Institut Eurécom
Networking and Security Department
2229, route des Crêtes
B.P. 193
06904 Sophia-Antipolis
FRANCE

Research Report RR-08-224

# On the Variance of the Least Attained Service Policy and its Use in Multiple Bottleneck Networks

May 15$^{th}$, 2008

Matthias Auchmann*, Guillaume Urvoy-Keller+
* Technische Universität Wien
+ Institut Eurecom

Email : m.auchmann@artech.at, Guillaume.Urvoy@eurecom.fr

[1]Institut Eurécom's research is partially supported by its industrial members: BMW Group Research & Technology - BMW Group Company, Bouygues Télécom, Cisco Systems, France Télécom, Hitachi Europe, SFR, Sharp, STMicroelectronics, Swisscom, Thales.

# On the Variance of the Least Attained Service Policy and its Use in Multiple Bottleneck Networks

Matthias Auchmann and Guillaume Urvoy-Keller

**Abstract**

Size-based scheduling has proved to be effective in a lot of scenarios involving Internet traffic. In this work, we focus on the Least Attained Service Policy, a popular size-based scheduling policy. We tackle two issues that have not received much attention so far. First, the variance of the conditional response time. We prove that the classification proposed by Wierman et al. [11], which classifies LAS has an always unpredictable policy, is overly pessimistic. We illustrate the latter by focusing on the M/M/1/LAS queue. Then we consider LAS queues in tandem. We provide preliminary results concerning the characterization of the output process of an M/M/1/LAS queue and the conditional average response time and its variance for LAS queues in tandem.

# Contents

## List of Figures

# 1 Introduction

Size-based scheduling has proved to be very effective in increasing performance in a lot of scenarios: Web servers [3], Internet traffic [8] or 3G networks [5]. The key idea behind size-based scheduling is to favor short jobs while ensuring that large jobs do not starve. The net result is better interactivity from the user point of view as short jobs correspond to interactive applications when considering Internet traffic. The extent to which large jobs suffer depends on the statistical characteristics of the job size distribution and especially how the mass is distributed among short and large jobs. Broadly speaking, the larger the mass carried by the large flows, the smaller the penalty since short flows, that have the highest priority, can not monopolize the server. Heavy-tailed distributions, which have often been observed in the Internet [2], feature such a property.

In this paper, we consider the Least Attained Service (LAS) policy, a.k.a the Foreground-Background policy [7]. LAS has been initially proposed and studied in the context of time-sharing computers in the late 60s [10]. Under LAS, priority is given to the job that has received the least amount of service. In case of ties, jobs share the server in a round-robin manner. A salient feature of LAS is that it has no internal parameter to tune.

Our focus in this work is twofold. First, we focus on the conditional variance of LAS. Few results exist concerning the variance of LAS [11, 1]. In [1], the asymptotic conditional variance is considered. In [11], the authors propose a classification of scheduling policies based on their variance. In particular, they propose to classify scheduling policies as: (i) always predictable, (ii) sometimes predictable or (iii) always unpredictable - precise definitions will be given in Section 2. LAS falls in the latter category, which seems at first sight disappointing. Indeed, LAS does a good job at providing low response time to small jobs but despite this nice property, the results in [11] seem to restrict the interest of LAS.
We revisit the variance of LAS by considering a worst case scenario : the M/M/1/LAS queue. We consider the M/M/1/LAS queue as a worst case scenario as (i) empirical distributions of flow size for Internet traffic have a much higher variability than the one of an exponential distribution[1] (ii) the performance of LAS and especially the fraction of flows that receive a better service under LAS than under PS is known to increase with the variability of the job size distribution [8].
Considering the M/M/1/LAS queue, we analytically bound the fraction of flows that are treated in a predictable manner. We obtain that at least 75% of flows are treated predictably, irrespective of the load. Numerical studies further demonstrate that the actual fraction of such flows should be closer to 95%.

Then, we focus on the problem of using LAS queues in tandem. The motivation behind this scenario is to determine the benefits that could be obtained with LAS in a multiple bottlenecks scenario. A typical example is wireless mesh networks

[1] Typical candidate distributions to model flow sizes are Pareto, lognormal or Weibul distribution with large coefficient of variations, e.g. [4].

based on the 802.11 protocol where the available bandwidth is known to be highly varying, as has been exemplified by the roofnet experiment (http://pdos.csail.mit.edu/roofnet/doku.php). In such a situation, LAS could be highly benefical as it allows to maintain a minimum level of interactivity, even when congestion is high. However, using LAS at multiple queues in tandem can also be detrimental to large flows that could be penalized multiple times.

