Proceedings of the 2021 Winter Simulation Conference
S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, eds.

MEASURING THE OVERLAP WITH OTHER CUSTOMERS IN THE SINGLE SERVER QUEUE

Sergio Palomo

Systems Engineering
Cornell University
Ithaca, NY 14850, USA

Jamol Pender

Operations Research and Information Engineering
Cornell University
Ithaca, NY 14850, USA

ABSTRACT

The single server queue is one of the most basic queueing systems to model stochastic waiting dynamics. Most work involving the single server queue only analyzes the customer or agent behavior. However, in this work, we are inspired by COVID-19 applications and are interested in the interaction between customers and more specifically, the time that adjacent customers overlap in the queue. To this end, we derive a new recursion for this overlap time and study the steady state behavior of the overlap time via simulation and probabilistic analysis. We find that the overlap time between adjacent customers in the M/M/1 queue has a conditional distribution that is given by an exponential distribution, however, as the distance between customers grows, the probability of a non-negative overlap time decreases geometrically. We also find via simulation that the exponential distribution still holds when the distribution is non-exponential, hinting at a more general result.

1 INTRODUCTION

COVID-19 has had a large impact on our society to date. To understand the impact of exposure of people lining up in person, we study the single server queue, but analyze the times that adjacent customers overlap in the queue. This overlap time can be used a proxy for understanding the exposure of infected individuals might have with others while waiting in an in-person queue. Much of the queueing literature is concerned with understanding the number of customers in the queue and how much time an individual customer spends in the queue. However, much of the work does not explore how much time customers overlap with other customers in the queue, but there is one notable exception Kang et al. (2021) where they use a spatial SIR model along with queueing theory to compute overlap between an infectious person and non-infectious persons. Thus, understanding how much time people overlap with one another is a very important performance measure to understand given the COVID-19 pandemic or any epidemic for that matter. Many government organizations are attempting to do contact tracing to see where infected people overlap with non-infected people. We analyze a stylized version of this this problem by analyzing the overlap dynamics of the M/M/1 queue and use simulation to analyze the G/G/1 queue with log-normal arrival and service times.

In this work, we consider the M/M/1 queueing system model where the customer inter-arrival distribution is general, the service distribution is general, and there is a single server. This model has been studied for many years and one can obtain a thorough description of this queueing system in (Shortle, Thompson, Gross, and Harris 2018). One of the hallmarks of the G/G/1 queue is the Lindley recursion, which describes the waiting time of the n^{th} customer in terms of the previous customer's waiting time, see for example (Lindley 1952). Thus, if we define W_n to be the waiting time for that customer, A_n to be the inter-arrival time between the n^{th} and the $(n+1)^{th}$ customers and S_n to be the n^{th} customer's service time, then Lindley's recursion can be represented by the following recursive equation

978-1-6654-3311-2/21/$31.00 ©2021 IEEE

$$W_{n+1} \quad = \quad \max\left(W_n + S_n - A_n, 0\right).$$

Previous analysis such as (Prabhu 1974) studies the steady state waiting time of the G/G/1 queue and shows that the steady state waiting time satisfies the following integral equation, which is known as Lindley's integral equation

$$F(x) = \int_0^\infty K(x-y)F(\mathrm{d}y) \quad x \geq 0$$

where $K(x)$ is the distribution function of the random variable denoting the difference between the service time of the n^{th} customer and the inter-arrival time of the n^{th} and $n+1^{th}$ customer. In fact the Wiener–Hopf method can be used to solve this integral equation in closed form (Prabhu 1974). In addition to the steady state waiting time, one can also calculate the conditional waiting time of the $(n+1)^{th}$ customer given the waiting time of the n^{th} customer, see for example Palomo and Pender (2020). Despite significant analysis of the Lindley recursion, much of the current research does focus on the interactions between customers. In the remaining of this paper, we will consider measuring the overlap time between customers in the M/M/1 queue and connect it to the Lindley recursion in a precise mathematical sense. We also explore the overlap in the G/G/1 queue where the arrival and service times are log-normal random variables.

1.1 Contributions of Our Work

In this paper, we make the following contributions to the literature:

- We derive the steady state distribution of the overlap time between customers in the Markovian case and show that it is given by an exponential distribution.
- We derive new conditional moments for the overlap time when the information given is the waiting time or the previous overlap time.

1.2 Organization of the Paper

The remainder of the paper is organized as follows. Section 2 introduces the overlap time and a recursion for computing. In Section 3, we present our simulation results for computing the distribution of the overlap time of customers in the G/G/1 queue. We also explain how the inter-arrival and service distributions affect the distribution of the overlap time process. Finally, a conclusion and ideas for future work is given in Section 4.

2 The Overlap Time

In this section, we describe the overlap time between nearby customers. The overlap time represents the time that nearby customers are both jointly in the queue at the same time. This time is important from a COVID-19 perspective since the Center for Disease Control (CDC) guidelines on their website at (cdc.gov) suggest that not being near someone for more than 15 minutes with a mask can help stop the spread of COVID-19. However, when social distancing is not preventable, we are interested in how long nearby customers are waiting in the same space at the same time. To gain some intuition about the overlap time, we first analyze the overlap time of back to back customers to get a bit of intuition for the overlap time process in the general case. Suppose that we want to compute the overlap time of the n^{th} and the $(n+1)^{th}$ customer, then we have the following recursion formula

$$O_{n,n+1} \quad = \quad \max\left(D_n - \sum_{i=1}^{n} A_i, 0\right). \tag{1}$$

where D_n is defined as the departure time of the n^{th} customer and $O_{n,n+j}$ is the overlap time of the n^{th} and $(n+j)^{th}$ customers. Moreover, by manipulating the departure time of the n^{th} customer we can show that the overlap between back to back customer is precisely the Lindley recursion and the waiting time of the $(n+1)^{th}$ customer i.e.

