Dynamic Scheduling in Switched Queueing Networks with Heavy-Tailed Traffic

Mihalis G. Markakis[1]

We consider the problem of scheduling in a single-hop switched queueing network. This setting is often used to model the dynamics and decisions arising in several communication networks, such as input-queued switches, uplinks and downlinks of cellular networks, and wireless ad-hoc networks. Motivated by empirical evidence of self-similarity and long-range dependence in Internet traffic traces (see Park & Willinger, 2000), we assume that the network receives a mix of heavy-tailed and light-tailed traffic. Our goal is to analyze the impact of heavy-tailed traffic on the performance of a widely studied class of dynamic scheduling policies, the Max-Weight policies, known for their throughput optimality property (see Tassiulas & Ephremides, 1992). As performance metric we use delay stability: a queue is delay stable if its expected steady-state delay is finite, and delay unstable otherwise.

We first show that a queue receiving heavy-tailed traffic is delay unstable under any scheduling policy (a direct consequence of the Pollaczek-Khinchine formula). Then, we focus on the celebrated Max-Weight scheduling policy, and show that a light-tailed queue that conflicts (i.e., cannot be served simultaneously) with a heavy-tailed queue is also delay unstable. This is true irrespective of the rate or the tail distribution of the light-tailed traffic, or other scheduling constraints in the network. Surprisingly, we present a simple queueing system where a light-tailed queue that does not conflict with a heavy-tailed queue can be delay unstable as well. Delay stability in this case depends on the arrival rates.

As a first step towards generalizing this result, we establish a connection between delay instability of queues in the stochastic network, and the evolution of the corresponding fluid model from certain initial conditions. This connection provides a *systematic and practical methodology for identifying delay unstable queues* in any given switched network. Moreover, it allows us to reach an interesting conclusion for the special case of switched networks with disjoint schedules: there exist arrival rate vectors in the stability region for which all queues are delay unstable.

Finally, we turn our attention to the class of Max-Weight-α scheduling policies. We show that if the α-parameters are chosen suitably, then the sum of the α-moments of steady-state queue lengths is finite. We provide an explicit upper bound on the latter quantity, which we use to derive results related to the delay stability of queues, and the optimal scaling of higher order moments of steady-state queue lengths with traffic intensity.

---

[1] Joint work with Eytan Modiano and John N. Tsitsiklis. The authors are affiliated with the Laboratory for Information and Decision Systems, at the Massachusetts Institute of Technology, Cambridge, MA, USA.