# Fluid and Diffusion Limits for Transient Sojourn Times of Processor Sharing Queues with Time Varying Rates 

Robert C. Hampshire<br>Department of Operations Research and Financial Engineering Princeton University<br>rhampshi@princeton.edu<br>Mor Harchol-Balter<br>School of Computer Science<br>Carnegie Mellon University<br>harchol@cs.cmu.edu<br>William A. Massey<br>Department of Operations Research and Financial Engineering<br>Princeton University<br>wmassey@princeton.edu

December 16, 2005


#### Abstract

We provide an approximate analysis of the transient sojourn time for a processor sharing queue with time varying arrival and service rates, where the load can vary over time, including periods of overload. Using the same asymptotic technique as uniform acceleration as demonstrated in [12] and [13], we obtain fluid and diffusion limits for the sojourn time of the $M_{t} / M_{t} / 1$ processor-sharing queue. Our analysis is enabled by the introduction of a "virtual customer" which differs from the notion of a "tagged customer" in that the former has no effect on the processing time of the other customers in the system. Our analysis generalizes to non-exponential service and interarrival times, when the fluid and diffusion limits for the queueing process are known.


Keywords: processor sharing, fluid limits, diffusion limits, transient behavior, timevarying queues, uniform acceleration, sojourn times, virtual customers.

## 1 Introduction

The processor sharing discipline has been used to model many aspects of computer systems, including the quantum-based time sharing of the CPU by computer operating systems (see Kleinrock [10]) and (elastic) traffic modeling in communication networks (see Nunez-Queija [16], Roberts [18] and Bonald and Proutiere [2]) and scheduling in Web servers [6].

Under processor sharing (PS), the service capacity is at all times equally shared among all the jobs present. If there are $n$ jobs present, each one receives a fraction $1 / n$ of the total service capacity. This scheduling policy induces simple formulas in the case of an underloaded (stable) M/G/1/PS queue. For example, in Kleinrock [10] it is shown that

$$
\lim _{t \rightarrow \infty} \mathrm{P}(Q(t)=n)=\left\{\begin{array}{cl}
(1-\rho) \rho^{n} & \text { if } \rho<1  \tag{1.1}\\
0 & \text { if } \rho \geq 1
\end{array}\right.
$$

where $Q(t)$ denotes the number of jobs in the system at time $t$. Moreover, $\rho=\lambda \cdot E[S]$ where $S$ denotes the random size (service requirement) of a job and $\lambda$ is the mean Poisson customer arrival rate. Furthermore, the expected sojourn time for a job with size (service requirement) $x$ is known to be

$$
E[T(x)]=\frac{x}{1-\rho} .
$$

The popularity of the PS queue is due in large part to Kleinrock [10] who uses processorsharing as an approximation of round-robin quantum-based scheduling. This is work that primarily deals with a stationary queue having load $\rho<1$. The survey papers of Yashkov [23, 24] provide a detailed overview of the results on stationary PS queues, including results by Coffman, Muntz and Trotter [4], Morrison [15], Guillemin and Boyer [7]. All these papers deal with various aspects of the sojourn time distribution for the $M / M / 1 / P S$. These sojourn time analyses are often based on the notion of tracking a "tagged customer" who arrives into the system and interacts with the customers there. Masuyama and Takine [14] obtain similar results for a $M A P / M / 1 / P S$ queue. In the early 1980's Ott [17], Schassberger [19] and Yashkov [22] all derive independently the Laplace-Stieltjes transform (LST) for the sojourn time in the $M / G / 1 / P S$ queue. Zwart and Boxma [25] show how to eliminate the contour integrals in the above results and obtain a more explicit formula, which they use to efficiently compute the moments of sojourn time. They also prove that for heavy-tailed service demand distributions, the heavy traffic limiting distribution of sojourn time has the same tail index as the service distribution.

While stationary behavior of the $\mathrm{M} / \mathrm{G} / 1 / \mathrm{PS}$ queue is well-studied, it is important to understand the transient behavior of the $\mathrm{M} / \mathrm{G} / 1 / \mathrm{PS}$ queue as well. When a job arrives into the system, the job finds a specific number of existing jobs with existing remaining sizes (remaining service requirements), not a steady-state distribution. Initial work on the M/M/l/PS transient queue was done by Sengupta and Jagerman [21] who produce Laplace transforms for the transient behavior. Similar work on the M/G/l/PS transient queue was done later by Kitaev [9]. This was followed by a key paper in the area of transient analysis of the M/G/1/PS queue by Jean-Marie and Robert [8] who provide a fluid approximation for the case of fixed load $\rho>1$. For this overloaded regime, they derive the asymptotic growth rate for the number of customers after $t$ time units of overload (for large $t$ ), where this rate is the solution to a simple integral equation. They also derive the asymptotic behavior of residual service times. These results were generalized by Chen, Kella, and Weiss [3] who also develop a fluid approximation, based on a time and space scaling. They too consider fixed load $\rho$ and
examine three regimes: $\rho<1, \rho=1$, and $\rho>1$. They also allow for general conditions on what the tagged arrival sees in terms of the number of jobs found in the system with their residual job sizes.

There are two areas that the prior work on transient analysis does not address, and these form the primary contributions of this paper. First, the prior work all deals with a fixed load $\rho$. In practice there are short-term fluctuations in load, which have a dramatic impact on the sojourn time, and are not captured by the existing constant-rate models. To motivate the importance of capturing load fluctuations, consider the performance graph of an Apache web server shown in Figure 1, taken from [20]. In this figure, instantaneous load fluctuates between 1.2 and 0.2 , where the time-average load is 0.7 . As the authors in [20] point out, steadystate queueing formulas for load 0.7 result in a very poor prediction of mean sojourn time (or equivalently, the mean response time). The analysis techniques that we introduce in this paper result in simple formulas that capture the effect of fluctuating loads and rates on transient sojourn time. Second, the prior work provides only a first-order (mean) approximation of sojourn time. Since the fluid approximations correspond to a strong law of large numbers, it is natural to extend these results to diffusion limits, which correspond to central limit theorems. We thus strengthen the prior work by using diffusion limits, which we rigorously establish for exponential service distributions. This provides us with estimations of the standard deviations about the mean and enables us to estimate the distribution of sojourn time, rather than just its mean.


Figure 1: Response time (averaged over 1 second intervals) for an Apache web server servicing HTTP requests, where load fluctuates between $\rho=1.2$ and $\rho=0.2$ (for 25 seconds each).

In this paper, we provide fluid and diffusion limits for the transient sojourn time of an $M_{t} / M_{t} / 1 / P S$ queue. Using this extended Kendall notation, the $M_{t}$ for an arrival process denotes a non-homogeneous Poisson process. Similarly, the $M_{t}$ for a service time distribution denotes the times between jumps for a non-homogeneous Poisson process. Our primary technique is uniform acceleration $[13,12]$, however we use it differently from how it has been used in the past. First and foremost, we are for the first time applying uniform acceleration to the sojourn time of a processor-sharing queue. Second, practical considerations have motivated us to extend the traditional uniform acceleration analysis and allow for general starting conditions in terms of the number of jobs seen by an arrival. Third, we introduce the notion of a "virtual customer," which differs from the traditional "tagged customer" in that the virtual
customer has no effect on the experience of the other customers in the system.
The sojourn time $T(x)$ of the virtual customer with a job of size $x$ can be determined from the following formula

$$
\begin{equation*}
x=\int_{0}^{T(x)} \frac{d t}{1+Q(t)} \tag{1.2}
\end{equation*}
$$

This would be the exact sojourn time of a "real customer" if the effect of the virtual customer on $Q(t)$ were taken into account. Such a job arriving at time 0 increases the number of jobs in the system by 1 , and is served at the instantaneous rate $1 /(1+Q(t))$ at time $t$.

