



UvA-DARE (Digital Academic Repository)

Markovian queues with Poisson control

Núñez-Queija, R.; Prabhu, B.J.; Resing, J.A.C.

DOI

[10.1016/j.indag.2023.03.002](https://doi.org/10.1016/j.indag.2023.03.002)

Publication date

2023

Document Version

Final published version

Published in

Indagationes Mathematicae

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Núñez-Queija, R., Prabhu, B. J., & Resing, J. A. C. (2023). Markovian queues with Poisson control. *Indagationes Mathematicae*, 34(5), 990-1013.
<https://doi.org/10.1016/j.indag.2023.03.002>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.



ELSEVIER



Available online at www.sciencedirect.com

ScienceDirect

indagationes
mathematicae

Indagationes Mathematicae 34 (2023) 990–1013

www.elsevier.com/locate/indag

Special Issue dedicated to the memory of J.W. Cohen

Markovian queues with Poisson control[☆]

R. Núñez-Queija^a, B.J. Prabhu^b, J.A.C. Resing^{c,*}

^a University of Amsterdam, Amsterdam, The Netherlands

^b LAAS-CNRS, Université de Toulouse, CNRS, INSA, Toulouse, France

^c Eindhoven University of Technology, Eindhoven, The Netherlands

Abstract

We investigate Markovian queues that are examined by a controller at random times determined by a Poisson process. Upon examination, the controller sets the service speed to be equal to the minimum of the current number of customers in the queue and a certain maximum service speed; this service speed prevails until the next examination time. We study the resulting two-dimensional Markov process of queue length and server speed, in particular two regimes with time scale separation, specifically for infinitely frequent and infinitely long examination times. In the intermediate regime the analysis proves to be extremely challenging. To gain further insight into the model dynamics we then analyse two variants of the model in which the controller is just an observer and does not change the speed of the server. © 2023 The Author(s). Published by Elsevier B.V. on behalf of Royal Dutch Mathematical Society (KWG). This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Keywords: Markovian queues; Variable service speeds; Asymptotics; Time-scale separation

1. Introduction

Consider a dynamically controlled queue in which the speed of the server can be varied in order to attain a given objective. The optimal policy for such problems usually requires adjustments of the speed based on the current queue-length. For example, it has been shown that the right trade-off between energy consumption and mean sojourn time is achieved by setting the speed to be proportional to the number of customers to the power of a constant [2,3]. These adjustments to the speed have to be made at arbitrary times which may depend upon the dynamics of the queue-length (for example, at arrival instants or departure instants). However, in practice communicating real-time state information of the queue may not always be feasible or desirable due to communication overheads. Nevertheless, it may be possible to control the

[☆] A short letter mainly devoted to Conjecture 1 of this paper appeared in Núñez-Queija et al. (2022) [11].

* Corresponding author.

E-mail addresses: nunezqueija@uva.nl (R. Núñez-Queija), balakrishna.prabhu@laas.fr (B.J. Prabhu), j.a.c.resing@tue.nl (J.A.C. Resing).

<https://doi.org/10.1016/j.indag.2023.03.002>

0019-3577/© 2023 The Author(s). Published by Elsevier B.V. on behalf of Royal Dutch Mathematical Society (KWG). This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

speed at, say, Poisson instants which are independent of the state of the queue. These control instants can be seen as moments at which a controller, oblivious of the state, decides to measure the state and take appropriate action.

In this paper, we analyse a queueing model in which the server speed is modified at Poisson instants of rate ν that are independent of the queue length evolution. Inspired from the policy in [2,3], it will be assumed that the speed is set to $\min(q, \bar{s})$ when q is the observed queue length and \bar{s} the maximum server speed. Our results can be used to compute the sub-optimality induced in the performance metrics due to control instants being different from the ones prescribed by the optimal policy. We shall call this model the *Poisson controller* model.

Our main focus is on the model with $\bar{s} = \infty$. The analysis of that model is not trivial as we show in the paper and for that reason we use the finite speed model to gain insight into the reasons for that complexity.

1.1. Contributions

The Poisson controller model can be described by a two-dimensional Markov process with state given by $(Q(t), S(t))$ where $Q(t)$ denotes the number of customers in the system at time t and $S(t)$ denotes the speed of the server at t .

For the analysis, we shall separate the two cases $\bar{s} < \infty$ and $\bar{s} = \infty$ since their treatment is based on different techniques. Moreover, as will be seen later, they give rise to different asymptotic results which are easier to present separately.

For the infinite maximum speed case (Section 2), that is $\bar{s} = \infty$, we determine the functional equation whose solution leads to the joint generating function of the steady-state process. The steady-state distribution of the queue-length conditional on the speed is shown to be equivalent to the transient distribution of a queue whose state is reset to the given speed at the control instants. The latter queue was analysed in [7] in which he gave the generating function for the steady-state queue-length. From this conditional generating function as well as the observation that the marginal distribution of Q is the same as that of S , we obtain a system of linear equations to compute this marginal distribution. Finally, we investigate the limiting behaviour of the steady-state distribution when the rate of the control process, ν , goes to 0 and to ∞ . For $\nu \rightarrow 0$, it is possible that the queue is unstable for some of the speeds. The queue-length can thus live on two different scales: (i) the fluid scale when the sampled speed is less than the arrival rate; (ii) the normal scale when the sampled speed is greater than the arrival rate.

For finite maximum speed \bar{s} (Section 3), we resort to the matrix geometric method [8] to analyse the system. Using a probabilistic interpretation, we obtain an explicit expression for the R matrix. As an illustration, we derive the steady-state probabilities for $\bar{s} = 1$. For the general case, we show how to obtain the joint generating function when $\nu \rightarrow 0$ and $\nu \rightarrow \infty$.

We end the paper with two variants of the model, the $M/M/1$ model and the $M/M/\infty$ model, in which the controller is just an observer and does not change the speed of the server. The reason that we look at these variants is that in these models, like in the model with Poisson controller, the two-dimensional Markov process also has jumps to diagonal states. However, contrary to the model with Poisson controller, for these variants an explicit expression of the joint generating function of the steady-state process, keeping track of the current queue length and the last observed queue length in this case, can be obtained by solving the corresponding functional equation.

1.2. Related work

Our model is related to queueing systems or Markov chains in random environments [9,10]. In the cited models, the arrival rate or the service rate of the queue depends upon the state of the environment which is a random process. As opposed to this, in our model, the environment (which is the speed observed by the controller) itself depends upon the queue length since it is set to the value measured at the control instants. The two variables – the queue-length and the state – thus influence the dynamics of each other in our model whereas in the classical random environment model it is only the environment that influences the dynamics of queue-length. We also mention [6] in which a random environment model with both unstable and stable speeds was analysed. They obtained the conditional generating function for the queue-length for the process on both the fluid as well as the normal scale. We obtain this type of results for the Poisson controller model.

Another related model is the model with workload dependent arrival and service rates in [4]. In that work, the rates change instantaneously with the state which is not the case for us. The finite speed Poisson controller model with only two speeds (in Section 3) is a special case of the queue with service speed adaptations analysed in [5].

2. Poisson controller with infinite maximum speed

Consider a single-server queue to which arrivals occur according to a Poisson process of rate λ . Each arrival brings with it an exponentially distributed service requirement with mean $1/\mu$. The speed of the server can be dynamically assigned values in the set $\{0, 1, 2, \dots\}$. Adjustments to the speed can be made only at control instants which are assumed to occur according to a Poisson process of rate ν independently of the arrivals and the departures. At a control instant, the speed of the server is set equal to the number of customers observed at that instant. Between any two consecutive control instants, the speed of the server remains constant at the value chosen at the earlier control instant.

Let $Q(t)$ denote the number of customers in the system at time t and $S(t) = Q(\tau_t)$, with τ_t the last control instant at or before time t . Thus, the speed at time t is maintained at $Q(\tau_t)$ until the next control instant. We do not specify the initial state $(Q(0), S(0))$ since we are interested in the steady-state behaviour of the system which is independent of the initial state. Furthermore, as our focus is on the number of customers in the system and service times are exponential, the service discipline can be any discipline in which the service order is independent of the actual service times of the customers (e.g., FCFS, LCFS or ROS).

