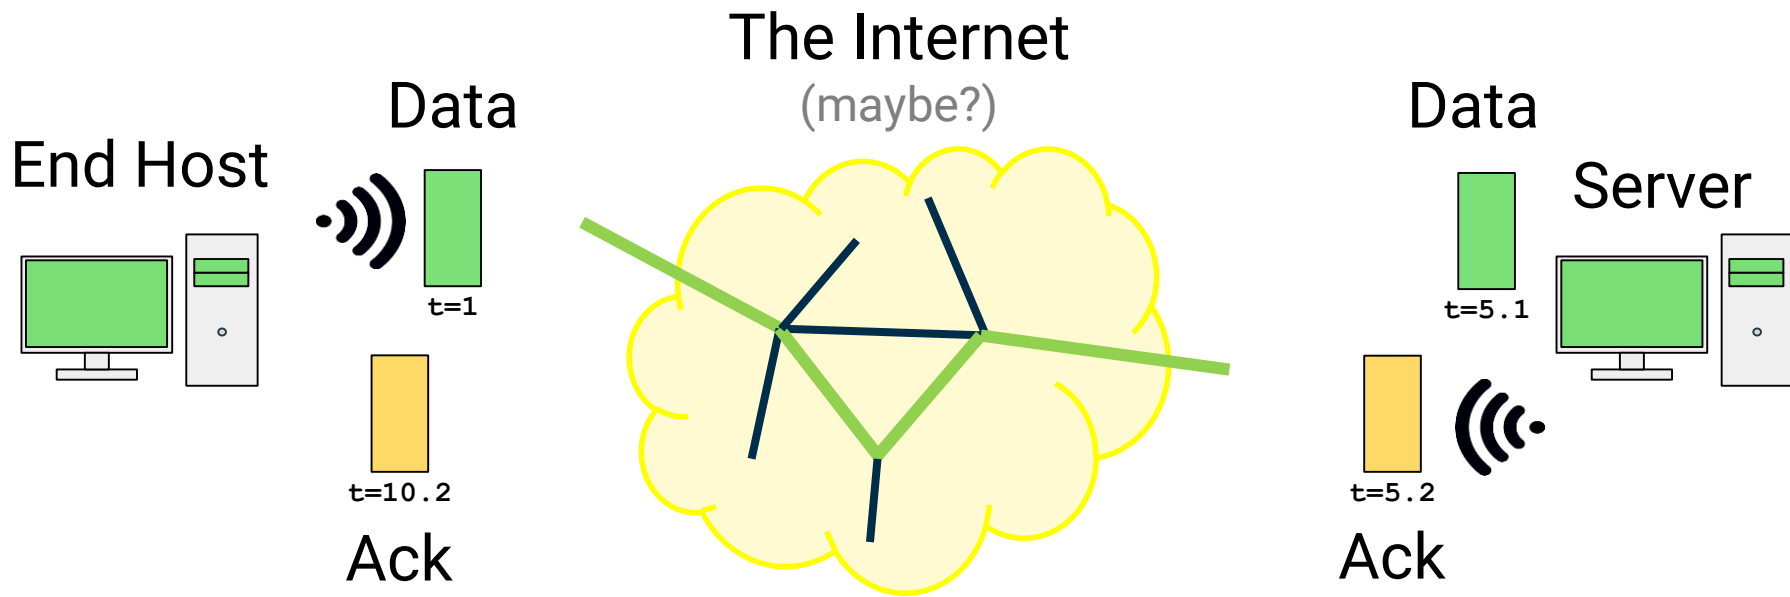


Server






# Internet Congestion Control



# Internet Congestion Control


Latency Trace of Internet Path\*



\*from pantheon.stanford.edu

# Internet Congestion Control


Latency Trace of Internet Path\*



\*from pantheon.stanford.edu

# Internet Congestion Control

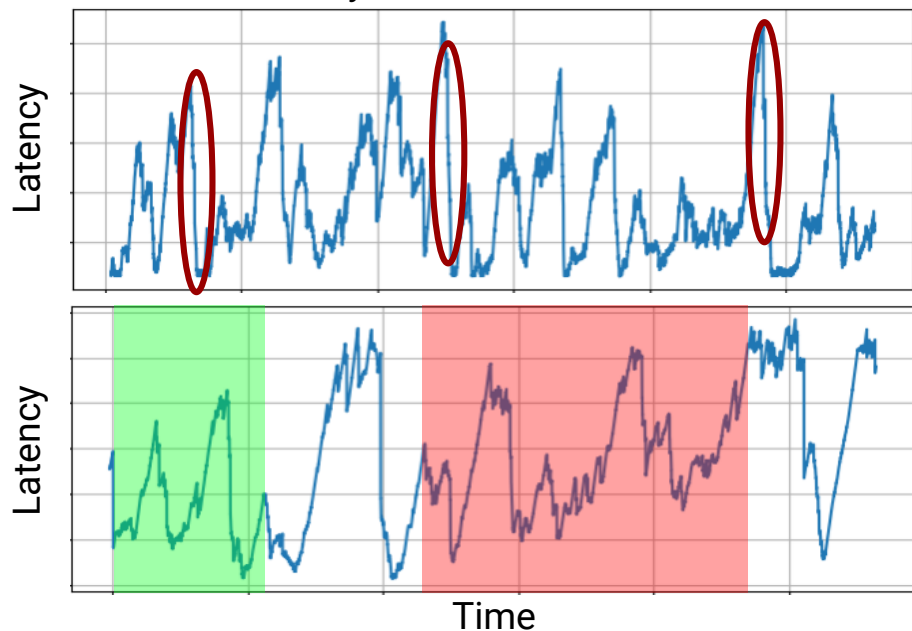
