Heavy-traffic Analysis of Mean Response Time Under Shortest Remaining Processing Time*

Minghong Lin, Adam Wierman, and Bert Zwart †

DRAFT

Abstract

SRPT has long been known to optimize the queue length distribution and the mean response time (a.k.a. sojourn time). As such, it has been the focus of a wide body of analysis. However, results about the heavy-traffic behavior of SRPT have only recently started to emerge and hold only under a few specific workloads. In this work, we tightly characterize the growth rate of the mean response time under SRPT in the M/GI/1 system under general job size distributions. Our results illustrate the relationship between the job size tail and the heavy traffic growth rate of mean response time. Further, we show that the heavy traffic growth rate can be used to provide an accurate approximation for mean response time outside of heavy traffic.

Key words: queueing; scheduling; SRPT; heavy traffic; response time; sojourn time.

1 Introduction

Shortest Remaining Processing Time (SRPT) has long been known to optimize the mean response time (a.k.a. flow time, sojourn time) in a single server model [26]. As a result, there has been extensive research studying SRPT in a wide variety of models over the last 50 years. For example, SRPT has been studied in the M/GI/1 model [25], with MAP arrivals [20], with setup times [12], and in a variety of other settings [21, 27, 19]. Further, there has been renewed interest in SRPT recently as a result of a number of computer system designs based on SRPT-like policies, e.g., web servers [13, 24], routers [22, 23], wireless networks [14], and beyond. This renewed interest has led

*This work was supported by grants from NSF CCF 0830511, Microsoft Research, and the Lee Center for Advanced Networking.

†M. Lin is a doctoral student at California Institute of Technology, M/C 136-93, Pasadena, CA 91125, mhlin@caltech.edu. A. Wierman is an assistant professor with the Department of Computer Science at California Institute of Technology, M/C 136-93, Pasadena, CA 91125, adamw@caltech.edu. B. Zwart has a joint affiliation at CWI Amsterdam, VU University Amsterdam, EURANDOM, and Georgia Tech, bertz@cwi.nl.

to new results studying the tail behavior [17, 18, 7], the heavy-traffic behavior [3, 2, 11], and the fairness of SRPT [29, 30].

However, despite this large literature, there are some simple properties of SRPT that are still not well understood. One such property is the focus of this paper: *How does the mean response time of SRPT scale with load?*

It is perhaps surprising that this question is still not understood given its fundamental nature, especially since the the mean response time, $E[T]$, was derived for the first time by Schrage & Miller [25] in 1966. However, the formula for the mean response time is complicated enough that the dependence of it on load, ρ, is not well understood. Specifically, there are a few papers that have derived the heavy-traffic growth rate of SRPT under specific job size distributions, i.e., under Exponential job sizes [2] and under Pareto job sizes [3]. Further, for general job sizes, Down et al. derive a process-level heavy-traffic limit for the conditional response time of a tagged job [11]. However, no results are known about the growth rate of $E[T]$ under general job sizes.

The contribution of this paper is to derive the growth rate of $E[T]$ as a function of ρ in heavy-traffic, i.e., as $\rho \to 1$ (see Theorems 1 and 2). This characterization of the growth rate highlights the relationship between the growth rate and the tail of the job size distribution. Specifically, the heavy-traffic growth rate is shown to depend on the tail of a measure $G(x)$, which characterizes the truncated load, and the Matuszewska index of the job size distribution, which relates to the moment conditions of the job sizes. The results illustrate SRPT provides an order of magnitude of improvement over other common scheduling policies such as Processor Sharing (PS) and First Come First Served (FCFS) if and only if the job size distribution is unbounded. Further, the results illustrate that a heavier-tail implies a slower growth rate. Additionally, once the tail is "heavy-enough" (i.e. has a large enough Matuszewska index) the growth rate becomes (up to a constant) independent of the job size distribution.

In addition to the insight provided by the heavy-traffic growth rate of SRPT, we illustrate that the heavy-traffic analysis can be used to provide a simple approximation of $E[T]$ SRPT, which is accurate even outside of heavy-traffic. This simple approximation is very useful when analyzing more complex models which have pieces that use SRPT. For example, this approximation has already been applied to attain results for a multi-queue load balancing model [8] and a power management model [1]. In each case, analyzing the system using the exact form of $E[T]$ under SRPT would have resulted in only numeric results, but using the heavy-traffic approximation led to analytic results providing new insights into the models.

Another important application of the results in this paper is to allow the competitive analysis of scheduling policies in the M/GI/1 model. Competitive analysis depends on tight bounds on the performance of the optimal algorithm. Our results provide simple, tight analytic forms for the optimal algorithm, which can be used as a baseline for comparison of other algorithms. We will discuss this further in the concluding remarks (Section 5).

The remainder of the paper is organized as follows. In section 2, we introduce our notation and provide some background on Matuszewska index. In section 3, we summarize our main results and relate them to prior work. Additionally, we illustrate the application of our main results to two specific distributions: the Pareto and the Weibull in order to highlight the usefulness of the heavy-traffic results as an approximation outside of heavy-traffic. In section 4 we show the details of the proofs. Finally, section 5 provides some concluding remarks.

2 Preliminaries

In this paper we study the performance of SRPT in an M/GI/1 queue in heavy-traffic. We assume the c.d.f. of job sizes, $F(x)$, is continuous. Denote by $E[T]$ the mean response time (a.k.a. sojourn time) under SRPT, which is the time from when a job enters the system until it completes service. Let $\bar{F}(x) = 1 - F(x)$, λ denote the arrival rate, and $\rho = \lambda E[X]$ be the load. Define $\rho(x) = \lambda \int_0^x t dF(t)$ and $m_2(x) = \int_0^x t^2 dF(t)$. Here $\rho(x)$ can be interpreted as the load made up by jobs with size less than x (ignoring all job with size greater than x), $m_2(x)$ is the second moment of the jobs with size $< x$.