Uniform acceleration is an asymptotic analysis method where we scale the arrival and service rates by a factor $\eta$. In the context of formula (1.2), this means that $\lambda(t) \Longrightarrow \eta \lambda(t)$ and $\mu(t) \Longrightarrow \eta \mu(t)$ and the corresponding queueing process is referred to as $Q^{\eta}(t)$, where we now take the limit as $\eta \rightarrow \infty$.

It is not clear why uniform acceleration should tell us anything about the original unscaled system. For example, in the case of a stationary queue, scaling the arrival and service rates each by a factor of $\eta$ should lead to a drop in the mean sojourn time by a factor of $\eta$. By contrast, for the case of transient (nonstationary) queues, we prove in Theorem 2.4, that we can induce an asymptotic analysis on $T(x)$ in terms of our asymptotic analysis of $Q^{\eta}(t)$. We also provide some underlying motivation for why the sojourn time obtained in the accelerated regime is indicative of what happens in the original (non-accelerated) system for the transient queue.

Our derivation of the sojourn time behavior of the $M_{t} / M_{t} / 1 / P S$ queue leads to some surprising results. First, we find that under systems with time-fluctuating load, the expected slowdown experienced by a job (its sojourn time divided by its size) is no longer constant. This is in sharp contrast to the classical M/G/1/PS queue with fixed load $\rho<1$, which is characterized by constant slowdown (which is also referred to as "fairness" [1]). Second, in studying the sojourn time distribution, we observe a point mass. This indicates that for a given fixed job size, we can explicitly determine a point at which the distribution has positive mass.

We end the paper with an exploration of numerical examples that show how well the asymptotic results for the $M_{t} / M_{t} / 1 / P S$ queue approximate the mean, variance, and distribution of these sojourn times.

## 2 Uniform Acceleration for the $M_{t} / M_{t} / 1 / P S$ Queue

Since the behavior of the $M_{t} / M_{t} / 1$ queueing process is independent of any work conserving queueing discipline, the sample path behavior of the $M_{t} / M_{t} / 1 / F I F O$ and $M_{t} / M_{t} / 1 / P S$ queues are identical. We can apply the asymptotic results to this system as found in Massey [13]. We use the asymptotic analysis of uniform acceleration, whereby we scale the time dependent arrival and service rates, $\lambda(t)$ and $\mu(t)$ respectively, by the same parameter $\eta>0$. This results in an $M_{t} / M_{t} / 1$ queueing process with arrival and service rates $\eta \lambda(t)$ and $\eta \mu(t)$, respectively, and we denote this queueing process as $\left\{Q^{\eta}(t) \mid t \geq 0\right\}$. We then analyze the behavior of this process asymptotically by letting our scale factor $\eta$ become very large.

We start by recalling the following asymptotic analysis for the transition probabilities of the $M_{t} / M_{t} / 1$ queueing process:

Theorem 2.1 (Massey, 1985) If $\lambda$ and $\mu$ are continuously differentiable functions of $t$, we have $\rho^{*}(t)<1$ where

$$
\begin{equation*}
\rho^{*}(t) \equiv \sup _{0 \leq s<t} \frac{\int_{s}^{t} \lambda(r) d r}{\int_{s}^{t} \mu(r) d r}, \tag{2.1}
\end{equation*}
$$

and $\rho(t)=\lambda(t) / \mu(t)$, then

$$
\begin{equation*}
\mathrm{P}\left(Q^{\eta}(t)=n\right)=(1-\rho(t)) \rho(t)^{n}+\frac{\rho^{\prime}(t)}{\eta \mu(t)}\left(\frac{\rho(t)}{(1-\rho(t))^{2}}-\frac{n(n+1)}{2}\right) \rho(t)^{n-1}+O\left(\frac{1}{\eta^{2}}\right), \tag{2.2}
\end{equation*}
$$

as $\eta \rightarrow \infty$. Moreover, if $\rho^{*}(t) \geq 1$, then

$$
\begin{equation*}
\lim _{\eta \rightarrow \infty} \mathrm{P}\left(Q^{\eta}(t)=n\right)=0 \tag{2.3}
\end{equation*}
$$

for all $t>0$.
From this analysis there are three different regimes of asymptotic behavior that are labeled as follows: underloaded $\left(\rho^{*}(t)<1\right)$, critically loaded $\left(\rho^{*}(t)=1\right)$ and overloaded $\left(\rho^{*}(t)>\right.$ 1). Observe that the parameter that determines these asymptotic regimes is not given by $\rho(t)$. Also, observe that the results of this theorem are non-trivial (non-zero) only for the underloaded case.

Mandelbaum and Massey [12] apply the analysis of uniform acceleration directly to the random sample path behavior of the $M_{t} / M_{t} / 1$ queue.

Theorem 2.2 (Mandelbaum and Massey, 1995) If $\lambda$ and $\mu$ are locally integrable functions and $Q^{\eta}(0)=Q(0)$ for all $\eta>0$, then $\lim _{\eta \rightarrow \infty} Q^{\eta}(t) / \eta=Q^{(0)}(t)$ a.s. uniformly on compact sets, where

$$
\begin{equation*}
Q^{(0)}(t)=\int_{0}^{t}(\lambda(s)-\mu(s)) d s-\inf _{0 \leq s \leq t} \int_{0}^{s}(\lambda(r)-\mu(r)) d r . \tag{2.4}
\end{equation*}
$$

Moreover, $\lim _{\eta \rightarrow \infty} \sqrt{\eta}\left(Q^{\eta}(t) / \eta-Q^{(0)}(t)\right) \stackrel{d}{=} Q^{(1)}(t)$, where

$$
\begin{equation*}
Q^{(1)}(t)=W\left(\int_{0}^{t}(\lambda(s)+\mu(s)) d s\right)-\inf _{s \in \Phi(t)} W\left(\int_{0}^{s}(\lambda(r)+\mu(r)) d r\right), \tag{2.5}
\end{equation*}
$$

$\{W(t) \mid t \geq 0\}$ is standard (mean 0, variance $t$ ) Brownian motion, and finally

$$
\begin{equation*}
\Phi(t)=\left\{s \mid \int_{s}^{t}(\lambda(r)-\mu(r)) d r=Q^{(0)}(t) \text { and } 0 \leq s \leq t\right\} . \tag{2.6}
\end{equation*}
$$

The results of Theorem 2.1 are non-zero precisely when the results of Theorem 2.2 are zero. Conversely, the results of Theorem 2.2 are non-zero when the results of Theorem 2.1 are zero. Thus these two types of uniform acceleration analysis literally complement each other.