Latency Trace of Internet Path\*



\*from pantheon.stanford.edu

# Internet Congestion Control


Latency Trace of Internet Path\*



\*from pantheon.stanford.edu

## Underlying Complexity:

- Enormous, dynamic network




- Massive agent churn  
~80,000 agents/second

**You**Tube

- Very little information

# Revisiting Congestion Control

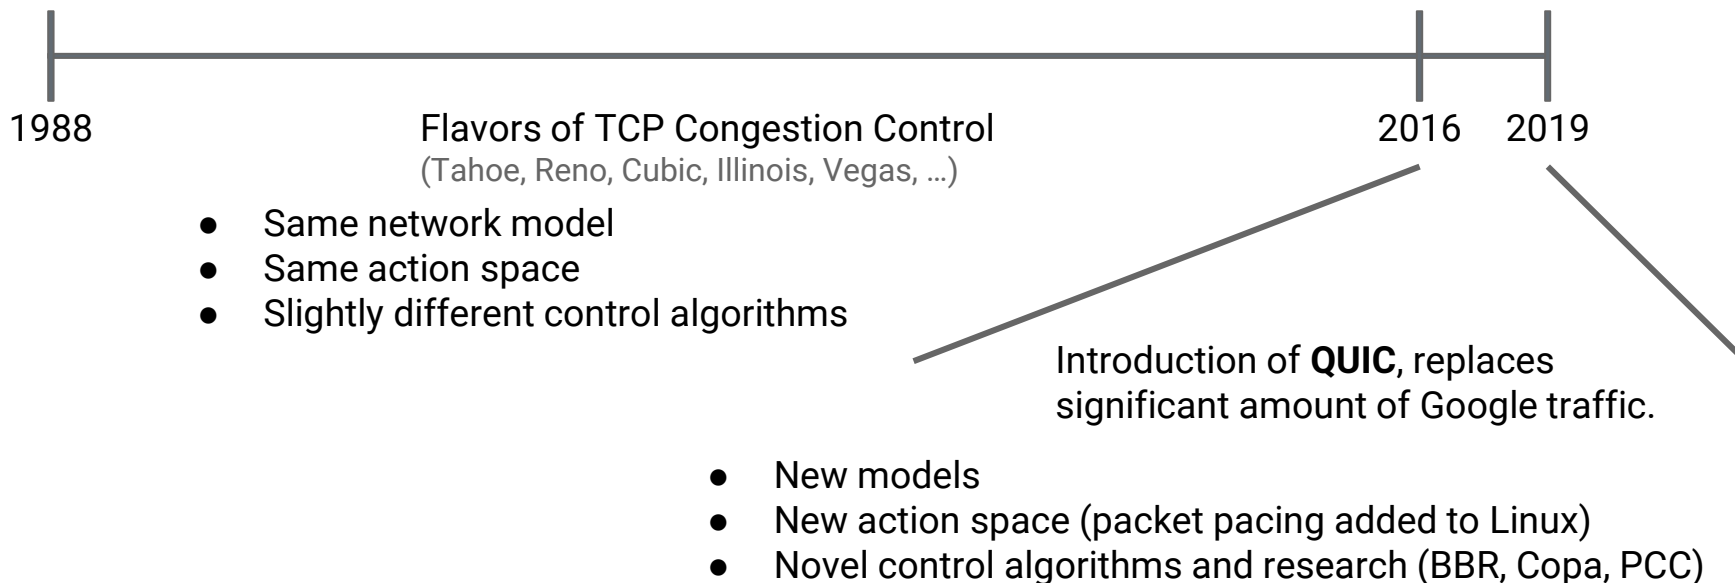
## Congestion Control Timeline



- Same network model
- Same action space
- Slightly different control algorithms


# Revisiting Congestion Control

## Congestion Control Timeline



# Reward-based architecture: PCC


## Observations



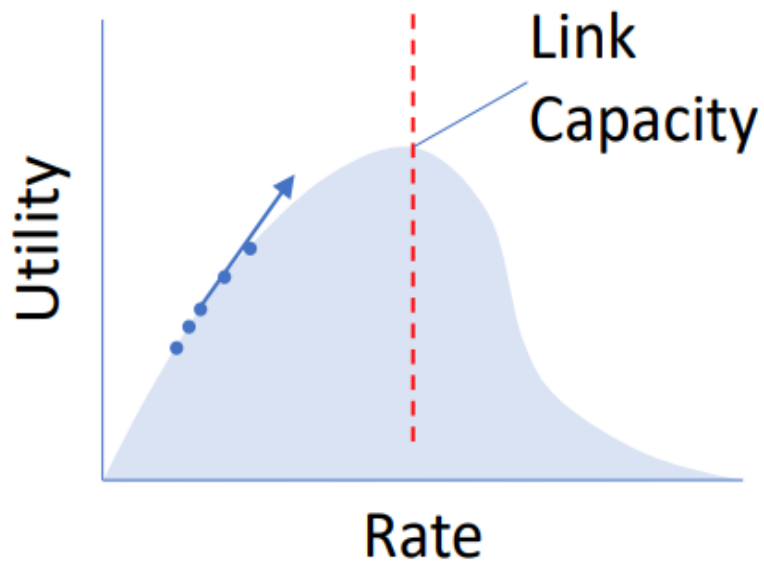


