UDC 519.872, 519.217

Bounding Moments of Sojourn Time in $M/G/1$ FCFS Queue with Inaccurate Job Size Information and Additive Error: Some Observations from Numerical Experiments

Tatiana A. Milovanova[*],
Lusine A. Meykhanadzhyan[†], **Rostislav V. Razumchik**[‡*]

[*] *Department of Applied Probability and Informatics*
Peoples' Friendship University of Russia (RUDN University)
6 Miklukho-Maklaya str., Moscow, 117198, Russian Federation
[†] *Department of Data Analysis, Decision-Making and Financial Technology*
Financial University under the Government of the Russian Federation
49 Leningradsky Prospekt, Moscow, 125993, Russian Federation
[‡] *Institute of Informatics Problems of the Federal Research Center "Computer Science and Control" of the Russian Academy of Sciences*
44-2 Vavilova Str., Moscow, 119333, Russian Federation

Email: milovanova_ta@rudn.university,lamejkhanadzhyan@fa.ru,rrazumchik@ipiran.ru,razumchik_rv@rudn.university

In this paper consideration is given to the $M/G/1$ FCFS system in which the service time distribution is not fully known. There is an unknown "theoretical distribution" of the actual service times. But those service times which are known, follow a different distribution, obtained by adding theoretical services times with the error term drawn from a left-truncated normal distribution. The goal is to derive bounds on the response time in the $M/G/1$ FCFS queue with the unknown "theoretical distribution" that are better than simply using the known service time distribution. In [1] it is shown that in the case when theoretical service times are multiplied by a log-normally distributed error the $M/G/1$ LCFS queue with resampling gives better upper bounds on the mean response time in the $M/G/1$ FCFS queue with the "theoretical distribution" of the service times. Here we present some numerical results, which show that once the error becomes additive the result of [1] is not valid any more. In the calculations it was assumed that the unknown "theoretical distribution" is left-truncated Weibull. The new modification of the LCFS resampling policy is suggested, which leads to the lower bounds for the unknown first and second moments of the response time in the $M/G/1$ FCFS system with the "theoretical distribution" of the service times. Behaviour of mean response times under other service policies is briefly discussed.

Key words and phrases: left-truncated Weibull distribution, inaccurate service time, additive error, additive noise, queueing system, mean sojourn time, sojourn time variance.

In: K. E. Samouylov, L. A. Sevastianov, D. S. Kulyabov (eds.): Selected Papers of the 1st Workshop (Summer Session) in the framework of the Conference "Information and Telecommunication Technologies and Mathematical Modeling of High-Tech Systems", Tampere, Finland, 20–23 August, 2018, published at http://ceur-ws.org

1. Introduction

In this paper a step forward is made in the problem of analyzing the queues with the inaccurate job size information. We understand the inaccurate job size information as described in [2]. Assume that at a top-level some technical system (e.g. a data-intensive execution engine) may be modelled by a $M/G/1$ queue. Suppose a job arrives at the system. Upon arrival its (future) service time, say $\hat{S}$, is sampled from the known distribution, say $\hat{B}(x)$, and the value of $\hat{S}$ is used for scheduling the job. After the job has received service, it turned out that its service time was $S \neq \hat{S}$. The same happens with the next job etc. It means that the (true) service time is sampled from another distribution, say $B(x)$, which is unknown to the scheduler. Thus the system's performance characteristics based on $\hat{B}(x)$ are biased and may differ significantly from those which are based on $B(x)$. The general question is: is it possible to tune the mathematical model so as to reduce the bias? In [1] it is shown that such tuning is possible in $M/G/1$ Processor Sharing (PS) queue with $\hat{S} = S \times X$ with $\hat{S}$ being long-tailed, S and X being independent, X having a log-normal distribution with with parameters 0 and σ. Under these conditions the mean response time in the $M/G/1$ LCFS queue with resampling may be a good upper bound[1] for the mean response time in the $M/G/1$ PS queue with service times drawn from $B(x)$. Since it is more common in the mathematical models of various phenomena that an error term X (usually normally distributed) enters the model in the additive (not multiplicative) way, we find it important to understand whether the results of [1] still hold in such case. The main contribution of this paper is the observation (made from the series of numerical experiments) that the results of [1] do not hold any more once the additive error (instead of multiplicative) is assumed. The PS (and the preemptive LCFS) policy filters out the additive error and thus the mean service time in the $M/G/1$ PS queue based on $\hat{B}(x)$ coincides with the mean service time based on $B(x)$. Yet for other policies (like FCFS, LCFS, RANDOM) under which the mean sojourn time depends on higher moments of the service time distribution (or its other characteristics) the situation is more complex. And for such cases some meaningful results for the additive model can be gained by developing the ideas from [1] (see sections 2 and 3 for the details). The end of this section is devoted to the discussion of the additive model.