The conditional mean response time for a job of size x, $E[T(x)]$, under $M/GI/1/SRPT$ was first derived by Schrage & Miller [25] and is equal to

$$E[T(x)] \quad = \quad \int_0^x \frac{dt}{1-\rho(t)} + \frac{\lambda x^2 \bar{F}(x)}{2(1-\rho(x))^2} + \frac{\lambda m_2(x)}{2(1-\rho(x))^2} \tag{1}$$

with $E[T] = E[E[T(X)]]$.

Despite the existence of this result, due to its complexity, understanding the behavior of $E[T]$ under SRPT is difficult. For example, it is hard to determine the impact of job size variability and load on this formula. Further, calculating $E[T]$ numerically is non-trivial. The goal of this paper is to provide insight into the behavior of $E[T]$ by studying SRPT in heavy traffic.

Our results illustrated that $E[T]$ under SRPT in heavy traffic depends on a measure $G(x) = \rho(x)/\rho = \int_0^x tf(t)dt/E[X]$ and the Matuszewska index [15] of $F(x)$.

Definition 1. *Let $f(\cdot)$ be positive,*

- *Its* **upper Matuszewska index** $\alpha(f)$ *is the infimum of those α for which there exists a constant $C = C(\alpha)$ such that for each $\Lambda > 1$,*

$$f(\lambda x)/f(x) \le C\{1 + o(1)\}\lambda^{\alpha} \quad (x \to \infty) \text{ uniformly in } \lambda \in [1, \Lambda];$$

- *Its* **lower Matuszewska index** $\beta(f)$ *is the supremum of those β for which, for some $D = D(\beta) > 0$ and all $\Lambda > 1$,*

$$f(\lambda x)/f(x) \ge D\{1 + o(1)\}\lambda^{\beta} \quad (x \to \infty) \text{ uniformly in } \lambda \in [1, \Lambda];$$

Intuitively, the Matuszewska index is a measure to bound the function $f(x)$ by functions of the form $g(x) = Cx^\mu$. A function $f(x)$ with upper Matuszewska index α and lower Matuszewska index β means that $f(x)$ lies roughly between $C_1 x^\alpha$ and $C_2 x^\beta$ as $x \to \infty$. A few properties of the Matuszewska index are listed in the Appendix. Some important results about the Matuszewska index that we use in the proofs can be found in the appendix. For more details about the Matuszewska index, please refer to [6].

Our main results are described in asymptotic notation. We say $f(x) = O(g(x))$ as $x \to a$, if and only if $\lim_{x\to a} |f(x)/g(x)| < \infty$. Similarly, $f(x) = o(g(x))$ denotes $\lim_{x\to a} |f(x)/g(x)| = 0$, and $f(x) = \Theta(g(x))$ denotes $0 < \lim_{x\to a} |f(x)/g(x)| < \infty$. To denote more accurate asymptotics, we write $f(x) \sim g(x)$ iff $\lim_{x\to a} |f(x)/g(x)| = 1$, and $f(x) \preceq g(x)$ iff $\lim_{x\to a} |f(x)/g(x)| \leq 1$. Other notations $\succeq$, $\prec$ and $\succ$ are defined in a similar way. In this paper, this notation is always used to describe asymptotic behavior as $\rho \to 1$.

3 Results and Discussion

In this section, we summarize and discuss our main results. Further, we illustrate the results numerically using two common job size distributions. The proofs of the results are provided in Section 4.

Our first result is a loose characterization of the growth rate of SRPT, which illustrates when the mean response time of SRPT is an order of magnitude better than the growth rate of PS and FCFS. Recall that, in an $M/G/1$ system, the mean response time for PS (FCFS) scales with $\frac{1}{1-\rho}$, i.e., $E[T] = \Theta(\frac{1}{1-\rho})$, regardless of the job size distribution as long as it has finite first (second) moment. Theorem 1 shows that the growth rate of $E[T]$ for SRPT is strictly slower than that of PS and FCFS as $\rho \to 1$ if and only if the job size distribution is unbounded.

Theorem 1. *In an $M/GI/1$ SRPT queue:*

(i) If $F(x)$ has bounded support ($\exists x_m < \infty$ s.t. $F(x_m) = 1$), then

$$\frac{E[X^2]}{2x_m}\frac{1}{1-\rho} \preceq E[T] = \Theta\left(\frac{1}{1-\rho}\right) \quad \textit{as } \rho \to 1.$$

(ii) If $F(x)$ has unbounded support, then

$$E[T] = o\left(\frac{1}{1-\rho}\right) \quad \textit{as } \rho \to 1.$$

Theorem 1 shows that not only does SRPT minimize $E[T]$, it also provides an order of magnitude improvement in many cases. Our next result characterizes the growth rate of SRPT explicitly. It illustrates that, not only does $E[T]$ scale more slowly than PS and FCFS when the job size distribution is unbounded, but the improvement can be exponential.

Theorem 2. *In an $M/GI/1$ SRPT queue:*

(i) If the upper Matuszewska index $\alpha(\bar{F}) < -2$, then $E[X^2] < \infty$, and

$$\frac{1}{4}\frac{E[X^2]}{(1-\rho)G^{-1}(\rho)} \preceq E[T] = \Theta\left(\frac{E[X^2]}{(1-\rho)G^{-1}(\rho)}\right) \quad as\ \rho \to 1.$$

(ii) If the lower Matuszewska index $\beta(\bar{F}) > -2$, then $E[X^2] = \infty$, and

$$\frac{E[X]}{2}\log\left(\frac{1}{1-\rho}\right) < E[T] = \Theta\left(E[X]\log\left(\frac{1}{1-\rho}\right)\right) \quad as\ \rho \to 1.$$

This result generalizes recent results from Bansal [4] and Bansal & Gamarnik [5] who each analyze special cases. Bansal [4] derived the growth rate of $E[T]$ under SRPT in $M/M/1$. Bansal & Gamarnik [5] derived the growth rate of $E[T]$ under SRPT with Pareto Job size distribution. We will discuss these special cases later when we illustrate the theorem for a few examples.

