Operations Research Letters 41 (2013) 317–320

Contents lists available at SciVerse ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

$M/G/k$ with staggered setup

Anshul Gandhi [*], Mor Harchol-Balter

Carnegie Mellon University, Pittsburgh, PA, 15213, USA

ARTICLE INFO

Article history:
Received 1 August 2012
Received in revised form
5 March 2013
Accepted 14 March 2013
Available online 25 March 2013

Keywords:
$M/G/k$
Setup times
Decomposition

ABSTRACT

We consider the $M/G/k$/staggered-setup, where idle servers are turned off to save cost, necessitating a setup time for turning a server back on; however, at most one server may be in setup mode at any time. We show that, for exponentially distributed setup times, the response time of an $M/G/k$/staggered-setup approximately decomposes into the sum of the response time for an $M/G/k$ and the setup time, where the approximation is nearly exact. This generalizes a prior decomposition result for an $M/M/k$/staggered-setup.

1. Introduction

1.1. Motivation

We consider an $M/G/k$ system, where there is a *setup cost* required to turn on a server which is currently turned off. Setup costs are common in manufacturing systems [1], where there is often a "warmup" time needed to get a machine running, or a "transport time" needed when calling a staff member into work. Setup costs are also common in data centers [4], where there is a "boot up" time needed to turn on an off server.

In data centers, idle servers are often turned off to save power. However, turning on an off server requires a setup cost. Setup costs are wasteful in two ways: (i) they impose a *setup time*, which is the time it takes the server to turn on, thereby increasing overall mean response time, (ii) they waste *power*, since peak power is consumed during the setup time, although no work is being done. To save on power, the number of servers that can be in setup at any time is often purposely limited. In the *staggered setup* model [6,1], at most one server can be in setup at any time.

1.2. Model

We consider an $M/G/k$ system where jobs arrive according to a Poisson process with rate λ and are served at rate $\mu = \frac{1}{\mathbb{E}[S]}$, where S denotes the job size. For stability, we assume that $k \cdot \mu > \lambda$.

Each of the k servers is in one of three states: off, on (being used to serve a job), or setup (undergoing the setup cost). In the $M/G/k/STAG$ model that we consider, we allow at most one server to be in setup at any given time. When servers are not in use, they are immediately switched to the off state. When a new job arrives, if there is already a server in the setup state, then the job simply joins the queue, otherwise the job picks an off server (assuming there is one) and switches it into the setup state. We use I to denote the setup times, with $\mathbb{E}[I] = \frac{1}{\alpha}$. Unless stated otherwise, we assume that the setup times are exponentially distributed, with rate α. When a job completes service at a server, j, the job at the head of the queue is moved to server j, without the need for setup, since server j is already on. Note that even if the job at the head of the queue was already waiting on another server i in setup mode, the job at the head of the queue is still directed to server j. At this point, if there is another job in the queue, then server i continues to be in setup for this job. If no such job exists in the queue, then server i is turned off.

1.3. Prior work

While there has been a lot of work on single-server systems with setup costs, there has been very little work on multi-server systems with setup costs. For the single-server, [5], in 1986, showed that the distribution of response time for an $M/G/1$ system with setup times, referred to as $M/G/1/Setup$, has the following decomposition property:

$$T_{M/G/1/Setup} \stackrel{d}{=} I + T_{M/G/1}, \tag{1}$$

where I denotes the setup time, and I is exponentially distributed. Note that in the case of a single-server system, $M/G/1/Setup$ is

* Corresponding author.
E-mail addresses: anshulg@cs.cmu.edu (A. Gandhi), harchol@cs.cmu.edu (M. Harchol-Balter).

0167-6377/$ – see front matter
http://dx.doi.org/10.1016/j.orl.2013.03.006

Table 1
Matrix analytic results for the $M/H_2/k/STAG$. In all cases, $\mathbb{E}[S] = 1$. The gray shaded results indicate expected values derived using matrix analytic results for the $M/H_2/k$ and assuming the decomposition property in Eq. (3). The white shaded results indicate matrix analytic results for the $M/H_2/k/STAG$.