The process $(Q(t), S(t))_{t \geq 0}$ is a Markov process with transition rates

$$(Q(t), S(t)) \rightarrow \begin{cases} (Q(t) + 1, S(t)) & \text{with rate } \lambda; \\ (Q(t) - 1, S(t)) & \text{with rate } \mu S(t); \\ (Q(t), Q(t)) & \text{with rate } \nu. \end{cases} \tag{1}$$

Fig. 1 shows the rate diagram of this Poisson controller model.

The process $(Q(t), S(t))_{t \geq 0}$ is ergodic for all possible combinations $\lambda > 0$, $\mu > 0$ and $\nu > 0$ (see also Remark 4 in Section 3) and we are interested in the steady-state behaviour of this two-dimensional Markov process. Denote with $\pi_{i,j} = \lim_{t \rightarrow \infty} P(Q(t) = i, S(t) = j)$ the steady-state probabilities and let

$$P(x, y) = \sum_{i \geq 0, j \geq 0} \pi_{i,j} x^i y^j \tag{2}$$




Fig. 1. Rate diagram of the model with Poisson controller with infinite maximum speed.

be the corresponding joint probability generating function. By (Q, S) we denote a pair of random variables with this joint probability generating function.

Proposition 1. $P(x, y)$ is the solution of the functional equation

$$(v + \lambda(1 - x))P(x, y) + \mu y \left(1 - \frac{1}{x}\right) \frac{\partial}{\partial y} [P(x, y) - P(0, y)] = vP(xy, 1). \tag{3}$$

Proof. From the transition in Fig. 1, the balance equations for the stationary probabilities can be seen to be, for $i \geq 1, j \geq 0$:

$$(\lambda + j\mu + v)\pi_{i,j} = \lambda\pi_{i-1,j} + j\mu\pi_{i+1,j} + \mathbb{1}_{\{i=j\}}v \sum_k \pi_{j,k} \tag{4}$$

and, for $i = 0, j \geq 0$:

$$(\lambda + v)\pi_{i,j} = j\mu\pi_{i+1,j} + \mathbb{1}_{\{i=j\}}v \sum_k \pi_{j,k}. \tag{5}$$

Multiplying both sides of the rate balance equations by $x^i y^j$ and summing over all possible i and j immediately leads to the equation

$$\begin{aligned} & (v + \lambda)P(x, y) + \mu y \frac{\partial}{\partial y} [P(x, y) - P(0, y)] \\ &= \lambda x P(x, y) + \frac{\mu y}{x} \frac{\partial}{\partial y} [P(x, y) - P(0, y)] + vP(xy, 1), \end{aligned}$$

which can be alternatively written as (3). \square

Corollary 1. *The marginal distribution of the speed is the same as the marginal distribution of the number of customers in the system. That is,*

$$P(1, y) = P(y, 1).$$

Proof. Substituting $x = 1$ in Proposition 1 leads to the above corollary. \square

An explanation for Corollary 1 is that the controller sees the marginal distribution of Q due to the PASTA property.¹

Since S is set to Q at the control instants, the marginal distribution of S just after these instants is the same as the marginal distribution of Q . Because the time until the next control instant is independent of S at a control instant (this time is exponentially distributed with parameter ν), the marginal distribution of S at an arbitrary instant is the same as the marginal distribution of S just after a control instant and hence also the same as the marginal distribution of Q at an arbitrary instant.

Let $\gamma_i = \sum_j \pi_{i,j}$ be the marginal distribution of the stationary queue-length. A second consequence of Proposition 1 is the following local balance when there are i customers in the queue.

Corollary 2.

$$\lambda \gamma_i = \sum_j j \mu \pi_{i+1,j}, \quad i \geq 0. \tag{6}$$

Proof. Substituting $y = 1$ in Proposition 1 gives

$$x \lambda P(x, 1) = \mu \frac{\partial}{\partial y} [P(x, y) - P(0, y)] \Big|_{y=1}. \tag{7}$$

Comparing the coefficient of x^i on either side we get the claimed result. \square

An alternative proof of this result can be obtained using the following up- and downcrossings argument. The rate at which the number of customers in the system goes from i to $i + 1$ equals the rate at which the number of customers goes from $i + 1$ to i .

Finally, we can also conclude that the expected queue-length is larger than in an $M/M/\infty$ queue.

Corollary 3.

$$\mathbb{E}[Q] > \rho, \quad \text{with } \rho = \lambda/\mu.$$

Proof. Substituting $x = 1$ in (7) yields the equation

$$\lambda = \mu (E(S) - E(S \cdot 1[Q = 0])),$$

from which we conclude that

$$E(S \cdot 1[Q > 0]) = \rho,$$

and, hence,

$$E[Q] = E[S] = E(S \cdot 1[Q > 0]) + E(S \cdot 1[Q = 0]) > \rho. \quad \square$$

¹ We use that ‘‘PASTA’’ holds for any sequence of Poisson-generated events that may or may not change the system state (see [12]). In the literature this property has been coined PASTA because it is mostly used for queueing systems at arrival instants.

We were unable to obtain an explicit solution to the functional equation (3). Next, we provide an alternative way to compute the generating function. This method is not explicit but is amenable to numerical computations. For this, we first compute the distribution of the queue-length conditioned on the speed.

2.1. The conditional distribution

Define $\sigma_j = \sum_i \pi_{i,j}$ to be the marginal distribution of the stationary server speed. Let

$$p_j(i) = \frac{\pi_{i,j}}{\sigma_j}$$

be the stationary conditional probability of having i customers in the system when the service rate is $j\mu$ and let $f_j(z)$ be the generating function of this stationary conditional distribution. Furthermore, define

$$\beta_j(v) = \frac{\lambda + j\mu + v - \sqrt{(\lambda + j\mu + v)^2 - 4\lambda(j\mu)}}{2\lambda} \tag{8}$$

Proposition 2. For the generating function $f_j(z)$ of the stationary conditional distribution we have

$$f_0(z) = \frac{v}{v + \lambda(1 - z)} \tag{9}$$

and, for $j > 0$,

$$f_j(z) = \frac{v\tilde{\beta}_j(v) \sum_{k=0}^{\infty} c_{k,j} z^k}{\lambda(1 - \tilde{\beta}_j(v)z)}, \tag{10}$$

where

$$\tilde{\beta}_j(v) = \frac{\lambda\beta_j(v)}{j\mu}, \tag{11}$$

and

$$c_{k,j} = \begin{cases} (1 - \beta_j(v))^{-1} \beta_j(v)^j & k = 0; \\ \beta_j(v)^{j-k} & k = 1, \dots, j; \\ 0 & k > j. \end{cases} \tag{12}$$

Proof. For the conditional process the following balance equations can be obtained by dividing (4) and (5) by σ_j and by noting that the $\sum_k \pi_{j,k} = \sigma_j$ (from Corollary 1 the marginal distribution of the queue-length is the same as the marginal distribution of the speed):

$$\begin{aligned} (\lambda + v)p_j(0) &= j\mu p_j(1), \\ (\lambda + v + j\mu)p_j(i) &= \lambda p_j(i - 1) + j\mu p_j(i + 1), \quad i > 0, \quad i \neq j; \\ (\lambda + v + j\mu)p_j(j) &= \lambda p_j(j - 1) + j\mu p_j(j + 1) + v. \end{aligned}$$

Hence the generating function satisfies

$$(\lambda + v + j\mu)f_j(z) - j\mu p_j(0) = \lambda z f_j(z) + v z^j + \frac{j\mu}{z} [f_j(z) - p_j(0)],$$

which leads to the equation

$$[\lambda z^2 - (\lambda + v + j\mu)z + j\mu] f_j(z) = j\mu(1 - z)p_j(0) - v z^{j+1}. \tag{13}$$




Fig. 2. Rate diagram of the conditional process.