Theorem 2 illustrates the precise impact of the job size distribution on the growth rate of $E[T]$. It turns out that the growth rate is determined by the measure $G(\cdot)$ and the Matuszewska index. Both of which are related to the tail of the job size distribution. It shows that $E[T]$ for distributions with heavier tail increases slower as $\rho \to 1$. And, once the tail is heavy enough, the growth rate becomes $\Theta\left(\log\left(\frac{1}{1-\rho}\right)\right)$, which is "independent" of the job size distribution (though the constant may depend on the distribution). Additionally, the lower bounds in Theorem 2 holds for general distributions with finite (infinite) second moment accordingly. Please refer to Equation (4) for details. However, the upper bounds are more complicated since the exact constant for $E[T]$ can be arbitrary large and depends on the distribution. Thus, one cannot hope to provide a general upper bound constant without using explicit information about the job size distribution.

Examples

To get more intuitive understanding of the theorems, we apply them in the case of a few common job size distributions. This allows us to rederive many of the known results in the literature as well as derive the growth rate in some novel cases. Further, we will illustrate, using numeric examples, the accuracy of the heavy traffic growth rate when it is used as an approximation outside of heavy traffic.

The two distributions we use as examples are the Pareto and Weibull distributions. The Pareto distribution is probably the most popular heavy-tailed distribution and is widely used as a model for the tails of other distributions. The Weibull distribution is often used in the field of life data analysis due to its flexibility. It can mimic the behavior of other statistical distributions such as the normal, the exponential and other distributions with increasing or decreasing failure rate. Note that the heavy traffic behavior of $E[T]$ under SRPT for Pareto distribution has been studied in [5], while

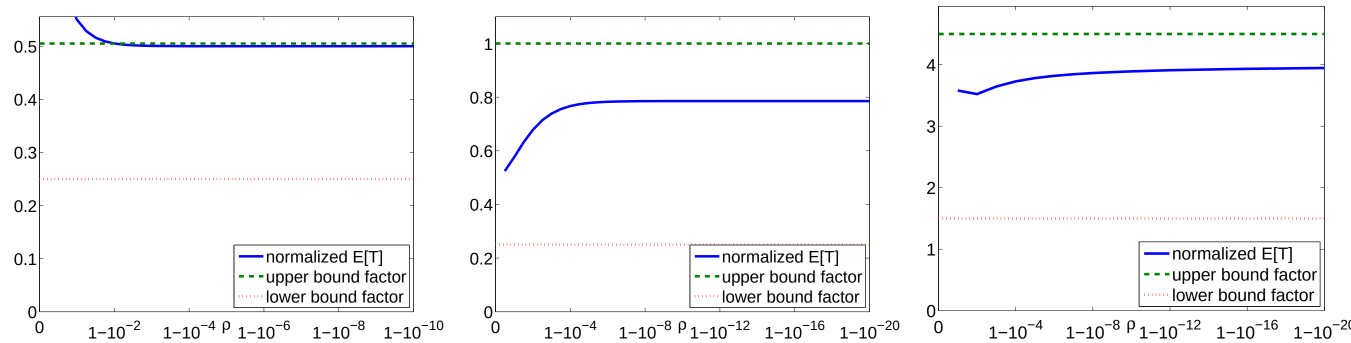

(a) $\alpha = 100$ (b) $\alpha = 3$ (c) $\alpha = 1.5$

Figure 1: Normalized $E[T]$ and the bounds for Pareto as $\rho \to 1$

the heavy traffic behavior for Weibull distribution is novel.

Example 1: Pareto Distribution

Our first example is the case of Pareto job sizes.

Corollary 1. *For Pareto Distribution with* $\bar{F}(x) = \left(\frac{x}{x_m}\right)^{-\alpha}$, *as* $\rho \to 1$,

$$\frac{E[X]}{2}\log\left(\frac{1}{1-\rho}\right) \preceq E[T] \preceq \left(\frac{(3-\alpha)E[X]}{4-2\alpha}\right)\log\left(\frac{1}{1-\rho}\right) \qquad (1<\alpha<2)$$

$$\frac{1}{4x_m}\frac{E[X^2]}{(1-\rho)^{\frac{\alpha-2}{\alpha-1}}} \preceq E[T] \preceq \frac{\alpha-1}{2(\alpha-2)x_m}\frac{E[X^2]}{(1-\rho)^{\frac{\alpha-2}{\alpha-1}}} \qquad (\alpha>2)$$

This result is a generalization of earlier work by Bansal & Gamarnik [5], who show that $E[T]$ under SRPT with Pareto Distribution has growth rate of $O((1-\rho)^{-(\alpha-2)/(\alpha-1)})(\alpha > 2)$ and $O(\log(1/(1-\rho)))(1<\alpha<2)$ by proving that Preemptive Shortest Job First (PSJF) has this upper bound and $E[T]$ under SRPT is less than $E[T]$ under PSJF. Pechinkin provides the exact constant which is very complicated for $E[T]$ under SRPT with Pareto Distribution as $\rho \to 1$.

A numeric illustration of $E[T]$ and our upper/lower bound for Pareto distributions with various parameter ($\alpha = 100, 3, 1.5$) is shown in Figure 1 and Figure 2. Figure 1 shows that the exact constant of $E[T]$(normalized by the $\Theta(\cdot)$ formula) is not very far away from the upper bound constant as $\rho \to 1$, especially when α is large. Figure 2 shows the accuracy of using formula $C_1 * growth\ rate + C_2$ as an approximation of $E[T]$ for smaller ρ (The constants C_1, C_2 are chosen to get the best fitting). The percentage error of this approximation is shown. We can see that the percentage errors are less than 6% for $\rho > 0.7$ for all three cases, i.e., this formula is a very good approximation, especially when α is large.