Parameters	$\mathbb{E}[T]$	$\mathbb{E}[T^2]$	$\mathbb{E}[T^3]$	$\mathbb{E}[T^4]$
$k = 2, \lambda = 1.9$	$1.9253 \cdot 10^2$	$8.2402 \cdot 10^4$	$5.3089 \cdot 10^7$	$4.5609 \cdot 10^{10}$
$\mathbb{E}[I] = 5, C^2 = 40$	$1.9220 \cdot 10^2$	$8.2264 \cdot 10^4$	$5.3000 \cdot 10^7$	$4.5532 \cdot 10^{10}$
$k = 4, \lambda = 1$	$1.0011 \cdot 10^3$	$2.0022 \cdot 10^6$	$6.0067 \cdot 10^9$	$2.4027 \cdot 10^{13}$
$\mathbb{E}[I] = 1000, C^2 = 7$	$1.0010 \cdot 10^3$	$2.0020 \cdot 10^6$	$6.0061 \cdot 10^9$	$2.4024 \cdot 10^{13}$
$k = 6, \lambda = 3$	$2.0167 \cdot 10^2$	$8.0678 \cdot 10^4$	$4.8391 \cdot 10^7$	$3.8702 \cdot 10^{10}$
$\mathbb{E}[I] = 200, C^2 = 15$	$2.0113 \cdot 10^2$	$8.0469 \cdot 10^4$	$4.8282 \cdot 10^7$	$3.8626 \cdot 10^{10}$
$k = 8, \lambda = 5$	$2.1917 \cdot 10^0$	$1.6463 \cdot 10^1$	$3.9238 \cdot 10^2$	$1.5723 \cdot 10^4$
$\mathbb{E}[I] = 1, C^2 = 10$	$2.1823 \cdot 10^0$	$1.6393 \cdot 10^1$	$3.9159 \cdot 10^2$	$1.5704 \cdot 10^4$
$k = 10, \lambda = 6$	$1.0154 \cdot 10^2$	$2.0326 \cdot 10^4$	$6.0943 \cdot 10^6$	$2.4361 \cdot 10^9$
$\mathbb{E}[I] = 100, C^2 = 20$	$1.0111 \cdot 10^2$	$2.0245 \cdot 10^4$	$6.0747 \cdot 10^6$	$2.4300 \cdot 10^9$

the same as $M/G/1/STAG$. For multi-server queues, only recently in 2010, [4] showed that the distribution of response time for an $M/M/k/STAG$ with exponential setup time, I, has the following decomposition property:

$$T_{M/M/k/STAG} \stackrel{d}{=} I + T_{M/M/k}. \quad (2)$$

However, no results exist for the $M/G/k/STAG$ system.

1.4. Our results

In this paper, we present results suggesting that the decomposition property in Eq. (2) provides a very good approximation for the $M/G/k/STAG$ system, with exponential setup times. That is, the distribution of response time for an $M/G/k/STAG$ can be well approximated as:

$$T_{M/G/k/STAG} \stackrel{d}{\approx} I + T_{M/G/k}. \quad (3)$$

In Section 2, we prove that the decomposition property, as in Eq. (3) above, holds *exactly* for the $M/H_2^*/k/STAG$, where the job size distribution, H_2^*, is a degenerate exponential. Then, in Section 3, we present matrix analytic results suggesting that the decomposition property provides a very good approximation (nearly exact) for the $M/H_2/k/STAG$ and the $M/E_2/k/STAG$, where H_2 and E_2 are the hyper-exponential and the hypo-exponential job size distributions respectively. In Section 4, we present simulation results suggesting that the decomposition property provides a very good approximation for various job size distributions including Deterministic, Uniform, Weibull, Bounded Exponential, and Bounded Pareto. Finally, we conclude in Section 5 with a discussion of the $G/G/k/STAG$ system, and the $M/G/k$ system with non-exponential setup times and non-staggered setup.

2. $M/H_2^*/k/STAG$

The H_2^* distribution is the degenerate exponential distribution, whereby with probability p, the job size is zero, and with probability $(1-p)$, the job size is exponentially distributed with mean $\frac{\mathbb{E}[S]}{1-p}$. Thus, the overall mean job size is $\mathbb{E}[S]$. The squared coefficient of variation for the H_2^* is $C^2 = 1 + \frac{2p}{1-p}$. The H_2^* is an important distribution in queuing theory because its C^2 value spans the entire range from 1 to ∞, allowing it to represent a variety of job size distributions. We now prove the decomposition property for the $M/H_2^*/k/STAG$.

