UNIVERSITY OF SUSSEX

Data Transport for the Orbiting Internet

Aiden Valentine
School of Engineering and Informatics
University of Sussex

Introduction

- *Internet from space* is widely adopted
- SpaceX, Amazon, Telesat have been and continue to deploy low earth orbit (LEO) satellite constellations
  - … competing with/complementing terrestrial networks
- 1000s of satellites in multiple orbital shells and planes per shell
- Inter-satellite and ground station to satellite links

LEO networks are the future

- Provide internet connection to remote communities
- Have been used in the aftermath of natural disasters
- Have been used in warzones

in places where terrestrial

networks are damaged

UNIVERSITY OF SUSSEX

LEO Satellite Network Characteristics

- Varying RTT over time and shorter paths will change the base RTT
  - Every 15 seconds the Starlink network reconfigures.
- The interruptions lead to loss, posing a challenge for loss-based protocols.
- A single path may encounter multiple bottlenecks
- Non-congestive loss due to weather interference

Network Dynamics

UNIVERSITY
OF SUSSEX

The Study

- How does congestion control handle the dynamic topology in terms of
  - Responsiveness in capturing bandwidth
  - Fairness
  - Latency inflation

Selected approaches

- Cubic, SaTCP, BBRv3, Sage, Astraea, Vivace
- State of the art schemes
- Interpretable human derived schemes and Reinforcement learning

Methodology

- Mininet based emulation through LeoEM and our mininet test bed looking at …
  - Goodput
  - Intra RTT Fairness
- We have conducted a systematic study comprised of 1000s of individual experiments looking at ...
  - Responsiveness
  - Inter RTT fairness
  - Efficiency

UNIVERSITY OF SUSSEX

Results of LEO emulation
Goodput

- Paths experience various levels of dynamics
- Hard to interpret, why do the RL schemes underperform ?