The deterministic process $\left\{Q^{(0)}(t) \mid t \geq 0\right\}$ is called the $M_{t} / M_{t} / 1$ queueing fluid limit. Observe that $Q^{(0)}(t)>0$ if and only if $\rho^{*}(t)>1$. The fluid limit is then an estimate of the backlog for the original (non accelerated, $\eta=1$ ) queueing process. Notice that since we always have $\rho(t) \leq \rho^{*}(t)$, it is possible to have $\rho(t)<1$ (or $\lambda(t)<\mu(t)$ ), but still have $\rho^{*}(t) \geq 1$. This
is due to a backlog of jobs, acquired during a period of overloading in the past, that have not been flushed out of the queue by time $t$.

The random process $\left\{Q^{(1)}(t) \mid t \geq 0\right\}$ is the first order "correction term" to the fluid limit and gives a sense of how the original queueing process deviates from the fluid model. For simplicity, we refer to it as the $M_{t} / M_{t} / 1$ queueing diffusion limit. Technically, it may not be a diffusion since the sample paths at fixed time points may have a non-zero probability of a discontinuity. However, these discontinuities only occur during transitions from the overloaded regime to the underloaded one.

Now we generalize these limit theorems to the extended case of a uniformly accelerated initial load where we set $Q^{\eta}(0)=\eta \cdot Q(0)$. Below we have our extended fluid and diffusion limits.

Theorem 2.3 If $Q^{\eta}(0)=\eta \cdot Q(0)$, then

$$
\begin{equation*}
Q^{(0)}(t)=Q(0)+\int_{0}^{t}(\lambda(s)-\mu(s)) d s-\inf _{0 \leq s \leq t}\left(Q(0)+\int_{0}^{s}(\lambda(r)-\mu(r)) d r\right) \wedge 0 \tag{2.7}
\end{equation*}
$$

and

$$
Q^{(1)}(t)= \begin{cases}Q_{0}^{(1)}(t) & \text { if } Q(0)<\sup _{0 \leq s \leq t} \int_{0}^{s} \mu(r)-\lambda(r) d r  \tag{2.8}\\ Q_{0}^{(1)}(t) \vee W\left(\int_{0}^{t}(\lambda(s)+\mu(s)) d s\right) & \text { if } Q(0)=\sup _{0 \leq s \leq t} \int_{0}^{s} \mu(r)-\lambda(r) d r \\ W\left(\int_{0}^{t}(\lambda(s)+\mu(s)) d s\right) & \text { if } Q(0)>\sup _{0 \leq s \leq t} \int_{0}^{s} \mu(r)-\lambda(r) d r\end{cases}
$$

where $\left\{Q_{0}^{(1)}(t) \mid t \geq 0\right\}$ is the $M_{t} / M_{t} / 1$ queueing diffusion limit given by (2.5) when $Q(0)=0$.
Proof: Both limits follow from the fact that the process $\left\{Q^{\eta}(t) \mid t \geq 0\right\}$ can be written as

$$
\begin{equation*}
Q^{\eta}(t)=Z^{\eta}(t)-\inf _{0 \leq s \leq t} Z^{\eta}(s) \wedge 0=\max \left(Q_{0}^{\eta}(t), Z^{\eta}(t)\right) \tag{2.9}
\end{equation*}
$$

where $\left\{Q_{0}^{\eta}(t) \mid t \geq 0\right\}$ is another $M_{t} / M_{t} / 1$ queueing process with the same arrival and service rates but with $Q(0)=0$, and

$$
\begin{equation*}
Z^{\eta}(t)=\eta \cdot Q(0)+\Pi_{1}\left(\int_{0}^{t} \eta \lambda(s) d s\right)-\Pi_{2}\left(\int_{0}^{t} \eta \mu(s) d s\right) \tag{2.10}
\end{equation*}
$$

where $\left\{\Pi_{i}(t) \mid t \geq 0\right\}$ for $i=1,2$ are two independent, standard (rate 1) Poisson processes that are used to construct both queueing processes $Q^{\eta}$ and $Q_{0}^{\eta}$. Using the theory of strong approximations for Poisson processes, we have

$$
\begin{equation*}
\lim _{\eta \rightarrow \infty} \sup _{0 \leq s \leq t}\left|\frac{1}{\eta} Z^{\eta}(s)-Q(0)-\int_{0}^{s} \lambda(r)-\mu(r) d r\right|=0 \quad \text { a.s. } \tag{2.11}
\end{equation*}
$$

and we can construct a Brownian motion, $\{W(t) \mid t \geq 0\}$, such that

$$
\begin{equation*}
\lim _{\eta \rightarrow \infty} \sup _{0 \leq s \leq t}\left|\sqrt{\eta}\left(\frac{1}{\eta} Z^{\eta}(s)-Q(0)-\int_{0}^{s} \lambda(r)-\mu(r) d r\right)-W\left(\int_{0}^{s} \lambda(r)+\mu(r) d r\right)\right| \stackrel{d}{=} 0 \tag{2.12}
\end{equation*}
$$

Applying the asymptotics of (2.11), (2.12) and Theorem 2.2 to (2.9) gives us the desired result.

Suppose that we have a single virtual customer with a job of size $x$ sharing a unit processing rate with $Q(t)$ customers. Combining our notion of a virtual customer with uniform acceleration, we transform this system into $\eta$ virtual customers with jobs of size $x / \eta$ sharing a unit processing rate with $Q^{\eta}(t)$ customers. This gives us

$$
\begin{equation*}
x / \eta=\int_{0}^{T^{\eta}(x)} \frac{d t}{\eta+Q^{\eta}(t)} \quad \Longrightarrow \quad x=\int_{0}^{T^{\eta}(x)} \frac{d t}{1+Q^{\eta}(t) / \eta} \tag{2.13}
\end{equation*}
$$

Now we can state our main result and its proof.
Theorem 2.4 Given an $M_{t} / M_{t} / 1 / P S$ queue and for all $x \geq 0$, we have the strong law of large numbers limit

$$
\begin{equation*}
\lim _{\eta \rightarrow \infty} T^{\eta}(x)=T^{(0)}(x) \text { a.s. } \quad \text { where } \quad x=\int_{0}^{T^{(0)}(x)} \frac{d t}{1+Q^{(0)}(t)} \tag{2.14}
\end{equation*}
$$

and the central limit theorem

$$
\begin{equation*}
\lim _{\eta \rightarrow \infty} \sqrt{\eta}\left(T^{\eta}(x)-T^{(0)}(x)\right) \stackrel{d}{=} T^{(1)}(x) \equiv T^{(0)^{\prime}}(x) \cdot \int_{0}^{T^{(0)}(x)} \frac{Q^{(1)}(t) d t}{\left(1+Q^{(0)}(t)\right)^{2}} \tag{2.15}
\end{equation*}
$$

Proof: To prove the strong law limit result (2.14), it suffices to show that every convergent subsequence $\left\{T^{\eta(k)}(x) \mid k=1,2, \ldots\right\}$, where $\eta(k) \rightarrow \infty$, converges to $T^{(0)}(x)$.