For $j = 0$, (13) leads to

$$f_0(z) = \frac{-\nu z}{\lambda z^2 - (\lambda + \nu)z} = \frac{\nu}{\lambda + \nu - \lambda z} = \frac{\nu}{\nu + \lambda(1 - z)}.$$

For $j > 0$, let $z_1 < 1$ and $z_2 > 1$ be the zeros of the polynomial in the lefthandside of (13). For $z = z_1$ also the righthandside of (13) should be zero, hence

$$p_j(0) = \frac{\nu z_1^{j+1}}{j\mu(1 - z_1)}.$$

So we conclude that

$$\begin{aligned} f_j(z) &= \frac{\nu(1 - z)z_1^{j+1} - \nu(1 - z_1)z^{j+1}}{\lambda(z - z_1)(z - z_2)(1 - z_1)} \\ &= \frac{(\nu/z_2)}{\lambda(1 - z/z_2)} \cdot \frac{(1 - z_1)z^{j+1} - (1 - z)z_1^{j+1}}{(z - z_1)(1 - z_1)} \\ &= \frac{(\nu/z_2)}{\lambda(1 - z/z_2)} \cdot \left(\sum_{k=1}^j z_1^{j-k} z^k + \frac{z_1^j}{1 - z_1} \right). \end{aligned}$$

Now, clearly, $z_1 = \beta_j(\nu)$ and $\lambda z_1 z_2 = j\mu$. Hence, $z_2 = (j\mu)/(\lambda z_1) = 1/\tilde{\beta}_j(\nu)$ and the result follows. \square

Remark 1. The stationary queue-length process conditional on $S = j$ can readily be seen to have the stationary distribution of a process $Q_j(\cdot)$, obtained from the ordinary queue length process after eliminating time intervals for which the speed $\neq j$. Then, $Q_j(\cdot)$ behaves as the number of customers in an $M/M/1$ queue with arrival rate λ , service rate $j\mu$, and restarting in state j after $\text{exp}(\nu)$ distributed time periods. It is a one-dimensional Markov chain with rate diagram shown in Fig. 2.

Further, the stationary distribution of Q_j is the same as the transient distribution after an exponentially distributed time period with parameter ν in an $M/M/1$ queue with arrival rate λ , service rate $j\mu$ and starting at time 0 in state j . Hence, Eq. (10) can be alternatively derived using results on the transient distribution in the $M/M/1$ queue (see Section I.4.4 and in particular formula (4.27) in [7]).

Remark 2. From Proposition 2 it follows that, for $j > 0$, we have that

$$(Q | S = j) \stackrel{d}{=} [j - A_j]_+ + B_j,$$

where A_j and B_j are independent random variables with A_j geometrically distributed with parameter $\beta_j(v)$ and B_j geometrically distributed with parameter $\tilde{\beta}_j(v)$.

From the conditional generating function we can obtain the following expression for the stationary conditional distribution.

Corollary 4. For $j > 0$,

$$f_j(z) = \frac{v\tilde{\beta}_j(v)}{\lambda} \sum_{l \geq 0} \left(\sum_{k=0}^l c_{k,j} \tilde{\beta}_j(v)^{l-k} \right) z^l, \tag{14}$$

and hence

$$p_j(l) = \frac{v}{\lambda} \sum_{k=0}^l c_{k,j} \tilde{\beta}_j(v)^{l-k+1}. \tag{15}$$

The mean conditional queue-length can also be expressed explicitly.

Corollary 5.

$$\mathbb{E}[Q_j] = f'_j(1) = \frac{v\tilde{\beta}_j(v)}{\lambda} \left(\frac{\tilde{\beta}_j(v) \sum_k c_{k,j}}{(1 - \tilde{\beta}_j(v))^2} + \frac{\sum_k k c_{k,j}}{1 - \tilde{\beta}_j(v)} \right) \tag{16}$$

with

$$\sum_k c_{k,j} = (1 - \beta_j(v))^{-1}, \tag{17}$$

and

$$\sum_k k c_{k,j} = \frac{j - (j + 1)\beta_j(v) + \beta_j(v)^{j+1}}{(1 - \beta_j(v))^2} \tag{18}$$

Proof. The expression in (16) follows directly. For (17), from (12) for $c_{k,j}$ we get

$$\sum_k c_{k,j} = (1 - \beta_j(v))^{-1}. \tag{19}$$

Alternatively, since $f_j(1) = 1$, we have

$$1 = \frac{v\tilde{\beta}_j(v) \sum_k c_{k,j}}{\lambda(1 - \tilde{\beta}_j(v))} \tag{20}$$

and hence

$$\sum_k c_{k,j} = \frac{\lambda(1 - \tilde{\beta}_j(v))}{v\tilde{\beta}_j(v)} = \frac{j\mu - \lambda\beta_j(v)}{v\beta_j(v)}. \tag{21}$$

These two expressions are the same because $\lambda\beta_j(v)^2 - (\lambda + v + j\mu)\beta_j(v) + j\mu = 0$. For (18), we have

$$\sum_k k c_{k,j} = \sum_{k=1}^j k \beta_j(v)^{j-k} = \frac{j - (j + 1)\beta_j(v) + \beta_j(v)^{j+1}}{(1 - \beta_j(v))^2} \quad \square \tag{22}$$

Proposition 3.

$$\sigma_l = \sum_j \sigma_j \frac{\nu \tilde{\beta}_j(\nu)}{\lambda} \left(\sum_{k=0}^l c_{k,j} \tilde{\beta}_j(\nu)^{l-k} \right), \forall l \geq 0. \tag{23}$$

Proof. Observe that we can rewrite $P(x, y)$ in terms of $f_j(x)$ in the following way.

$$P(x, y) = \sum_j \sigma_j f_j(x) y^j. \tag{24}$$

Now use $P(1, y) = P(y, 1)$ to obtain

$$\sum_j \sigma_j f_j(1) y^j = \sum_j \sigma_j f_j(y). \tag{25}$$

Substituting (14) in (25) and comparing the coefficient of y^l on either side, we get the desired system of linear equations for σ_j . \square

Remark 3. Eq. (23) can be interpreted as the balance equation of the embedded Markov chain of the server speed at control instants (see also (15)).

Since (23) is an infinite set of linear equations, it is not straightforward to solve them. Therefore, we look at the asymptotics of $P(x, y)$ when the rate of control ν goes to either ∞ or to 0.

2.2. Asymptotics for $\nu \rightarrow \infty$

When the rate of control $\nu \rightarrow \infty$, intuitively one expects the speed to be the same as the queue length since measurements are being made at a much faster rate compared to rates of variations in the queue-length. The two-dimensional process will live mainly on the diagonal states. The following result formalizes this intuition.

Proposition 4. If $\nu \rightarrow \infty$, then $P(x, y) \rightarrow e^{\rho(xy-1)}$.

Proof. First of all, remark that, when $\nu \rightarrow \infty$, functional equation (3) reduces to $P(x, y) = P(xy, 1)$ from which we conclude that $\pi_{i,j} \rightarrow 0$ for $i \neq j$. Next, assume that $\pi_{i,j}$ has the following analytic expansion

$$\pi_{i,j} = \sum_m \pi_{i,j}^{(m)} \nu^{-m}. \tag{26}$$

Then we have that $\pi_{i,j}^{(0)} = 0$ for $i \neq j$ and furthermore $\pi_{i,j}^{(1)} = 0$ for $|i - j| \geq 2$. From the balance equation for state $(i, i - 1)$ we obtain $\pi^{(1)}(i, i - 1) = \lambda \pi^{(0)}(i - 1, i - 1)$. From the balance equation for state $(i - 1, i)$ we obtain $\pi_{i-1,i}^{(1)} = i \mu \pi_{i,i}^{(0)}$. Furthermore, from a cut between speed level $i - 1$ and speed level i we obtain $\pi_{i-1,i}^{(1)} = \pi_{i,i-1}^{(1)}$. Combining these three equations leads to the relation $\pi_{i,i}^{(0)} = (\lambda/i\mu) \pi_{i-1,i-1}^{(0)}$ which together with the normalization equation $\sum_i \pi_{i,i}^{(0)} = 1$ leads to the solution $\pi_{i,i}^{(0)} = (\rho^i/i!) e^{-\rho}$ and hence to the result that $\lim_{\nu \rightarrow \infty} P(x, y) = e^{\rho(xy-1)}$. \square




Fig. 3. Trajectories of the rescaled queue-length and the rescaled server-speed in the limit $\nu \rightarrow 0$. The vertical dashed lines are control instants.

