State-Dependent Response Times via Fluid Limits in Shortest Remaining Processing Time Queues

Douglas G. Down
Department of Computing and Software
McMaster University
downd@mcmaster.ca

H. Christian Gromoll
Department of Mathematics
University of Virginia
gromoll@virginia.edu

Amber L. Puha
Department of Mathematics
California State University, San Marcos
apuha@csusm.edu

ABSTRACT

We consider a single server queue with renewal arrivals and i.i.d. service times, in which the server employs the Shortest Remaining Processing Time (SRPT) policy. We provide a fluid model (or formal law of large numbers approximation) for this system. The foremost payoff of our fluid model is a fluid level approximation for the state-dependent response time of a job of arbitrary size, that is, the amount of time it spends in the system, given an arbitrary system configuration at the time of its arrival.

1. INTRODUCTION

Interest in the SRPT policy stretches back to the first optimality result due to Schrage [4], who showed that SRPT minimizes the number of jobs in the system at any point in time. Expressions for the mean response time for a single server M/G/1/SRPT queue were earlier developed by Schrage and Miller [5]. This expression depends on the entire service time distribution through nested integrals and thus is somewhat difficult to work with, particularly if one wishes to make comparisons with other policies. Recently, there has been renewed interest in the SRPT policy. To cite but one work, Bansal and Harchol-Balter [1] are interested in the issue of fairness for SRPT. There has also been a recent body of work on the tail behavior of single server queues under SRPT; see for example Nuyens et al. [3]. They discuss the advisability of implementing SRPT using large deviations techniques. In particular, they show that SRPT is effective in the heavy-tailed service time setting, but may not be effective in the light-tailed service time setting.

In this paper, we discuss fluid limits (functional law of large numbers approximations) for single server SRPT queues, as well as implications for analyzing state-dependent response times. That is, we obtain from our fluid model a fluid level approximation of the amount of time a job of a given size spends in the system, given an arbitrary system configuration at the time of its arrival. This paper is an extended abstract of results in [2]. In [2], a more thorough introduction is provided, as well as all proofs and a detailed analysis of the fluid model behavior. There is also a rigorous justification of the fluid model as an approximation to the underlying stochastic model. Here we briefly summarize the stochastic and fluid models and develop some examples that highlight the implications of the results in [2].

2. STOCHASTIC MODEL

Consider a single server queue operating under the SRPT scheduling policy. The SRPT scheduling policy gives preemptive priority to the job in the system with the shortest remaining processing time. Note that to implement this policy, it is assumed that the service times of jobs are known upon arrival. We assume that the arrival process is a delayed renewal process, with rate α. Service times are independent and identically distributed, with distribution ν. We assume that ν is continuous with unbounded support and has mean $1/\alpha$. In particular, we are assuming that the queue is critically loaded. In [2], we significantly relax these conditions, namely we examine queues with arbitrary loads, as well as allowing for ν to have atoms and bounded support.

3. FLUID MODEL

Let $\mathbf{M}$ denote the set of finite nonnegative Borel measures on $\mathbb{R}_+$. For $\xi \in \mathbf{M}$ and a Borel measurable function g on $\mathbb{R}_+$, we define $\langle g, \xi \rangle = \int_{\mathbb{R}_+} g(x)\xi(dx)$, when the integral exists. Also, let $\chi(x) = x$, for $x \in \mathbb{R}_+$. Given a measure-valued function ζ of time taking values in $\mathbf{M}$, for $t \geq 0$ let

$$l(t) = \sup\{x \in \mathbb{R}_+ : \langle 1_{[0,x)}, \zeta(t)\rangle = 0\},$$

which is the infimum of the support of $\zeta(t)$. We refer to $l(\cdot)$ as the *left edge* of the measure-valued function $\zeta(\cdot)$.

We now define our fluid model. To do so, we need the fluid analog, ξ, of the initial condition. Note that the initial condition must contain the remaining service time for each job initially in the system. The fluid analog ξ is a an element of $\mathbf{M}$ that describes the initial mass associated with jobs that have a particular remaining service time. Here for simplicity we assume that ξ is absolutely continuous with respect to Lebesgue measure and also that $\langle \chi, \xi \rangle < \infty$.

DEFINITION 3.1. *A measure-valued function ζ is a critical fluid model solution if each of the following hold:*

(C1) $\zeta(\cdot)$ is right continuous;

(C2) for all $t \in [0, \infty)$, $\langle \chi, \zeta(t)\rangle = \langle \chi, \xi \rangle$;

(C3) for all $t \in [0, \infty)$ and for all continuous, nonnegative, bounded, real functions g, $\langle g, \zeta(t)\rangle = \langle g 1_{[l(t),\infty)}, \xi + \alpha t \nu\rangle$.

Condition (C1) reflects the right continuous nature of the the state descriptor for the stochastic model. Condition (C2) reflects the work conserving nature of SRPT. Condition (C3) is specific to SRPT. It implies that, for each $t \in [0, \infty)$, $\zeta(t)$ has no support below $l(t)$ and agrees with the measure $\xi + \alpha t \nu$ at and above $l(t)$. If we intuitively regard the fluid model as a deterministic system that receives αt units of mass during each time interval $(0, t]$ which, as it arrives, is instantaneously distributed over $\mathbb{R}_+$ according to the distribution ν, and is processed according to the SRPT discipline, then (C3) can be interpreted as follows. Mass arriving below level $l(t)$ at time t is instantaneously flushed out of the system, while mass arriving above $l(t)$ by time t receives no processing by time t. Hence, the mass that is at $l(t)$ at time t is being processed at time t. This reflects the fact that in an SRPT queue, jobs with the shortest remaining processing time are served first.