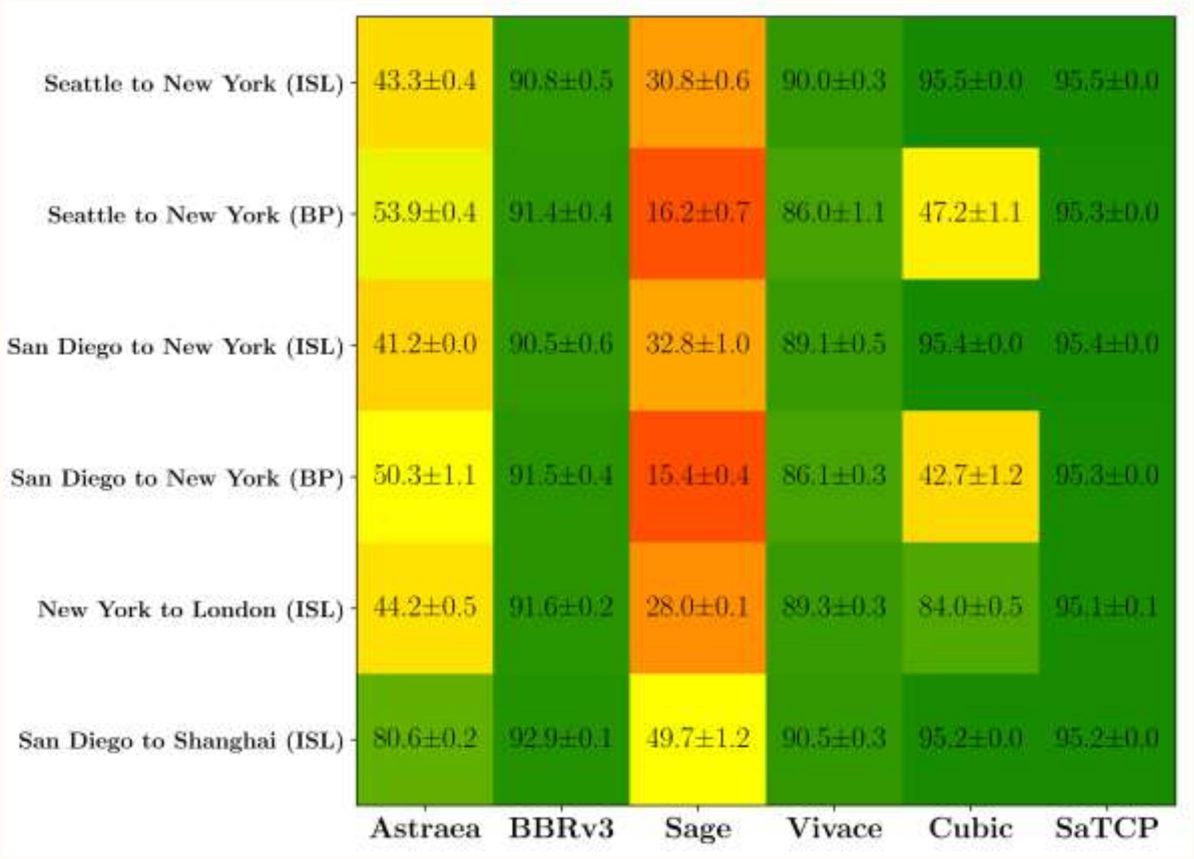

	Astraea	BBRv3	Sage	Vivace	Cubic	SaTCP
Seattle to New York (ISL)	43.3±0.4	90.8±0.5	30.8±0.6	90.0±0.3	95.5±0.0	95.5±0.0
Seattle to New York (BP)	53.9±0.4	91.4±0.4	16.2±0.7	86.0±1.1	47.2±1.1	95.3±0.0
San Diego to New York (ISL)	41.2±0.0	90.5±0.6	32.8±1.0	89.1±0.5	95.4±0.0	95.4±0.0
San Diego to New York (BP)	50.3±1.1	91.5±0.4	15.4±0.4	86.1±0.3	42.7±1.2	95.3±0.0
New York to London (ISL)	44.2±0.5	91.6±0.2	28.0±0.1	89.3±0.3	84.0±0.5	95.1±0.1
San Diego to Shanghai (ISL)	80.6±0.2	92.9±0.1	49.7±1.2	90.5±0.3	95.2±0.0	95.2±0.0