2.3. Asymptotics for $\nu \rightarrow 0$

On the other extreme, when the measurements are performed at a much slower rate, time-scale separation between the queue-length and the server-speed occurs. Since the server-speed does not change between two control instants, the queue-length evolves on a faster time-scale compared to the server-speed. In the spatial dimension, both the queue-length and the server-speed can evolve on two scales: fluid scale on which they are $O(\nu^{-1})$ and the normal scale on which they are $O(1)$. The trajectories on these processes rescaled by ν will be cyclic as shown in Fig. 3.

Assume time has been rescaled by a factor ν so that control instants form a Poisson process of rate 1. A cycle begins when the measured queue-length is 0. For this speed, the queue-length process being unstable, it grows linearly on the fluid scale at rate λ until the next measurement instant. That is, during this period, the limiting process $\lim_{\nu \rightarrow 0} \nu Q(t)$ will grow linearly. The queue-length being on the fluid scale, the next measurement will set the speed to a value $O(\nu^{-1})$. This will bring the queue-length to 0 instantaneously in time $O(\nu)$. The speed is now $O(\nu^{-1})$ whereas the arrival rate is λ . So, the queue-length will remain at 0 (on the normal scale and not just on the fluid scale) until the next measurement instant at which point the speed will be set to 0. At this point, a new cycle will begin. We were unable to formalize this intuition and leave it a conjecture.

Define $\hat{P}(x, y) := \lim_{\nu \rightarrow 0} P(x^\nu, y^\nu)$ to be the generating function of the scaled pair of random variables $(\nu Q, \nu S)$ in the limit $\nu \rightarrow 0$.

Conjecture 1.

$$\hat{P}(x, y) = \frac{1}{2} \frac{1}{1 - \lambda \log(x)} + \frac{1}{2} \frac{1}{1 - \lambda \log(y)}. \tag{27}$$

As $\nu \rightarrow 0$, the probability mass of the joint process concentrates around the two axes $Q = 0$ and $S = 0$. When Q is on the fluid scale, $S = 0$ while when $Q = 0$, S is on the fluid scale.




Fig. 4. Rate diagram of the model with $\bar{s} = 2$.

than the expected duration of the corresponding excursion in the model with maximal speed $\bar{s} < \infty$ (essentially because during excursions outside the set U the server in the model with infinite maximal speed $\bar{s} = \infty$ works always at least as fast as in the model with finite maximal speed $\bar{s} < \infty$). As in the latter model this expected duration will be finite, because $\lambda < \bar{s}\mu$, the expected duration will also be finite in the model with maximal speed $\bar{s} = \infty$ and hence the system will be ergodic.

The vector $\pi_n = [\pi_{n,0}, \pi_{n,1}, \dots, \pi_{n,\bar{s}}]$ containing the stationary probabilities for the different states with n customers in the system, are, for $n \geq \bar{s} - 1$, of the form

$$\pi_n = \pi_{\bar{s}-1} R^{n-\bar{s}+1}, \tag{30}$$

where R is the minimal non-negative solution of the matrix equation

$$A_0 + RA_1 + R^2A_2 = 0,$$

where $A_1 = A_{1,\bar{s}}$.

The matrix R can be explicitly calculated for this model. It is a diagonal matrix with additional non-zero values in the last column.

Proposition 5. For the matrix R we have that

$$R = \begin{pmatrix} \frac{\lambda}{\lambda+\nu} & & & & \frac{\lambda}{\bar{s}\mu} \\ & \frac{\lambda\beta_1(\nu)}{\mu} & & & \frac{\lambda(1-\beta_1(\nu))}{\bar{s}\mu} \\ & & \ddots & & \vdots \\ & & & \frac{\lambda\beta_i(\nu)}{i\mu} & \frac{\lambda(1-\beta_i(\nu))}{\bar{s}\mu} \\ & & & & \vdots \\ & & & & \frac{\lambda}{\bar{s}\mu} \end{pmatrix}. \tag{31}$$

Proof. We use the fact that $R = A_0N = \lambda IN = \lambda N$, where the matrix N contains, for an arbitrary $n \geq \bar{s}$, the expected sojourn times in the different states on level n , starting from the different states on level n , before the first visit to level $n - 1$ (see Theorem 6.4.1 in [8]). Now,

clearly $N_{i,j} = 0$ for all $i \neq j$ and $j \neq \bar{s}$ because for these values of i and j you can only reach state (n, j) starting from state (n, i) via level $n - 1$. Furthermore, $N_{0,0} = \frac{1}{\lambda + \nu}$, because if you start in state $(n, 0)$ you stay there an exponentially distributed time with parameter $\lambda + \nu$ and will never return there before the first visit to level $n - 1$. The value $N_{\bar{s},\bar{s}}$ is equal to the expected time there is one customer in the system during a busy period of an $M/M/1$ queue with arrival rate λ and service rate $\bar{s}\mu$. Hence, $N_{\bar{s},\bar{s}} = \frac{1}{\lambda + \bar{s}\mu} + \frac{\lambda}{\lambda + \bar{s}\mu} N_{\bar{s},\bar{s}}$ and so $N_{\bar{s},\bar{s}} = \frac{1}{\bar{s}\mu}$. Furthermore, we have that $N_{0,\bar{s}} = N_{\bar{s},\bar{s}}$ because starting from state $(n, 0)$ we certainly reach state (n, \bar{s}) before the first visit to level $n - 1$. For $i \neq 0$ and $i \neq \bar{s}$ we have that $N_{i,\bar{s}} = (1 - \beta_i(\nu))N_{\bar{s},\bar{s}}$, because $\beta_i(\nu)$ is the probability that the busy period in an $M/M/1$ queue with arrival rate λ and service rate $i\mu$ is smaller than an independent exponential random variable with parameter ν . Hence $1 - \beta_i(\nu)$ is the probability that we reach state (n, \bar{s}) before the first visit to level $n - 1$ when we start in a state (n, i) with $i \neq 0$ and $i \neq \bar{s}$.

Finally, the quantity $N_{j,j}$ for $j \neq 0$ and $j \neq \bar{s}$ is equal to the expected time there is one customer in the system during a busy period of an $M/M/1$ queue with arrival rate λ and service rate $j\mu$ and in which a disaster occurs removing all the customers in the system after an exponentially distributed time with parameter ν . We have that $N_{j,j} = \frac{1}{\lambda + \nu + j\mu} + \frac{\lambda}{\lambda + \nu + j\mu} \beta_j(\nu)N_{j,j}$ and so $N_{j,j} = \frac{1}{\lambda(1 - \beta_j(\nu)) + \nu + j\mu} = \frac{\beta_j(\nu)}{j\mu}$. \square

Since R is of the above form, R^n has the following simple formula which can be used for computing π_n in (30).

Corollary 6. *Let $R_{i,j}$ be the element (i,j) of R . Then,*

$$R^n = \begin{pmatrix} R_{0,0}^n & & & R_{0,\bar{s}} \frac{(R_{\bar{s},\bar{s}}^n - R_{0,0}^n)}{R_{\bar{s},\bar{s}} - R_{0,0}} \\ & \ddots & & \vdots \\ & & R_{i,i}^n & R_{i,\bar{s}} \frac{(R_{\bar{s},\bar{s}}^n - R_{i,i}^n)}{R_{\bar{s},\bar{s}} - R_{i,i}} \\ & & & \ddots \\ & & & R_{\bar{s},\bar{s}}^n \end{pmatrix}. \tag{32}$$

From the above corollary, it can be inferred that the joint probability vector of level $n \geq \bar{s}$ and speeds other than \bar{s} is a geometric term with rate $R_{i,i}$. For level $n \geq \bar{s}$ and speed \bar{s} , the probability is a mixture of geometric terms: terms with rate $R_{i,i}$, $i < \bar{s}$, and one with rate $R_{\bar{s},\bar{s}}$.