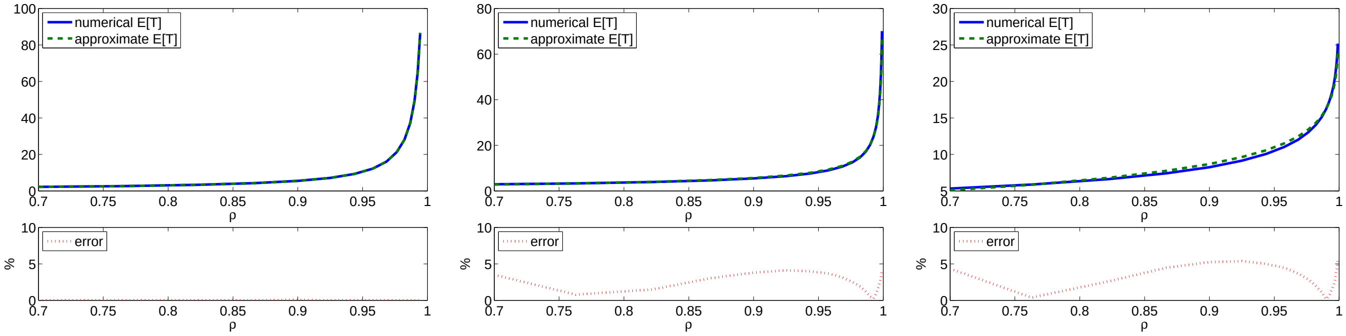

(a) $\alpha = 100$ (b) $\alpha = 3$ (c) $\alpha = 1.5$

Figure 2: The accuracy of using $C_1 * growth\ rate + C_2$ as an approximation of $E[T]$ for Pareto

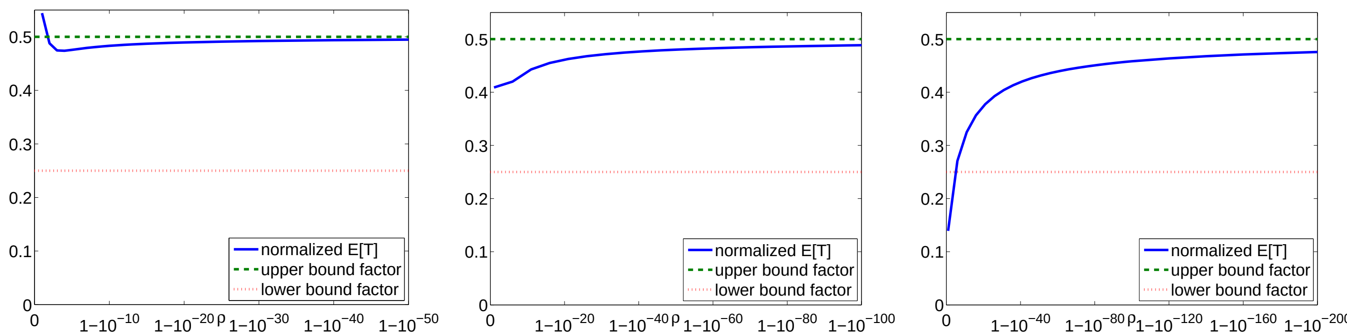

(a) $\alpha = 2$ (b) $\alpha = 1$ (c) $\alpha = 0.5$

Figure 3: Normalized $E[T]$ and the bounds for Weibull as $\rho \to 1$

Example 2: Weibull Distribution

We now move to the case of Weibull job sizes.

Corollary 2. *For Weibull Distribution with $\bar{F}(x) = e^{-\mu x^{\alpha}}$ $(\alpha > 0)$, as $\rho \to 1$, we have*

$$\frac{E[X^2]}{4} \frac{1}{(1-\rho) \cdot \mu^{-1/\alpha} \log(\frac{1}{1-\rho})^{1/\alpha}} \preceq E[T] \preceq \frac{E[X^2]}{2} \frac{1}{(1-\rho) \cdot \mu^{-1/\alpha} \log(\frac{1}{1-\rho})^{1/\alpha}}.$$

Note that $M/M/1$ is a special case of Weibull Distribution by setting $\alpha = 1$; thus this result generalizes earlier work by Bansal [4], who proves that for $\rho \in [2/3, 1)$ in $M/M/1$, $\frac{1/(18e)}{\mu(1-\rho)\log(1/(1-\rho))} \leq E[T] \leq \frac{7}{\mu(1-\rho)\log(1/(1-\rho))}$.

A numeric illustration of $E[T]$ and our bounds for Weibull distributions with various parameter $(\alpha = 2, 1, 0.5)$ is shown in Figure 3 and Figure 4. Again, the approximation seems very accurate even when ρ is not very large.

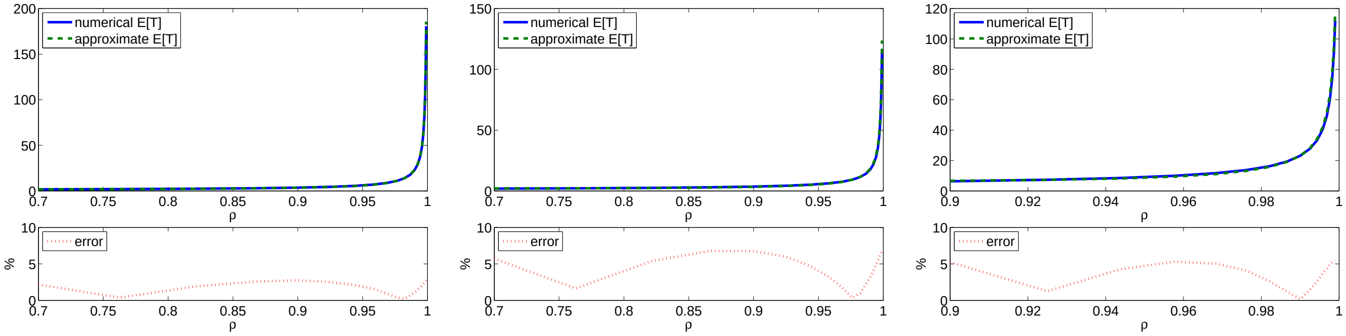

(a) $\alpha = 2$ (b) $\alpha = 1$ (c) $\alpha = 0.5$

Figure 4: The accuracy of using $C_1 * growth\ rate + C_2$ as an approximation of $E[T]$ for Weibull

