Scheduling in Multiserver Systems: Approaches and Open Problems (Part II)

Mor Harchol-Balter
Computer Science Department, Carnegie Mellon University

In Part II we move on to looking at server farm models that are representative of Web server farms. Here the individual servers all employ Processor-Sharing (PS) service order, which is a preemptive time-sharing scheduling policy. Examples of such server farms include the Cisco Local Director product, the IBM Network Dispatcher, Microsoft SharePoint, and F5 Labs BIG/IP.

We first show that the desired routing/dispatching policy for minimizing mean response time in the case of PS server farms can be very different from that for FCFS server farms. We then focus on the Join-the-Shortest-Queue routing policy and discuss some existing approximations in the literature (e.g., [5, 6]) and some new approximations that apply to the case of PS server farms [2].

Finally, we turn to the question of what server farm architectures are optimal for minimizing mean response time. Here we consider server farms where the individual servers employ Shortest-Remaining-Processing-Time (SRPT) scheduling, or there is a central SRPT queue. Such models are very difficult to analyze stochastically. The closest stochastic result only manages a server farm with a central priority queue [3]. Primary work on SRPT server farms is dominated by the STOC/FOCS/SPAA community, which uses competitive ratios as its metric. We will describe server farm architectures that appear to be optimal, but aren't, and discuss their competitive ratios, both for the case of a central queue model [4] and an immediate-dispatch model [1]. These results motivate future research directions for researchers in the stochastic community.

References

[1] N. Avrahami and Y. Azar. Minimizing total flow time and total completion time with immediate dispatching. In *ACM Symposium on Parallel Algorithms and Architectures (SPAA 1997)*, pages 11–18, 2003.

[2] V. Gupta, M. Harchol-Balter, K. Sigman, and W. Whitt. Analysis of Join the Shortest Queue for Processor-Sharing Server Farms. Work in progress 2006.

[3] M. Harchol-Balter, T. Osogami, A. Scheller-Wolf, and A. Wierman. Multi-server queueing systems with multiple priority classes. *Queueing Systems: Theory and Applications*, 51(3-4):331–360, 2005.

[4] S. Leonardi and D. Raz. Approximating total flow time on parallel machines. In *ACM Symposium on Theory of Computing*, pages 110–119, 1997.

[5] R. D. Nelson and T. K. Philips. An Approximation to the Response Time for Shortest Queue Routing. In *ACM SIGMETRICS Conference*, pages 181–189, May 1989.

[6] R. D. Nelson and T. K. Philips. An Approximation for the Mean Response Time for Shortest Queue Routing with General Interarrival and Service Times. *Performance Evaluation*, 17:123–139, 1993.