# Reward-based architecture: PCC

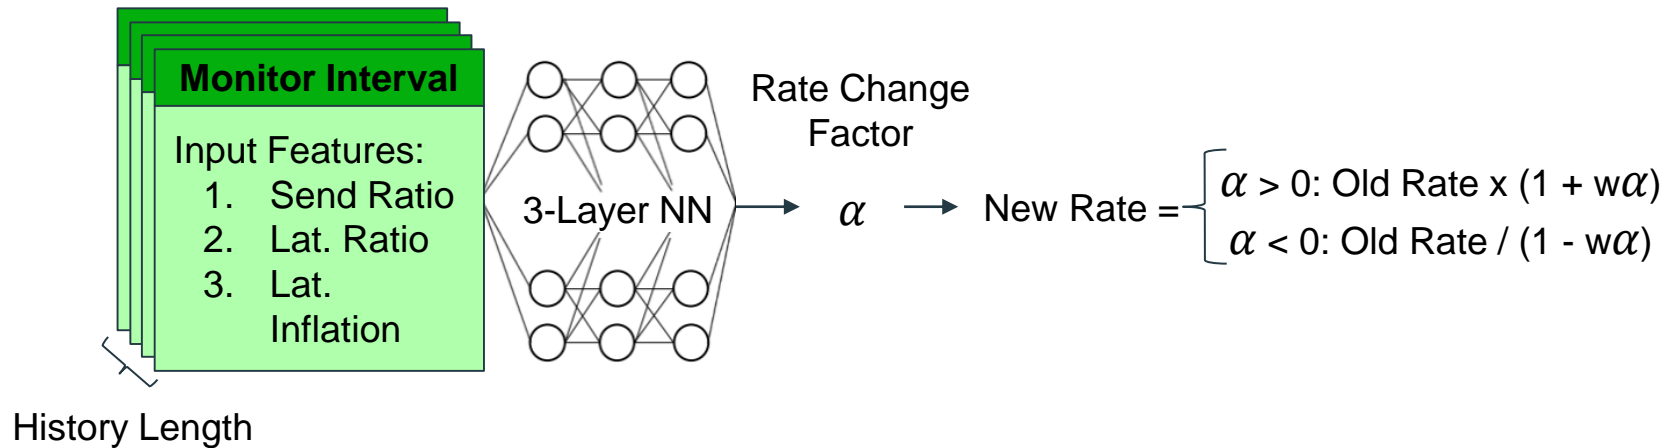
## Observations



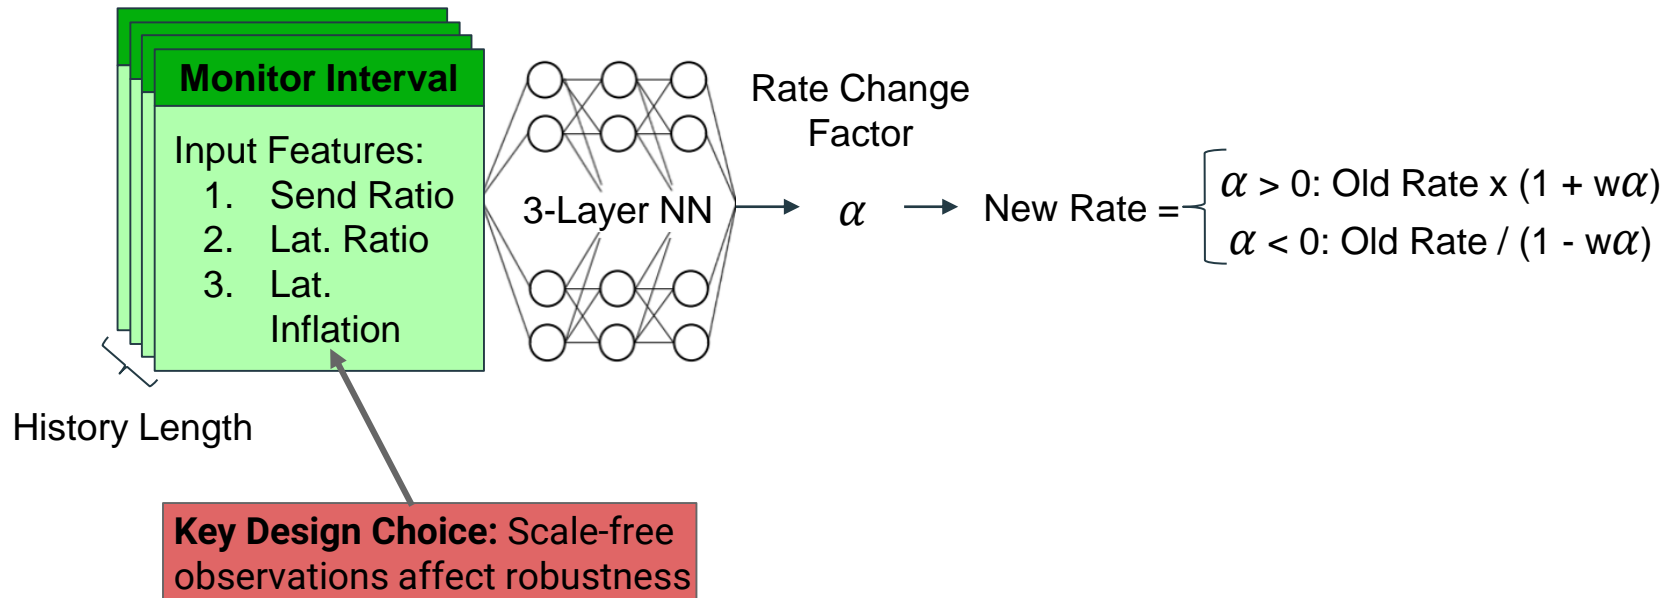
## Actions



# Agent Architecture



# Agent Architecture



# Training/Testing Environment

## Training Environment:

- Simulated network
- Each episode chooses link parameters from a range:

Capacity	Latency	Loss	Queue
1 - 6mbps	50 - 500ms	0 - 5%	1 - ~3000pkt

- **Standard gym** at [github.com/PCCProject/PCC-RL](https://github.com/PCCProject/PCC-RL)

# Training/Testing Environment

## Training Environment:

- Simulated network
- Each episode chooses link parameters from a range:

Capacity	Latency	Loss	Queue
1 - 6mbps	50 - 500ms	0 - 5%	1 - ~3000pkt

- **Standard gym** at [github.com/PCCProject/PCC-RL](https://github.com/PCCProject/PCC-RL)

## Testing Environment:

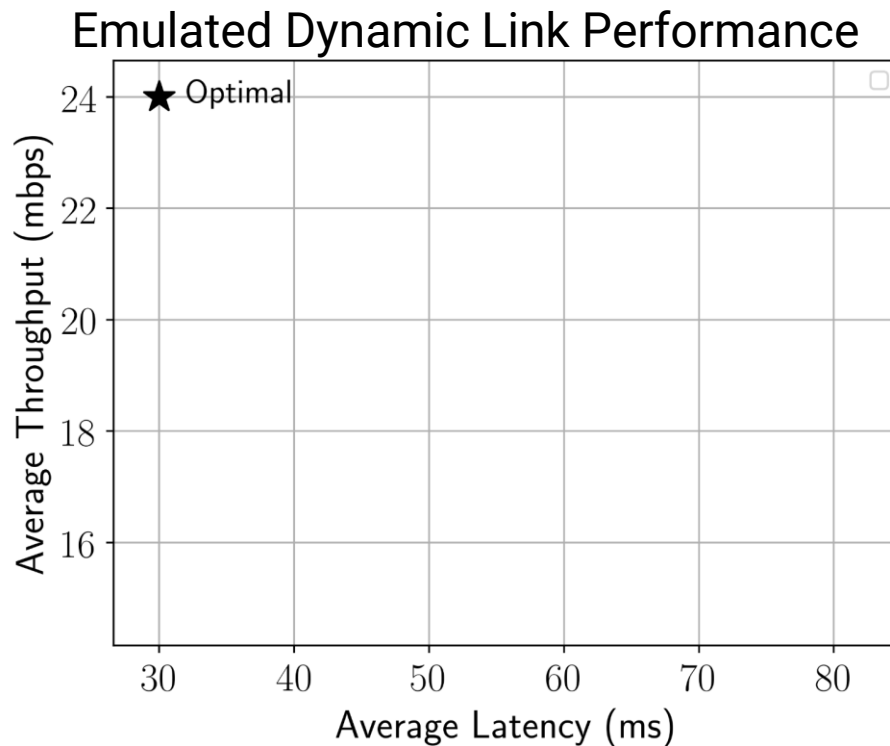
- Real packets in Linux kernel network emulation
- Much wider testing range:

Capacity	Latency	Loss	Queue
1 - 128mbps	1 - 512ms	0 - 20%	1 - 10000pkt

# State-of-the-art Results

## Test Description:


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



# State-of-the-art Results

## Test Description:


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



# State-of-the-art Results

## Test Description:

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

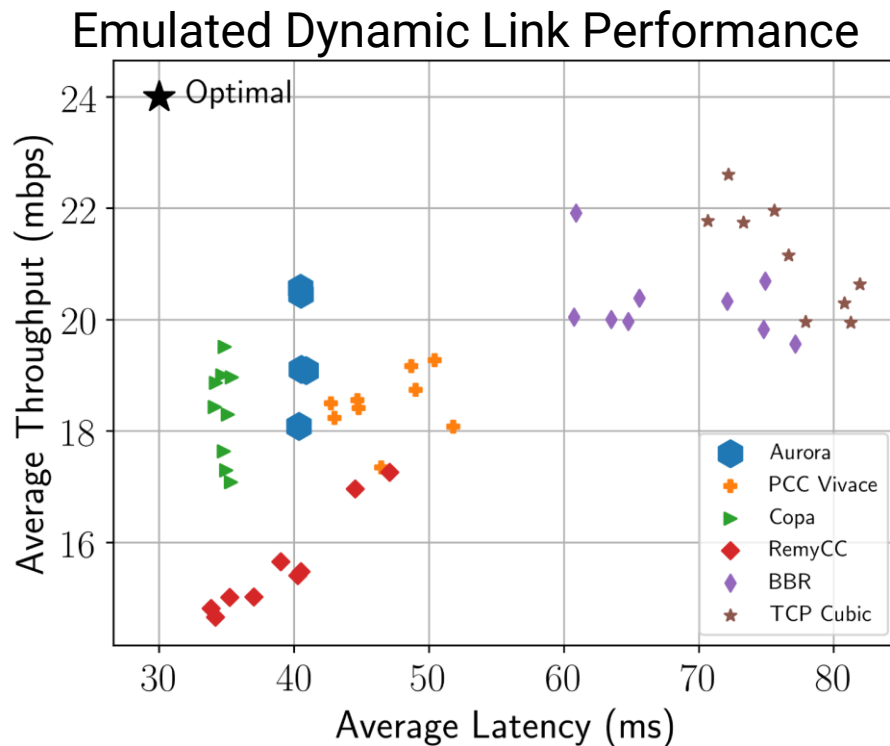




# State-of-the-art Results

## Test Description:

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

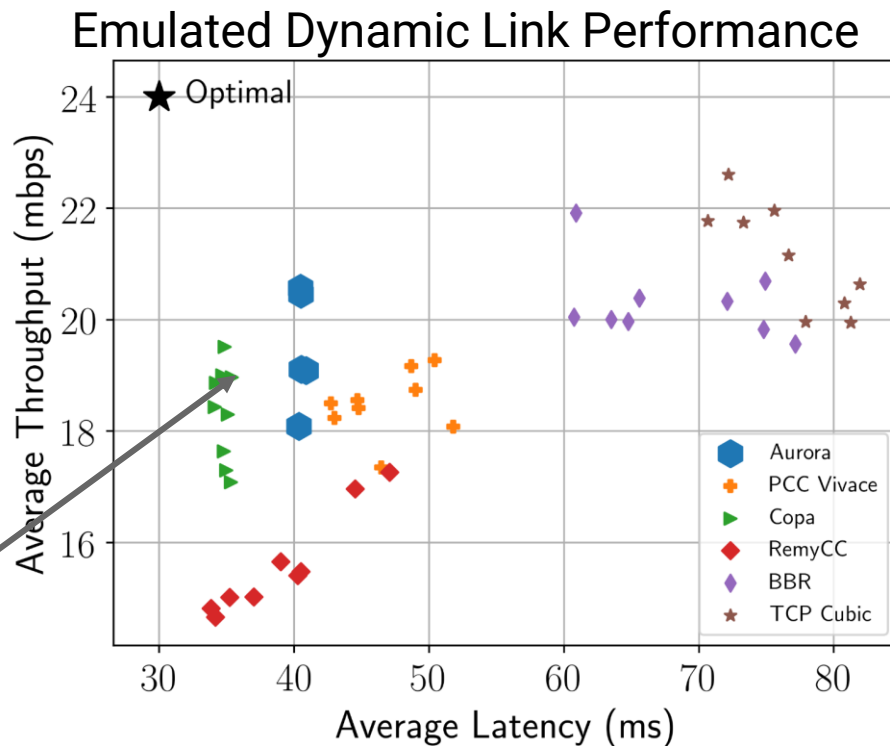


# State-of-the-art Results

## Test Description:


- 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



See us at:

Poster #45

6:30pm - 9:00pm

Pacific Ballroom

Code available at

[github.com/PCCProject/PCC-RL](https://github.com/PCCProject/PCC-RL)