The probability vectors π_n , $n \leq \bar{s} - 1$, can now be computed using a system of linear equations and the normalization equation (see, e.g., [8]). It then follows that the joint generating function can be expressed as

Proposition 6.

$$P(x, y) = \sum_{n < \bar{s} - 1, j} \pi_{n,j} x^n y^j + \pi_{\bar{s}-1} (I - Rx)^{-1} x^{\bar{s}-1} \cdot \mathbf{y}^T, \tag{33}$$

with $\mathbf{y} = [1, y, y^2, \dots, y^{\bar{s}}]$

Observe that from the particular form of R in Proposition 5, $(I - Rx)^{-1}$ can be explicitly computed:

$$(I - Rx)^{-1} = \begin{pmatrix} \frac{1}{1-R_{0,0}x} & & & \frac{R_{0,\bar{s}}x}{(1-R_{\bar{s},\bar{s}}x)(1-R_{0,0}x)} \\ & \ddots & & \vdots \\ & & \frac{1}{1-R_{i,i}x} & \frac{R_{i,\bar{s}}x}{(1-R_{\bar{s},\bar{s}}x)(1-R_{i,i}x)} \\ & & & \ddots \\ & & & & \frac{1}{1-R_{\bar{s},\bar{s}}x} \end{pmatrix}. \tag{34}$$

As an illustration of this analysis, we obtain the steady-state probabilities for $\bar{s} = 1$. In this case, our model is a special case of the model with two different service speeds discussed in [5].

For $\bar{s} = 1$, (30) and Proposition 5 reduce to

$$[\pi_{n,0}, \pi_{n,1}] = [\pi_{0,0}, \pi_{0,1}]R^n, \tag{35}$$

with

$$R = \begin{pmatrix} \frac{\lambda}{\lambda+\nu} & \frac{\lambda}{\mu} \\ 0 & \frac{\lambda}{\mu} \end{pmatrix}. \tag{36}$$

Furthermore,

$$\pi_{0,1} = \frac{\lambda}{\nu}\pi_{0,0} \tag{37}$$

and

$$R^n = \begin{pmatrix} \left(\frac{\lambda}{\lambda+\nu}\right)^n & \frac{\lambda+\nu}{\lambda+\nu-\mu} \left[\left(\frac{\lambda}{\mu}\right)^n - \left(\frac{\lambda}{\lambda+\nu}\right)^n \right] \\ 0 & \left(\frac{\lambda}{\mu}\right)^n \end{pmatrix}. \tag{38}$$

Hence, we obtain

$$\pi_{n,0} = \left(\frac{\lambda}{\lambda+\nu}\right)^n \pi_{0,0}, \tag{39}$$

$$\pi_{n,1} = \left[\left(\frac{\lambda}{\nu} + \frac{\lambda+\nu}{\lambda+\nu-\mu}\right) \left(\frac{\lambda}{\mu}\right)^n - \frac{\lambda+\nu}{\lambda+\nu-\mu} \left(\frac{\lambda}{\lambda+\nu}\right)^n \right] \pi_{0,0}. \tag{40}$$

The constant $\pi_{0,0}$ follows of course from the normalization equation, and takes the value

$$\pi_{0,0} = \frac{\nu(\mu-\lambda)}{\mu(2\lambda+\nu)}. \tag{41}$$

Remark that

$$P(S = 0) = \frac{(\mu-\lambda)(\lambda+\nu)}{\mu(2\lambda+\nu)}, \quad P(S = 1) = \frac{\lambda(\lambda+\mu+\nu)}{\mu(2\lambda+\nu)} \tag{42}$$

and that

$$P(Q = 0, S = 1) = \pi_{0,1} = \frac{\lambda(\mu-\lambda)}{\mu(2\lambda+\nu)} \tag{43}$$

so that

$$P(Q > 0, S = 1) = \frac{\lambda(\lambda+\mu+\nu)-\lambda(\mu-\lambda)}{\mu(2\lambda+\nu)} = \frac{\lambda(2\lambda+\nu)}{\mu(2\lambda+\nu)} = \frac{\lambda}{\mu} \tag{44}$$

as it should be.

For the joint probability generating function we have

$$P(x, y) = \sum_{n=0}^{\infty} \pi_{n,0} x^n + \sum_{n=0}^{\infty} \pi_{n,1} x^n y \tag{45}$$

$$= \pi_{0,0} \left[\frac{\lambda + \nu}{\nu + \lambda(1-x)} - \frac{\lambda + \nu}{\lambda + \nu - \mu} \cdot \frac{\lambda + \nu}{\nu + \lambda(1-x)} \cdot y + \left(\frac{\lambda}{\nu} + \frac{\lambda + \nu}{\lambda + \nu - \mu} \right) \cdot \frac{\mu}{\mu - \lambda x} \cdot y \right] \tag{46}$$

3.1. Asymptotics for $\nu \rightarrow \infty$

In this asymptotic regime, intuitively, as in Section 2, the joint probability will be non-zero only on states of the type (q, q) for $q < \bar{s} - 1$ and (q, \bar{s}) for $q \geq \bar{s}$. The stationary probability of seeing queue-length q will be that of an $M/M/\bar{s}$ queue. Formally,

Proposition 7.

$$\lim_{\nu \rightarrow \infty} \pi_{q,j} = \begin{cases} \pi_{0,0} \frac{\rho^q}{j! \bar{s}^{(q-j)}}, & j = \min(q, \bar{s}); \\ 0, & \text{otherwise.} \end{cases} \tag{47}$$

The proof is based on arguments from singular perturbation theory [1]. Here, we give a sketch of the proof. Rescale time by ν so that the control rate is 1, the arrival rate is $\lambda \nu^{-1}$ and the service requirement has rate $\mu \nu^{-1}$. With this transformation, the generator in (29) can be written as

$$G = G_0 + \epsilon G_1 \tag{48}$$

with $\epsilon = \nu^{-1}$. Note that G_0 is a block diagonal matrix. That is, upon setting $\epsilon = 0$, each block becomes a separate ergodic class. Thus, this chain is singularly perturbed for $\epsilon > 0$. The limiting distribution for $\epsilon \rightarrow 0$ can be computed from the stationary distribution of each block of G_0 as well as that of the aggregated chain (see Sec. 4 in [1]). It can be easily checked that the stationary distribution of block q of G_0 is $\mathbf{e}_{\min(q, \bar{s})}$, where \mathbf{e}_i is the unit vector with 1 in the i th column. This is true since any observation in block q sets the speed to $\min(q, \bar{s})$. Moreover, the aggregated chain has the dynamics of an $M/M/\bar{s}$ queue with arrival rate λ and service requirement μ^{-1} . Combining these two observations, the proof follows.

3.2. Asymptotics for $\nu \rightarrow 0$

We shall use similar arguments as in Section 2.3 to obtain the generating function for the finite maximum speed case as $\nu \rightarrow 0$. In the spatial dimension, the queue-length process can be either on the fluid scale or on the normal scale depending on whether the speed of the server is smaller or larger than λ . Unlike in the infinite maximum speed case, here the speed always remains on the normal scale.

In the following, we shall assume that time has been rescaled by a factor ν so that the control instants happen according to a Poisson process of rate 1. In order to derive various quantities related to the queue-length and the speed, we shall use an intuitive argument based on time-scale separation to first compute the marginal distribution of the server-speed. From the marginal distribution, the generating function of the joint process will be obtained using (24).




Fig. 5. Trajectories of the queue-length on the fluid scale and the server speed in the limit $\nu \rightarrow 0$. The vertical dashed lines are control instants. Trajectories that occur on the normal scale all collapse to 0.

Define

$$S^+ = \{j : j > \rho\}, \quad S^- = \{j : j \leq \rho\}. \tag{49}$$

The set S^+ contains the speeds for which the queue is positive recurrent while S^- is its complement set.

The marginal distribution of the speed will be computed from the stationary distribution of the Markov chain embedded at the control instants (and one more point, which will be explained later) and the renewal reward theorem. Let a cycle denote the time between two consecutive observations in S^- . An illustration of trajectories of the *queue-length on the fluid scale* and the speed is shown in Fig. 5. Until the first control instant after the start of a cycle, the queue-length grows linearly on the fluid scale with rate $\lambda - j\mu$ where j is the queue-length sampled at the start of the cycle. Hence, at the first control instant the speed will necessarily be set to \bar{s} , which is in S^+ . Since the queue-length process is now stable, it will decrease with rate $\bar{s}\mu - \lambda$ for one or more control instants until the fluid hits 0. (Notice that control instants do not affect the server working at constant speed \bar{s} while the queue is at the fluid level, ensuring that level 0 will be reached.) The queue-length process will now evolve on the normal scale for one or more control instants until a speed from S^- is sampled. At this point, a new cycle will begin. A cycle can thus be decomposed into three phases: the *fluid unstable* phase during which the fluid grows; the *fluid stable* phase during which the fluid drains; and the *normal* phase during which the queue length lives on the normal scale.