$$\begin{aligned} O_{n,n+1} &= \max\left(D_n - \sum_{i=1}^{n} A_i, 0\right) \\ &= \max\left(W_n + S_n + \sum_{i=1}^{n-1} A_i - \sum_{i=1}^{n} A_i, 0\right) \\ &= \max\left(W_n + S_n - A_n, 0\right). \end{aligned}$$

Th recursion given in Equation 1 says that the overlap time between the n^{th} and the $(n+1)^{th}$ customers is exactly the maximum of zero and the departure time of the n^{th} customer minus the arrival time of the $(n+1)^{th}$ customer. Thus, if the n^{th} customer departs before the $(n+1)^{th}$ customer arrives, then there is no overlap. However, if the n^{th} customer departs after the $(n+1)^{th}$ customer arrives, the overlap time is positive and is equal to the difference of the two times. Using this idea, we can derive a recursion for the overlap time between non-adjacent customers as well. For a pair of customers that are exactly k spaces apart, we have the following recursion for their overlap time:

$$\begin{aligned} O_{n,n+k} &= \max\left(D_n - \sum_{i=1}^{n+k-1} A_i, 0\right) \\ &= \max\left(W_n + S_n + \sum_{i=1}^{n-1} A_i - \sum_{i=1}^{n+k-1} A_i, 0\right) \\ &= \max\left(W_n + S_n - \sum_{i=n}^{n+k-1} A_i, 0\right). \end{aligned}$$

In addition to deriving a recursion for the overlap time of the n^{th} and $(n+k)^{th}$ customer in terms of the waiting time of the n^{th} customer, we can also derive a recursion in terms of the previous overlap time. This recursion is given by

$$O_{n,n+1} = \max\left(O_{n,n} - A_n, 0\right)$$

where $O_{n,n}$ is defined to be $O_{n,n} = W_n + S_n$. Moreover, this recursive equation can be iterated k times fo obtain the overlap time of the n^{th} and $(n+k)^{th}$ customer i.e.

$$O_{n,n+k} = \max\left(O_{n,n+k-1} - A_{n+k-1}, 0\right). \tag{2}$$

This observation in Equation 2 shows that the overlap between customer n and customer $n+k$ must be smaller than the overlap between customer n any customer that comes before $n+k$, but after customer n. Thus, there is a natural ordering between the overlap times i.e.

$$O_{n,n+k} \leq O_{n,n+k-1} \leq \ldots \leq O_{n,n+1} \leq O_{n,n} = W_n + S_n.$$

However, why should we care about $O_{n,n+k}$ and not just worry about $O_{n,n}$ or $O_{n,n+1}$? In the context of COVID-19, it is apparent that it is not sufficient to just know whether or not the customer next to you is infectious or not. Perhaps even adjacent customers might transmit air particles that could infect anyone. Thus, the overlap between any two customers is a compelling problem to study. We outline this analysis in the sequel.

2.1 Overlap Distribution for the M/M/1 Queue

Now that we have a good idea of how to construct the overlap times of different customers in the queue, we are interested in understand the distribution of those overlap times as a function of the numerical spacing between the customers in line. In what follows, we prove that in the case of exponential inter-arrival times and exponential service times, the overlap time between customers follows an exponential distribution.

Theorem 1 Let O_k be the steady state overlap time of two customers who are exactly k customers apart in an stable M/M/1 queue where the arrival rate is λ and the service rate is μ. Then, in steady state, the conditional overlap time distribution is given by an exponential distribution i.e.

$$\mathbb{P}\left(O_k > t | O_k > 0\right) = e^{-(\mu-\lambda)t} \tag{3}$$

and the unconditional distribution is

$$\mathbb{P}\left(O_k > t\right) = \rho^k e^{-(\mu-\lambda)t} \tag{4}$$

Proof.

$$\begin{aligned}
\mathbb{P}\left(O_k > t | O_k > 0\right) &= \sum_{j=0}^{\infty} \mathbb{P}\left(O_k > t \middle| \{O_k > 0\} \cap \{Q(t_{n+k}) = j\}\right) \cdot \mathbb{P}\left(Q(t_{n+k}) = j\right) \\
&= \sum_{j=0}^{\infty} \mathbb{P}\left(O_k > t \middle| \{O_k > 0\} \cap \{Q(t_{n+k}) = j\}\right) \cdot (1-\rho) \cdot \rho^j \\
&= \sum_{j=k+1}^{\infty} \mathbb{P}\left(\sum_{l=1}^{(j-k)} S_l > t\right) \cdot (1-\rho) \cdot \rho^j \\
&= \sum_{j=k+1}^{\infty} \left(\sum_{n=0}^{j-k-1} \frac{1}{n!} e^{-\mu t} \cdot (\mu t)^n\right) \cdot (1-\rho) \cdot \rho^j \\
&= e^{-\mu t} \sum_{j=k+1}^{\infty} \left(\sum_{n=0}^{j-k-1} \frac{(\mu t)^n}{n!}\right) \cdot (1-\rho) \cdot \rho^j \\
&= e^{-\mu t} \sum_{n=0}^{\infty} \sum_{j=k}^{\infty} \frac{(\mu t)^n}{n!} \cdot (1-\rho) \cdot \rho^j \\
&= e^{-\mu t} \sum_{n=0}^{\infty} \frac{(\mu \rho t)^n}{n!} \\
&= e^{-(\mu-\lambda)t}.
\end{aligned}$$