The major drawback of the additive model is that in its basic form i.e. when the inaccurate job size $\hat{S}$ is modelled by $\hat{S} = S + X$, where S and X are independent, X is normally distributed with mean μ and standard deviation σ, it is meaningless. Indeed, the true service times S are always positive and since X is defined on the real line, $\hat{S}$ may become negative, which never happens in practice. The most straightforward way out is to assume truncated distributions for S and X. The use of truncated distributions for S conforms the experimental data (see, for example, data for the actual service times of Mapreduce jobs from Facebook in [3, 4]). But we are unaware of practical use-cases of the truncated normal (or other) distribution for an additive error X in the service times (in the queueing-theory context). Thus we assume[2] that the components of the inaccurate job size $\hat{S}$ are independent random variables $S > a$ and $X > b$, for some constants $a > 0$ and $b < 0$, and $\hat{S} = S + X$. We also assume that S has the left-truncated Weibull distribution with parameters k and α and its probability density $p_S(x)$ has the form

$$p_S(x) = \frac{k}{\alpha}\left(\frac{x}{\alpha}\right)^{k-1} e^{-\left(\frac{x}{\alpha}\right)^k + \left(\frac{a}{\alpha}\right)^k}, \ x > a, \ k > 0, \ \alpha > 0. \quad (1)$$

[1] It is always good if S is exponentially distributed.

[2] It is worth noticing here that, on the contrary, for the multiplicative model there are statistical evidences from practice. See [2].

The mean $\mathbf{E}S$ and the variance $Var\, S$ can be derived in a straightforward manner (for the details one can refer, for example, to [5]). The probability density $p_X(x)$ of the left-truncated normal distribution is

$$p_X(x) = \frac{\frac{1}{\sigma}\phi\left(\frac{x-\mu}{\sigma}\right)}{1-\Phi\left(\frac{b-\mu}{\sigma}\right)}, x > b, \sigma > 0, \tag{2}$$

where $\phi(x)$ and $\Phi(x)$ are the probability density and the cumulative function of the standard normal distribution correspondingly. The closed-form expressions for the mean $\mathbf{E}X$ and the variance $Var\, X$ are well-known and are thus omitted.

Estimation of the unknown parameters a, b, k, α, μ and σ from the given trace may not be trivial (or not possible at all). The constants a and b, which enter the additive model, make the situation more complicated. This is due to the fact that in practice from the available data it is only possible to estimate the value of $a + b$ (which is the minimal value in the data set), but not the individual values of a and b. Finally, in the additive models it is commonly assumed that no systematic errors are present, which means that $\mathbf{E}\hat{S} = \mathbf{E}(S + X) = \mathbf{E}S$ i.e. $\mathbf{E}X \equiv 0$. Thus the parameters μ, σ and b must satisfy the equation

$$\mu + \sigma\frac{\phi\left(\frac{b-\mu}{\sigma}\right)}{1-\Phi\left(\frac{b-\mu}{\sigma}\right)} = 0.$$

For example, for $b = -2$, $\sigma = 1$ the value of μ is -0.0627. But it can be shown that this equation may not have a solution at all (see remarks in the section 4).