Let \hat{S}_n be the server-speed process embedded just after control instants except in the *fluid stable* state when the process is embedded at instants when the fluid hits 0. Observe that \bar{s} is the only state in which the queue-length process can be in both the fluid scale as well as the normal scale. For the computation of the generating function, it will be convenient to compute

the marginal probability of the server-speed being \bar{s} separately for each of the two scales of the queue-length process. For this, define the server-speed \bar{s}_f which indicates the maximum speed when the queue-length is on the fluid scale. We shall use \bar{s} as before for the maximum speed when the queue-length is on normal scale. \hat{S}_n is a discrete-time Markov chain taking values in $\mathcal{S} \cup \{\bar{s}_f\}$, with transition probability matrix

$$M = \begin{matrix} & \mathcal{S}^- & \mathcal{S}^+ & \bar{s}_f \\ \begin{matrix} \mathcal{S}^- \\ \mathcal{S}^+ \\ \bar{s}_f \end{matrix} & \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{1} \\ V_0 & V & \mathbf{0} \\ \mathbf{0} & \mathbf{b} & 0 \end{bmatrix} \end{matrix}. \tag{50}$$

The elements (i, j) of matrix M for $i \in \mathcal{S}^+$ and $j \in \mathcal{S}^- \cup \mathcal{S}^+$ (these are all entries of the matrices V_0 and V) can be deduced directly: $M_{i,j}$ is the probability of setting the speed to $j \in \mathcal{S}^- \cup \mathcal{S}^+$ at the next control instant given that the current speed is $i \in \mathcal{S}^+$. Because of time-scale separation, between two control instants in the normal phase, given the server-speed to be $i \in \mathcal{S}^+$, the queue-length process is a stationary $M/M/1$ queue with arrival rate λ and service rate $i\mu$. Thus,

$$M_{i,j} = \begin{cases} (1 - \rho_i)\rho_i^j, & j \neq \bar{s}; \\ \rho_i^{\bar{s}}, & j = \bar{s}. \end{cases} \tag{51}$$

with $\rho_i = \lambda/(i\mu)$. The vector $\mathbf{b} = [0, 0, \dots, 0, 1]$ since from \bar{s}_f the chain can only go to \bar{s} (at the additional embedded time instant when the fluid hits 0).

Let ψ be the stationary distribution of \hat{S}_n . Then, by the renewal reward theorem, the marginal distribution of the speed in the limit $\nu \rightarrow 0$ can be written in terms of ψ as follows:

$$\sigma_j = \frac{\tau_j \psi_j}{\sum_{k \in \mathcal{S} \cup \{\bar{s}_f\}} \tau_k \psi_k}, \quad j \in \mathcal{S} \cup \{\bar{s}_f\}, \tag{52}$$

where τ_j is the average time spent in state j before the next jump. We remind the reader that the state \bar{s} has been split in two: \bar{s}_f refers to the maximum speed when the queue-length is on the fluid scale while \bar{s} refers to maximum speed when the queue-length is on the normal scale. Note that for all states except \bar{s}_f , the jumps occur after an exponentially distributed time of rate 1. The only unknown quantity is thus $\tau_{\bar{s}_f}$, which can be written as

$$\tau_{\bar{s}_f} = \frac{\sum_{j \in \mathcal{S}^-} \tau_{\bar{s}_f,j} \psi_j}{\sum_{j \in \mathcal{S}^-} \psi_j}. \tag{53}$$

where $\tau_{\bar{s}_f,j}$ is the expected time spent in \bar{s}_f conditioned on the speed being j in the preceding *fluid unstable* phase. It can be seen that, given speed j in the *fluid unstable* phase, the amount of fluid at the end of this phase is distributed as $\exp((\lambda - j\mu)^{-1})$. The expected amount of fluid at the start of *fluid stable* phase is thus $\lambda - j\mu$. The speed in the *fluid stable* phase being \bar{s} , the expected time it takes to drain this fluid is

$$\tau_{\bar{s}_f,j} = \frac{\lambda - j\mu}{\bar{s}\mu - \lambda}. \tag{54}$$

Let us decompose the vector ψ in the same manner as the matrix M : $\psi = [\psi^-, \psi^+, \psi_{\bar{s}_f}]$. The component ψ^- thus contains the stationary probabilities of the states in \mathcal{S}^- . The other components are similarly defined. The following result presents formulas for ψ and $\tau_{\bar{s}_f}$. Its proof is by a simple check and is therefore omitted.

Proposition 8. Let $\theta = \mathbf{b}(I - V)^{-1}$ and $\kappa = \theta \cdot \mathbf{1}^T$. Then,

$$\psi^- = \frac{\theta V_0}{2 + \kappa}, \quad \psi^+ = \frac{\theta}{2 + \kappa}, \quad \bar{\psi}_{\bar{s}_f} = \frac{1}{2 + \kappa}, \tag{55}$$

and

$$\tau_{\bar{s}_f} = \frac{1}{\bar{s}\mu - \lambda} (\lambda - \theta V_0 J^T \mu) \tag{56}$$

with $J = [0, 1, \dots, |\mathcal{S}^-|]$.

Intuitively, θ computes the number of visits of \hat{S}_n to states in \mathcal{S}^+ starting in state \bar{s}_f while remaining in \mathcal{S}^+ . Upon leaving \mathcal{S}^+ , \hat{S}_n enters states in \mathcal{S}^- with probability distribution θV_0 , initiating a *fluid unstable* phase. It leaves \mathcal{S}^- after one epoch (corresponding to the first next control instant) and visits \bar{s}_f where it stays for one epoch before entering \mathcal{S}^+ through the state \bar{s} (which corresponds to the fluid hitting 0).

We now have all the ingredients to compute the generating function $P(x, y)$. Define $\hat{P}(x, y) = \lim_{\nu \rightarrow 0} P(x^\nu, y)$ for the generating function that captures the queue-length on the fluid scale and $\tilde{P}(x, y) = \lim_{\nu \rightarrow 0} P(x, y)$ to be the one that captures the normal scale queue-length. From (24) and (52), we get

$$\hat{P}(x, y) = \sum_{j \in \mathcal{S}^-} \frac{\sigma_j}{1 - (\lambda - j\mu) \log(x)} \left(y^j + \frac{\lambda - j\mu}{\bar{s}\mu - \lambda} y^{\bar{s}} \right) + \sum_{j \in \mathcal{S}^+} y^j \sigma_j, \tag{57}$$

$$\tilde{P}(x, y) = \sum_{j \in \mathcal{S}^+} \frac{1 - \rho_j}{1 - \rho_j x} y^j \sigma_j. \tag{58}$$

On the fluid scale, conditional on speed being $j \in \mathcal{S}^-$, the queue-length is an exponential random variable of rate $\lambda - j\mu$. This is also true in the fluid stable phase. The main difference between the two phases is the time spent in there. Since for a *fluid unstable* phase of length 1, the corresponding length of the *fluid stable* phase is $\tau_{\bar{s}_f}$, we get the terms for $j \in \mathcal{S}^-$ in (57). The terms for \mathcal{S}^+ are a direct consequence of the fact that on the fluid scale, a stable queue is always of length 0 and only the marginal distribution of the speeds appears. On the normal scale, the queue-length is just a stationary $M/M/1$ queue which gives the coefficient of $y^j \sigma_j$ in (58).

Remark 5. The above asymptotic analysis cannot be used to obtain the results in Section 2.3 for the infinite speed case by taking $\bar{s} \rightarrow \infty$. The difficulty comes from the fact that when $\bar{s} = \infty$, the speed sampled after the fluid (unstable) phase is an exponential random variable. On the other hand, in the finite \bar{s} case, this speed is always a fixed value. The limit of this sequence of deterministic values is unable to capture the exponential random variable at $\bar{s} = \infty$. One will have to scale \bar{s} as ν^{-1} to get back the results of Section 2.3.

Next we apply the above arguments to the special case $\bar{s} = 1$ and show that we indeed obtain the same expression for the generating function as the one in (46) with $\nu \rightarrow 0$.