To the best of our knowledge, no work has tackled the problem of studying LAS in tandem queues. We rely on numerical evalutions to address this issue as the analytical approach seems too complex at the moment, even for the case of M/M/1/LAS queues. We make the following contributions: (i) We demonstrate that while the departure process of an M/M/1/LAS queue is apparently a Poisson process, the LAS scheduling policy introduces a negative correlation between departure times and job sizes; (ii) We evalutate the impact of the above correlation on the conditional average response times and the conditional variance of the response time of LAS queues in tandem.

## 2 Conditional Variance of an M/M/1/LAS queue

The variance of the conditional response time for an $M/G/1/LAS$ queue is given by [12]:

$$Var[T(x)]^{LAS} = \frac{\lambda x \tilde{m}_2(x)}{(1-\tilde{\rho}(x))^3} + \frac{\lambda \tilde{m}_3(x)}{3(1-\tilde{\rho}(x))^3} + \frac{3}{4}\Big(\frac{\lambda \tilde{m}_2(x)}{(1-\tilde{\rho}(x))^2}\Big)^2 \qquad (2.1)$$

The $\tilde{m}_i(x)$ are the truncated moments, $\tilde{\rho}(x) = \lambda \tilde{m}_1(x)$ is the truncated rate. Since we focus on the exponential distribution we introduce the truncated moments for $Exp(\mu)$ :

$$\tilde{m}_i(x) = \int_0^x y^i \mu e^{-\mu y} dy + x^i e^{-\mu x}$$

In [11], scheduling policies were classified based on the variance of conditional response times as always predictable, always unpredictable or sometimes predictable. For a policy P, jobs of size $x$ are treated predictably if:

$$\frac{Var[T(x)]^P}{x} \leq \frac{\lambda m_2}{(1-\rho)^3} \qquad (2.2)$$

Otherwise jobs of size x are treated unpredictably. See [11] for a justification of the right side term in Eq. (2.2). More generally, a scheduling policy P is:

- **Always predictable** if it is predictable under all loads and service distributions;
- **Sometimes predictable** if it is predictable under some loads and service distributions, and unpredictable under others;

- **Always unpredictable** if it is unpredictable regardless of service distribution and load.

In [11] a variance bound was derived to show that LAS is always unpredictable. In contrast, PS is shown to be always predictable. Comparison of LAS to PS is important as it is shown in [9] that the M/G/1/LAS queue is an accurate model for a LAS router while the M/G/1/PS queue is an accurate model for a FIFO router.

Our initial motivation in this paper is to show that LAS outperforms PS not only in terms of conditional response time offered to a majority of short jobs [8], but also in terms of conditional variance. Figure 1 illustrates the relative performance of LAS and PS for an M/M/1 queue with $\lambda = 1$ and $\mu = 1.25$, i.e. a load $\rho = 0.8$. We observe that while LAS offers both low average and variance of response time for about 70% of the jobs[2], its performance becomes enventually worse as compared to PS for the largest jobs. Note that in a more realistic case of Internet traffic, LAS achieves better results as compared to PS due to the way the mass is distributed among short and large flows.

The large variance observed for the large jobs in Figure 1 illustrates why it is classified as always unpredictable, as it is due to the largest jobs only that Eq. (2.2) is violated. The question we address here is to determine the fraction of jobs that are treated predictably under LAS for the case of an M/M/1/LAS queue.