2. Problem statement

Let us consider two queueing systems operating independently in parallel. The first (denote it system I) is the $M/G/1$ FCFS queue with the arrival rate λ and the service times distributed as S with the probability density (1). Denote the mean sojourn time in this system by v_S. The second (denote it system II) is the $M/G/1$ FCFS queue with the same arrival rate λ but with the service times distributed as $\hat{S} = S + X$, where the density of X is given by (2) and the parameters μ, σ and b are such that $\mathbf{E}X = 0$. Denote the mean sojourn time in this system by $v_{\hat{S}}$. For any $\sigma > 0$ it holds that

$$\mathbf{E}\hat{S} = \mathbf{E}S + \mathbf{E}X = \mathbf{E}S, \quad \mathbf{E}\hat{S}^2 = \mathbf{E}S^2 + \mathbf{E}X^2.$$

From the Pollaczek-Khintchine formula it follows that $v_S < v_{\hat{S}}$ for any value of λ such that $0 < \lambda < (\mathbf{E}S)^{-1}$. One of the questions we are interested in is the following: is it possible to tune the mathematical model of system II in such a way that it gives the value of the mean sojourn time, say v^*, satisfying $v_S < v^* < v_{\hat{S}}$?

In [7–9] there was introduced the service policy — LCFS with resampling — which, when applied to system II, may give such value of v^*. But only if the error is multiplicative and the service time distribution is long-tailed. With the additive error, as it is shown by the numerical examples in the next section, this result does not hold any more[3] . Yet the LCFS policy with resampling may be useful for finding the lower bound for v_S. If we alter it in such a way that the resampling customer occupies a place in the queue with a certain probability[4] , say θ, then by making the value of θ dependent on the service time distribution $\hat{B}(x)$ (or only its moments), the mean sojourn time $v^*(\theta)$ in the $M/G/1$ LCFS with θ-resampling may provide a lower bound for v_S. The numerical examples in the next section give the first impressions about this effect.

[3] And we don't have any suggestion for a modification which leads to an upper bound better than $v_{\hat{S}}$.

[4] Note that if $\theta = 1$ then θ-resampling policy is the simple resampling policy.

3. Some observations from the numerical experiments

Even though some analytic analysis of the interplay between v_S, $v_{\hat{S}}$ and $v^*(\theta)$ is possible under the assumptions made above, we will restrict ourselves here only to some observations made from the numerical experiments.

Let us assume that the distributions of S and X (and thus $\hat{S}$) are known and given by (1) and (2). Then the numerical computation of v_S, $v_{\hat{S}}$ and $v^*(\theta)$ is possible. We consider the following three cases for the distribution of S: (i) S has a constant hazard rate, (ii) S has a decreasing hazard rate, (iii) S has an increasing hazard rate. The basic data for the distributions of S and X are given in the Table 1. The only idea for picking up such values was to keep the variance of X comparable with the variance of S; in the rest the values of the parameters were chosen arbitrarily[5] .

Table 1

Basic data of the distributions of S, X and $\hat{S}$.

	Case (i)	Case (ii)	Case (iii)
X	$b = -20.358$, $\mu = -10$, $\sigma = 20$		
$\mathbf{E}X$	0		
$VarX$	196.417		
S	$a = 20.358$		
	$k = 1$, $\alpha = 10$	$k = 0.762$, $\alpha = 5$	$k = 1.263$, $\alpha = 15$
	hazard rate const.	hazard rate $\downarrow$	hazard rate $\uparrow$
$\mathbf{E}S$	30.358	30.358	30.358
$VarS$	100	115.241	87.473
$\hat{S}$	hazard rate $\uparrow$	hazard rate $\uparrow\downarrow$	hazard rate $\uparrow$
$\mathbf{E}\hat{S}$	30.358	30.358	30.358
$Var\hat{S}$	296.417	311.654	283.890