UNIVERSITY OF SUSSEX

Results of LEO emulation Fairness over

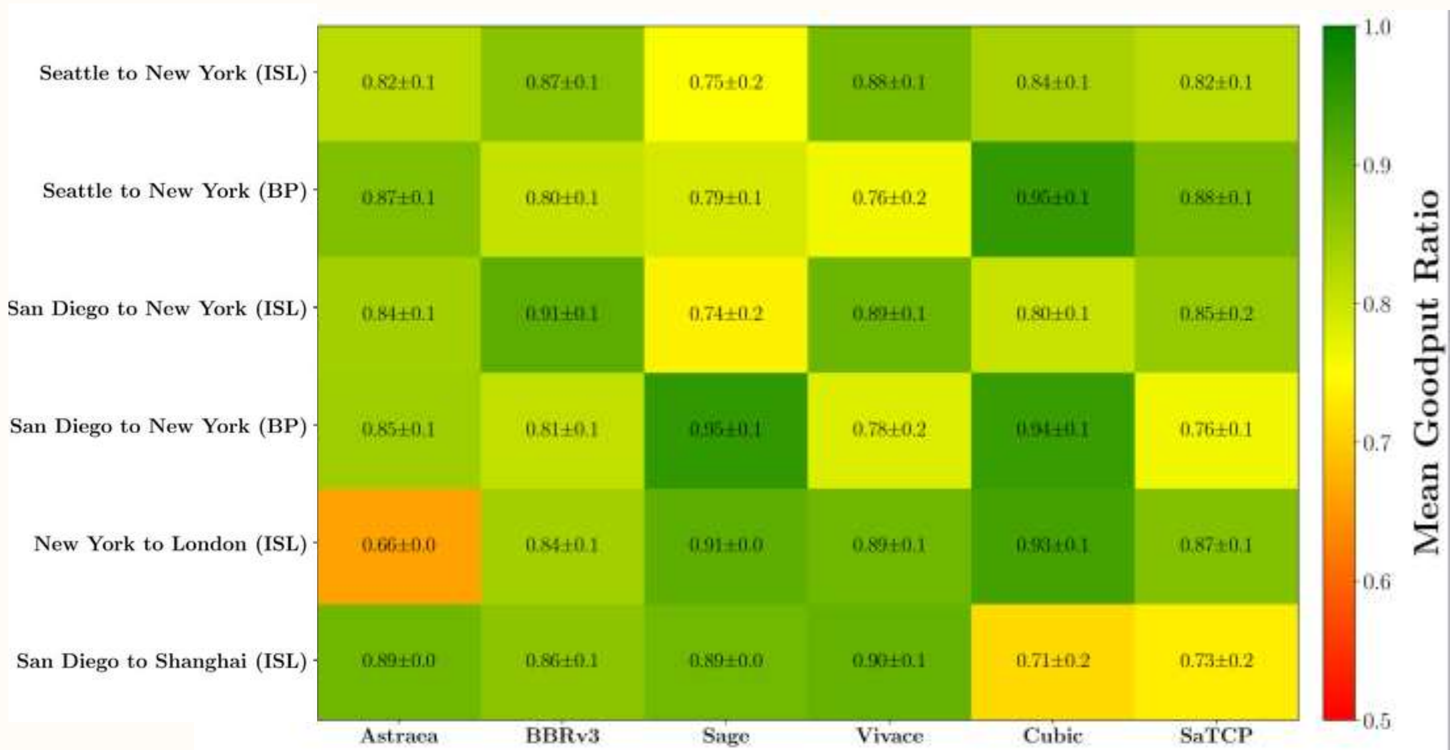

(a) Buffer Size: 1×BDP

- Fairness appears good? But is it under-utilisation ?
- Hard to interpret fairness due to the RTT variation and lossy handovers

UNIVERSITY OF SUSSEX

Microbenchmark - Responsiveness

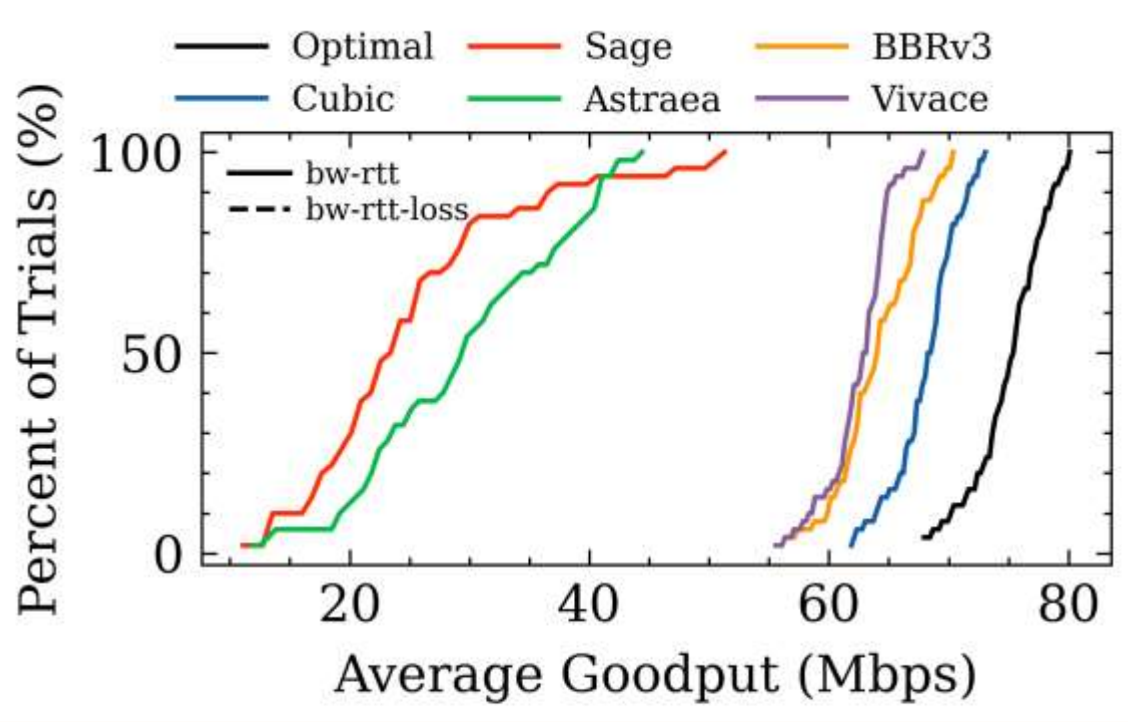

- Human derived schemes are more responsive
- The RL protocols struggle with responsiveness