Finally, the unconditional probability follows easily from conditioning that at least $k+1$ customers are present in the system. □

Theorem 1 shows that the overlap time distribution is given by an exponential distribution. This is consistent with the conditional waiting time distribution of the M/M/1 queue. We observe in Figures 1 - 3 that the exponential distribution given by simulation is consistent with our theoretical analysis. In all of the Figures in this paper, the number of total customers used was 100,000. We see in Figure 1 that the exponential distribution is quite good when the traffic is light and we see in Figure 2 that the exponential distribution is also good when we consider a heavier traffic setting. We also increase the rates of traffic in Figure 3 and also observe that the exponential distribution fits the data quite well. We will explore non-exponential settings in the sequel.

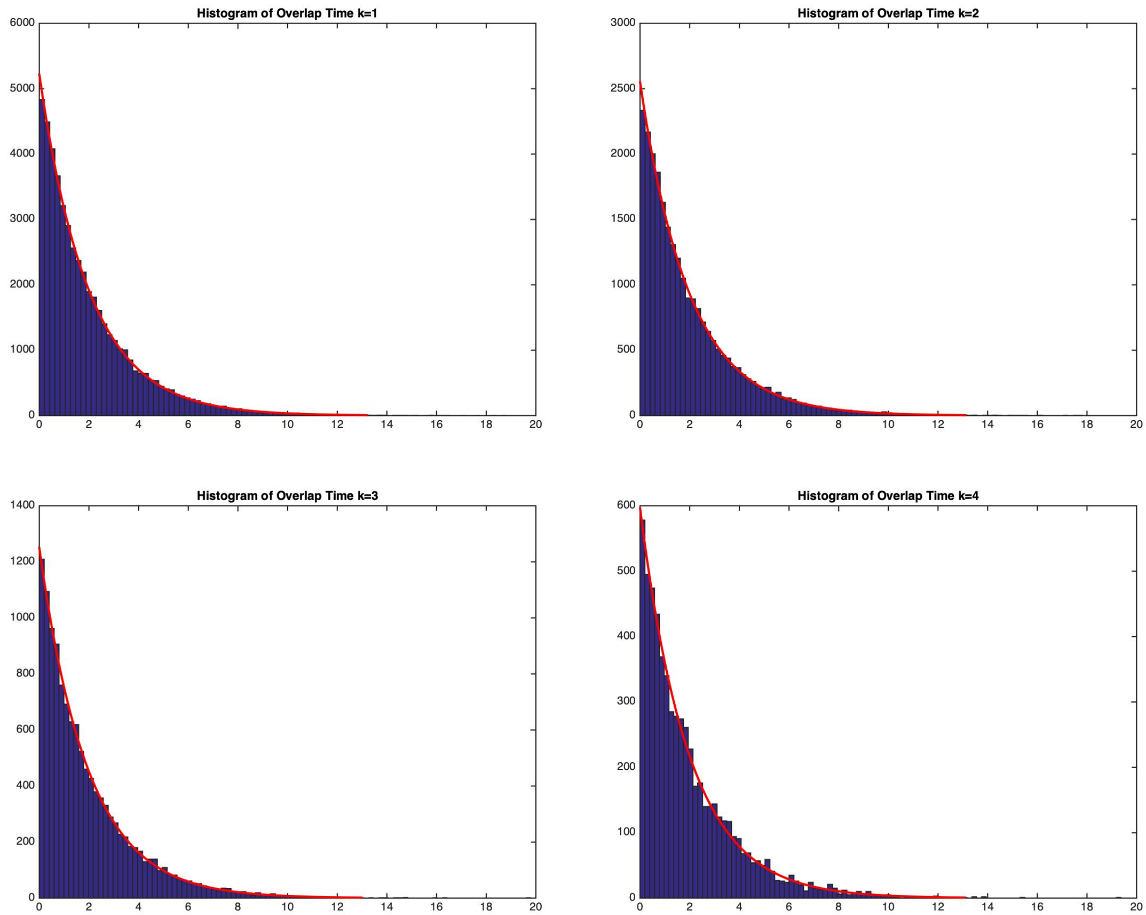

Figure 1: Histograms for the Overlap Times Exponential Distributions $(\lambda = .5, \mu = 1)$. Upper Left (k=1), Upper Right (k=2), Lower Left (k=3), Lower Right (k=4).

2.2 Conditional Mean and Variance Formulas

Proposition 1 Let X be an exponential random variable with rate μ and Y be an Erlang(λ, k) random variable. Then, the probability density function of $Z = X - Y$ is equal to

$$f_Z(z) = \begin{cases} 1 - \frac{\mu\lambda^k}{(\lambda+\mu)^k} e^{-\lambda z} \sum_{j=0}^{k-1} \frac{[(\lambda+\mu)z]^j}{j!}, & \text{for } z \leq 0 \\ \frac{\mu\lambda^k}{(\lambda+\mu)^k} e^{-\mu z} & \text{for } z > 0 \end{cases} \tag{5}$$