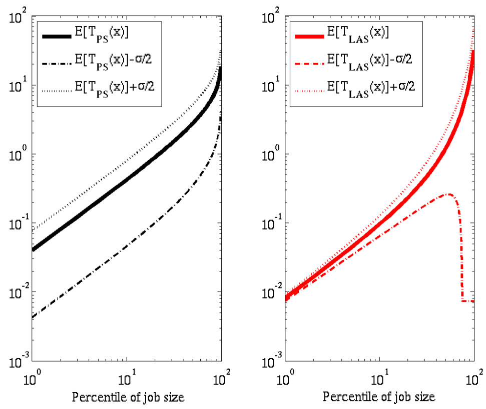


Figure 1: M/M/1/LAS against M/M/1/PS: mean and variance of response time

## 2.1 Preliminary Results

**Lemma 1:** $\tilde{m}_3(x) \leq \frac{6}{e^2} x \frac{2}{\mu^2}$

[2] We consider here the intersection between the curves $E[T_{PS}(x)] + \sigma/2$ and $E[T_{LAS}(x)] + \sigma/2$ as an increase in response time is more important than a decrease in response time

*Proof.* $\tilde{m}_3(x) = \int_0^x y^3 \mu e^{-\mu y} dy + x^3 e^{-\mu x} \overset{z=\mu y}{=} \int_0^{\mu x} (\frac{z}{\mu})^3 \mu e^{-z} \frac{dy}{\mu} + x^3 e^{-\mu x} = \frac{1}{\mu^3}\gamma(4, \mu x) + x^3 e^{-\mu x} = \frac{1}{\mu^3} 3! \Big(1 - e^{-\mu x} \sum_{k=0}^{3} \frac{(\mu x)^k}{k!}\Big) + x^3 e^{-\mu x} = \frac{1}{\mu^3}\Big(6 - 6e^{-\mu x}(1 + \mu x + \frac{1}{2}(\mu x)^2 + \frac{1}{6}(\mu x)^3\Big) + x^3 e^{-\mu x} \overset{?}{\leq} \frac{6}{e^2} x \frac{2}{\mu^2}$
At 0, both sides are 0. Deriving the previous inequation gives

$$3e^{-\mu x} x^2 \leq \frac{12}{\mu^2 e^2} \Leftrightarrow \underbrace{\mu^2 x^2 e^{-\mu x}}_{findmax} \leq \frac{4}{e^2} \tag{2.3}$$

Deriving and setting 0 we easily find the mode of the lefthand side being $\frac{2}{\mu}$. Verifying the equation by substituting the mode in equation (2.3) we get

$$\mu^2 (\frac{2}{\mu})^2 e^{-\mu \frac{2}{\mu}} = 4e^{-2} \tag{2.4}$$

This ends the proof. $\square$

**Lemma 2:** $\tilde{m}_2(x) \leq \frac{2}{\mu e} x$

*Proof.* Analogous to Lemma 1, after evaluating $\tilde{m}_2(x)$ we need to show

$$\frac{1}{\mu^2}\Big(2 - 2e^{-\mu x}(1 + \mu x + \frac{1}{2}(\mu x)^2)\Big) + x^2 e^{-\mu x} = \tag{2.5}$$

$$\frac{1}{\mu^2}\Big(2 - 2e^{-\mu x}(1 + \mu x)\Big) \leq \frac{2}{\mu e} x \tag{2.6}$$

Again, the inequality holds for 0, and we derive both sides:

$$\underbrace{2e^{-\mu x} x}_{findmax} \leq \frac{2}{\mu e} \tag{2.7}$$

The mode is easily shown to be $\frac{1}{\mu}$; substituting gives $\frac{2}{\mu e}$ as the maximum, thus ending the proof. $\square$

**Lemma 3:** $m_{i+1}(x) \leq x m_i(x)$

*Proof.* $m_{i+1}(x) = \int_0^x t^{i+1} f(t) dt \leq x \int_0^x t^i f(t) dt = x m_i(x)$ $\square$

**Lemma 4:** $\tilde{m}_{i+1}(x) \leq x \tilde{m}_i(x)$

*Proof.* $\tilde{m}_{i+1}(x) = m_{i+1}(x) + x^{i+1}\overline{F}(x) \le x(m_i(x) + x^i\overline{F}(x)) = x\tilde{m}_i(x)$ □

**Lemma 5:** $\tilde{m}_i(x) \le m_i$

*Proof.* $\tilde{m}_i(x) = \int_0^x y^i \mu e^{-\mu y} dy + x^i e^{-\mu x} \le \int_0^\infty y^i \mu e^{-\mu y} dy = m_i$ □

**Lemma 6:** $\tilde{\rho}(x) \le \rho$

*Proof.* $\tilde{\rho}(x) = \lambda \tilde{m}_1(x) \le \lambda m_1 = \rho$ □

## 2.2 M/M/1/LAS variance bounds

In this section, we first present two bounds on the variance of the conditional response time for a M/M/1/LAS queue. We next derive corresponding bounds on the fraction of jobs that are treated predictably under LAS.

