Performance Improvement Through Active Idleness
José A. Moreira, Carlos F. Bispo

Abstract

We illustrate by an example the potential of distributed scheduling with Active Idleness to improve the performance of multi-class queueing networks, which are originally controlled by non-idling, or work-conserving, policies. The queuing network we use in our simulation experiments is due to Dai, [1]. This particular network was shown to be unstable if FIFO is the scheduling policy. However, for the *Last Buffer First Serve* policy, [2], it is stable. Under this setting we show that forcing inactivity during some periods of time in the presence of customers may result in significant performance gains.

Keywords: Queuing networks, Distributed Scheduling, Stability.

You may find the original 2003 document by following the link
http://users.isr.ist.utl.pt/∼cfb/paper11.pdf