Proof. Apply convolution. □

Proposition 2 Let Z be a random variable with probability density function given in Equation 5. Then, we define the following conditional moments as

$$\begin{aligned} \alpha_m(x) &= \mathbb{E}\left[Z^m \cdot \{x + Z > 0\}\right] & (6) \\ &= \frac{(-x)^{m+1}}{m+1} - \frac{\mu\lambda^k}{(\lambda+\mu)^k} \sum_{j=0}^{k-1} \frac{(\lambda+\mu)^j}{j!} \left(\frac{\Gamma(m+j+1, -\lambda x) - \Gamma(m+j+1)}{(-\lambda)^{m+j+1}} \right) + \frac{\mu\lambda^k}{(\lambda+\mu)^k} \frac{\Gamma(m+1)}{\mu^{m+1}} & (7) \end{aligned}$$

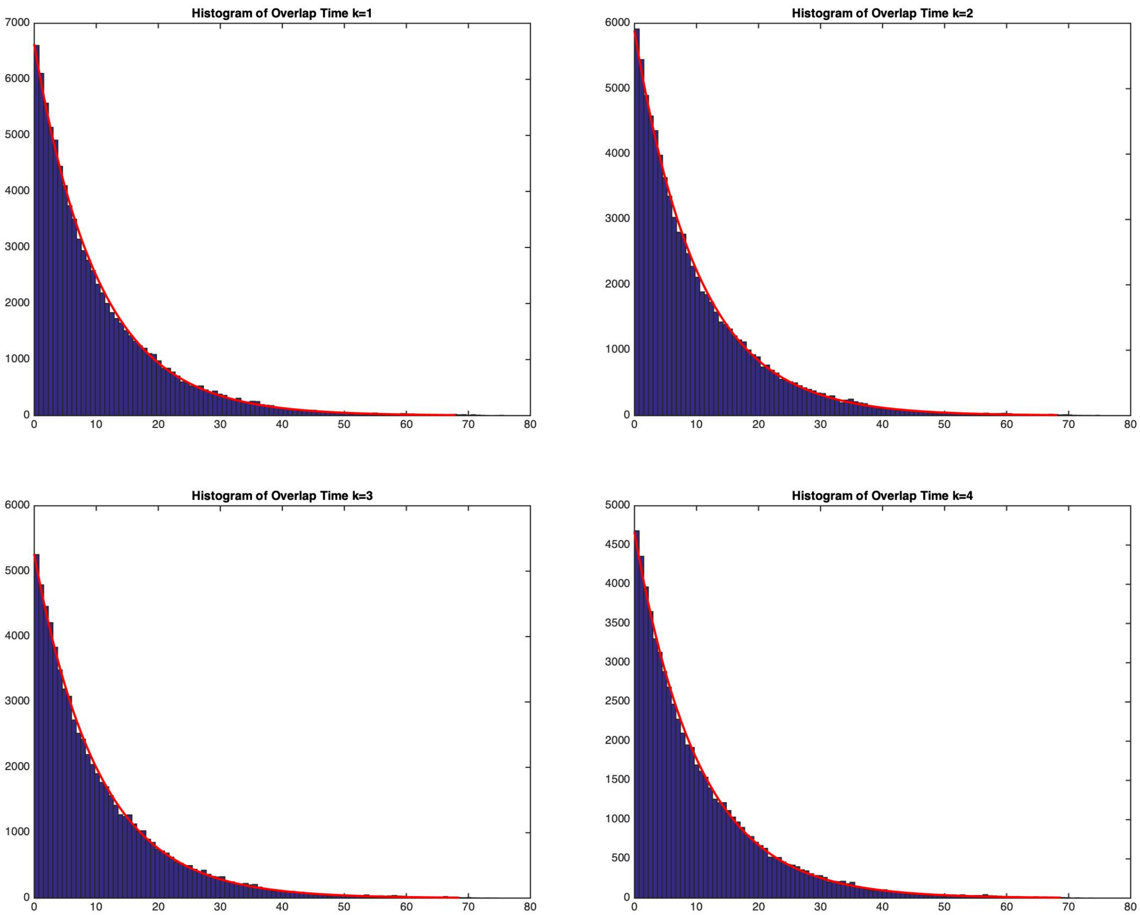

Figure 2: Histograms for the Overlap Times Exponential Distributions $(\lambda = .9, \mu = 1)$. Upper Left (k=1), Upper Right (k=2), Lower Left (k=3), Lower Right (k=4).

Proof.

$$
\begin{aligned}
\alpha_m(x) &= \mathbb{E}\left[Z^m \cdot \{x + Z > 0\}\right] = \int_{-x}^{\infty} z^m f_Z(z) dz \\
&= \int_{-x}^{0} z^m \left(1 - \frac{\mu\lambda^k}{(\lambda+\mu)^k} e^{-\lambda z} \sum_{j=0}^{k-1} \frac{[(\lambda+\mu)z]^j}{j!}\right) dz + \int_0^{\infty} z^m \frac{\mu\lambda^k}{(\lambda+\mu)^k} e^{-\mu z} dz \\
&= \int_{-x}^{0} z^m \left(1 - \frac{\mu\lambda^k}{(\lambda+\mu)^k} e^{-\lambda z} \sum_{j=0}^{k-1} \frac{[(\lambda+\mu)z]^j}{j!}\right) dz + \frac{\mu\lambda^k}{(\lambda+\mu)^k} \frac{\Gamma(m+1)}{\mu^{m+1}} \\
&= \int_{-x}^{0} z^m dz - \frac{\mu\lambda^k}{(\lambda+\mu)^k} \sum_{j=0}^{k-1} \frac{(\lambda+\mu)^j}{j!} \int_{-x}^{0} \left(z^{m+j} e^{-\lambda z}\right) dz + \frac{\mu\lambda^k}{(\lambda+\mu)^k} \frac{\Gamma(m+1)}{\mu^{m+1}} \\
&= \frac{(-x)^{m+1}}{m+1} - \frac{\mu\lambda^k}{(\lambda+\mu)^k} \sum_{j=0}^{k-1} \frac{(\lambda+\mu)^j}{j!} \left(\frac{\Gamma(m+j+1, -\lambda x) - \Gamma(m+j+1)}{(-\lambda)^{m+j+1}}\right) + \frac{\mu\lambda^k}{(\lambda+\mu)^k} \frac{\Gamma(m+1)}{\mu^{m+1}}
\end{aligned}
$$

