



HAL
open science

Markov-modulated M/G/1 type queue in heavy traffic and its application to time-sharing disciplines

Halldora Thorsdottir, Ina Maria Maaïke Verloop

► **To cite this version:**

Halldora Thorsdottir, Ina Maria Maaïke Verloop. Markov-modulated M/G/1 type queue in heavy traffic and its application to time-sharing disciplines. *Queueing Systems*, 2016, 83, pp.29-55. 10.1007/s11134-016-9477-y . hal-02640380

HAL Id: hal-02640380

<https://hal.science/hal-02640380>

Submitted on 28 May 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Open Archive Toulouse Archive Ouverte


OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is an author's version published in: <http://oatao.univ-toulouse.fr/22224>

Official URL:

<https://doi-org.gorgone.univ-toulouse.fr/10.1007/s11134-016-9477-y>

To cite this version:

Thorsdottir, Halldora and Verloop, Maaïke  *Markov-modulated M/G/1 type queue in heavy traffic and its application to time-sharing disciplines.* (2016) *Queueing Systems*, 83. 29-55. ISSN 0257-0130.

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

Markov-modulated M/G/1-type queue in heavy traffic and its application to time-sharing disciplines

H. Thorsdottir¹ · I. M. Verloop^{2,3}

Abstract This paper deals with a single-server queue with modulated arrivals, service requirements and service capacity. In our first result, we derive the mean of the total workload assuming generally distributed service requirements and any service discipline which does not depend on the modulating environment. We then show that the workload is exponentially distributed under heavy-traffic scaling. In our second result, we focus on the discriminatory processor sharing (DPS) discipline. Assuming exponential, class-dependent service requirements, we show that the joint distribution of the queue lengths of different customer classes under DPS undergoes a state-space collapse when subject to heavy-traffic scaling. That is, the limiting distribution of the queue-length vector is shown to be exponential, times a deterministic vector. The distribution of the scaled workload, as derived for general service disciplines, is a key quantity in the proof of the state-space collapse.

Keywords Markov-modulation · Heavy traffic · Discriminatory processor sharing · Single-server queue

✉ H. Thorsdottir
halldora@cwil.nl

I. M. Verloop
verloop@irit.fr

CWI, Science Park 123, 1098 XG Amsterdam, The Netherlands

² CNRS, IRIT, 2 rue C. Camichel, 31071 Toulouse, France

Université de Toulouse, INP, 31071 Toulouse, France

1 Introduction

Markov-modulation is a way to formalize the embedding of queues in a random environment. The parameters of the queue in question, typically arrival rates, service requirements or both, are governed by an external Markov chain, thereby creating an extra layer of randomness around the stochastic queueing process. For classical results on Markov-modulated single-server queues with the first-come-first-serve (FCFS) discipline see, for example, [26,28,30]. Recent work on systems in a Markov-modulated environment can be found in, for example, [17,22,35].

In this paper we will analyse a modulated queue under a heavy-traffic scaling, that is, evaluate the system at its critical load. It is a well-known result from the literature on single-server queues that, under fairly general conditions [23], the steady-state distributions of the appropriately scaled queue length and workload become exponential when the critical load is approached. This property has been seen to carry over to certain systems where arrivals and service times are modulated by an external Markov process; see [3,12,15]. In fact, [3] establishes an even stronger result: convergence of the queue-length *process* to a reflected Brownian motion. *Multi-class single-server queues* under a heavy-traffic scaling have been studied in, for example, [21] for FCFS with feedback routing, [20] for the discriminatory random order of service discipline and [14,32] for the discriminatory processor sharing (DPS) policy. In particular, [20,32] show that the steady-state queue-length vector undergoes a so-called *state-space collapse* and converges to an exponentially distributed variable, times a deterministic vector. The cited multi-class results under heavy-traffic scaling are all for non-modulated systems. In light of this, we will in this paper put special emphasis on a modulated multi-class single-server queue, and the limiting steady-state queue-length distribution is derived.

