

Recent Advances in Average-Reward Restless Bandits

Monday, June 17, 2024 4:00 PM (30 minutes)

We consider the infinite-horizon, average reward restless bandit problem. For this problem, a central challenge is to find asymptotically optimal policies in a computationally efficient manner in the regime where the number of arms, N , grows large. Existing policies, including the renowned Whittle index policy, all rely on a uniform global attractor property (UGAP) assumption to achieve asymptotic optimality, which is a complex and difficult-to-verify assumption. In this talk, I will present new, sampling-based policy designs for restless bandits. One of our proposed policies breaks the long-standing UGAP assumption for the first time; then our subsequent policies eliminate the need for the UGAP assumption to achieve asymptotic optimality entirely. Our techniques offer new insights into guaranteeing convergence (avoiding undesirable attractors or cycles) in large stochastic systems.

Primary authors: HONG, Yige (Carnegie Mellon University); XIE, Qiaomin (University of Wisconsin–Madison); CHEN, Yudong (University of Wisconsin–Madison); WANG, Weina (Carnegie Mellon University)

Presenter: WANG, Weina (Carnegie Mellon University)

Session Classification: Parallel session: Online learning in stochastic networks