UNIVERSITY OF SUSSEX

Microbenchmark - Responsiveness

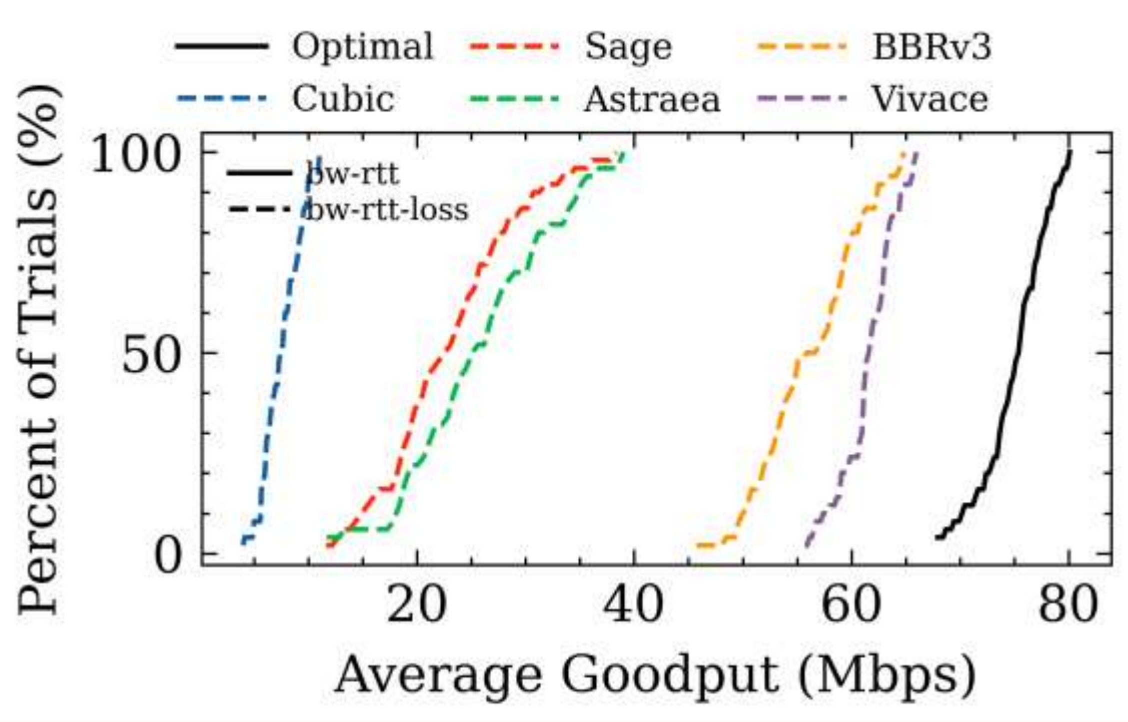

- Cubic doing the worst
- The RL schemes are resilient to non-congestive loss

UNIVERSITY OF SUSSEX

Microbenchmark – Individual run