Theorem 1.

$$T_{M/H_2^*/k/STAG} \stackrel{d}{=} I + T_{M/H_2^*/k}.$$

Proof. The time in queue for an $M/H_2^*/k/STAG$ system, $T^Q_{M/H_2^*/k/STAG}$, is the same as the time in queue for an $M'/M'/k/STAG$ system, where the arrival process is Poisson with rate $\lambda \cdot (1-p)$, and the job size distribution, S', is exponential with mean $\frac{\mathbb{E}[S]}{1-p}$. This is because the zero-sized jobs in the $M/H_2^*/k/STAG$ system do not contribute to the queuing time. Thus:

$$\begin{aligned} T_{M/H_2^*/k/STAG} &\stackrel{d}{=} T^Q_{M/H_2^*/k/STAG} + S \\ &= T^Q_{M'/M'/k/STAG} + S \\ &= (T_{M'/M'/k/STAG} - S') + S \\ &= (T_{M'/M'/k} + I - S') + S \quad \text{(from Eq. (2))} \\ &= I + (T_{M'/M'/k} - S' + S) \\ &= I + (T^Q_{M'/M'/k} + S) \\ &= I + (T^Q_{M/H_2^*/k} + S) \\ &= I + T_{M/H_2^*/k}. \quad \square \end{aligned}$$

3. $M/H_2/k/STAG$ and $M/E_2/k/STAG$

The H_2 is the hyper-exponential job size distribution, whereby with probability p, the job size is of type I (exponential with mean $\frac{1}{\mu_1}$), and with probability $(1-p)$, the job size is of type II (exponential with mean $\frac{1}{\mu_2}$). The H_2 distribution is far broader than the H_2^* distribution.

The $M/H_2/k/STAG$ can be analyzed numerically, via matrix analytic methods. While the Markov chain is complex (and is thus omitted due to lack of space), it is tractable via matrix analytic methods due to its regular repeating structure. Likewise, the $M/H_2/k$ without setup times can be analyzed via matrix analytic methods. A subset of our results for the $M/H_2/k/STAG$ are shown in Table 1 (additional results can be found in [3]). In our results, $\mathbb{E}[S] = \frac{p}{\mu_1} + \frac{1-p}{\mu_2}$ is the mean job size, and $C^2 = \frac{\text{Var}(S)}{\mathbb{E}^2[S]}$, is the squared coefficient of variation. From Table 1, it appears that the decomposition property provides a very good approximation for the $M/H_2/k/STAG$ for different parameter values for the first four moments of response time. The maximum difference between the $M/H_2/k/STAG$ results and the decomposition results among all the cases in Table 1 is 0.43%. Note that we derived the higher moments of response time from the higher moments of number of jobs in system using the distributional Little's Law [2].

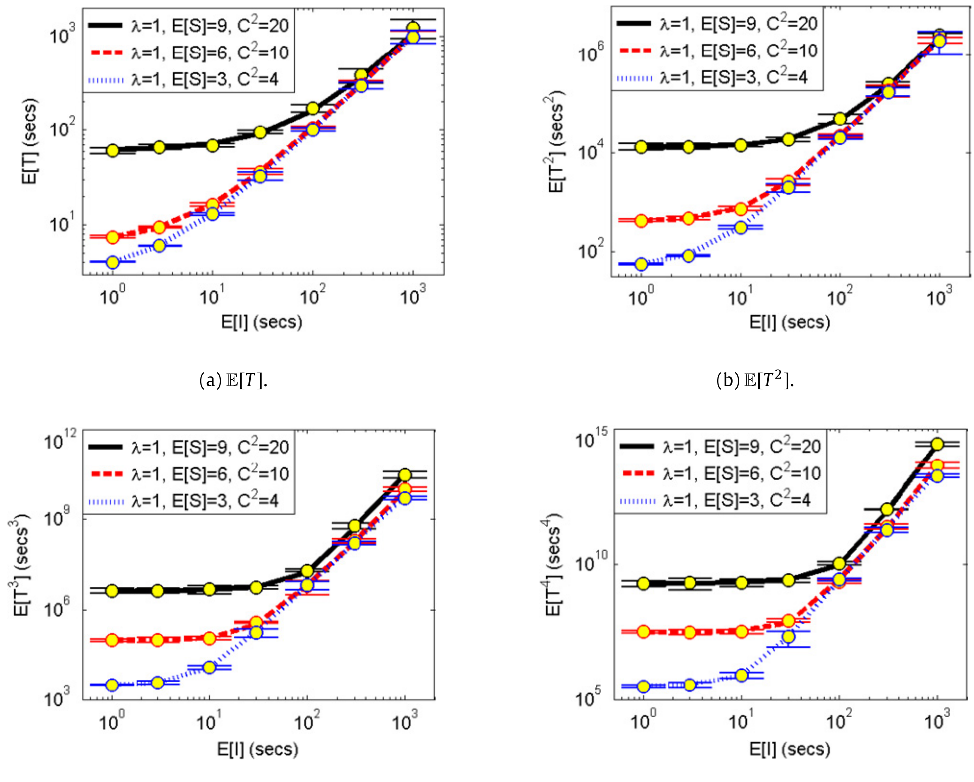

(a) $\mathbb{E}[T]$.

(b) $\mathbb{E}[T^2]$.