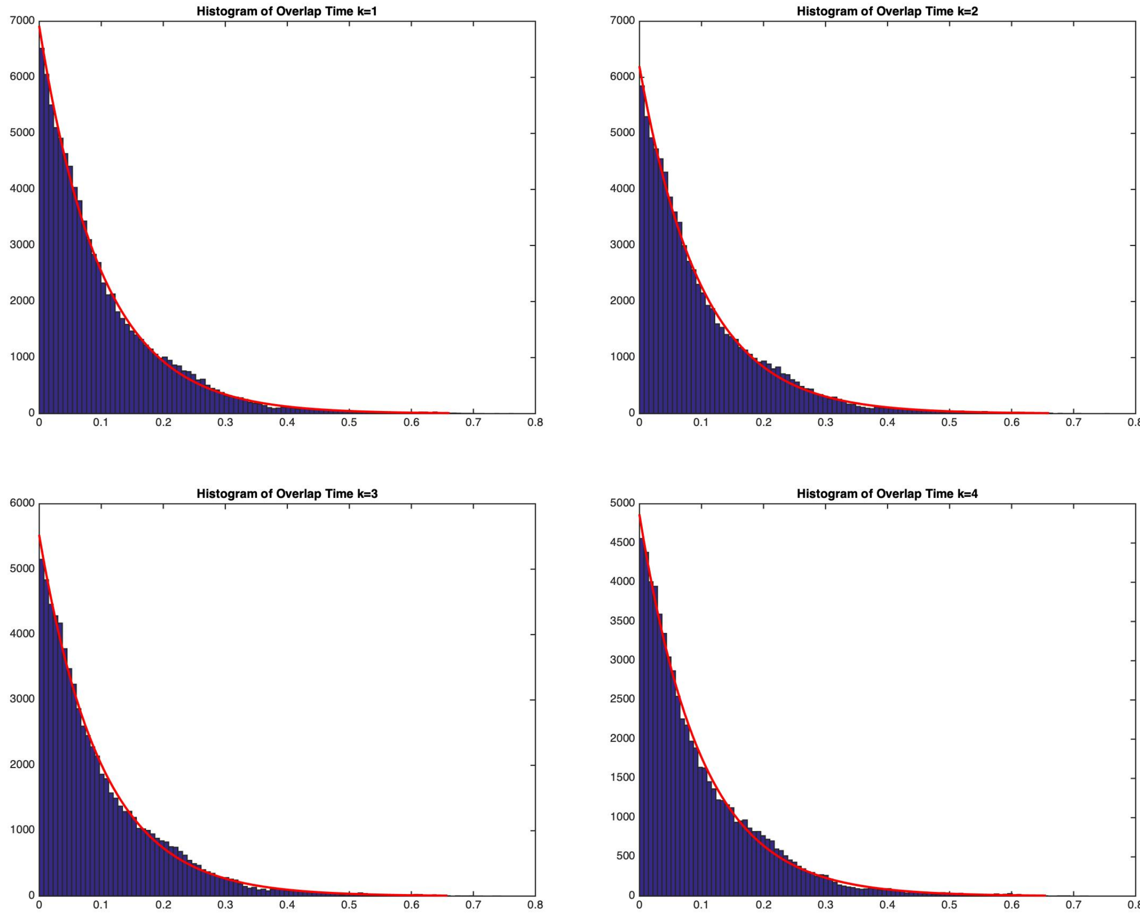

Figure 3: Histograms for the Overlap Times Exponential Distributions $(\lambda = 90, \mu = 100)$. Upper Left (k=1), Upper Right (k=2), Lower Left (k=3), Lower Right (k=4).

where

$$\Gamma(m,x) = \int_x^\infty y^{m-1} e^{-y} dy \quad \text{and} \quad \Gamma(m) = \int_0^\infty y^{m-1} e^{-y} dy$$

are defined as the incomplete gamma function and the gamma function respectively. □

Theorem 2 The conditional m^{th} moment of the overlap time for the k^{th} adjacent customer in the M/M/1 queue is given by the following formula

$$\mathbb{E}[O^m_{n,n+k} | W_n] = \sum_{j=0}^{m} W_n^j \binom{m}{j} \alpha_{m-j}(W_n). \tag{8}$$

Proof.

$$\begin{aligned}
\mathbb{E}[O_{n,n+k}^m | W_n] &= \mathbb{E}\left[\left(\left(W_n + S_n - \sum_{i=n+1}^{n+k} A_i\right)^+\right)^m \middle| W_n\right] && (9) \\
&= \mathbb{E}\left[\left((W_n + Z)^+\right)^m \middle| W_n\right] && (10) \\
&= \mathbb{E}\left[(W_n + Z)^m \cdot \{W_n + Z > 0\} \middle| W_n\right] && (11) \\
&= \sum_{j=0}^{m} W_n^j \binom{m}{j} \alpha_{m-j}(W_n). && (12)
\end{aligned}$$