While there is little ambiguity in how arrival rates are modulated, there are in the literature typically two ways in which to modulate the service rates. One can (i) let the departure rate be *continuously* modulated throughout a customer's service, and the other approach is to (ii) let a customer's service requirement distribution be based on the state of the environment when it arrives and remain the same until the customer departs. We note that by adapting the number of different customer classes, the fixed service requirements of case (ii) can be seen as a special case of the continuously modulated requirements (i); we further elaborate on this later in the paper in Remark 1 in Sect. 2.

The way the load or traffic intensity for modulated queues is defined goes hand in hand with the way the service rates are affected by the environment. In case (i), the load is typically the average of arrival rates (where the averaging is with respect to the equilibrium distribution of the environment), say λ_∞ , divided by the average of service rates, say μ_∞ (see, for example, [26]). In case (ii), the load is taken as the average over the arrival rate times the mean service requirement, say, λ_d/μ_d per state d of the environment (see for example, [12,30]). The two load definitions represent different scenarios. In particular, when load (ii) is equal to 1 (the critical load) it means that in at least one state of the environment the total load over all classes must exceed 1, i.e. for at least one state we must have overload. This is true only for special cases of definition (i).

In this paper, special focus will be given to a multi-class single-server queue under the DPS discipline. The DPS discipline was first introduced by Kleinrock in [24] as an extension of the well-known egalitarian processor sharing (PS) discipline and has turned out to be very suitable to model the simultaneous parsing of diverse tasks, such as processing network data. Under this service discipline, the service capacity is divided between all present customers in proportion to their prescribed weights. Due to the challenging nature of DPS systems, most available results are in terms of limit theorems and moments. Fayolle et al. [16] established the mean sojourn time conditioned on the service requirement. That analysis also yielded the mean queue lengths of the different classes, which were shown to depend on the entire service requirement distributions of all classes. The DPS model has finite mean queue lengths irrespective of any higher-order characteristics of the service distribution; see Avrachenkov et al. [5]. This is an extension of a result for the Processor Sharing (PS) system, which holds while the queue is stable. DPS under a heavy-traffic regime was analysed in [14] assuming finite second moments of the service requirement distributions. Assuming exponential service requirement distributions, a direct approach to establish a heavy-traffic limit for the joint queue-length distribution was described in [29] and extended to phase-type distributions in [32]. Combining light- and heavy-traffic limits, in [19] an interpolation approximation is derived for the steady-state distribution of the queue length and waiting time of DPS. The performance of DPS in overload is considered in [1]. Asymptotics of the sojourn time have received attention in [8,9]. Game-theoretic aspects of DPS have been studied in [18,33]. A thorough overview of DPS results can be found in [2].

We are not aware of work analysing a DPS system under modulation. We refer to [27] where the processor sharing discipline (DPS discipline with unit weights) was analysed in a Markovian random environment. Multi-class queues in a random environment have been studied for different models in [11,31]. In [31], a modulated system is studied where arrivals can only occur at transition epochs of the modulating process but service requirements are class dependent and generally distributed. Using a time-changing argument, the waiting time distribution is derived under the FCFS discipline. In [11], the authors derive a Brownian control problem to establish a form of the $c\mu$ scheduling rule in heavy traffic under continuously modulated service requirements. By using a particular scaling, the time-scale separation of the external environment and the queue-length process is exploited. Similar scaling of a modulated queue can also be seen in results on the Markov-modulated infinite server queue in, for example, [7].

The system we analyse in this paper is a single-server queue where the arrival rates, service requirements and service capacity are modulated. We focus on the setting where a customer's service requirement distribution is based on the state of the environment when it arrives and does not change throughout its service. The service capacity is, however, continuously modulated. This assumption is in line with the literature for various types of modulated queues; see [10,13,25,31]. In Remark 3 in Sect. 6, we discuss how part of our results can be extended to a more general model with continuously modulated service requirements. We derive the distribution of the workload under a heavy-traffic scaling for generally distributed service requirements and any service discipline which does not depend on the environment. We then turn our attention to the DPS discipline in a multi-class queue, which is a particular case

of the general modulated system as described above. The weights of the DPS system, determining the service proportion, depend on a customer's class and do not change with the environment, which means that the workload result remains valid. An important finding in the present paper is that the queue-length vector under DPS becomes independent of the modulating process in the heavy-traffic limit, which is consistent with the modulated M/G/1 queue in, for example, [3, 12]. This, together with the obtained result on the workload, allows us to derive the distribution of the queue-length vector under DPS and to show that it undergoes a state-space collapse.