4 Proofs

We now prove the results described in Section 3. Before the main proof, it is useful to decompose $E[T]$ under SRPT as follows. Denote

$$\begin{aligned} \tilde{R}(x) &= \int_0^x \frac{dt}{1-\rho(t)} + \frac{\lambda x^2 \bar{F}(x)}{2(1-\rho(x))^2} \\ \tilde{W}(x) &= \frac{\lambda m_2(x)}{2(1-\rho(x))^2} \end{aligned}$$

Then $E[T] = E[E[T(x)]] = E[\tilde{R}] + E[\tilde{W}]$. We can derive the bounds for $E[\tilde{R}]$ and $E[\tilde{W}]$ separately, and combine them to get the bounds for $E[T]$.

$$E[\tilde{R}] = \int_0^\infty \tilde{R}(x)dF(x) = \int_0^\infty \frac{\bar{F}(x)}{1-\rho(x)}dx + \frac{\lambda}{2}\int_0^\infty \frac{x^2\bar{F}(x)}{(1-\rho(x))^2}dF(x)$$

Consider

$$\log\frac{1}{1-\rho} = \lambda\int_0^\infty \frac{xdF(x)}{1-\rho(x)} = \lambda\int_0^\infty \frac{\bar{F}(x)}{1-\rho(x)}dx + \lambda^2\int_0^\infty \frac{x^2\bar{F}(x)}{(1-\rho(x))^2}dF(x)$$

Thus

$$\frac{1}{2\lambda}\log\frac{1}{1-\rho} < E[\tilde{R}] < \frac{1}{\lambda}\log\frac{1}{1-\rho}$$

As $\lambda E[X] = \rho \to 1$, we have,

$$\frac{E[X]}{2}\log\frac{1}{1-\rho} < E[\tilde{R}] \preceq E[X]\log\frac{1}{1-\rho}$$

For $E[\tilde{W}]$,

$$E[\tilde{W}] = \int_0^\infty W(x)dF(x) = \int_0^\infty \frac{\lambda m_2(x)}{2(1-\rho(x))^2}dF(x) = \int_0^\infty \frac{m_2(x)}{2(1-\rho(x))^2 x}d\rho(x)$$

We now begin by proving Theorem 1 as follows.

Proof of Theorem 1. If $F(x)$ has bounded support, we have $x \in [0, x_m]$, $m_2(x) \in [0, E[X^2]]$. Pick x_p so that $m_2(x_p) = E[X^2] - \epsilon$. Then

$$\begin{aligned} E[\tilde{W}] &= \frac{1}{2}\int_0^{x_m} \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x}dx \\ &\geq \frac{1}{2}\int_{x_p}^{x_m} \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x_p)}{x_m}dx \\ &= \frac{m_2(x_p)}{2x_m}\frac{1}{1-\rho(x)}|_{x_p}^{x_m} \end{aligned}$$

We can pick x_p to make $m_2(x_p)$ arbitrarily close to $E[X^2]$. Thus as $\rho \to 1$, we have

$$E[\tilde{W}] \succeq \frac{E[X^2]}{2x_m}\frac{1}{1-\rho}$$

On the other hand,

$$\begin{aligned} E[\tilde{W}] &= C_1 + \frac{1}{2}\int_{x_p}^{x_m} \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x}dx \\ &< C_1 + \frac{E[X^2]}{2x_p}\frac{1}{1-\rho(x)}|_{x_p}^{x_m} \\ &\sim \frac{E[X^2]}{2x_p}\frac{1}{1-\rho} \qquad as\ \rho \to 1 \end{aligned}$$

Thus we get $E[\tilde{W}] = \Theta\left(\frac{1}{1-\rho}\right)$, so $E[T] = E[\tilde{R}] + E[\tilde{W}] = \Theta\left(\frac{1}{1-\rho}\right)$ for $F(x)$ with bounded support. If $F(x)$ has unbounded support, for any given $\epsilon > 0$, we can find an x_0 so that $\bar{G}(x_0) < \epsilon$.

$$\lambda m_2(x) = x\rho(x) - \int_0^x \rho(t)dt = \int_0^x (\rho - \rho(t))dt - x(\rho - \rho(x)) \leq \rho \int_0^x \bar{G}(t)dt$$

Thus

$$\begin{aligned} \lim_{x\to\infty} \frac{m_2(x)}{x} &\leq \lim_{x\to\infty} \frac{E[X]\int_0^x \bar{G}(t)dt}{x} \\ &= \lim_{x\to\infty} \frac{C_1 + E[X]\int_{x_0}^x \bar{G}(t)dt}{x} \\ &\leq \lim_{x\to\infty} \frac{C_1 + E[X]\epsilon(x - x_0)}{x} \\ &\leq E[X]\epsilon \end{aligned}$$

By definition, we can find a x_p so that for all $x > x_p$, $m_2(x)/x \leq E[X]\epsilon + \epsilon$.

$$
\begin{aligned}
E[\tilde{W}] &= \frac{1}{2}\int_0^\infty \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x}dx \\
&= C_2 + \frac{1}{2}\int_{x_p}^\infty \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x}dx \\
&\leq C_2 + \frac{E[X]\epsilon+\epsilon}{2}\int_{x_p}^\infty \frac{\rho'(x)}{(1-\rho(x))^2}dx \\
&\sim \frac{E[X]\epsilon+\epsilon}{2}\frac{1}{1-\rho}
\end{aligned}
$$

in which ϵ could be arbitrary small. Thus $E[\tilde{W}] = o\left(\frac{1}{1-\rho}\right)$ for $F(x)$ with unbounded support. $\square$

To prove Theorem 2, we will prove a series of lemmas.