□

Corollary 3 The conditional mean and variance of the overlap time for the k^{th} adjacent customer in the M/M/1 queue is given by the following expressions

$$\mathbb{E}[O_{n,n+k}|W_n] = W_n - W_n \left(\frac{\mu}{\mu+\lambda}\right)^k e^{-\lambda W_n} + \frac{1}{\lambda}\left(\frac{\mu}{\mu+\lambda}\right)^k e^{-\lambda W_n}(1 - \lambda W_n) + \frac{k}{\mu} - \frac{1}{\lambda} \quad (13)$$

$$\mathrm{Var}[O_{n,n+k}|W_n] = W_n^2 \cdot \alpha_0(W_n) \cdot (1 - \alpha_0(W_n)) + 2W_n \cdot \alpha_1(W_n) \cdot (1 - \alpha_0(W_n)) - \alpha_1(W_n)^2 + \alpha_2(W_n) \quad (14)$$

Proposition 3 Let Y be an $Erlang(k,\lambda)$ random variable. Then, we define the following conditional moments as

$$\beta_m(x) = \mathbb{E}\left[Y^m \cdot \{x - Y > 0\}\right] = \frac{\Gamma(m+k, \lambda x)}{\Gamma(k)\lambda^m}$$

Proof.

$$\begin{aligned}
\beta_m(x) &= \mathbb{E}\left[Y^m \cdot \{x - Y > 0\}\right] \\
&= \int_x^\infty y^m e^{-\lambda y} \frac{\lambda^k y^{k-1}}{\Gamma(k)} dy \\
&= \frac{\lambda^k}{\Gamma(k)} \int_x^\infty y^{m+k-1} e^{-\lambda y} dy \\
&= \frac{\Gamma(m+k, \lambda x)}{\Gamma(k)\lambda^m}.
\end{aligned}$$

□

Theorem 4 The conditional m^{th} moment of the overlap time for the k^{th} adjacent customer in the M/M/1 queue is given by the following formula

$$\mathbb{E}[O_{n,n+k}^m | O_n] = \sum_{j=0}^{m} O_n^j \binom{m}{j} \frac{\Gamma(m-j+k, \lambda O_n)}{\Gamma(k)\lambda^{m-j}} \cdot (-1)^{m-j}$$

Proof.

$$\begin{aligned}
\mathbb{E}[O_{n,n+k}|O_n] &= \mathbb{E}\left[\left(\left(O_n - \sum_{i=n}^{n+k-1} A_i\right)^+\right)^m \middle| O_n\right] \\
&= \mathbb{E}\left[(O_n - Y)^m \cdot \{O_n - Y > 0\} \middle| O_n\right] \\
&= \sum_{j=0}^{m} O_n^j \binom{m}{j} \mathbb{E}\left[(-Y)^{m-j} \cdot \{O_n - Y > 0\} \middle| O_n\right] \\
&= \sum_{j=0}^{m} O_n^j \binom{m}{j} \frac{\Gamma(m-j+k, \lambda O_n)}{\Gamma(k)\lambda^{m-j}} \cdot (-1)^{m-j}.
\end{aligned}$$

□

Corollary 5 The conditional mean and variance of the overlap time for the k^{th} adjacent customer in the M/M/1 queue is given by the following expressions

$$\begin{aligned}
\mathbb{E}[O_{n,n+k}|O_n] &= O_n \cdot \beta_0(O_n) - \beta_1(O_n) \\
\mathrm{Var}[O_{n,n+k}|O_n] &= O_n^2 \cdot \beta_0(O_n) \cdot (1 - \beta_0(O_n)) - 2O_n \cdot \beta_1(O_n) \cdot (1 - \beta_0(O_n)) - \beta_1(O_n)^2 + \beta_2(O_n)
\end{aligned}$$

3 ADDITIONAL NUMERICAL RESULTS

In addition to exploring the Markovian setting, we also explored the setting where the inter-arrival times and service times are not given by exponential random variables. We use LogNormal random variables for the inter-arrival and service times and explore the impact on the distribution. We should mention that we set the mean and variance of the LogNormal random variables to be the values we want and compute the associated parameters. In Figures 4 - 6, we observe that the overlap distribution is close to an exponential. However, it does not appear to be the same as the one given by the exponential parameter setting. Thus, we believe a result similar to the heavy traffic limit result holds for the overlap times as well, which means that the exponential distribution will depend on the mean and variance of the inter-arrival and service time distributions. We hope to prove a result like this in the future.

Finally, we should mention we implemented two goodness of fit tests for each figure in this paper. The first goodness of fit test was based on the chi-square distribution and second was based on the Kolmogorov-Smirnov test. In all cases, we rejected the null hypothesis that the data came from an exponential distribution that was fit from the data i.e sample mean. We would like to emphasize that this was true even for the M/M/1 setting, which we proved to be exponential. This is generally well known in statistics and happens with large datasets where any deviation can cause a rejection from the null.

