

Designing M/G/1 Scheduling Policies from Job Size Samples

Tuesday, June 18, 2024 5:00 PM (30 minutes)

The Gittins policy is known to minimize mean response time in the M/G/1 with unknown job sizes, provided the job size distribution is known. In practice, however, one does not have direct access to the job size distribution. Motivated by this, we consider the problem of designing a scheduling policy based on a finite number of samples from the job size distribution. This is an open problem. A natural idea is a policy one might call *empirical Gittins*, which constructs a policy using the empirical distribution of the samples. But it is unclear how empirical Gittins compares to the true Gittins policy (constructed from the true job size distribution).

This talk will present initial results on M/G/1 scheduling with only sample access to the job size distribution. We show that empirical Gittins is a good proxy for true Gittins in the finite support case. Specifically, we show that empirical Gittins is a $(1 + \epsilon)$ -approximation for mean response time with probability $1 - \delta$, provided at least $O\left(\frac{1}{\epsilon^2} \log \frac{1}{\delta}\right)$ samples. Our results are based on applications of the recently derived “WINE” queueing identity. We will also present initial progress towards the general case, which combine WINE with the “SOAP” analysis of M/G/1 scheduling policies.

Primary author: RAMAKRISHNA, Shefali (Cornell University)

Co-author: Dr SCULLY, Ziv (Cornell University)

Presenter: RAMAKRISHNA, Shefali (Cornell University)

Session Classification: Parallel session: Reinforcement learning and queueing I