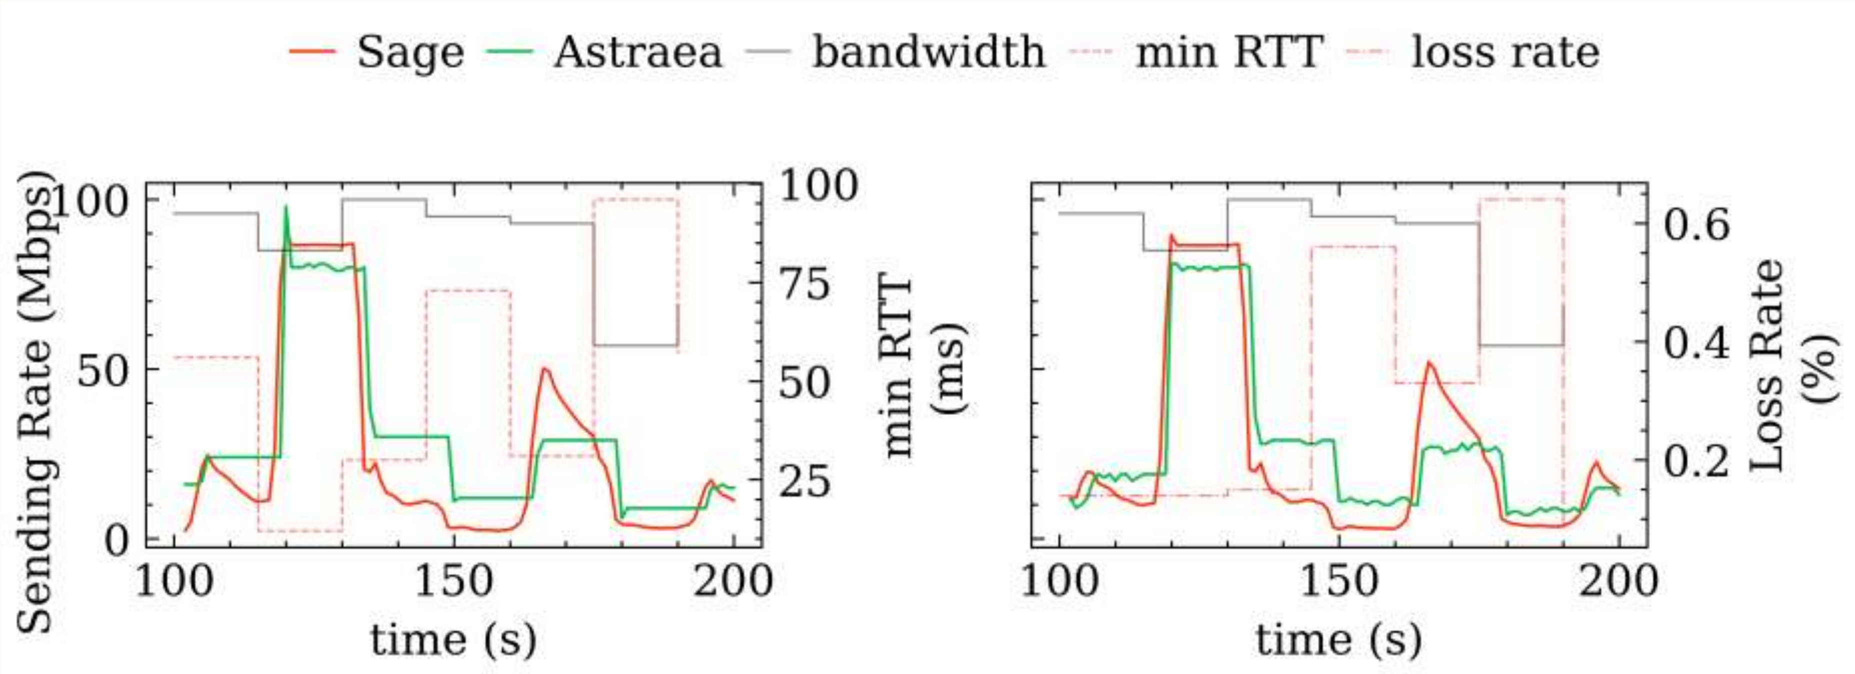

Add bbr

Microbenchmark - Fairness Inter-RTT

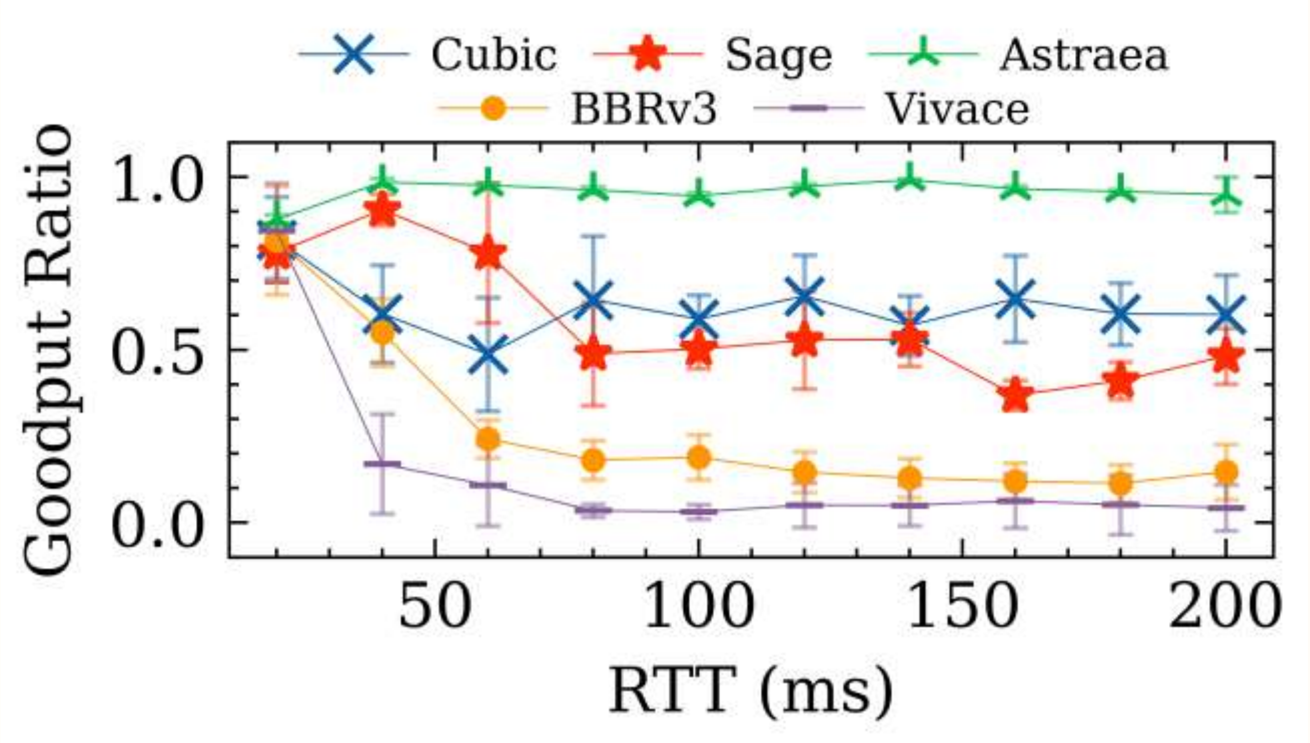

- We have not found a heuristic for RTT fairness
- Embedding fairness in the reward