Lemma 1. *For $M/GI/1$ under SRPT, if the job size distribution $F(x)$ has $E[X^2] < \infty$, then*

$$\frac{E[X^2]}{4(1-\rho)G^{-1}(\rho)} \preceq E[\tilde{W}] \preceq \frac{E[X^2]}{2}\int_0^\rho \frac{1}{(1-y)^2G^{-1}(y)}dy \quad \text{as } \rho \to 1.$$

Proof. Begin by calculating as follows

$$
\begin{aligned}
E[\tilde{W}] &= \frac{1}{2}\int_0^\infty \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x}dx \\
&> \frac{1}{2}\int_{x_0}^\infty \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x}dx \\
&> \frac{m_2(x_0)}{2}\int_{x_0}^\infty \frac{\rho'(x)}{(1-\rho(x))^2x}dx
\end{aligned}
$$

Denote $y = \rho(x)$, we have

$$
\begin{aligned}
E[\tilde{W}] &> \frac{m_2(x_0)}{2}\int_{\rho_0}^\rho \frac{1}{(1-y)^2\rho^{-1}(y)}dy \\
&\geq \frac{m_2(x_0)}{2}\int_{\rho_0}^{\rho^2} \frac{1}{(1-y)^2\rho^{-1}(y)}dy
\end{aligned}
$$

Remember that $G(x) = \rho(x)/\rho$, which implies $\rho^{-1}(z) = G^{-1}(z/\rho)$.

$$\begin{aligned} E[\tilde{W}] &> \frac{m_2(x_0)}{2}\int_{\rho_0}^{\rho^2} \frac{1}{(1-y)^2 G^{-1}(\frac{y}{\rho})} dy \\ &> \frac{m_2(x_0)}{2G^{-1}(\rho)}\int_{\rho_0}^{\rho^2} \frac{1}{(1-y)^2} dy \\ &= \frac{m_2(x_0)}{2G^{-1}(\rho)}\left(\frac{1}{1+\rho}\cdot\frac{1}{1-\rho} - \frac{1}{1-\rho_0}\right) \end{aligned}$$

We can pick a x_0 large enough to make $m_2(x_0)$ arbitrarily close to $E[X^2]$. Thus we have

$$E[\tilde{W}] \succeq \frac{E[X^2]}{4(1-\rho)G^{-1}(\rho)}$$

For the other inequality:

$$\begin{aligned} E[\tilde{W}] &= \frac{1}{2}\int_0^\infty \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x} dx \\ &< \frac{E[X^2]}{2}\int_0^\infty \frac{\rho'(x)}{(1-\rho(x))^2 x} dx \end{aligned}$$

Denote $y = \rho(x)$, we have

$$\begin{aligned} E[\tilde{W}] &< \frac{E[X^2]}{2}\int_0^\rho \frac{1}{(1-y)^2\rho^{-1}(y)} dy \\ &\leq \frac{E[X^2]}{2}\int_0^\rho \frac{1}{(1-y)^2\rho^{-1}(\rho y)} dy \end{aligned}$$

Remember that $\rho^{-1}(z) = G^{-1}(z/\rho)$:

$$E[\tilde{W}] < \frac{E[X^2]}{2}\int_0^\rho \frac{1}{(1-y)^2 G^{-1}(y)} dy$$

$\square$

Lemma 2. *$\alpha(\bar{G}(x)) \leq \alpha(\bar{F}(x)) + 1$; $\beta(\bar{G}(x)) \geq \beta(\bar{F}(x)) + 1$.*

Proof. Assume $\alpha(\bar{F}(x)) = \mu$, by definition we have $\bar{F}(\lambda x)/\bar{F}(x) \leq C\{1+o(1)\}\lambda^\mu$ $(x \to \infty)$ uniformly in $\lambda \in [1, \Lambda]$. Thus

$$\begin{aligned}
\bar{G}(\lambda x) &= \int_{\lambda x}^{\infty} \bar{F}(t)dt + \lambda x \bar{F}(\lambda x) \\
&= \lambda \int_{x}^{\infty} \bar{F}(\lambda u)du + \lambda x \bar{F}(\lambda x) \\
&\leq C\{1+o(1)\}\lambda^{\mu+1} \left(\int_{x}^{\infty} \bar{F}(u)du + x\bar{F}(x) \right) \\
&= C\{1+o(1)\}\lambda^{\mu+1}\bar{G}(x) \qquad as\ x \to \infty
\end{aligned}$$

Thus $\alpha(\bar{G}(x)) \leq \alpha(\bar{F}(x)) + 1$. Similarly we can prove $\beta(\bar{G}(x)) \geq \beta(\bar{F}(x)) + 1$. $\square$

Proof of Theorem 2 for $\alpha(\bar{F}) < -2$. By Lemma 2, $\alpha(\bar{G}(x)) \leq -1$, thus $\alpha(x\bar{G}(x)) \leq 0$. By Lemma 5, $\int_0^\infty \bar{G}(t)dt < \infty$.

$$\lambda m_2(x) = x\rho(x) - \int_0^x \rho(t)dt = \int_0^x (\rho - \rho(t))dt - x(\rho - \rho(x)) \leq \rho \int_0^x \bar{G}(t)dt.$$

Which implies $E[X^2] < \infty$. By Lemma 1, the lower bound is easily followed. For the upper bound, we will prove that $E[\tilde{W}]$ satisfies this bound, and $E[\tilde{R}]$ is dominated by $E[\tilde{W}]$ as $\rho \to 1$.

Assume $\alpha(\bar{G}(x)) = \mu < -1$, by Lemma 3, $\beta(1/\bar{G}(x)) = -\mu > 1$. By Lemma 6, $\alpha(\bar{G}^{-1}(1/x)) = -1/\mu$. By Lemma 3, $\beta(1/\bar{G}^{-1}(1/x)) = 1/\mu$. Thus $\beta(x/\bar{G}^{-1}(1/x)) = 1 + 1/\mu > 0$. By Lemma 4, $\liminf_{z\to\infty} \frac{z/\bar{G}^{-1}(1/z)}{\int_1^z 1/\bar{G}^{-1}(\frac{1}{u})du} > 0$, which implies $\limsup_{z\to\infty} \int_1^z \frac{\bar{G}^{-1}(\frac{1}{z})}{z\bar{G}^{-1}(\frac{1}{u})}du < \infty$. Denote $u = 1/(1-y)$, $z = 1/(1-\rho)$, we get

$$\begin{aligned}
\lim_{\rho\to 1} \int_0^\rho \frac{(1-\rho)G^{-1}(\rho)}{(1-y)^2 G^{-1}(y)}dy &= \lim_{\rho\to 1} \int_1^{\frac{1}{1-\rho}} \frac{(1-\rho)G^{-1}(\rho)}{\bar{G}^{-1}(1/u)}du \\
&= \lim_{z\to\infty} \int_1^z \frac{\bar{G}^{-1}(\frac{1}{z})}{z\bar{G}^{-1}(\frac{1}{u})}du \\
&< \infty
\end{aligned}$$