For $x \in \mathbb{R}_+$, let

$$s(x) = \frac{\langle \chi 1_{[0,x]}, \xi \rangle}{1 - \alpha \langle \chi 1_{[0,x]}, \nu \rangle},$$

and let $s_r^{-1}(\cdot)$ be the right continuous inverse of $s(\cdot)$. The following result characterizes the left edge dynamics of fluid model solutions.

THEOREM 3.2. *If ζ is a fluid model solution, then $l(t) = s_r^{-1}(t)$ for all $t \in [0, \infty)$.*

In light of Theorem 3.2, $s(x)$ can be viewed as the fluid analog of the waiting time for a job of size x that is in the system at time 0. Furthermore, since service times become negligible on fluid scale, the fluid analog of the waiting time is synonymous with the fluid analog of the response time. Therefore, $s(\cdot)$ can also be viewed as the fluid analog of the response time.

4. EXAMPLES AND DISCUSSION

EXAMPLE 4.1. *Suppose that ν is an exponential distribution with rate α. Then, for $x \in \mathbb{R}_+$,*

$$s(x) = \langle \chi 1_{[0,x]}, \xi \rangle \frac{e^{\alpha x}}{\alpha x + 1}.$$

EXAMPLE 4.2. *Suppose that ν has Pareto density $f(x) = (k+1)b^{k+1}x^{-(k+2)}$ for $x \geq b$, where $b, k > 0$ are such that $k/(k+1)b = \alpha$. Then, for $x \geq b$,*

$$s(x) = \langle \chi 1_{[0,x]}, \xi \rangle \left(\frac{x}{b}\right)^k.$$

EXAMPLE 4.3. *Suppose that ν has Weibull density $f(x) = a\lambda^a x^{a-1} e^{-(\lambda x)^a}$ for $x \geq 0$, where $\lambda, a > 0$ are such that $\lambda/\Gamma(1 + \frac{1}{a}) = \alpha$. Then, for $x \in \mathbb{R}_+$,*

$$s(x) = \langle \chi 1_{[0,x]}, \xi \rangle \frac{\Gamma(1 + \frac{1}{a})}{\lambda} e^{(\lambda x)^a}.$$

Finally, we give an example where $s(x)$ cannot be given explicitly, but must be numerically computed.

EXAMPLE 4.4. *Suppose that ν has lognormal density $f(x) = (1/\sqrt{2\pi}\sigma x)e^{-(\ln(x)-\mu)^2/(2\sigma^2)}$ for $x \geq 0$, where $\mu, \sigma > 0$ are such that $e^{-(\mu+\sigma^2/2)} = \alpha$. Then for $x \in \mathbb{R}_+$,*

$$s(x) = \frac{\langle \chi 1_{[0,x]}, \xi \rangle}{1 - e^{-(\mu+\sigma^2/2)} \int_0^x \frac{1}{\sqrt{2\pi}\sigma} e^{-(\ln(u)-\mu)^2/(2\sigma^2)} du}.$$

Note that when the service time distributions have an exponentially decaying tail, the behavior, as demonstrated by the approximate response time, $s(x)$, is quite different from

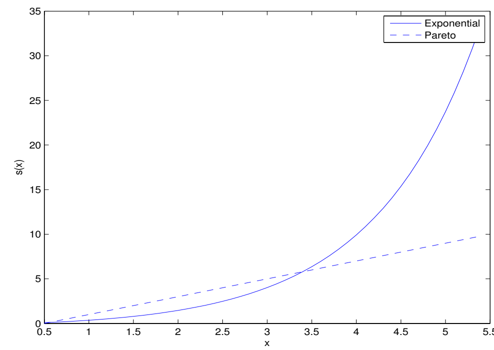

Figure 1: $s(x)$ for exponential and Pareto distributions

the case when the service time distributions have subexponential tails. In the first case, we see that $s(x)$ grows exponentially in the job size, x, while for the Pareto distribution the growth is polynomial. This is consistent with the observation in Nuyens et al. [3] that SRPT (as a member of the class of policies they call SMART policies) may exhibit poor performance for large jobs in a light tailed service setting, while it performs much better for large jobs in the heavy tailed setting. Suppose that we let $\alpha = 1$ and take for the initial measure $\xi = \nu$. For the Pareto distribution, set $k = 1$ and $b = 1/2$. The fact that large jobs are not treated well under the exponential distribution is shown clearly in Figure 1. For the Pareto distribution, the fluid analog of the response time is proportional to the job size.

Given an initial job distribution, it can be seen from the examples that it is easy to calculate a fluid model that describes the dynamics of the system for all job sizes. In other words, conditional on the initial job distribution, we provide detailed dynamics. As such, our work can be seen as a complement to much of the existing work on SRPT, which is concerned with steady-state behavior or tail asymptotics of the waiting time distribution.

5. REFERENCES

[1] N. Bansal and M. Harchol-Balter. Analysis of SRPT scheduling: investigating unfairness. *SIGMETRICS 2001.*

[2] D.G. Down, H.C. Gromoll and A.L. Puha. Fluid limits for shortest remaining processing time queues. *Mathematics of Operations Research*, forthcoming.

[3] M. Nuyens, A. Wierman and B. Zwart. Preventing large sojourn times using SMART scheduling. *Operations Research*, 56:88-101, 2008.

[4] L. Schrage. A proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 16:687-690, 1968.

[5] L.E. Schrage and L.W. Miller. The queue M/G/1 with the shortest remaining processing time discipline. *Operations Research*, 14:670-684, 1966.