In the sub-figures of the Figure 1 one can see the values of v_S, $v_{\hat{S}}$ and $v^*(\theta)$ as the functions of the arrival rate λ for the FCFS policy. The uppermost sub-figure corresponds to the case (i) in the Table 1; the midmost — to the case (ii), the bottommost — to the case (iii). Since the stability condition[6] for the $M/G/1$ FCFS queue is $\lambda\mathbf{E}\hat{S} = \lambda\mathbf{E}S < 1$, then $\lambda = (\mathbf{E}X)^{-1} = 0.033$ is the asymptote of v_S and $v_{\hat{S}}$. Each sub-figure contains three graphs for $v^*(\theta)$ corresponding to the following values of θ: $\theta_1 = 1$, $\theta_2 = 0.5$, $\theta_3 = 1 - e^{-\lambda\sqrt{\mathbf{E}\hat{S}^2}}$.

The most important observation, which follows from the Figure 1, is the following. Irrespective of the tail of the distribution of $\hat{S}$, the simple resampling policy (the curve of $v^*(\theta_1)$) does not provide better than $v_{\hat{S}}$ estimate of v_S across all possible values of the load. This empirical fact is in the sharp contrast with the results for the multiplicative model: as shown in [1,6], with the multiplicative model if Z has a long-tail distribution, the resampling policy may provide better than $v_{\hat{S}}$ estimates for v_S (even though both S and X are unknown!). Surprisingly, even though the simple resampling policy provides a useless upper bound for the unknown mean sojourn time v_S, it can be modified in such a way that it provides the lower bound for the $v_{\hat{S}}$. This can be achieved by varying

[5] And $a = -b$ was chosen for simplicity.

[6] For the resampling queue the stability condition is different: $\beta(\lambda) > \theta(1+\theta)^{-1}$.

the value of θ. As can be seen from the Figure 1, the lower bound may be quite tight and may hold across the whole (or at least meaningful) range of load.

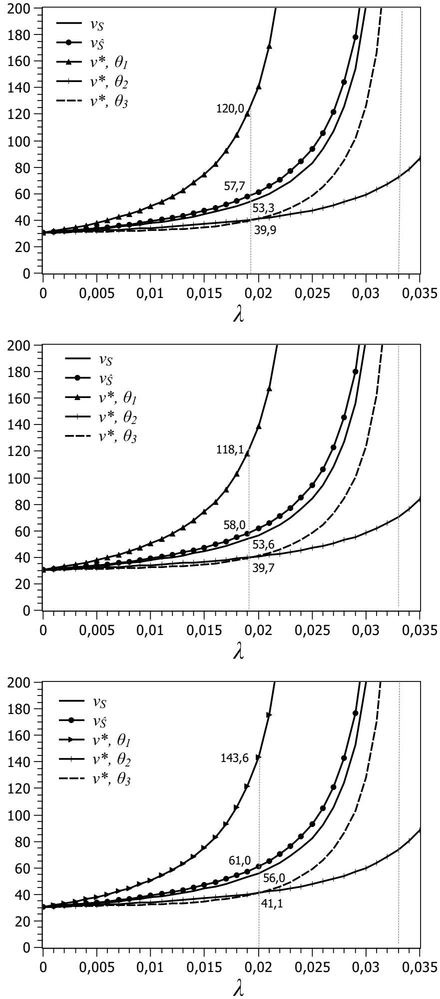

Figure 1. The values of v_S and $v_{\hat S}$ for the FCFS policy and $v^*(\theta)$ for the resampling policy versus the arrival rate. Here $\theta_1 = 1$, $\theta_2 = 0.5$ and $\theta_3 = 1 - e^{-\lambda\sqrt{\mathbf{E}\hat S^2}}$. Uppermost figure — case (i), midmost figure — case (ii), bottommost figure — case (iii).

It is worth also mentioning that if we switch from the FCFS policy to the other policy (like Shortest-Job-First or SRPT) under which the mean sojourn time depends not only on $\mathbf{E}\hat{S}$ but on some other characteristics of the service time distribution, the situation becomes more complicated and requires a special, delicate treatment[7] .