Let $U=\lim _{k \rightarrow \infty} T^{\eta(k)}(x)$. Since it follows that

$$
\begin{equation*}
\left|\int_{0}^{T^{\eta(k)}(x)} \frac{d t}{1+Q^{\eta(k)}(t) / \eta(k)}-\int_{0}^{U} \frac{d t}{1+Q^{(0)}(t)}\right| \leq\left|T^{\eta(k)}(x)-U\right| \tag{2.16}
\end{equation*}
$$

we then have

$$
\begin{equation*}
x=\lim _{k \rightarrow \infty} \int_{0}^{T^{\eta(k)}(x)} \frac{d t}{1+Q^{\eta(k)}(t) / \eta(k)}=\int_{0}^{U} \frac{d t}{1+Q^{(0)}(t)} . \tag{2.17}
\end{equation*}
$$

Since $1 /\left(1+Q^{(0)}(t)\right)$ is never zero, we must then have $U=T^{(0)}(x)$.
To prove the central limit theorem result (2.15), we first observe that

$$
\begin{align*}
\int_{T^{(0)}(x)}^{T^{\eta}(x)} \frac{d t}{1+Q^{(0)}(t)} & =\int_{0}^{T^{\eta}(x)} \frac{d t}{1+Q^{(0)}(t)}-\int_{0}^{T^{(0)}(x)} \frac{d t}{1+Q^{(0)}(t)}  \tag{2.18}\\
& =\int_{0}^{T^{\eta}(x)} \frac{d t}{1+Q^{(0)}(t)}-\int_{0}^{T^{\eta}(x)} \frac{d t}{1+Q^{\eta}(t) / \eta}  \tag{2.19}\\
& =\int_{0}^{T^{\eta}(x)} \frac{\left(Q^{\eta}(t) / \eta-Q^{(0)}(t)\right) d t}{\left(1+Q^{(0)}(t)\right)\left(1+Q^{\eta}(t) / \eta\right)} \tag{2.20}
\end{align*}
$$

Now we obtain the identity

$$
\begin{align*}
\sqrt{\eta}\left(T^{\eta}(x)-T^{(0)}(x)\right)= & \left(\frac{1}{T^{\eta}(x)-T^{(0)}(x)} \int_{T^{(0)}(x)}^{T^{\eta}(x)} \frac{d t}{1+Q^{(0)}(t)}\right)^{-1} \\
& \times \int_{0}^{T^{\eta}(x)} \frac{\sqrt{\eta}\left(Q^{\eta}(t) / \eta-Q^{(0)}(t)\right) d t}{\left(1+Q^{(0)}(t)\right)\left(1+Q^{\eta}(t) / \eta\right)} . \tag{2.21}
\end{align*}
$$

The rest follows since $\sqrt{\eta}\left(Q^{\eta}(t) / \eta-Q^{(0)}(t)\right)$ converges in distribution to a random process and all the other limits converge to constants.
It should be pointed out that the proof given here holds for any processor sharing queueing system where the queueing process for the number of customers in the system has fluid and diffusion limits.

We note that Theorem 2.4 applies to the transient sojourn times of the $M_{t} / M_{t} / 1 / P S$ queue. We emphasize the word transient to stress the fact that we are not analyzing sojourn times where the initial queue has a steady state distribution. The results that we discuss here do not necessarily have any implications for constant rate queues in steady state.

Figure 2 illustrates the heuristic idea behind Theorem 2.4. The original queue is represented in the top part of the figure. To illustrate our scale transformation, the jobs are now "quantized" into a series of shaded squares that represent unit processing times that can be viewed as "CPU time units". The bottom part of Figure 2 represents an accelerated queue with scale factor 2 , where both the number of new jobs and the processing rates of these jobs are doubled. Observe that this is equivalent to leaving the processing rate untouched but breaking each job in the original queue into two identical "subjobs", that are each half the size of the original total job.

Figure 2 illustrates a specific example where the job sizes in the original queue are all multiples of the scale factor 2 and all jobs are initially present. In this specific case shown, the completion times of the final subjob for the accelerated queue are identical to the corresponding completion times of the jobs for the original queue. We refer to this property as "scale invariance" for the processor sharing service discipline. While it is clear that scale invariance holds for the specific example illustrated in Figure 2, and probably generalizes to other cases of deterministic job sizes with appropriately sized jobs, this argument cannot be generalized indefinitely. Yet we hope that this provides some intuition behind Theorem 2.4, specifically why the accelerated system allows us to understand the behavior of the nonaccelerated system.

For deterministic job sizes and arrival processes, we can then define a precise notion of "scale invariance" for the processor sharing service discipline and a limited version of round robin. For the round robin case, we are limited to the case of the job sizes in the original queue initially being integer multiples of the scale factor 2 times the CPU time unit and all jobs are present initially. In either case, the completion time of the final subjobs for the accelerated queue is identical to the completion times of the corresponding total jobs for the original queue. All these results generalize to any scale factor of integer size $\eta$.

For a random processor sharing queue, we define a stochastic generalization of this type of acceleration. For the example of a scale factor equal to 2 , we still assume that the initial load of jobs in the system is deterministic and we double this number. For the job arrival process however, we replace this notion of doubling (scaling by 2 ) by the superposition of two i.i.d. replicas of the original arrival process. Note that for deterministic processes, this is the same as doubling. The sizes of these arriving subjobs are i.i.d. random variables but each one has the same distribution as one half the size of a random job (half as many CPU time units) in the original queue. For the case of the $M_{t} / M_{t} / 1 / P S$ queue, this notion of "stochastic acceleration" is identical to uniform acceleration with scale factor $\eta=2$. Now we extend the definition of uniform acceleration to the case of the scale factor $\eta$ being any integer. Instead of scale invariance, we have the fluid and diffusion limits of Theorem 2.4 as $\eta \rightarrow \infty$.

In the next section, we reduce these formulas to the constant rate case. We can obtain more explicit formulas that give us insight into how the sojourn time responds to periods of underloading and overloading.


Figure 2: Scale Invariance for Processor Sharing and Limited Round Robin

## 3 Exact Formulas for the Constant Rate Case

We have the following formula for the fluid limit of the $M / M / 1 / P S$ queue, which is a special case of Chen, Kella and Weiss [3].
Theorem 3.1 Given an $M / M / 1 / P S$ queue with arrival rate $\lambda$ and service rate $\mu$, we have

$$
\begin{equation*}
Q^{(0)}(t)=(Q(0)+(\lambda-\mu) t)^{+}, \tag{3.1}
\end{equation*}
$$

for all $t \geq 0$. Moreover, for all job sizes $x \geq 0$, we have

$$
\begin{equation*}
T^{(0)}(x)=x+\int_{0}^{x}\left((Q(0)+1) \cdot e^{(\lambda-\mu) y}-1\right)^{+} d y \tag{3.2}
\end{equation*}
$$

Finally, for all job sizes $x \geq 0$, we have

$$
\begin{equation*}
Q^{(0)}\left(T^{(0)}(x)\right)=\left((Q(0)+1) \cdot e^{(\lambda-\mu) x}-1\right)^{+} . \tag{3.3}
\end{equation*}
$$

For the case of $\lambda<\mu$, let $t_{*}$ equal the first time that the extended fluid limit process empties. Similarly, let $x_{*}$ equal the size of the smallest virtual job starting at time 0 that finishes just as the fluid limit first empties. We can write them explicitly as

$$
\begin{equation*}
t_{*} \equiv \frac{Q(0)}{\mu-\lambda} \quad \text { and } \quad x_{*} \equiv \frac{\log (Q(0)+1)}{\mu-\lambda} . \tag{3.4}
\end{equation*}
$$