The remainder of the paper is organized as follows. In Sect. 2 we describe the model. In Sect. 3 we study the workload of a single-server queue with modulated arrivals, service capacity and service requirement distribution, establishing its distribution in heavy traffic. From Sect. 4 onwards the focus is on DPS. In Sect. 4 we derive some basic properties of the queue-length distribution, obtain a rate conservation law and derive an equation for the moments of the queue lengths weighted with the modulated service capacity. Section 5 is devoted to the heavy-traffic scaling; there we show that the distribution of the environment becomes independent of that of the queue-length vector, in addition to deriving two technical lemmas. The exponential limiting distribution of the joint queue length in heavy traffic follows in Sect. 6. The result is shown in two steps; Sect. 6.1 for the state-space collapse and Sect. 6.2 for the exact limiting distribution, where we rely on the workload result of Sect. 3. We conclude with a summary and some open questions in Sect. 7.

2 Model

We analyse a single-server queue modulated by an independent external environment, which is formalized by an irreducible continuous time Markov chain on a finite state-space $\{1, \dots, D\}$. The modulating process is denoted by Z and is governed by an infinitesimal generator matrix $Q = (q_{d\ell})_{d,\ell=1}^D$ with an invariant distribution $\boldsymbol{\pi} = (\pi_1, \dots, \pi_D)$. In what follows, vectors are generally denoted in bold. New customers arrive according to a Poisson distribution with rate λ_d when the environment is in state d . A customer arriving in state d has a service requirement distribution given by a function $H_d(\cdot)$ with Laplace–Stieltjes transform (LST) $h_d(\cdot)$. The service requirement does not change further with the environment. The first and second moments are given by h_{d1} and h_{d2} , respectively. In addition we let the service capacity be scaled by a factor c_d during the environment's stay in state d ; this can thus change during the service of the customers. The traffic intensity will be measured as

$$\rho_\infty = \sum_d \pi_d \lambda_d h_{d1} / c_\infty, \quad (1)$$

with $c_\infty := \sum_d \pi_d c_d$ being the service capacity averaged over the environment. The workload is defined as the time it takes to empty the system at an arbitrary moment in time given the observed environment and is denoted by W . In Sect. 3 we study the workload and the environment as a two-dimensional process (W, Z) , under any service discipline that is independent of the environment.

The first main result of this paper concerns the distribution of the workload when the traffic intensity approaches its critical point. The system is said to be in heavy traffic when ρ_∞ approaches 1. Let $N > 0$ and define the following parametrization:

$$\lambda_d^{(N)} := \frac{\lambda_d}{\rho_\infty} (1 - 1/N) \rightarrow \frac{\lambda_d}{\rho_\infty} =: \hat{\lambda}_d, \quad \text{as } N \rightarrow \infty, \quad (2)$$

where ρ_∞ is based on the unscaled parameters. Prelimit quantities will be denoted with a superscript (N) ; the prelimit traffic intensity is thus $\rho_\infty^{(N)}$ and is equal to $1 - 1/N$. Limiting quantities will have a $\hat{\cdot}$; in heavy traffic the traffic intensity is denoted $\hat{\rho}_\infty$ and is equal to 1.