**Bound 1:** For M/M/1/LAS,

$$Var[T(x)]^{LAS} \le \frac{\lambda x m_2}{(1-\tilde{\rho}(x))^4}\Big(1 + \frac{2}{e^2} - (\frac{1}{4} + \frac{2}{e^2})\tilde{\rho}(x)\Big) \tag{2.8}$$

*Proof.* Lemma 5 easily gets us the first fraction in (2.1):

$$\frac{\lambda x \tilde{m}_2(x)}{(1-\tilde{\rho}(x))^3} \le \frac{\lambda x m_2}{(1-\tilde{\rho}(x))^3} = \frac{\lambda x m_2}{(1-\tilde{\rho}(x))^4}(1-\tilde{\rho}(x)) \tag{2.9}$$

Applying Lemma 1 to the second fraction gives:

$$\frac{\lambda \tilde{m}_3(x)}{3(1-\tilde{\rho}(x))^3} \le \frac{\lambda \frac{6}{e^2} x m_2}{3(1-\tilde{\rho}(x))^3} = \frac{\lambda x m_2}{(1-\tilde{\rho}(x))^4}\Big(\frac{2}{e^2} - \frac{2}{e^2}\tilde{\rho}(x)\Big) \tag{2.10}$$

We will now bound the last term in (2.1) with Lemma 4 and 5:

$$\frac{3}{4}\Big(\frac{\lambda \tilde{m}_2(x)}{(1-\tilde{\rho}(x))^2}\Big)^2 \le \frac{3}{4}\frac{\lambda x \tilde{m}_1(x) \lambda m_2}{(1-\tilde{\rho}(x))^4} = \frac{\lambda x m_2}{(1-\tilde{\rho}(x))^4}\Big(\frac{3}{4}\tilde{\rho}(x)\Big) \tag{2.11}$$

$$\begin{aligned} Var[T(x)]^{LAS} &= \underbrace{\frac{\lambda x \tilde{m}_2(x)}{(1-\tilde{\rho}(x))^3}}_{(2.9)} + \underbrace{\frac{\lambda \tilde{m}_3(x)}{3(1-\tilde{\rho}(x))^3}}_{(2.10)} + \underbrace{\frac{3}{4}\Big(\frac{\lambda \tilde{m}_2(x)}{(1-\tilde{\rho}(x))^2}\Big)^2}_{(2.11)} \\ &\le \frac{\lambda x m_2}{(1-\tilde{\rho}(x))^4}\Big(1 + \frac{2}{e^2} - (\frac{1}{4} + \frac{2}{e^2})\tilde{\rho}(x)\Big) \end{aligned}$$

□

**Bound 2:** For M/M/1/LAS,

$$Var[T(x)]^{LAS} \le \frac{8e\rho + 9\rho^2 - 8e\rho^2}{3e^2(1-\rho)^4} x^2 \tag{2.12}$$

*Proof.* By applying Lemma 4 to the second term of (2.1), and Lemma 6 to the denominators, we get

$$Var[T(x)]^{LAS} \leq \frac{4}{3}\frac{\lambda x \tilde{m}_2(x)}{(1-\rho)^3} + \frac{3}{4}\Big(\frac{\lambda \tilde{m}_2(x)}{(1-\rho)^2}\Big)^2 \tag{2.13}$$

Now bounding $\tilde{m}_2(x)$ by $\frac{2}{e\mu}x$ (Lemma 2), and substituting $\frac{\lambda}{\mu}$ by $\rho$, rearranging the terms ends the proof. □

Note that Bound 1 and Bound 2 complement each other well in two ways. For Bound 1, $\tilde{\rho}(x)$ tends to $\rho$ as $x$ tends to $\infty$ and $m_2$ was used to bound $\tilde{m}_2$. Using $m_2$ as a bound is more accurate for large jobs. In addition the difference between the variance divided by $x$ and Bound 1 divided by $x$ (both expressions we will use in a moment) tend to 0 as $x$ tends to $\infty$. Thus Bound 1 should be more accurate for large values of $\rho$ and for larger jobs.