function during training yields better fairness

UNIVERSITY OF SUSSEX

Microbenchmark - Efficiency

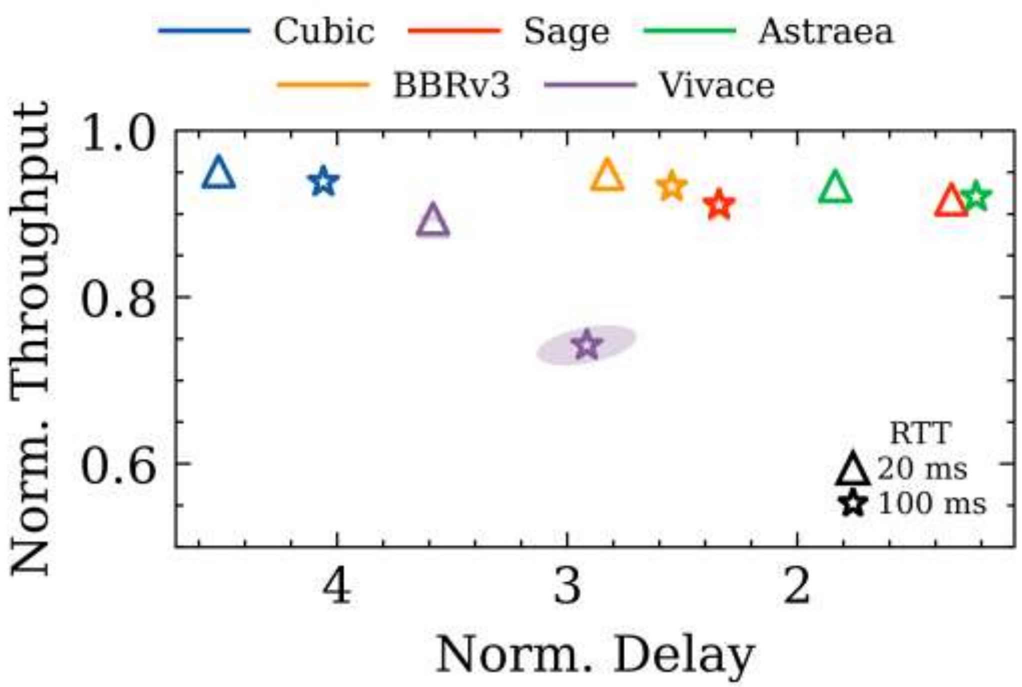

- Sage has learned a policy that outperforms delay based schemes

UNIVERSITY OF SUSSEX