Using $t_{*}$, the extended diffusion limit for the queueing process can be written as follows:
Theorem 3.2 If $\lambda>\mu$ or $\lambda=\mu$ and $Q(0)>0$, then

$$
\begin{equation*}
\left\{Q^{(1)}(t) \mid t \geq 0\right\} \stackrel{d}{=}\{W((\lambda+\mu) t) \mid t \geq 0\} . \tag{3.5}
\end{equation*}
$$

If $\lambda=\mu$ and $Q(0)=0$, then

$$
\begin{equation*}
\left\{Q^{(1)}(t) \mid t \geq 0\right\} \stackrel{d}{=}\left\{W(2 \lambda t)-\inf _{0 \leq s \leq t} W(2 \lambda s) \mid t \geq 0\right\} . \tag{3.6}
\end{equation*}
$$

Finally, if $\lambda<\mu$ and $Q(0) \geq 0$, then

$$
Q^{(1)}(t) \stackrel{d}{=} \begin{cases}W((\lambda+\mu) t) & \text { if } t<t_{*},  \tag{3.7}\\ W((\lambda+\mu) t)^{+} & \text {if } t=t_{*}, \\ 0 & \text { if } t>t_{*} .\end{cases}
$$

Note that for the case of $\lambda<\mu$ and $Q(0)>0$, almost all the sample paths of $\left\{Q^{(1)}(t) \mid t \geq 0\right\}$ are discontinuous at $t_{*}$. Probabilistically speaking, half of them are only left continuous (i.e. $\left.W\left((\lambda+\mu) t_{*}\right)>0\right)$ and the other half are only right continuous at $t_{*}$.

Excluding the critical loading case of $\lambda=\mu$ and $Q(0)=0$, we can totally characterize the distribution of $T^{(1)}(x)$.
Theorem 3.3 If $Q(0)>0$ or $\lambda \neq \mu$, then $T^{(1)}(x)$ is a Gaussian random variable with $\mathrm{E}\left[T^{(1)}(x)\right]=0$. Moreover, we have

$$
\begin{equation*}
\operatorname{Var}\left[T^{(1)}(x)\right]=\frac{(1+Q(0))(\lambda+\mu)}{(\lambda-\mu)^{3}}\left[e^{2(\lambda-\mu) x}-2(\lambda-\mu) x e^{(\lambda-\mu) x}-1\right] \tag{3.8}
\end{equation*}
$$

provided that either $\lambda>\mu$ or the conjunction of $\lambda<\mu$ and $x<x_{*}$ holds. When $\lambda=\mu$, this formula reduces by L'Hopital's rule to

$$
\begin{equation*}
\operatorname{Var}\left[T^{(1)}(x)\right]=(1+Q(0)) \frac{2 \lambda}{3} x^{3} \tag{3.9}
\end{equation*}
$$

Otherwise, when $\lambda<\mu$ and $x \geq x_{*}$, we have

$$
\begin{equation*}
\operatorname{Var}\left[T^{(1)}(x)\right]=\frac{\lambda+\mu}{(\mu-\lambda)^{3}}\left[Q(0)+1-2 \cdot \log (Q(0)+1)-\frac{1}{Q(0)+1}\right] \tag{3.10}
\end{equation*}
$$

Proof of Theorem 3.1: The extended constant rate fluid limit (3.1) follows immediately from equation (2.7). The constant rate formula given by (3.3) follows from differentiating $T^{(0)}(x)$ and subtracting one from it. It remains to derive a closed form solution to the extended fluid limit of the sojourn time.

If $\lambda=\mu$, then $Q^{(0)}(t)=Q(0)$ for all $t \geq 0$, and so $T^{(0)}(x)=(Q(0)+1) \cdot x$.
If $\lambda>\mu$, then

$$
\begin{equation*}
x=\int_{0}^{T^{(0)}(x)} \frac{d t}{1+Q(0)+(\lambda-\mu) t}=\frac{1}{\lambda-\mu} \log \left(1+\frac{\lambda-\mu}{Q(0)+1} T^{(0)}(x)\right) \tag{3.11}
\end{equation*}
$$

Solving for $T^{(0)}(x)$ gives us

$$
\begin{equation*}
T^{(0)}(x)=(Q(0)+1) \cdot \frac{e^{(\lambda-\mu) x}-1}{\lambda-\mu}=x+\int_{0}^{x}\left((Q(0)+1) \cdot e^{(\lambda-\mu) y}-1\right)^{+} d y \tag{3.12}
\end{equation*}
$$

The last step follows from $e^{(\lambda-\mu) x}$ being an increasing function of $x$ and greater than one when $\lambda>\mu$ and $x>0$.

Now let $\lambda<\mu$. Since $t_{*}=\inf \left\{t \mid Q^{(0)}(t)=0\right\}$, we have the two cases of $T^{(0)}(x)<t_{*}$ and $T^{(0)}(x) \leq t_{*}$. If $T^{(0)}(x)<t_{*}$, then we have the same equation for $T^{(0)}(x)$ as (3.11) for the case of $\lambda>\mu$. By continuity, we then have $T^{(0)}\left(x_{*}\right)=t_{*}$.

If $T^{(0)}(x) \geq t_{*}$, then $x \geq x_{*}$, and so we have $x-x_{*}=T^{(0)}(x)-t_{*}$. This equation simply states the fact that $x_{*}$ is the amount of the fluid model job of size $x$ that was processed with the initial load of fluid model jobs. Since $\lambda<\mu$, the fluid level in zero for all time after $t_{*}$. This means that the remaining job amount of $x-x_{*}$ has the server all to itself. Consequently,

$$
\begin{equation*}
T^{(0)}(x)=x-x_{*}+t_{*} \tag{3.13}
\end{equation*}
$$

which gives us

$$
\begin{equation*}
T^{(0)}(x)=x+\int_{0}^{x_{*}}\left((Q(0)+1) e^{(\lambda-\mu) y}-1\right) d y=x+\int_{0}^{x}\left((Q(0)+1) e^{(\lambda-\mu) y}-1\right)^{+} d y \tag{3.14}
\end{equation*}
$$

The last step follows from $e^{(\lambda-\mu) x}$ being a decreasing function of $x$ when $\lambda<\mu$. This argument also proves that this last formula is true for the previous case of $T^{(0)}(x)<t_{*}$.

