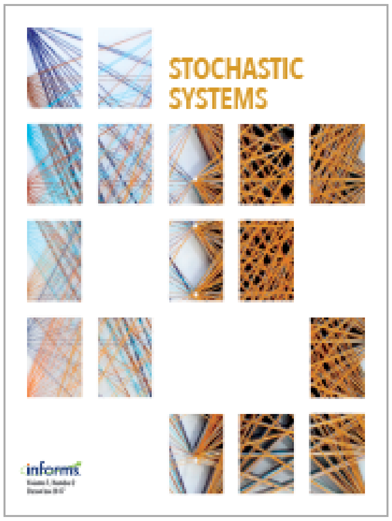

Stochastic Systems

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Open Problem—M/G/k/SRPT Under Medium Load

Isaac Grosof

To cite this article:
Isaac Grosof (2019) Open Problem—M/G/k/SRPT Under Medium Load. Stochastic Systems 9(3):297-298. https://doi.org/10.1287/stsy.2019.0042

informs
http://pubsonline.informs.org/journal/stsy
STOCHASTIC SYSTEMS
Vol. 9, No. 3, September 2019, pp. 297–298
ISSN 1946-5238 (online)

Special Section: Open Problems in Applied Probability

Open Problem—M/G/k/SRPT Under Medium Load

Isaac Grosof[a]

[a] Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
Contact: igrosof@cs.cmu.edu, https://orcid.org/0000-0001-6205-8652 (IG)

Published Online in Articles in Advance: September 18, 2019

https://doi.org/10.1287/stsy.2019.0042

History: This paper was accepted for the *Stochastic Systems* Special Section on Open Problems in Applied Probability, presented at the 2018 INFORMS Annual Meeting in Phoenix, Arizona, November 4–7, 2018.

Funding: I. Grosof was supported by the National Science Foundation [Grants CSR-1763701 and XPS-1629444].

Keywords: stochastic networks • control • optimization

The shortest remaining processing time (SRPT) scheduling policy has been deployed in many computer systems, such as web servers (Harchol-Balter et al. 2003), networks (Montazeri et al. 2018), databases (Guirguis et al. 2009), and operating systems (Bunt 1976). SRPT has also received extensive study by queueing theorists. In 1966, the mean response time for SRPT was first derived (Schrage and Miller 1966), and in 1968, SRPT was shown to minimize mean response time (Schrage 1968). However, these results are only known for *single-server* systems. Much less is known for *multiserver* systems, such as the M/G/k.

To our knowledge, the first analytic results on the M/G/k/SRPT were given by Grosof et al. (2018). They proved analytic bounds on the mean response time of the M/G/k/SRPT. They showed these bounds are tight in the heavy-traffic limit, namely as load approaches capacity. However, their bounds leave open the mean response time under *medium* load.

The medium-load regime has deep importance to the design of computer systems. When designing a multiserver computer system, we want to achieve high utilization while also keeping mean response time low. On a graph of mean response time $E[T]$ by load ρ, the region of importance is the one known as the "knee" of the curve. The knee refers to the load at which mean response time begins to increase above its minimum, the mean job size. This motivates the following question: in an M/G/k with SRPT scheduling, where is the knee of the mean response time curve?

To formalize this question, we pose the following open problem: For a given $\epsilon > 0$, what is the maximum load $\rho_H(\epsilon)$ such that

$$E[T]^{\mathrm{M/G/k/SRPT}} \leq (1+\epsilon)E[S], \tag{1}$$

where T denotes the response time distribution and S denotes the job size distribution? We are particularly interested in this question for ϵ in the range 0.1 to 1.

Prior Work: M/G/k/FCFS

Note that the equivalent question for a first-come, first-served system (FCFS), the M/G/k/FCFS, is far better understood. Early analysis by Kingman (1970) and later by Daley (1997) bounded mean response time based on the first two moments of the job size distribution. Harchol-Balter et al. (2005) exactly analyzed mean response time in an M/Ph/k/FCFS system. Because an arbitrary distribution can be approximated by a phase-type distribution, this analysis allows an approximation of mean response time in the M/G/k/FCFS. Most recently, Gupta and Osogami (2011) showed how to use the entire sequence of moments of a job size distribution to give tight bounds on mean response time. These results allow us to tightly bound mean response time in the M/G/k/FCFS at all loads and answer questions such as (1) for the M/G/k/FCFS.