Bound 2 is more accurate for small jobs than Bound 1, because in Bound 2 we basically used a Taylor expansion around 0 to bound $\tilde{m}_2$; thus this estimate is more accurate for smaller job sizes.

With these two bounds we will now express the percentage of jobs which are treated predictably in a M/M/1/LAS queue. Note that for the case of an exponential service distribution, Equation ((2.2)) simplifies to:

$$\frac{Var[T(x)]^{LAS}}{x} \leq \frac{\lambda\frac{2}{\mu^2}}{(1-\rho)^3} = \frac{\frac{2\rho}{\mu}}{(1-\rho)^3} \tag{2.14}$$

**Bound 3:** For M/M/1/LAS and a given $\rho$, at least a fraction of $\frac{1-\left((1-\rho)^3(1+\frac{2}{e^2})\right)^{1/4}}{\rho}$ of the jobs is treated predictably

*Proof.* Dropping the negative $(\frac{1}{4}+\frac{2}{e^2})\tilde{\rho}(x)$ term in (2.8) and dividing by x, and then equating to (2.2) leaves

$$\frac{\lambda m_2}{(1-\tilde{\rho}(x))^4}(1+\frac{2}{e^2}) = \frac{\lambda m_2}{(1-\rho)^3} \tag{2.15}$$

for us to solve. Cancelling out $\lambda$ and $m_2$ and rearranging the terms gives

$$\tilde{\rho}(x) = 1 - \left((1-\rho)^3(1+\frac{2}{e^2})\right)^{1/4} \tag{2.16}$$

Since $\tilde{\rho}(x) = \lambda\frac{1-e^{-\mu x}}{\mu} = \rho(1-e^{-\mu x})$, and $1-e^{-\mu x}$ is exactly the wanted percentage, dividing by $\rho$ completes the proof. □

**Bound 4:** For M/M/1/LAS and a given $\rho$, at least a fraction of $1-e^{-\frac{6e^2(1-\rho)}{8e+9\rho-8e\rho}}$ of the jobs is treated predictably

*Proof.* Solving for x in

$$\frac{Var[T(x)]^{LAS}}{x} \overset{(2.12)}{\leq} \frac{8e\rho + 9\rho^2 - 8e\rho^2}{3e^2(1-\rho)^4} x = \frac{2\rho}{\mu(1-\rho)^3} \tag{2.17}$$

gives

$$x = \frac{6e^2(1-\rho)}{\mu(8e + 9\rho - 8e\rho)} \tag{2.18}$$

Substituting this into the distribution function $1 - e^{-\mu x}$ ends the proof. □

Note that bounds 3 and 4 depend on $\rho$ only, not on $\lambda$ and $\mu$. Bounds 3 and 4 show that the majority of jobs are in fact treated predictably in a $M/M/1/LAS$ queue, as can be seen in Figure 2 where we plot the maximum of the two bounds (solid line). The minimum job percentage that is guaranteed to be treated predictably in this case, is about 75 percent.

As the above approach relies on lower bounding the number of jobs that are treated unpredictably under LAS, we further evaluated using simulations the percentage of jobs treated predictably. Figure 2 reports the results obtained with a simulator written in Matlab. Each simulation involved more than 500,000 jobs, and 4 different factorizations of $\rho$ by using four different $(\lambda, \mu)$ pairs, which is why there are 4 points for each $\rho$ value. We observe from Figure 2 that the fraction of flows that are treated pedictably seems to depend on $\rho$ only, not on the exact factorization of $\rho$. In addition, one can see that the actual percentage of jobs being treated predictably is very high: at least 90 percent of the jobs are treated predictably, regardless of the actual load value.

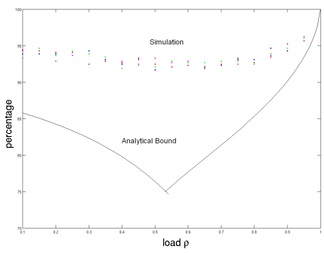


Figure 2: Bounds 3 and 4 (solid line) and simulations (dots)

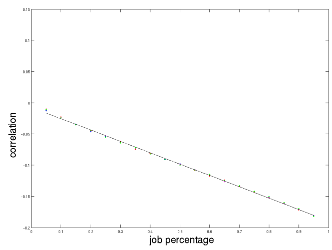


Figure 3: Correlation of the two interdeparture times and jobsize for LAS

# 3 LAS in tandem queues

In this section, we focus on LAS in tandem queues where fresh arrivals at each queue follow a Poisson process and service requirements follow an exponential distribution. We first study the output process of an isolated LAS queue and then, we will discuss the impact of the characteristics of the output process to the total response time of two LAS queues in tandem. Note that, to the best of our knowledge, no work has tackled this issue so far.