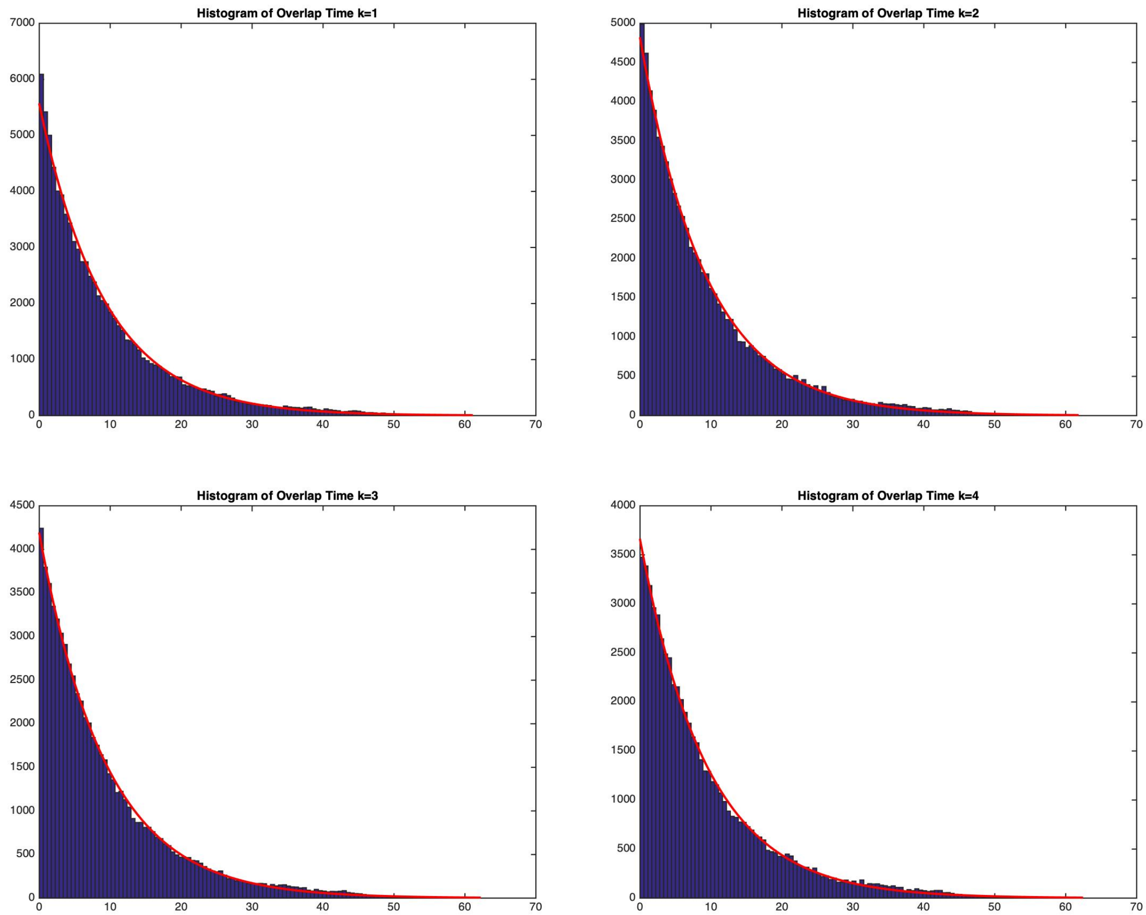

Figure 4: Histograms for the Overlap Times (LogNormal Distribution). (Inter-Arrival (Mean = 1/.9, Variance = 1)), (Service (Mean = 1, Variance = 1)). Upper Left (k=1), Upper Right (k=2), Lower Left (k=3), Lower Right (k=4).

4 CONCLUSION AND FUTURE WORK

In this paper, we consider the overlap time process of consecutive customers in the G/G/1 queue. We prove in the Markovian case that the overlap time distribution is given by an exponential distribution that decays geometrical as the distance between customers increases. Moreover, we compute the conditional moments of the overlap time process by conditioning on previous overlap times or previous waiting times. These formulas would be useful for predicting overlap times when given some information about the system. We also observe that the steady state distribution also appears to be exponential even in the case of non-exponential inter-arrival and service times, which hints at a stronger result about the waiting time distribution of the overlap time. We suspect that this result holds in heavy traffic like the waiting time result for the G/G/1 queue as we have tried several other distributions in heavy traffic and they all are very close to exponential in the heavy traffic regime.

In terms of future work, it would be great to consider the overlap time process for more general queueing models such as multi-server models, models with abandonment, and even infinite server models, see for example Massey and Pender (2018), Palomo et al. (2020), Koole and Mandelbaum (2002). Finally, it would be interesting to consider a more spatial queueing model like the one developed in Aldous (2017) and think about overlap in a true distance sense as well.

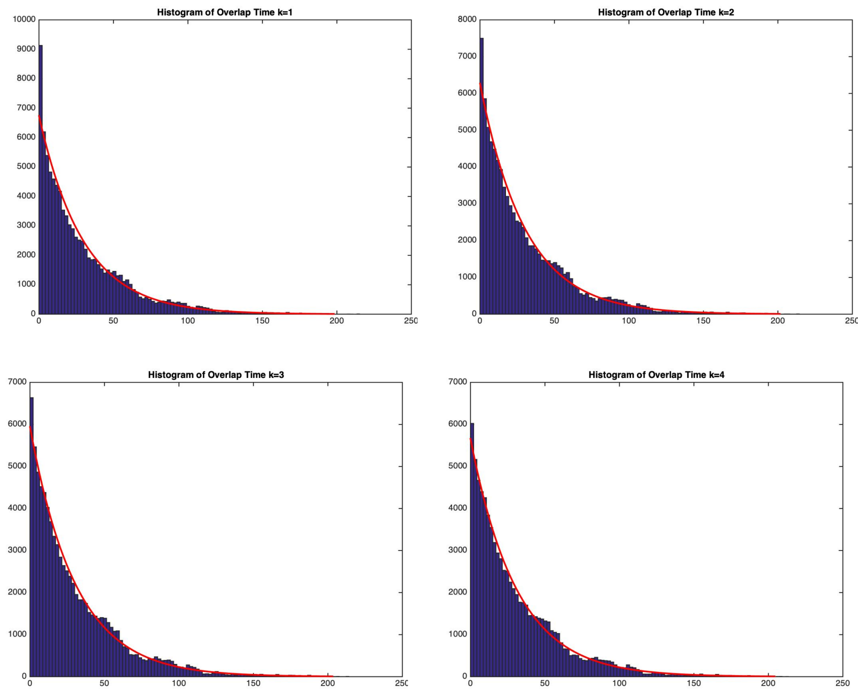

Figure 5: Histograms for the Overlap Times (LogNormal Distribution). (Inter-Arrival (Mean = 1/.9, Variance = 4)), (Service (Mean = 1, Variance = 4)). Upper Left (k=1), Upper Right (k=2), Lower Left (k=3), Lower Right (k=4).