Proof of Theorem 3.3: When $\lambda>\mu$, let $\alpha \equiv(\lambda-\mu) /(1+Q(0))$. Using the identities

$$
\begin{equation*}
\int_{0}^{t} \frac{s d s}{(1+\alpha s)^{2}}=\frac{1}{\alpha^{2}}\left[\log (1+\alpha t)+\frac{1}{1+\alpha t}-1\right] \tag{3.15}
\end{equation*}
$$

and

$$
\begin{equation*}
\int_{0}^{t} \frac{\log (1+\alpha s) d s}{(1+\alpha s)^{2}}=\frac{1}{\alpha}\left[1-\frac{\log (1+\alpha t)}{1+\alpha t}-\frac{1}{1+\alpha t}\right] \tag{3.16}
\end{equation*}
$$

gives us

$$
\begin{aligned}
\operatorname{Var}\left[T^{(1)}(x)\right]= & (1+Q(0))^{2} \cdot e^{2(\lambda-\mu) x} \\
& \times \int_{0}^{T^{(0)}(x)} \int_{0}^{T^{(0)}(x)} \frac{\operatorname{Cov}[W((\lambda+\mu) s), W((\lambda+\mu) t)] d s d t}{(1+Q(0)+(\lambda-\mu) s)^{2}(1+Q(0)+(\lambda-\mu) t)^{2}} \\
= & \frac{2(\lambda+\mu) e^{2(\lambda-\mu) x}}{(1+Q(0))^{2}} \cdot \int_{0}^{T^{(0)}(x)}\left(\int_{0}^{t} \frac{s d s}{(1+\alpha s)^{2}}\right) \frac{d t}{(1+\alpha t)^{2}} \\
= & \frac{2(\lambda+\mu) e^{2(\lambda-\mu) x}}{(1+Q(0))^{2} \alpha^{2}} \cdot \int_{0}^{T^{(0)}(x)} \frac{\log (1+\alpha t)+1 /(1+\alpha t)-1}{(1+\alpha t)^{2}} d t \\
= & \frac{(\lambda+\mu)(1+Q(0))}{(\lambda-\mu)^{3}} \cdot\left[e^{2(\lambda-\mu) x}-2(\lambda-\mu) x e^{(\lambda-\mu) x}-1\right] .
\end{aligned}
$$

For the case of $\lambda<\mu$, observe that when $x<x_{*}$, we have

$$
\begin{equation*}
T^{(1)}(x)=T^{(0)^{\prime}}(x) \cdot \int_{0}^{T^{(0)}(x)} \frac{W((\lambda+\mu) t) d t}{(1+Q(0)+(\lambda-\mu) t)^{2}} . \tag{3.17}
\end{equation*}
$$

It follows that the variance formula is the same here as for the case of $\lambda>\mu$.
For the case of $\lambda<\mu$ and $x \geq x_{*}$, we have $Q^{(0)}\left(T^{(0)}(x)\right)=Q^{(0)}\left(T^{(0)}\left(x_{*}\right)\right)=0$. Therefore,

$$
\begin{equation*}
T^{(1)}(x)=T^{(0)^{\prime}}\left(x_{*}\right) \cdot \int_{0}^{T^{(0)}\left(x_{*}\right)} \frac{W((\lambda+\mu) t) d t}{(1+Q(0)+(\lambda-\mu) t)^{2}} . \tag{3.18}
\end{equation*}
$$

If follows that the variance formula for $T^{(1)}(x)$ is computed here by applying the previous variance formula to $T^{(1)}\left(x_{*}\right)$.

## 4 Numerics for the Mean and Variance

The aim of this section is to compare our analytical results for these $M_{t} / M_{t} / 1 / P S$ sojourn times with results from simulation. The limit theorems of Section 2 suggest the following fluid and diffusion limit approximations for the mean and variance of the sojourn times,

$$
\begin{equation*}
\mathrm{E}[T(x)] \approx T^{(0)}(x)+\mathrm{E}\left[T^{(1)}(x)\right] \quad \text { and } \quad \operatorname{Var}[T(x)] \approx \operatorname{Var}\left[T^{(1)}(x)\right] \tag{4.1}
\end{equation*}
$$

where in general we have $\mathrm{E}\left[T^{(1)}(x)\right]=0$. Now we give some general conditions for computing $\operatorname{Var}\left[T^{(1)}(x)\right]$.

Suppose that the queueing process alternates between periods of underloading and overloading, where the times of critical loading are discrete, isolated points. If we set $\sigma_{n}$ and $\tau_{n}$ to be respectively, the starting time for the $n$-th period of overloading and the ending time for the $n$-th period of overloading, then we have

$$
\begin{equation*}
\int_{\sigma_{n}}^{\tau_{n}}(\lambda(s)-\mu(s)) d s=0 \tag{4.2}
\end{equation*}
$$

which implies

$$
Q^{(0)}(t)=\left\{\begin{array}{cc}
\int_{\sigma_{n}}^{t}(\lambda(s)-\mu(s)) d s & \text { if } \sigma_{n} \leq t<\tau_{n}  \tag{4.3}\\
0 & \text { otherwise }
\end{array}\right.
$$

and

$$
Q^{(1)}(t)=\left\{\begin{array}{cc}
W\left(\int_{0}^{t}(\lambda(s)+\mu(s)) d s\right)-W\left(\int_{0}^{\sigma_{n}}(\lambda(s)+\mu(s)) d s\right) & \text { if } \sigma_{n} \leq t<\tau_{n}  \tag{4.4}\\
\left(W\left(\int_{0}^{t}(\lambda(s)+\mu(s)) d s\right)-W\left(\int_{0}^{\sigma_{n}}(\lambda(s)+\mu(s)) d s\right)\right)^{+} & \text {if } t=\tau_{n} \\
0 & \text { otherwise }
\end{array}\right.
$$

Theorem 4.1 If we are given an $M_{t} / M_{t} / 1$ queue where critical loading only occurs at isolated time points, then $T^{(1)}(x)$ is a Gaussian random variable, with $\mathrm{E}\left[T^{(1)}(x)\right]=0$ and

$$
\begin{equation*}
\operatorname{Var}\left[T^{(1)}(x)\right]=T^{(0)^{\prime}}(x)^{2} \cdot \sum_{n=0}^{\infty} \int_{\sigma_{n}(x)}^{\tau_{n}(x)} \int_{\sigma_{n}(x)}^{\tau_{n}(x)} \frac{\left(\int_{\sigma_{n}(x)}^{s \wedge t}(\lambda(r)+\mu(r)) d r\right) d s d t}{\left(1+Q^{(0)}(s)\right)^{2}\left(1+Q^{(0)}(t)\right)^{2}}, \tag{4.5}
\end{equation*}
$$

with

$$
\begin{equation*}
\sigma_{n}(x) \equiv \sigma_{n} \wedge T^{(0)}(x) \quad \text { and } \quad \tau_{n}(x) \equiv \tau_{n} \wedge T^{(0)}(x) \tag{4.6}
\end{equation*}
$$

Proof: We have

$$
\begin{equation*}
T^{(1)}(x)=T^{(0)^{\prime}}(x)^{2} \cdot \int_{0}^{\infty} \frac{Q^{(1)}(t) d t}{\left(1+Q^{(0)}(t)\right)^{2}}=T^{(0)^{\prime}}(x)^{2} \cdot \sum_{n=0}^{\infty} \int_{\sigma_{n}(x)}^{\tau_{n}(x)} \frac{Q^{(1)}(t) d t}{\left(1+Q^{(0)}(t)\right)^{2}} \tag{4.7}
\end{equation*}
$$