## 3.1 Characterization of the output process of an M/M/1/LAS queue

### 3.1.1 Interdeparture time

First one might ask the question whether for LAS queueing disciplines Burke's theorem [6] remains valid. Burke's theorem states that if a FIFO single-server queuing process with exponential interarrival and service-time distributions is stationary, then it follows that the departure process is a homogeneous Poisson process with rate equal to the reciprocal of the mean interarrival time.

To check if the departure process of an M/M/1/LAS queue can be considered a Poisson process, we used the Kolmogorov-Smirnov statistical test to assess if the inter-departure times follow an exponential distribution. We also checked that inter-arrival times are uncorrelated. In Table 1, we report p-values for the KS test and the correlation coefficient for the interdeparture times. The factorization of $\rho$ is done with four different $(\lambda, \mu)$ pairs, and each time we use 5 runs with about 500,000 jobs each and average the results. We do conclude that the output process of an M/M/1/LAS queue is apparently a Poisson process.

| $\rho$ | Test agains Exp. | | | | Correlation Coefficient | | | |
|---|---|---|---|---|---|---|---|---|
| 0.05 | 0.7445 | 0.5469 | 0.5447 | 0.3876 | -0.0122 | 0.0076 | -0.0122 | 0.0006 |
| 0.15 | 0.5145 | 0.4331 | 0.6281 | 0.4594 | -0.0000 | -0.0013 | 0.0115 | 0.0035 |
| 0.25 | 0.6818 | 0.5520 | 0.6039 | 0.4695 | 0.0040 | 0.0053 | 0.0059 | 0.0045 |
| 0.35 | 0.5615 | 0.6113 | 0.5978 | 0.5508 | 0.0070 | 0.0134 | -0.0092 | -0.0026 |
| 0.45 | 0.5687 | 0.3485 | 0.7266 | 0.4194 | -0.0046 | 0.0097 | 0.0050 | 0.0013 |
| 0.55 | 0.6091 | 0.3382 | 0.6008 | 0.3410 | 0.0144 | 0.0065 | 0.0004 | -0.0057 |
| 0.65 | 0.3549 | 0.4652 | 0.6352 | 0.4439 | -0.0016 | 0.0020 | -0.0014 | 0.0060 |
| 0.75 | 0.2201 | 0.3547 | 0.4259 | 0.4020 | 0.0038 | -0.0080 | -0.0148 | -0.0075 |
| 0.85 | 0.7515 | 0.4955 | 0.4351 | 0.2496 | 0.0012 | -0.0066 | -0.0076 | 0.0054 |
| 0.95 | 0.5242 | 0.8156 | 0.4396 | 0.4803 | -0.0031 | -0.0086 | -0.0108 | 0.0037 |

Table 1: Interdeparture Time p-values for test against the exponential distribution and correlation coefficients.

### 3.1.2 Correlation of Interdeparture Time and Job size

To further characterize the output process of an M/M/1/LAS queue, we investigate whether the job size is independent of the interdeparture time. Figure 3 shows the correlation between the job sizes and the interdeparture times for LAS, plotted for different utilizations $\rho$. Again, different $(\lambda, \mu)$ pairs are used for the plots, but since curves overlap, the correlation seems to depend on $\rho$ only. We observe a negative correlation between inter-departure time and jobs sizes, which increases for increasing load. We believe that this correlation stems from the fact that as load increases, large flows are interrupted more frequently by shorter jobs. As a consequence, their remaining service requirement can reach very low values, lower than the service requirements of short jobs. When a large job leaves the queue, it is highly likely to be in a short period of time where the load is low. During this period, many large jobs that were stuck in the queue, with small remaining service times, leave the system. We believe this explains the negative correlation we observe.

Note that such a correlation does not seem to be present for the PS queueing discipline.