By Lemma 1, we have $E[\tilde{W}] = O\left(\frac{E[X^2]}{(1-\rho)G^{-1}(\rho)}\right)$.
Remember that $E[\tilde{R}] = \Theta(\log(\frac{1}{1-\rho}))$, thus

$$\lim_{\rho\to 1} \log(\frac{1}{1-\rho}) \Big/ \frac{E[X^2]}{(1-\rho)G^{-1}(\rho)} = \lim_{y\to\infty} \log y \Big/ \frac{E[X^2]y}{\bar{G}^{-1}(1/y)}$$

Denote $g(y) = \log y / \frac{E[X^2]y}{\bar{G}^{-1}(1/y)}$. We have shown that $\beta(x/\bar{G}^{-1}(1/x)) > 0$. And we have $\alpha(g(y)) < 0$ by definition. Based on Lemma 7, $g(y) \leq Cg(X)(y/X)^{\alpha(g(y))+\epsilon}$ for some constant C and X. Thus $\lim_{y\to\infty} g(y) = 0$, which means that $E[\tilde{R}]$ is dominated by $E[\tilde{W}]$. This completes our proof. $\square$

Proof of Theorem 2 for $\beta(\bar{F}) > -2$. The lower bound simply follows by seeing $E[T] > E[\tilde{R}]$. Now let us prove $E[X^2] = \infty$ by contradiction. Suppose $E[X^2] < \infty$, then $E[X^2] = \int_0^\infty t^2 f(t)dt =$

$2\int_0^\infty t\bar{F}(t)dt$, $m_2(x) = \int_0^x t^2 f(t)dt = 2\int_0^x t\bar{F}(t)dt - x^2\bar{F}(x)$. Thus

$$\begin{aligned} E[X^2] - m_2(x) &\geq x^2\bar{F}(x) \\ \Rightarrow \forall x \quad x^2\bar{F}(x) &< E[X^2] \end{aligned}$$

However, we know that $\beta(x^2\bar{F}(x)) > 0$, which means $\lim_{x\to\infty} x^2\bar{F}(x)$ is unbounded and proves that $E[X^2] = \infty$ by contradiction.

Next, we begin to calculate $E[\tilde{W}]$ by starting with $\lambda m_2(x)$.

$$\lambda m_2(x) = x\rho(x) - \int_0^x \rho(t)dt = \int_0^x (\rho - \rho(t))dt - x(\rho - \rho(x))$$

Building on the above calculation, we have

$$\begin{aligned} \frac{\lambda m_2(x)}{x(\rho - \rho(x))} &= \frac{\int_0^x (\rho - \rho(t))dt}{x(\rho - \rho(x))} - 1 \\ &= \frac{\int_0^x \bar{G}(t)dt}{x\bar{G}(x)} - 1 \end{aligned}$$

Since $\beta(\bar{G}(x)) > -1$, we have $\beta(x\bar{G}(x)) > 0$. By Lemma 4, $\liminf_{z\to\infty} \frac{z\bar{G}(z)}{\int_1^z \bar{G}(u)du} > 0$. Thus $\limsup_{x\to\infty} \frac{\lambda m_2(x)}{x(\rho-\rho(x))} < \infty$, Therefore

$$\begin{aligned} E[\tilde{W}] &= \frac{1}{2}\int_0^\infty \frac{\rho'(x)}{(1-\rho(x))^2}\frac{m_2(x)}{x}dx \\ &< \frac{1}{2\lambda}\int_0^\infty \frac{\rho'(x)}{(1-\rho(x))} \cdot \frac{\lambda m_2(x)}{x(\rho - \rho(x))}dx \\ &= \frac{1}{2\lambda} O\left(\log\frac{1}{1-\rho}\right) \end{aligned}$$

Finally, $E[T] = E[\tilde{R}] + E[\tilde{W}] = \Theta\left(E[X]\log\frac{1}{1-\rho}\right)$. □

5 Conclusion

SRPT has long been known to optimize the queue length distribution and the mean response time (a.k.a. sojourn time). As such, it has been the focus of a wide body of analysis. However, results about the heavy-traffic behavior of SRPT have only recently started to emerge and hold only under a few specific workloads. In this work, we tightly characterize the growth rate of $E[T]$ under SRPT in the M/GI/1 system under general job size distributions (Theorem 2). This provides some interesting insight into the behavior of SRPT. For example, SRPT provides an order of magnitude improvement over PS and FCFS if and only if the job size distribution is unbounded (Theorem 1). Further, if the distribution is unbounded the growth rate depends delicately on the tail of the job size distribution through the measure $G(x)$ when the upper Matuszewska index is less than -2. However, if the lower Matuszewska index is greater than -2, then the growth rate of $E[T]$ is independent (up to a constant) of the job size distribution.

To illustrate the heavy traffic results, we used numeric experiments in the cases of the Pareto and Weibul job sizes. These experiments illustrated, surprisingly, that the heavy-traffic growth rates can be used to provide quite accurate approximations of $E[T]$ even outside of heavy-traffic. Thus, the heavy-traffic results provide simple, accurate approximations for use in more complicated models. These results have already led to new approximate analyses of SRPT in load balancing [8] and power management [1] settings.

Finally, the characterization of the growth rate of SRPT is especially important because of the optimality of SRPT. These results provide a baseline with which to compare the performance of other policies. Without these results it has been at difficult to compare the performance of other scheduling disciplines with the optimal $E[T]$. As a result, there are almost no "competitive analysis" results in the M/GI/1 model. Using the results in this paper it is now possible to ask, and hopefully answer, a number of new interesting questions. For example, we have shown that PS and FCFS are constant-competitive for $E[T]$ if and only if the job size distribution is bounded. A similar question can be asked for many other scheduling policies. For example, under which distributions is it possible for a policy that is blind to job sizes to be constant competitive? Some partial results on this question can be found in [28, 16].