Since Brownian motion has the independent increment property, then by (4.4) $Q^{(1)}(s)$ and $Q^{(1)}(t)$ are independent random variables whenever $s$ and $t$ belong to disjoint periods of overloading. This means that

$$
\operatorname{Var}\left[T^{(1)}(x)\right]=T^{(0)^{\prime}}(x)^{2} \cdot \operatorname{Var}\left[\sum_{n=0}^{\infty} \int_{\sigma_{n}(x)}^{\tau_{n}(x)} \frac{Q^{(1)}(t) d t}{\left(1+Q^{(0)}(t)\right)^{2}}\right]
$$

$$
\begin{aligned}
& =T^{(0)^{\prime}}(x)^{2} \cdot \sum_{n=0}^{\infty} \operatorname{Var}\left[\int_{\sigma_{n}(x)}^{\tau_{n}(x)} \frac{Q^{(1)}(t) d t}{\left(1+Q^{(0)}(t)\right)^{2}}\right] \\
& =T^{(0)^{\prime}}(x)^{2} \cdot \sum_{n=0}^{\infty} \int_{\sigma_{n}(x)}^{\tau_{n}(x)} \int_{\sigma_{n}(x)}^{\tau_{n}(x)} \frac{\operatorname{Cov}\left[Q^{(1)}(s), Q^{(1)}(t)\right] d s d t}{\left(1+Q^{(0)}(s)\right)^{2}\left(1+Q^{(0)}(t)\right)^{2}} .
\end{aligned}
$$

The last step follows from setting

$$
\begin{equation*}
\operatorname{Cov}\left[Q^{(1)}(s), Q^{(1)}(t)\right]=\int_{\sigma_{n}(x)}^{s \wedge t}(\lambda(r)-\mu(r)) d r \tag{4.8}
\end{equation*}
$$

which completes the proof.
Figure 3 shows the approximations for mean and variance given by Equation (4.1) as compared with simulation, as well as the relative error. The left column of Figure 3 assumes a non-homogeneous Poisson arrival process with mean rate $\lambda(t)=1.2+0.2 * \sin (2 \pi * 0.2 * t)$, while the right column assumes a non-homogeneous Poisson arrival process with mean rate $\lambda(t)=1.2+0.2 * \sin (2 \pi * 10.0 * t)$. Throughout, all jobs have sizes that are exponentiallydistributed with mean 1, and the initial number of jobs in the system, $Q(0)$ is fixed at 10 . All the simulation results presented are derived from 10,000 realizations.

The maximum relative error of the fluid limit for the smaller frequency case (i.e. 0.2 versus 10.0 ) is less than the higher frequency case. This is consistent with results from the theory of uniform acceleration as seen in Theorem 3.1. The leading order terms of Equations 3.2, 3.3 and 3.4 do not depend on the rate of change in the offered load, $\rho^{\prime}(t)$. However their correction terms grow in magnitude as $\rho^{\prime}(t)$ becomes larger. Therefore it is reasonable to observe a smaller relative error in the slowly varying rate case.

## 5 Bimodality and Numerics for the Distribution

In this section, we first show that the sojourn time distribution always has a point mass distribution at the constant $(1+Q(0)) x$.

Theorem 5.1 Conditioned on $Q(0)$, we can define the following two independent events:

1. The first $Q(0)$ i.i.d. exponential service times are all larger than $x$.
2. The number of non-homogeneous Poisson arrivals for the next $(1+Q(0)) x$ time units is zero.

It then follows that the intersection of these two events 1 and 2 implies the event $\{T(x)=$ $(1+Q(0)) x\}$.

Moreover, if $T(x)$ is the sojourn time for the $M_{t} / M_{t} / 1 / P S$ queue with some initial load $Q(0)$, then we have

$$
\begin{equation*}
\mathrm{P}(T(x)=(1+Q(0)) x)=\exp \left(-\int_{0}^{(1+Q(0)) x} \lambda_{t} d t-Q(0) \cdot \int_{0}^{x} \mu_{t} d t\right), \tag{5.1}
\end{equation*}
$$

which equals the probability for the intersection of these two events 1 and 2.


Figure 3: Mean, standard deviation and relative errors for the sojourn times of $M_{t} / M / 1 / P S$ when $\lambda(t)=1.2+.2 * \sin 0.4 \pi t$ (left column) and $\lambda(t)=1.2+.2 * \sin 20 \pi t$ (right column), where $\mu=1$. Here, we assume throughout that $Q(0)=10$.

| $\beta(x)$ | $\lambda=2.0$ | $\lambda=0.5$ |
| :---: | :---: | :---: |
| $x=0.10536$ (10th percentile job size) | 3.37152 | 1.63308 |
| $x=0.69315$ (50th percentile job size) | 22.1808 | 10.743825 |
| $x=2.30259$ (90th percentile job size) | 73.68288 | 35.690145 |

Table 1: Values for $\beta(x)$ used in Figure 4.

Proof: When one of these two events does not happen, then there is some exponentially distributed random variable that is less than $(1+Q(0)) x$. This random variable, conditioned on being less than the constant $(1+Q(0)) x$, has a density. Since our underlying queueing process is Markovian, the remaining time until $T(x)$ occurs is independent of this conditioned random variable.

Now we use the fact that if $X$ and $Y$ are two independent random variables and if $X$ has a probability density, then so does $X+Y$. In fact the new density for this sum is the convolution of the density for $X$ with the probability distribution for $Y$.

If we use the notion of generalized functions, we can let $f_{T(x)}(t)$ denote the "density" of $T(x)$, where we may use delta functions to allow for the possibility of point mass distributions. Given our limit theorems, we know that this sojourn time distribution is asymptotically normal. This suggests that the actual sojourn time distribution is approximately bimodal with peaks about the values $(1+Q(0)) x$ and $T^{(0)}(x)$. Our density approximation formula is then

$$
\begin{equation*}
f_{T(x)}(t) \approx e^{-\beta(x)} \cdot \delta(t-(1+Q(0)) x)+\left(1-e^{-\beta(x)}\right) \cdot \frac{1}{\sqrt{2 \pi v(x)}} e^{-(t-m(x))^{2} /(2 \sqrt{v(x)})} \tag{5.2}
\end{equation*}
$$

where $\delta(\cdot)$ is the delta function, $m(x) \equiv T^{(0)}(x), v(x) \equiv \operatorname{Var}\left[T^{(1)}(x)\right]$ and

$$
\begin{equation*}
\beta(x) \equiv \int_{0}^{(1+Q(0)) x} \lambda(s) d s+Q(0) \int_{0}^{x} \mu(s) d s \tag{5.3}
\end{equation*}
$$

For the constant rate case, we have $\beta(x)=(\lambda \cdot(Q(0)+1)+\mu \cdot Q(0)) x$.
For the graphs of our numerical examples in Figure 4, we set $\mu=1, Q(0)=10.0$ and either $\lambda=2.0$ or $\lambda=0.5$. Given an exponentially distributed, service time with mean one, we have $x=0.10536$ for the 10 percentile job size, $x=0.69315$ for the 50 th percentile job size and $x=2.30259$ for the 90 th percentile job size. The values for $\beta(x)$ are given by Table 1.

In Figure 4 the simulated densities for the 10,50 and 90 -th percentile job sizes are compared with their respective normal approximations. The left column of Figure 4 displays this comparison for an overloaded $M / M / 1 / P S$ with arrival rate $\lambda=2$, while the right column considers an underloaded $M / M / 1 / P S$ with arrival rate $\lambda=0.5$.