For $\bar{s} = 1$, the transition matrix of the embedded chain becomes

$$M = \begin{matrix} & 0 & 1 & 1_f \\ \begin{matrix} 0 \\ 1 \\ 1_f \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 1 - \rho & \rho & 0 \\ 0 & 1 & 0 \end{bmatrix}, \end{matrix} \tag{59}$$

for which $\psi = (1+2(1-\rho))^{-1}[1-\rho, 1, 1-\rho]$, and $\tau_{1_f} = \rho(1-\rho)^{-1}$. The marginal distribution of the speed is thus

$$\sigma = \left[\frac{1-\rho}{2}, \frac{1}{2}, \frac{\rho}{2} \right]. \tag{60}$$

We now check that the above values are consistent with (46) with $\nu \rightarrow 0$. Setting x to x^ν and taking the limit $\nu \rightarrow 0$ in (46), we get

$$\hat{P}(x, y) = \frac{1-\rho}{2} \cdot \frac{1}{1-\lambda \log(x)} + \frac{\rho}{2} \cdot \frac{1}{1-\lambda \log(x)} \cdot y + \frac{1}{2} \cdot y, \tag{61}$$

which is the same expression as the one when (60) is substituted in (57).

The queue length process is with probability $\frac{1-\rho}{2}$ on the fluid scale with speed 0, it is with probability $\frac{\rho}{2}$ on the fluid scale with speed 1, and it is with probability $\frac{1}{2}$ on the normal scale with speed 1. Furthermore, remark that when the queue length process is on the normal scale, the number of customers in the system is geometrically distributed with parameter ρ and hence the probability that the server is working on this scale is ρ . In total, the fraction of time the server is working is $\frac{1-\rho}{2} \cdot 0 + \frac{\rho}{2} \cdot 1 + \frac{1}{2} \cdot \rho = \rho$ as it should be.

4. Other models with transitions to diagonal states

In this section, we analyse two models where the controller is just an observer and does not modify the speed of the server.

4.1. The M/M/1 queue with a Poisson observer

Consider the M/M/1 queue with arrival rate λ , exponentially distributed service times with parameter μ and with an observer arriving to the system according to a Poisson process with rate ν . The server always works with speed 1 (so no speed adaptations will occur) and assume $\lambda < \mu$.

In this case, the process $(Q(t), S(t))_{t \geq 0}$ is a Markov process with transitions

$$(Q(t), S(t)) \rightarrow \begin{cases} (Q(t) + 1, S(t)) & \text{with rate } \lambda; \\ (Q(t) - 1, S(t)) & \text{with rate } \mu; \\ (Q(t), Q(t)) & \text{with rate } \nu. \end{cases} \tag{62}$$

Fig. 6 shows the rate diagram of the M/M/1 queue with Poisson observer.

The balance equations for the stationary probabilities are

$$\begin{aligned} (\lambda + \mu + \nu)\pi_{i,j} &= \lambda\pi_{i-1,j} + \mu\pi_{i+1,j} + \mathbb{1}_{\{j=i\}}\nu \sum_k \pi_{i,k}, \quad i \geq 1, j \geq 0; \\ (\lambda + \nu)\pi_{0,j} &= \mu\pi_{1,j} + \mathbb{1}_{\{j=0\}}\nu \sum_k \pi_{0,k}, \quad j \geq 0, \end{aligned}$$

Proposition 9. *Generating function $P(x, y)$ satisfies the functional equation*

$$\left[\nu + \lambda(1-x) + \mu \left(1 - \frac{1}{x}\right) \right] P(x, y) = \mu \left(1 - \frac{1}{x}\right) P(0, y) + \nu P(xy, 1). \tag{63}$$

The solution of functional equation (63) is given by

$$P(x, y) = \frac{\frac{1-\frac{1}{x}}{1-x} \nu \frac{\mu-\lambda}{\mu-\lambda x y} + \nu \frac{\mu-\lambda}{\mu-\lambda x y}}{\nu + \lambda(1-x) + \mu \left(1 - \frac{1}{x}\right)} \tag{64}$$




Fig. 6. Rate diagram of the $M/M/1$ queue with Poisson observer.

with $0 < x_1 < 1$ solution of the equation $\lambda x^2 - (\lambda + \mu + \nu)x + \mu = 0$.

Proof. Multiplying both sides of the rate balance equations by $x^i y^j$ and summing over all possible i and j immediately leads to the equation

$$(v + \lambda)P(x, y) + \mu [P(x, y) - P(0, y)] = \lambda x P(x, y) + \frac{\mu}{x} [P(x, y) - P(0, y)] + \nu P(x, y),$$

which can be alternatively written as (63). The proof of (64) makes use of the fact that $P(x, y) = \frac{\mu - \lambda}{\mu - \lambda x y}$ and the fact that if $0 < x_1 < 1$ is a solution of the equation $\lambda x^2 - (\lambda + \mu + \nu)x + \mu = 0$ then

$$\mu \left(1 - \frac{1}{x_1}\right) P(0, y) + \nu \frac{\mu - \lambda}{\mu - \lambda x_1 y} = 0. \quad \square$$

Corollary 7 (Asymptotics when $\nu \rightarrow \infty$). From (64) and the fact that $x_1 \rightarrow 0$ if $\nu \rightarrow \infty$, it immediately follows that $P(x, y) \rightarrow \frac{\mu - \lambda}{\mu - \lambda x y}$ when $\nu \rightarrow \infty$.

Corollary 8 (Asymptotics when $\nu \rightarrow 0$). From (64) and the fact that $x_1 \rightarrow 1$ and $\nu / (\frac{1}{x_1} - 1) \rightarrow \mu - \lambda$, if $\nu \rightarrow 0$, it immediately follows that $P(x, y) \rightarrow \frac{\mu - \lambda}{\mu - \lambda x} \cdot \frac{\mu - \lambda}{\mu - \lambda y}$ when $\nu \rightarrow 0$.

The above results can be explained intuitively by the fact that the observer does not influence the queue-length. In the limit $\nu \rightarrow 0$, it just samples from the stationary distribution of an $M/M/1$ queue. Hence, in stationarity, the S and Q processes are distributionally equivalent to two independent $M/M/1$ queues.




Fig. 7. Rate diagram of the $M/M/\infty$ queue with Poisson observer.

Remark 6. The joint steady-state distribution of (Q_t, S_t) is equal to the joint steady-state distribution of $(Q_{\tau_t+Y}, Q_{\tau_t})$, where the random time $Y = t - \tau_t$ is exponentially distributed with parameter ν . So, $P(x, y)$ is the probability generating function of the joint steady-state distribution of the queue length in an $M/M/1$ queue at two different time points, where the distance between the two time points is exponentially distributed with parameter ν . Formula (64) can be alternatively derived using results on the transient distribution in the $M/M/1$ queue (see Section I.4.4 and in particular formula (4.27) in [7]).

4.2. The $M/M/\infty$ queue with an observer

Next, we consider the infinite server model with an observer arriving to the system according to a Poisson process with rate ν . As before, customers arrive according to a Poisson process with rate λ and they require exponentially distributed service times with parameter μ .

The process $(Q(t), S(t))_{t \geq 0}$ is a Markov process with transitions

$$(Q(t), S(t)) \rightarrow \begin{cases} (Q(t) + 1, S(t)) & \text{with rate } \lambda; \\ (Q(t) - 1, S(t)) & \text{with rate } \mu Q(t); \\ (Q(t), Q(t)) & \text{with rate } \nu. \end{cases} \tag{65}$$

Fig. 7 shows the rate diagram of the $M/M/\infty$ queue with Poisson observer.

