

**Title: Enabling fairness in multi-agent reinforcement learning for decision making.**

**Tian Lan**

George Washington University

Fairness is a crucial objective in many decision-making problems that involve multiple agents potentially competing for shared utilities/rewards. Examples range from resource allocation in networked systems to mission planning and automation. Recently, multi-agent reinforcement learning has been successfully applied to many decision-making problems, but often does not take fairness into account. In this talk, we will share some of our recent work on a utility-based framework for bringing the notion of fairness into multi-agent reinforcement learning, as well as its applications to network management problems.