Figure 1. Mean Response Time (E[T]) vs. Load Under SRPT and FCFS Scheduling

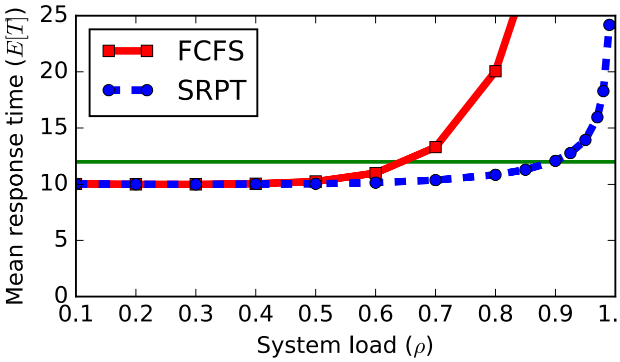

Notes. Load at which E[T] first exceeds 1.2 times mean job size E[S] is about 0.65 under FCFS and about 0.9 under SRPT. We use $k = 10$ servers. The job size distribution S is hyperexponential with $E[S] = 10$ and $C^2 = 10$.

Simulation

Although the M/G/k/FCFS system is better understood than the M/G/k/SRPT, simulations indicate that the SRPT system offers superior performance under medium load. SRPT's excellent empirical performance in this regime makes theoretical analysis especially important.

Figure 1 shows mean response time under FCFS and SRPT scheduling policies for an M/G/k system. The job size distribution is hyperexponential with mean 10 and $C^2 = 10$, and there are $k = 10$ servers in the system. Under FCFS scheduling, the load at which mean response time first exceeds mean job size by 20%, $\rho_H(0.2)$, is approximately 0.65. In contrast, under SRPT scheduling $\rho_H(0.2)$ is approximately 0.9. In this scenario, we see that SRPT scheduling allows much better server utilization while maintaining the same mean response time.

Although we empirically observe that SRPT scheduling yields relatively low mean response time under medium load in the M/G/k, we have no theoretical justification for this behavior. In particular, it remains to be shown whether this behavior persists across all job size distributions and any amount of servers.

Conclusion

We propose an important open problem: analyzing the mean response time in the M/G/k/SRPT under medium load. Although mean response time in the M/G/k/SRPT has recently been tightly bounded under heavy traffic, little to nothing is known under medium load. The excellent empirical mean response time of the M/G/k/SRPT under medium load motivates its further study.

References

Bunt RB (1976) Scheduling techniques for operating systems. *Comput.* 9(10):10–17.

Daley DJ (1997) Some results for the mean waiting-time and workload in GI/GI/K queues. Dshalalow JH, ed. *Frontiers in Queueing: Models and Applications in Science and Engineering* (CRC Press, New York), 35–59.

Grosof I, Scully Z, Harchol-Balter M (2018) SRPT for multiserver systems. *Performance Evaluation* 127–128:154–175.

Guirguis S, Sharaf MA, Chrysanthis PK, Labrinidis A, Pruhs K (2009) Adaptive scheduling of web transactions. Ioannidis Y, Lee D, Ng R, eds. *Data Engrg., 2009. ICDE'09. IEEE 25th Internat. Conf.* (IEEE, Pisacataway, NJ), 357–368.

Gupta V, Osogami T (2011) On Markov–Krein characterization of the mean waiting time in M/G/K and other queueing systems. *Queueing Systems* (68)3:339–352.

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

Harchol-Balter M, Schroeder B, Bansal N, Agrawal M (2003) Size-based scheduling to improve web performance. *ACM Trans. Comput. Systems* 21(2):207–233.

Kingman JFC (1970) Inequalities in the theory of queues. *J. Roy. Statist. Soc. Ser. B. Methodological* 32(1):102–110.

Montazeri B, Li Y, Alizadeh M, Ousterhout J (2018) Homa: A receiver-driven low-latency transport protocol using network priorities. *Proc. 2018 Conf. ACM Special Interest Group Data Comm.* (ACM, New York), 221–235.

Schrage L (1968) Letter to the editor—A proof of the optimality of the shortest remaining processing time discipline. *Oper. Res.* 16(3):687–690.

Schrage LE, Miller LW (1966) The queue M/G/1 with the shortest remaining processing time discipline. *Oper. Res.* 14(4):670–684.