The bimodal phenomenon discussed earlier in this section is readily apparent in the two graphs for the 10 percentile job size. Here, the point masses occur with probabilities 0.034 and 0.20. The bimodality is observable here because $\beta(x)$ is sufficiently small. Since $\beta$ increases linearly in $x$, the point mass probability decreases exponentially. Thus the bimodal behavior of the sojourn times of the 50 -th and 90 -th job size percentiles are not observable, since their point mass probabilities equal $2.33 \times 10^{-10}\left(50 \%\right.$-tile,left), $2.17 \times 10^{-5}(50 \%$-tile,right $)$, $9.83 \times 10^{-33}(90 \%$-tile,left $)$, and $3.13 \times 10^{-16}$ (90\%-tile,right) .


Figure 4: Empirical results. Density function for $M / M / 1 / P S$, restricted to jobs in (a) 10percentile, (b) 50-percentile, (c) 90-percentile. Left column shows overloaded case where $\lambda=2.0$ and $\mu=1.0$. Right column shows underloaded case where $\lambda=0.5$ and $\mu=1$. Here, we assume throughout that $Q(0)=10$.

## 6 Conclusion

We introduce the notion of a virtual customer for an approximate analysis of the sojourn time for a processor sharing queue. This creates a virtual job of a known size that is affected by the other jobs in the queue, but does not affect the response times of those other jobs.

We extend previous asymptotic results for the $M_{t} / M_{t} / 1 / P S$ queueing process to accommodate a non-zero, scaled initial load. These results can then be transformed into fluid and diffusion limits for the sojourn times. These fluid and diffusion limit, sojourn time formulas hold in general for any queueing process that has fluid and diffusion limits.

Our uniform acceleration scaling gives a simpler analysis of the mean and variance of the sojourn time, yet yields results that are a good approximation of the original stochastic model. We then obtain a time-varying analysis of the response time for systems that may experience alternating periods of underloading and overloading. Our numerical examples of overload behavior show that our approximations work well for a wide range of virtual job sizes. This type of behavior cannot be captured by steady-state models.

The density of the virtual response time is found to be well approximated by the convolution of a normal density and a point mass. Guided by simulation of the sojourn times of the $M_{t} / M_{t} / 1 / P S$ queue, we show that the sojourn time density may sometimes have a bimodal property. Even though we have convergence to the normal distribution, in practice the point mass contribution may not decay to zero quickly enough. In our numerical examples, we see when such a point mass term can contribute to our approximation of the density.

In future work we will extend this fluid and diffusion analysis to sojourn times for heterogeneous classes of customers with general job size (service) distributions, customer abandonment (having jobs "time out" after a given amount of time), and weighted processor sharing.

## References

[1] N. Bansal and M. Harchol-Balter. "Analysis of SRPT Scheduling: Investigating Unfairness." Proceedings of Association of Computing Machinery Sigmetrics 2001 Conference on Measurement and Modeling of Computer Systems. Cambridge, MA, 2001.
[2] T. Bonald and A. Proutiere. "On Performance Bounds for the Integration of Elastic and Adaptive Streaming Flows". Proceedings of Association of Computing Machinery Sigmetrics/Performance Joint International Conference on Measurement and Modeling of Computer Systems, 235-245, 2004.
[3] H. Chen, O. Kella, and G. Weiss. "Fluid Approximations for a Processor Sharing Queue". Queueing Systems: Theory and Applications, 27:99-125, 1997.
[4] E. G. Coffman, R. R. Muntz and H. Trotter. "Waiting Time Distribution for ProcessorSharing Systems." Journal of the Association of Computing Machinery, Vol 17, 123-130, 1970.
[5] F. Delcoigne, A. Proutiere and G. Regnie. "Modeling Integration of Streaming and Data Traffic". Performance Evaluation 55; 185-209; 2004.
[6] M. Harchol-Balter, B. Schroeder, N. Bansal and M. Agrawal. "Size-based Scheduling to Improve Web Performance." Association of Computing Machinery Transactions on Computer Systems, Vol. 21, No. 2, 207-233, May 2003.
[7] F. Guillemin and J. Boyer. "Analysis of the $M / M / 1$ Queue with Processor Sharing via Spectral Theory". Queueing Systems 39(4): 377-397; Dec 2001.
[8] A. Jean-Marie and P. Robert. "On the Transient Behavior of the Processor Sharing Queue." Queueing Systems 17, pp. 129-136, 1994.
[9] M. Yu. Kitaev. "The M/G/1 Processor-Sharing Model: Transient Behavior". Queueing Systems 14, pp. 239-273, 1993.
[10] L. Kleinrock. Queueing Systems, Volume II: Computer Applications, John Wiley \& Sons., 1976.
[11] J. D. C. Little. "A Proof of the Queueing Formula $L=\lambda W$ ", Operations Research, vol. 9, pp. 383-387, 1961.
[12] A. Mandelbaum and W. A. Massey. "Strong Approximations for Time Dependent Queues", Mathematics of Operations Research, 20:1, pp. 33-64, February 1995.
[13] W. A. Massey. "Asymptotic Analysis of the Time Dependent M/M/1 Queue", Mathematics of Operations Research, 10, pp. 305-327, May 1985.
[14] H. Masuyama and T. Takine. "Sojourn Time Distribution in a MAP/M/1 ProcessorSharing Queue." Operations Research Letters, vol. 31, no. 5, pp. 406-412, 2003.
[15] J. Morrison. "Response Time for a Processor-Sharing system". SIAM J. Appl. Math. 45(1), 152-167, 1985.
[16] R. Nunez-Queija. "Processor Sharing Models for Integrated Services Networks", Ph.D. Thesis, Eindhoven University of Technology, 2000.
[17] T. Ott "The Sojourn Time Distribution in the M/G/1 Queue with Processor Sharing", J. Appl. Prob.,pp. 360-378, 1984.
[18] J. W. Roberts. "Engineering for Quality of Service." K. Park, and W. Willinger (editors). Self-Similar Network Traffic and Performance Evaluation. Wiley, New York, pp. 401-420, 2000.
[19] R. Schassberger. "A New Approach to the M/G/1 Processor Sharing Queue", Adv. Appl. Prob., pp. 202-213, 1984.
[20] B. Schroeder and M. Harchol-Balter. "Web servers under overload: How scheduling can help." 18th International Teletraffic Congress. Berlin, Germany. September 2003.
[21] B. Sengupta and D. L. Jagerman. "A Conditional Response Time of the M/M/1 Processor-Sharing Queue." ATB T Technical Journal, Vol 64, pp. 409-421, 1985.
[22] S. F. Yashkov. "A Derivation of Response Time Distribution for a M/G/1 ProcessorSharing Queue." Problems of Control and Information Theory, Vol 12 (2), pp. 133-148, 1983.
[23] S. F. Yashkov. "Processor-Sharing Queues: Some Progress in Analysis", Queueing Systems, Vol 2, pp. 1-17, 1987.
[24] S. F. Yashkov. "Mathematical Problems in the Theory of Shared-Processor Systems", Journal of Soviet Mathematics, Vol 58, pp. 101-147, 1993.
[25] B. Zwart and O. J. Boxma. "Sojourn Time Asymptotics in the M/G/1 Processor Sharing queue" Queueing Systems, 35:141-166, 2000.