Coming back to the FCFS case, just for the illustrations purposes, let us briefly consider one real-life scenario. Specifically let us estimate the values of the parameters a, b, k, α, μ and σ based on one of the workloads generated by the tool SWIM (see [10]), which is said to be used (see [11, 12]) to test MapReduce systems. Specifically we use the trace FB09-1 from Facebook[8] , which contains a bunch of data on 6638 jobs. Since we are interested in the job's processing times only we used the methodology from [3, Section 2.2] to combine the available data to obtain the processing times (further denoted by t_i). The basic statistics for this trace are: the minimal processing time is $t_{min} = 1.3738 \times 10^{-8}$, the sample mean processing time is $t_{avg} \approx 11.8$, the sample second moment is $t_{sec.\,mom.} \approx 13574.8$.

The direct way to obtain the values of a, b, k, α, μ and σ would be to equate the first four theoretical moments $\mathbf{E}\hat{S}$, $\mathbf{E}\hat{S}^2$, $\mathbf{E}\hat{S}^3$ and $\mathbf{E}\hat{S}^4$ to the sample moments and solve the system under the natural constraints $a > 0$, $-a < b < 0$, $k > 0$, $\alpha > 0$, $\mu < 0$ and $\sigma > 0$. Since there are 6 unknowns, the other two equations are: $a + b = t_{min}$, $\mathbf{E}X = 0$. The problem is that this system may not have a solution and it is so with the trace[9] FB09-1. Thus in order to obtain some estimates we artificially limit the number of unknowns. Firstly we fix the value of $b = -2$ which immediately gives us also the value of $a = 2 + t_{min} = 2 + 1.3738 \times 10^{-8}$. Next we fix the value of $\sigma = 5$ and from the equation $\mathbf{E}X = 0$ we find the value of $\mu = -10.89427$. Finally by solving (numerically) the system of two equation $\mathbf{E}\hat{S} = t_{avg}$, $\mathbf{E}\hat{S}^2 = t_{sec.\,mom.}$, we find the values of $k \approx 0.07639$ and $\alpha \approx 8.1147477 \times 10^{-15}$. So the probability density functions of the true service time S and the error X have the form:

$$p_S(x) = 262511.47x^{-0.92361}e^{-11.92313\,x^{0.07639}}, \; x > 2.000000014,$$

$$p_X(x) = \frac{\phi(0.2x + 2.17895802)}{5\,(1 - \Phi(1.77895802))}, \; x > -2.$$

Numerical computations show that for these S and X the picture is qualitatively the same as in the Figure 1, case (ii).

Conclusion

The empirical observations show that in a case of inaccurate job size information, when the error is additive the simple resampling policy (i.e. when $\theta = 1$) does not lead to better estimates of the unknown mean sojourn time v_S. Yet there is a modification of the simple resampling policy (for example, with $\theta = \theta_3 = 1 - e^{-\lambda\sqrt{\mathbf{E}\hat{S}^2}}$) which may lead to quite tight lower bounds for the v_S. The usefulness of this empirical fact can be seen from the following example. According to the Figure 1 in the $M/G/1$ FCFS queue with the arrival rate λ and the service time $\hat{S} = S + X$ (where S and X are independent, given by (1) and (2); S is the true service time, X is the error) the true sojourn time v_S satisfies the inequality $v^*(\theta_3) < v_S$ for all $0 < \lambda < (ES)^{-1}$. Thus

[7] It may happen, for example, that $v_{\hat{S}} < v_S$ even though $\mathbf{E}\hat{S} = \mathbf{E}S$ and $\mathbf{E}\hat{S}^2 > \mathbf{E}S^2$.

[8] The trace can be downloaded from `https://raw.github.com/SWIMProjectUCB/SWIM/master/workloadSuite/FB-2009_samples_24_times_1hr_1.tsv`.