References

[1] L. Andrew, M. Lin, and A. Wierman. Robustness, fairness, and efficiency in speed scaling designs, 2009.

[2] N. Bansal. On the average sojourn time under M/M/1/SRPT. *Oper. Res. Letters*, 22(2):195–200, 2005.

[3] N. Bansal and D. Gamarnik. Handling load with less stress. *Queueing Systems*, 54(1):45–54, 2006.

[4] Nikhil Bansal. On the average sojourn time under m/m/1/srpt. *SIGMETRICS Perform. Eval. Rev.*, 31(2):34–35, 2003.

[5] Nikhil Bansal and David Gamarnik. Handling load with less stress. *Queueing Syst. Theory Appl.*, 54(1):45–54, 2006.

[6] N.H. Bingham, C.M Goldie, and J.L. Teugels. *Regular Variation.* Cambridge University Press, 1987.

[7] O. Boxma and B. Zwart. Tails in scheduling. *Perf. Eval. Rev.*, 34(4):13–20, 2007.

[8] Ho-Lin Chen, J. Marden, and A. Wierman. On the impact of heterogeneity and back-end scheduling in load balancing designs, 2009.

[9] L. de Haan and S.J. Resnick. positive increase and bounded increase for non-decreasing functions. *Unpublished appendix to [10]*, 1983.

[10] L. de Haan and S.J. Resnick. stochastic compactness and point processes. *Journal of the. Australian Mathematical Society*, 37:307–316, 1984.

[11] D. Down, H. C. Gromoll, and A. L. Puha. The fluid limit of an SRPT queue. Under submission.

[12] C. Goerg. Evaluation of the optimal SRPT strategy with overhead. *IEEE Trans. on Comm.*, 34(4):338–344, 1986.

[13] M. Harchol-Balter, B. Schroeder, M. Agrawal, and N. Bansal. Size-based scheduling to improve web performance. *ACM Transactions on Computer Systems*, 21(2), May 2003.

[14] M. Hu, J. Zhang, and J. Sadowsky. A size-aided opportunistic scheduling scheme in wireless networks. In *Globecom*, 2003.

[15] W. Matuszewska. On a generalization of regularly increasing functions. *Studia Math.*, 24:271–279, 1964.

[16] M. Nuyens and A. Wierman. The foreground-background queue: A survey. *Performance evaluation*, 65(3-4):286–307, 2008.

[17] M. Nuyens and B. Zwart. A large-deviations analysis of the GI/GI/1 SRPT queue. *Queueing Sys.*, 54(2):85–97, 2006.

[18] Misja Nuyens, Adam Wierman, and Bert Zwart. Preventing large sojourn times using SMART scheduling. *Oper. Res.*, 56(1):88–101, 2008.

[19] T.M. O'Donovan. Direct solutions of M/G/1 priority queueing models. *Revue Francaise d'Automatique Informatique Recherche Operationnelle*, 10:107–111, 1976.

[20] A. V. Pechinkin. The MAP/G/1 queue with SRPT service discipline. *Theory of Prob. and its App.*, 45(3):532–539, 2000.

[21] R. Perera. The variance of delay time in queueing system M/G/1 with optimal strategy SRPT. *Archiv fur Elektronik und Uebertragungstechnik*, 47:110–114, 1993.

[22] I. A. Rai, G. Urvoy-Keller, and E. Biersack. Analysis of FB scheduling for job size distributions with high variance. In *Proc. of ACM Sigmetrics*, 2003.

[23] I. A. Rai, G. Urvoy-Keller, M. Vernon, and E. W. Biersack. Performance modeling of FB based scheduling in packet switched networks. In *Proc. of ACM Sigmetrics-Performance*, 2004.

[24] M. Rawat and A. Kshemkalyani. SWIFT: Scheduling in web servers for fast response time. In *Symp. on Net. Comp. and App.*, 2003.

[25] 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.

[26] Linus E. Schrage. A proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 16:678–690, 1968.

[27] F. Schreiber. Properties and applications of the optimal queueing strategy SRPT - a survey. *Archiv fur Elektronik und Uebertragungstechnik*, 47:372–378, 1993.

[28] A. Wierman. *Scheduling for today's computer systems: Bridging theory and practice.* PhD thesis, Carnegie Mellon University, 2007.

[29] A. Wierman and M. Harchol-Balter. Classifying scheduling policies with respect to unfairness in an M/GI/1. In *Proc. of ACM Sigmetrics*, 2003.

[30] Adam Wierman. Fairness and classifications. *Perf. Eval. Rev.*, 34(4):4–12, 2007.

Appendix

In this section, we summarize some important, basic properties of the Matuszewska index that we apply in our proofs.

Lemma 3. $\beta(1/f) = -\alpha(f)$.

Lemma 4. *[6] Let f be positive and locally integrable on $[X, \infty)$, and set $\tilde{f}(x) = \int_X^x f(t)dt/t$. If $\beta(f) > 0$ then $\liminf_{x\to\infty} f(x)/\tilde{f}(x) > 0$. And $\limsup_{x\to\infty} f(x)/\tilde{f}(x) \geq \alpha(f)$.*

Lemma 5. *[6] Let f be positive and measurable, and set $\tilde{f}(x) = \int_x^\infty f(t)dt/t \leq \infty$. If $\alpha(f) < 0$ then $\tilde{f}(x) < \infty$ for all large x.*

Lemma 6. *[9] Let $f : (0, \infty) \to R$ be non-decreasing and unbounded above, $f^{\leftarrow}(x) = \inf\{y \in [X, \infty) : f(y) > x\}$. $\alpha(f) < \infty$ iff $\beta(f^{\leftarrow}) > 0$ and $\beta(f^{\leftarrow}) = 1/\alpha(f)$.*

Lemma 7. *[6] Let f be positive. If $\alpha(f) < \infty$ then for every $\alpha > \alpha(f)$ there exist positive constants C, X such that $f(y)/f(x) \leq C(y/x)^\alpha$ for $y \geq x \geq X$.*