(c) $\mathbb{E}[T^3]$.

(d) $\mathbb{E}[T^4]$.

Fig. 1. Simulation results showing the first four moments of response time for $M/G/10/STAG$ under the Bounded Pareto job size distribution. The lines indicate expected values derived using simulation results for the $M/G/10$ response time and assuming the decomposition property in Eq. (3). The solid circles indicate simulation results for $M/G/10/STAG$. Horizontal lines indicate 95% confidence intervals.

We also repeated the above analysis for the $M/E_2/k/STAG$, where E_2 represents the hypo-exponential distribution, for which C^2 ranges from 0 to 1. We find that the decomposition property also provides a very good approximation for the $M/E_2/k/STAG$ for different parameter values for the first four moments of response time.

4. $M/G/k/STAG$

In this section, we demonstrate via simulations that the decomposition property provides a very good approximation for the $M/G/k/STAG$, when G follows distributions other than the hyper-exponential and hypo-exponential. Each simulation run consists of 10^7 arrivals, and we average our results over multiple runs for each job size distribution.

Fig. 1 shows a subset of our simulation results for the first four moments of response time for an $M/G/10$ system (additional results can be found in [3]). In these results, the job size distribution is Bounded Pareto. The maximum difference between the $M/G/10/STAG$ simulation results and the decomposition results among all cases in Fig. 1 is less than 5%, which is reasonable considering the magnitude of the values for the higher moments. We also ran simulations where the job size distribution was Deterministic, Uniform, Bounded Exponential, and Weibull, and also tried different values of k. In all cases, we find that the response time of the $M/G/k/STAG$ (shown as solid circles) can be well approximated using the decomposition property as suggested by Eq. (3) (shown by the lines). Tight confidence intervals (horizontal lines) indicate a good agreement between the simulation results and Eq. (3).

5. Conclusion

In this paper, using analysis, matrix analytic methods, and simulations, we show that the response time distribution for an $M/G/k/STAG$ ($M/G/k$ with exponential staggered setup) can be well approximated using the sum of the setup time distribution and the response time distribution of an $M/G/k$ without setup, as in Eq. (3). For the case when $G \sim H_2^*$, we prove that this approximation is exact.

The fact that the setup time is exponentially distributed is important. We can prove [3] that the decomposition property for the mean response time of an $M/G/1/STAG$ only holds if the squared coefficient of variation for the setup time is $C_I^2 = 1$. If we use the decomposition result as an approximation for the mean response time of an $M/G/1/STAG$ with non-exponential setup time, the results can be very inaccurate. For example, in the case of an $M/M/1/STAG$ with deterministic setup times, the approximation (given by Eq. (2)) can be off by as much as 90% when compared to the exact value (given by Eq. (1)). We find similar results for the $M/M/k/STAG$ as well, by comparing the approximation to the exact numbers given by simulation. In general, we find that the approximation gets worse as the setup time increases.

The fact that at most one server is in setup at any point of time is also critical for the decomposition property. If we relax this requirement, and allow multiple servers to be in setup, as in [4], we find that the decomposition property does not provide a good approximation, even for the case of an $M/M/2$ system (see [3] for details). We further find that the decomposition property is a poor approximation for the $G/G/k/STAG$ with exponential setup times when the arrival process is not Poisson.

It is worth noting that the decomposition property stated in this paper is counter-intuitive. One might expect that the setup time impacts the response time of *some* of the jobs, but the decomposition property states that the setup time is approximately added, in its entirety, to the full distribution of response time. We have searched for a long time for some simple intuition which explains why the decomposition property provides such a good approximation for the $M/G/k/STAG$ with exponential setup time, but this intuition remains elusive.

References

[1] J.R. Artalejo, A. Economou, M.J. Lopez-Herrero, Analysis of a multiserver queue with setup times, Queueing Syst. Theory Appl. 51 (1–2) (2005) 53–76.

[2] Shelby L. Brumelle, A generalization of $L = \lambda W$ to moments of queue length and waiting times, Oper. Res. 20 (6) (1972) 1127–1136.

[3] Anshul Gandhi, Mor Harchol-Balter, $M/G/k$ with Exponential Setup, Technical Report CMU-CS-09-166, Carnegie Mellon University, 2009.

[4] Anshul Gandhi, Mor Harchol-Balter, Ivo Adan, Server farms with setup costs, Perform. Eval. 67 (2010) 1123–1138.

[5] Hanoch Levy, Leonard Kleinrock, A queue with starter and a queue with vacations: delay analysis by decomposition, Oper. Res. 34 (3) (1986) 426–436.

[6] Intel Corporation. Serial ATA Staggered Spin-Up (White paper), 2004.