### 3.2 Impact on tandem queue performance

The negative correlation observed above relates to the behavior of LAS that tends to sort jobs in the queue in ascending order, which means that short jobs tend to leave the LAS queue in group and the same for large jobs. When considering the case of a tandem queue, this sorting operated by the scheduler can have a detrimental impact on the performance of large flows. Indeed, when the large flows that enter the tandem queue at the first queue reach the second queue, they again have to compete with the large flows with whom they were in competition in the first queue.

We quantitatively evaluated the previous intuition by comparing the conditional average response time and its variance for two LAS queues in tandem with the sum of response time for the same system where we "re-draw" the job sizes of the job leaving the first queue using the initial distribution, thus wiping out the correlation observed in the previous section. To enable a meaningfull comparison, we first present the case of the PS discipline, for which the correlation between job sizes and inter-departure times is negligible. Then we contrast with the results for LAS.

Figures 4 and 5 illustrates the above scenario for the case where the load is 0.8. In the first row you see the actual measured response time of the tandem system versus the added theoretical values; in the first column you see the absolute values, in the second column the relative error of the actual measurement.

The second row plots twice the theoretical variance, four times the theoretical variance and the actual variance in the first column. This is done since for two random variables $X_1$, $X_2$,
$Var(X_1 + X_2) = Var(X_1) + 2\,Cov(X_1, X_2) + Var(X_2)$
and
$Cov(X_1, X_2) = \rho\sqrt{Var(X_1)}\sqrt{Var(X_2)} \leq \rho\, max(Var(X_1), Var(X_2))$.
Thus all measurements between twice the variance and four times the variance could be due to covariance, and the variance formulas for the second queue could still be valid.
To get a better picture of which proportion of the overall variance is due to covari-

ance, the third row shows theoretical variance for one queue and the measurements for the first and the second queue alone, as above as absolute and relative plots in the first and second column, respectively.

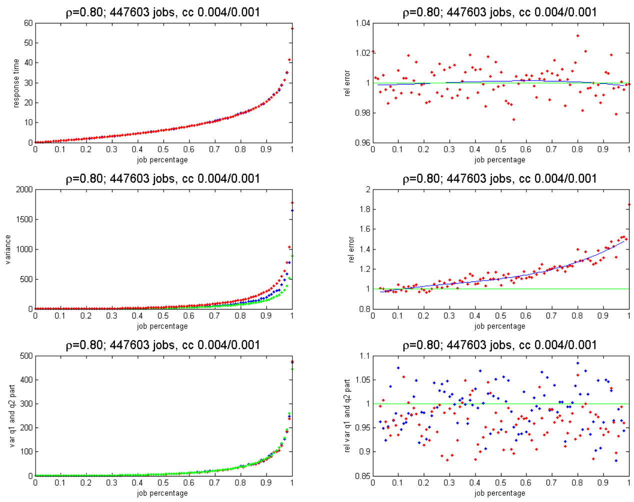


Figure 4: Tandem Queue simulation results for PS, $\rho = 0.8$

Note that the response times of the two queues do add nicely here, as it would be expected for classical, independent of each other queues. The overall variance is 2 $Var(X)$ for small jobs and goes up until $4\ Var(X)$ for the biggest jobs and is due to $Cov(X_1, X_2)$, since the variances of the two queues alone are perfectly in line with the theoretical value for one queue only. So for PS, the second queue behaves like the first queue in terms of mean and variance even with the output process of the first queue as its input process. The covariance increases for big jobs, thus increasing the variance of the overall system.
As we shall now see in Figure5, things are different for LAS.

One can observe from Figure 5 that the fact that the output process of the first queue is significantly different from the input process has a strong effect on the plots. While the response time is even better than the reference model for smaller sized jobs, LAS further penalizes big jobs in the tandem system. This is also true for variance, the variance of the second queue is up to three times higher than in formula (2.1) (with $G = M$).