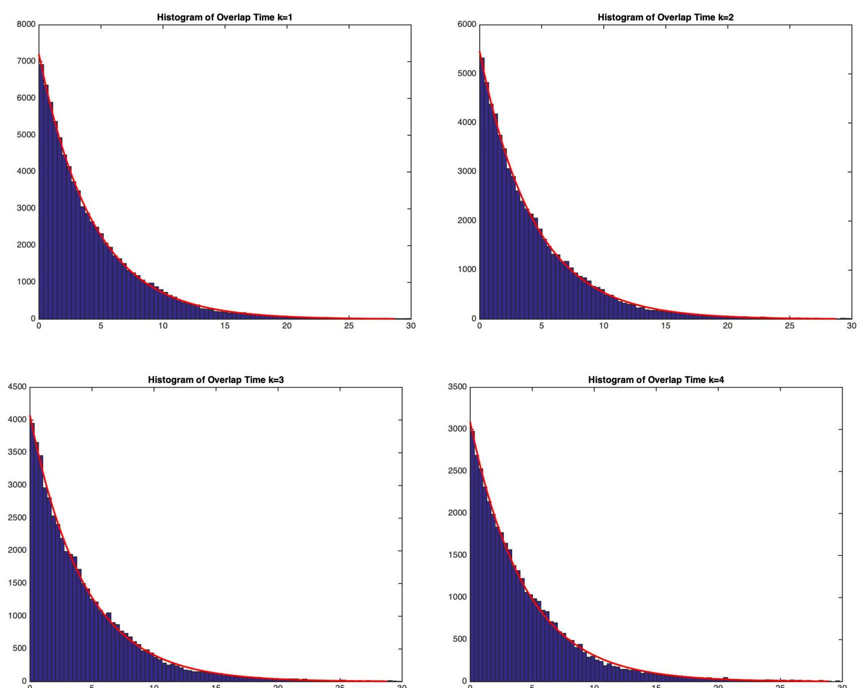

Figure 6: Histograms for the Overlap Times (LogNormal Distribution). (Inter-Arrival (Mean = 1/.9, Variance = .5)), (Service (Mean = 1, Variance = .5)). Upper Left (k=1), Upper Right (k=2), Lower Left (k=3), Lower Right (k=4).

ACKNOWLEDGEMENTS

We would like to thank Cornell University and the Sloan Foundation for providing Sergio Palomo with a Sloan Fellowship for his research. We also thank Sherwin Deroudi for early comments on this work.

REFERENCES

Aldous, D. 2017. "Waves in a spatial queue". *Stochastic Systems* 7(1):197–236.

Kang, K., S. Doroudi, M. Delasay, and A. Wickeham. 2021. "A Queueing-Theoretic Framework for Evaluating Transmission Risks in Service Facilities During a Pandemic". *arXiv preprint arXiv:2103.13441*.

Koole, G., and A. Mandelbaum. 2002. "Queueing models of call centers: An introduction". *Annals of Operations Research* 113(1-4):41–59.

Lindley, D. V. 1952. "The theory of queues with a single server". In *Mathematical Proceedings of the Cambridge Philosophical Society*, Volume 48, 277–289. Cambridge University Press.

Massey, W. A., and J. Pender. 2018. "Dynamic rate Erlang-A queues". *Queueing Systems* 89(1):127–164.

Palomo, S., and J. Pender. 2020. "Learning Lindley's Recursion". In *Proceedings of the 2020 Winter Simulation Conference*, edited by K. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, 2595–2605. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Palomo, S., J. Pender, W. Massey, and R. C. Hampshire. 2020. "Flattening the curve: Insights from queueing theory". *arXiv preprint arXiv:2004.09645*.

Prabhu, N. 1974. "Wiener-Hopf techniques in queueing theory". In *Mathematical methods in queueing theory*, 81–90. Springer.

Shortle, J. F., J. M. Thompson, D. Gross, and C. M. Harris. 2018. *Fundamentals of queueing theory*, Volume 399. John Wiley & Sons.

AUTHOR BIOGRAPHIES

SERGIO PALOMO is a PhD student in Systems Engineering at Cornell University. He holds a Masters in Mathematics from City College of New York and a Bachelors in Applied Mathematics from Stony Brook University. His research interests are in queueing theory, machine learning and stochastic simulation. His e-mail address is sdp85@cornell.edu.

JAMOL PENDER is an assistant professor in the School Operations Research and Information Engineering (ORIE) at Cornell University. He earned his PhD in the Department of Operations Research and Financial Engineering (ORFE) at Princeton University. His research interests include queueing theory, point processes, stochastic simulation, dynamical systems and applied probability. His e-mail address is jjp274@cornell.edu. His website is https://blogs.cornell.edu/jamolpender/.