[9] Indeed, due to the given values of t_{min} and t_{avg} the value $b \approx 0$ and for $\sigma > 1$ we have that $b/\sigma \approx 0$. Thus, by disregarding the term b/σ and by introducing the new variable $c = \mu/\sigma$, the equation $\mathbf{E}X = 0$ can be rewritten as $c + \phi(-c)(1 - \Phi(-c))^{-1} = 0$. It does not have a solution in $(-\infty, 0)$.

from the Pollaczek-Khintchine formula for v_S we get the following lower bound for the (unknown!) second moment of S:

$$\mathbf{E}S^2 \geq \max\left(0, \frac{2(1-\lambda\mathbf{E}S)}{\lambda}(v^*(\theta_3) - \mathbf{E}S)\right).$$

Finding conditions when this bound is trivial as well as the proofs of the presented empirical are yet to be done.

Acknowledgments

The publication has been prepared with the support of the "RUDN University Program 5-100". The work of L.A. Meykhanadzhyan and R.V. Razumchik (specifically, the idea to consider the inaccuracy model with the additive noise, the results and the analysis of the numerical experiments) was funded by RFBR according to the research project No. 18-37-00283.

References

1. L.A. Meykhanadzhyan, R.V. Razumchik. New Scheduling Policy For Estimation Of Stationary Performance Characteristics In Single Server Queues With Inaccurate Job Size Information. Proceedings 30th European Conference on Modelling and Simulation, 2016. doi:10.7148/2016-0710.
2. M. Pastorelli, A. Barbuzzi, D. Carra, M. Dell'Amico, P. Michiardi. HFSP: Bringing Size-Based Scheduling for Hadoop. IEEE Big Data, pp. 51–59, 2013. doi:10.1109/TCC.2015.2396056
3. M. Dell'Amico. A simulator for data-intensive job scheduling. arXiv, Tech. Rep., 2013, `arXiv:1306.6023`.
4. `raw.github.com/SWIMProjectUCB/SWIM/master/workloadSuite/FB-2009_samples_24_times_1hr_1.tsv` [Last access: 18.09.2018]
5. D.R. Wingo. The left-truncated Weibull distribution: theory and computation. *Statistical Papers*, 30 (1) (1989) 39–48. doi:10.1007/BF02924307
6. I.S. Horvath, R. Razumchik, M. Telek. Estimating mean sojourn time in the processor sharing $M/G/1$ queue with inaccurate job size information Proceedings of the International Scientific Conference Analytical And Computational Methods In Probability Theory And Its Applications, Moscow, Russia, 2017. pp. 741–744.
7. L.A. Meykhanadzhyan, T.A. Milovanova, R.V. Razumchik. Stationary waiting time in a queueing system with inverse service order and generalized probabilistic priority, Inform. Primen., 9 (2) (2015) 14–22. doi:10.14357/19922264150202
8. L.A. Meykhanadzhyan. Stationary characteristics of the finite capacity queueing system with inverse service order and generalized probabilistic priority, Inform. Primen., 10 (2) (2016) 123–131. doi:10.14357/19922264160214
9. L.A. Meykhanadzhyan, T.A. Milovanova, A.V. Pechinkin, R.V. Razumchik. Stationary distribution in a queueing system with inverse service order and generalized probabilistic priority, *Inform. Primen.*, 8 (3) (2014) 28–38. doi:10.14375/19922264140304
10. Y. Chen, S. Alspaugh, A. Ganapathi, R. Griffith, and R. Katz. Statistical workload injector for MapReduce, `https://github.com/SWIMProjectUCB/SWIM`.
11. M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker, I. Stoica. Delay scheduling: A simple technique for achieving locality and fairness in cluster scheduling. *Proceedings of the 5th European conference on Computer systems*, 265–278, 2010. doi:10.1145/1755913.1755940
12. Y. Chen, A. Ganapathi, R.Griffith, and R. Katz. The case for evaluating mapreduce performance using workload suites. MASCOTS 2011, IEEE, Singapore, 390–399, 2011. doi:10.1109/MASCOTS.2011.12.