2021 NSF CYBER-PHYSICAL SYSTEMS PRINCIPAL INVESTIGATORS' MEETING

Distributed Learning for Control of Cyber-Physical Systems

Michael M. Zavlanos
Duke University

Challenge:

A novel distributed RL framework for the control of CPS, so that it has performance guarantees under partial observations.

Project Overview

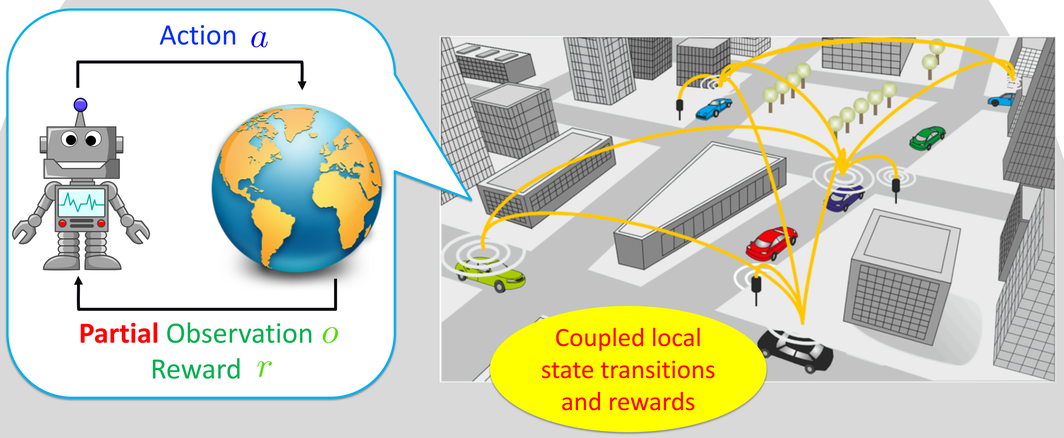

Solution:

- A new one-point residual-feedback policy gradient estimator, that has much lower variance than conventional one-point estimators.
- Decentralization of the policy gradient estimator using consensus to get local estimates of the global objective value.

Proposed Framework

Existing distributed Actor-Critic Methods require full state and action observation

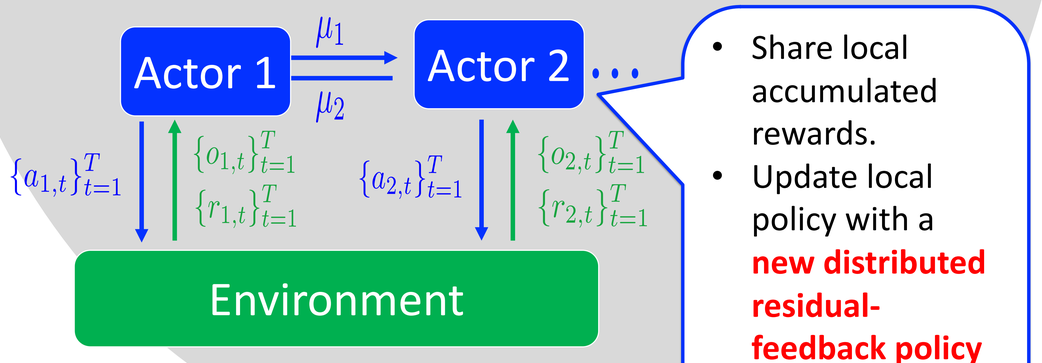

- Share local accumulated rewards.
- Update local policy with a **new distributed residual-feedback policy gradient.**

Scientific Impact:

- Application of residual-feedback gradient estimation to any black-box CPS learning and control problem.
- The first guarantees on sample complexity of distributed RL under partial observations.

Broader Impact:

- Preservation of user's privacy and robustness to single-node failure
- Applications to many domains, e.g., smart city, health care, etc.
- K-12, undergraduate, and graduate education
- Diversity

Award ID#: CNS #1932011, Award Date: 10/01/2019