The balance equations for the stationary probabilities are

$$(\lambda + i\mu + \nu)\pi_{i,j} = \lambda\pi_{i-1,j} + (i + 1)\mu\pi_{i+1,j} + \mathbb{1}_{\{i=j\}}\nu \sum_k \pi_{j,k}, \quad i \geq 1, j \geq 0;$$

$$(\lambda + \nu)\pi_{i,j} = (i + 1)\mu\pi_{i+1,j} + \mathbb{1}_{\{i=j\}}\nu \sum_k \pi_{j,k}, \quad i = 0, \quad j \geq 0,$$

Proposition 10. *Generating function $P(x, y)$ is solution of the functional equation*

$$(\nu + \lambda(1 - x))P(x, y) + \mu(x - 1)\frac{\partial}{\partial x}P(x, y) = \nu P(xy, 1). \tag{66}$$

The solution of functional equation (66) is given by

$$P(x, y) = e^{\rho(x-1)} \cdot e^{\rho(y-1)} \cdot \sum_{k=0}^{\infty} \frac{\nu}{\nu + k\mu} \frac{\rho^k}{k!} (x - 1)^k (y - 1)^k. \tag{67}$$

Proof. Multiplying both sides of the rate balance equations by $x^i y^j$ and summing over all possible i and j immediately leads to the equation

$$(\nu + \lambda)P(x, y) + \mu x \frac{\partial}{\partial x}P(x, y) = \lambda x P(x, y) + \mu \frac{\partial}{\partial x}P(x, y) + \nu P(xy, 1),$$

which can be alternatively written as (66).

Now, use the fact that

$$P(xy, 1) = e^{\rho(xy-1)} = e^{\rho(x-1)(y-1)} \cdot e^{\rho(x-1)} \cdot e^{\rho(y-1)},$$

and introduce the function $h(x, y)$ via

$$P(x, y) = h(x, y) \cdot e^{\rho(x-1)} \cdot e^{\rho(y-1)}.$$

Then, using (66), we obtain

$$\nu h(x, y) + \mu(x - 1)\frac{\partial}{\partial x}h(x, y) = \nu e^{\rho(x-1)(y-1)}. \tag{68}$$

If we now write

$$h(x, y) = \sum_{k=0}^{\infty} \sum_{\ell=0}^{\infty} h_{k,\ell} (x - 1)^k (y - 1)^\ell,$$

then from (68) we obtain

$$(\nu + \mu k)h_{k,\ell} = \begin{cases} \nu \frac{\rho^k}{k!} & \text{if } k = \ell, \\ 0 & \text{if } k \neq \ell, \end{cases}$$

and hence

$$h(x, y) = \sum_{k=0}^{\infty} \frac{\nu}{\nu + k\mu} \frac{\rho^k}{k!} (x - 1)^k (y - 1)^k. \quad \square$$

Corollary 9 (Asymptotics when $\nu \rightarrow \infty$ or $\nu \rightarrow 0$). *From (67) it follows that, if $\nu \rightarrow \infty$, then $P(x, y) \rightarrow e^{\rho(xy-1)}$ and, if $\nu \rightarrow 0$, then $P(x, y) \rightarrow e^{\rho(x-1)} \cdot e^{\rho(y-1)}$.*

Remark 7. $P(x, y)$ is the probability generating function of the joint steady-state distribution of the queue length in an $M/M/\infty$ queue at two different time points, where the distance between the two time points is exponentially distributed with parameter ν . Formula (67) can be alternatively derived using results on the transient distribution in the $M/M/\infty$ queue.

Remark 8. An alternative expression for $P(x, y)$ is the following:

$$P(x, y) = e^{\rho(x-1)} \cdot e^{\rho(y-1)} \cdot \int_{u=0}^1 \frac{\nu}{\mu} u^{\nu/\mu-1} e^{\rho u(y-1)(x-1)} du. \tag{69}$$

This can be derived from (66) in the following way. We have

$$\begin{aligned} \frac{\partial}{\partial x} (e^{-\rho x}(1-x)^{\nu/\mu} P(x, y)) &= e^{-\rho x}(1-x)^{\nu/\mu} \left[\frac{\partial}{\partial x} P(x, y) \right. \\ &\quad \left. + \frac{\nu + \lambda(1-x)}{\mu(x-1)} P(x, y) \right] \\ &= e^{-\rho x}(1-x)^{\nu/\mu} \frac{\nu P(xy, 1)}{\mu(x-1)} \\ &= e^{-\rho x}(1-x)^{\nu/\mu} \frac{\nu}{\mu(x-1)} e^{\rho(xy-1)}. \end{aligned}$$

Hence,

$$e^{-\rho x}(1-x)^{\nu/\mu} P(x, y) = \int_{z=1}^x \frac{\nu}{\mu(z-1)} (1-z)^{\nu/\mu} e^{\rho(z y-1-z)} dz, \tag{70}$$

which can alternatively be written as

$$P(x, y) = e^{\rho(x-1)} \int_{z=1}^x \frac{\nu}{\mu(z-1)} \left(\frac{1-z}{1-x} \right)^{\nu/\mu} e^{\rho z(y-1)} dz. \tag{71}$$

Now, substituting $u = \frac{z-1}{x-1}$ (and hence $z = (x-1)u + 1$) yields (69).

5. Discussion

We close the paper by briefly describing how our results help with decision-making in queueing systems. One application of this can be to model the trade-off between the monitoring costs and the suboptimality due to reduced monitoring. Control policies are sometimes obtained assuming changes can be made at arbitrary time instants (or equivalently assuming infinite monitoring frequency). In practice, however, monitoring the state incurs costs pushing one to monitor less frequently. Reducing the monitoring frequency lowers the measurement cost but also makes the policy more suboptimal. With the analysis in this paper, one can compute the performance obtained for a given monitoring frequency, ν , and then determine the appropriate value of ν that optimizes the objective that accounts for both the performance as well as the monitoring costs.

As an illustration, consider the problem of optimizing a linear combination of sojourn time and energy consumption that is mentioned in the introduction of the paper. The optimal policy is to set the speed to the number of jobs to the power of a coefficient (in this paper, we assume this coefficient to be 1). In Fig. 8, we plot the expected queue length as well as the expected speed as a function of ν , the monitoring frequency, for $\bar{s} = 2$. It can be seen that both decrease with ν which points to a trade-off since a higher ν will entail a higher monitoring cost.

We also remark that for the finite-speed case Proposition 6 can easily be generalized to arbitrary increasing function speed profiles s_i , where s_i is the value the speed is set to when state i is observed. For the above figures, it was assumed that $s_i = i\mu$.




Fig. 8. Expected queue length and expected speed as functions of the monitoring frequency (ν). $\bar{s} = 2$, $\mu = 1$.

References

- [1] E. Altman, K.E. Avrachenkov, R. Núñez Queija, Perturbation analysis for denumerable Markov chains with application to queueing models, *Adv. Appl. Probab.* 36 (3) (2004) 839–853.
- [2] L. Andrew, M. M. Lin, A. Wierman, Optimality, fairness, and robustness in speed scaling designs, in: *Proceedings of the ACM SIGMETRICS’10*, ACM, New York, USA, 2010, pp. 37–48.
- [3] N. Bansal, T. Kimbrel, K. Pruhs, Speed scaling to manage energy and temperature, *J. ACM* 54 (1) (2007) 3:1–3:39.
- [4] R. Bekker, S.C. Borst, O.J. Boxma, O. Kella, Queues with workload-dependent arrival and service rates, *Queueing Syst.* 46 (3–4) (2004) 537–556.
- [5] R. Bekker, O.J. Boxma, J.A.C. Resing, Queues with service speed adaptations, *Stat. Neerl.* 62 (2008) 441–457.
- [6] S. Cheung, R. Boucherie, R. Núñez Queija, Quasi-stationary analysis for queues with temporary overload, in: *2010 22nd International Teletraffic Congress (ITC 22)*, 2010, pp. 1–8, <http://dx.doi.org/10.1109/ITC.2010.5608739>.
- [7] J. Cohen, *The Single Server Queue*, Elsevier Science, 1982.
- [8] G. Latouche, V. Ramaswami, Introduction To Matrix Analytic Methods in Stochastic Modeling, in: *ASA-SIAM Series on Statistics and Applied Probability*, SIAM, 1999.
- [9] Y. Liu, H. Honnappa, S. Tindell, N.K. Yip, Infinite server queues in a random fast oscillatory environment, *Queueing Syst.* 98 (1) (2021) 145–179.
- [10] A. Lovas, M. Rásonyi, Markov chains in random environment with applications in queueing theory and machine learning, *Stochastic Process. Appl.* 137 (2021) 294–326.
- [11] R. Núñez-Queija, B. Prabhu, J. Resing, Queueing models with service speed adaptations at arrival instants of an external observer, *Queueing Syst.* 100 (3) (2022) 233–235.
- [12] R.W. Wolff, Poisson arrivals see time averages, *Oper. Res.* 30 (2) (1982) 223–231.