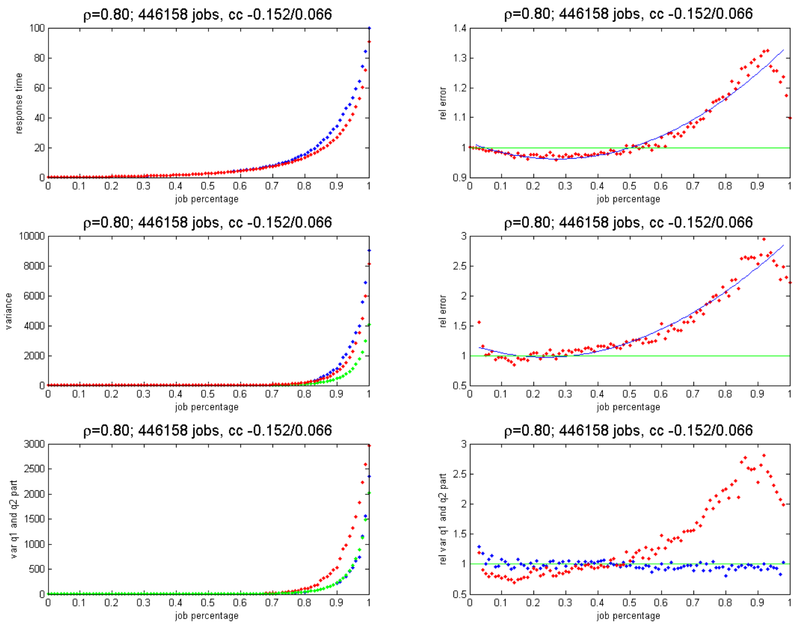


Figure 5: Tandem Queue simulation results for LAS, $\rho = 0.80$

# 4 Conclusions

In this paper, we have focused both one the variance of the conditional response time of an M/M/1/LAS queue and the use of LAS in tandem queues.

For the variance, we proved that the fraction of jobs that are treated predictably is very high even in the unfavorable case of exponential service requirements. The practical implication of this result is that it is likely that when LAS is used for real Internet traffic, the variance in response time it offers to the majority of the flows be small as compared to the legacy FIFO scheduling policy (to be modeled by a PS queue), which is the default scheduling policy in network equiments, ranging from routers to access points.

Concerning the use of LAS in tandem queues, we demonstrated that while the output process of an M/M/1/LAS queue is apparently Poisson, LAS tends to group together jobs of similar sizes, which results in a high penalty for the large jobs that cross the two queues. Practical implications of this result are less clear than for the variance case. Indeed routers work on a packet basis and do not output all the packets of a connection simultaneously as it happens in an M/M/1/LAS queue. This should dampen the effect we have observed.

As future work, we plan to continue working on the use of LAS in wireless mesh networks to precisely assess its performance. Note that a typical path in a

WMN from an end user to an Internet gateway should be quite small, no more than 4 hops. As a consequence, we can expect that the nice properties of LAS, namely its ability to maintain interactivity despite adversal load conditions, outweighs its side effect, namely the penalty experienced by large flows crossing multiple bottlenecks.

# References


[1] P. Brown. Comparing fb and ps scheduling policies. Technical report, Orange Labs, 2008.

[2] M. E. Crovella et al. *A Practical Guide to Heavy Tails*, chapter 3. Chapman and Hall, New-York, 1998.

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

[4] F. Hernandez-Campos, M. Karaliopoulos, M. Papadopouli, and H. Shen. Spatio-temporal modeling of traffic workload in a campus wlan. In *WICON '06*, page 1, 2006.

[5] M. Hu, J. Zhang, and J. Sadowsky. Size-aided opportunistic scheduling in wireless networks. In *GLOBECOM '03. IEEE*, 2003.

[6] L. Kleinrock. *Queuing Systems, Volume II: Computer Applications*. Wiley, New York, 1976.

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

[8] I. A. Rai, G. Urvoy-Keller, and E. W. Biersack. Analysis of las scheduling for job size distributions with high variance. In *Proc. ACM SIGMETRICS*, pages 218–228, June 2003.

[9] I. A. Rai, G. Urvoy-Keller, M. Vernon, and E. W. Biersack. Performance models for las-based scheduling disciplines in a packet switched network. In *ACM SIGMETRICS-Performance*, June 2004.

[10] L. E. Schrage. The queue m/g/1 with feedback to lower priority queues. *Management Science*, 13(7):466–474, 1967.

[11] A. Wierman and M. Harchol-Balter. Classifying scheduling policies with respect to higher moments of conditional response time. *ACM SIGMETRICS Performance Evaluation Review*, 33(1):229–240, June 2005.

[12] S. Yashkov. Processor-sharing queues: some progress in analysis. *Queueing Systems: Theory and Applications*, 2(1):1–17, June 1987.