In the remainder of the paper, starting in Sect. 4, we analyse a single-server queue with K customer classes under the discriminatory processor sharing policy; the queue is again embedded in a random environment. Let $\alpha_{k,d}$ be the probability that a customer, arriving while the environment is in state d , is of class k ; note that $\sum_k \alpha_{k,d} = 1$ for a given d . The Poisson arrival rate of a class- k customer is denoted by $\lambda_{k,d} := \alpha_{k,d} \lambda_d$ and for each class k it is assumed that $\lambda_{k,d} > 0$ for at least one state d . In the multi-class setting we assume that a class- k customer has an exponentially distributed service requirement with mean $1/\mu_k$. We believe that the results obtained in this paper can be extended to phase-type distributed service requirements, the latter being dense in the space of all distributions on $[0, \infty)$. For the non-modulated DPS queue, the phase-type analysis was performed in [32] using similar proof techniques. For ease of exposition, however, we focus here on the exponential case.

We no longer let the service requirement of a particular customer be environment dependent (although the distribution of an arbitrary customer is, as explained in Sect. 6.2). One can, however, retrieve the environment-dependent service requirements by introducing additional classes for each environment; see Remark 1. By referring to a class- k customer's service rate while in state d as $\mu_{k,d} := \mu_k c_d$, we take the modulated service capacity into account. Most of the results for the queue length can in fact be shown without assuming this product form, representing a system where the service requirements are continuously modulated; see Remark 3, Sect. 6.1, for further details.

We denote the average arrival rate of class- k customers by $\lambda_{k,\infty} := \sum_d \lambda_{k,d} \pi_d$; similarly we denote the average service rate for class k by $\mu_{k,\infty} := \sum_d \mu_{k,d} \pi_d = \mu_k c_\infty$ and $\rho_{k,\infty} := \lambda_{k,\infty} / \mu_{k,\infty}$. The aggregate traffic intensity for the multi-class model is defined as

$$\rho_\infty := \sum_{k=1}^K \rho_{k,\infty},$$

which is consistent with the definition in Eq. (1).

Let the state of the multi-class system be described by the vector of random variables $(M_1, \dots, M_K, Z) =: (\mathbf{M}, Z)$, where M_k is the number of class- k customers, for $k = 1, \dots, K$. As before, Z represents the state of the background process. In a DPS system, the random fraction of service given to a class- k customer is

Table 1 A multi-class system where the service requirement distribution is fixed upon arrival can be translated into one with non-modulated service requirements

Classes	Arrival rates		Serv. rate $d \in \{1, 2\}$	Weight
	$d = 1$	$d = 2$		
(1,1)	$\lambda_{1,1}$	0	$\mu_{1,1}$	g_1
(1,2)	0	$\lambda_{1,2}$	$\mu_{1,2}$	g_1
(2,1)	$\lambda_{2,1}$	0	$\mu_{2,1}$	g_2
(2,2)	0	$\lambda_{2,2}$	$\mu_{2,2}$	g_2

$$\frac{g_k}{\sum_j g_j M_j},$$

where g_k are weight parameters associated with each class k .

When in heavy traffic, we denote $\hat{\rho}_{k,\infty} := \hat{\lambda}_{k,\infty}/\mu_{k,\infty}$ and thus have for the multi-class system

$$\sum_k \hat{\rho}_{k,\infty} = \sum_k \frac{\hat{\lambda}_{k,\infty}}{\mu_{k,\infty}} = \frac{1}{\rho_\infty} \sum_k \frac{\lambda_{k,\infty}}{\mu_{k,\infty}} = \frac{1}{\rho_\infty} \sum_k \rho_{k,\infty} = 1.$$

Remark 1 (Modulated service requirements rewritten to classes) Any multi-class system where the service requirement distribution is determined by the modulating process at a customer's arrival can be written as a multi-class model with non-modulated, only class-dependent, service rates μ_k , as illustrated below: While in state d of the environment, class- k customers arrive with rate $\lambda_{k,d}$ and have exponential service requirement with mean $1/\mu_{k,d}$ and weight g_k . Such customers we refer to as being of class (k, d) and count with $M_{k,d}$; hence, we need to keep track of $K \cdot D$ different customer ‘‘classes’’. Arrivals to class (k, d) are inactive while not in state d .

From Table 1 one can easily see how a $K = D = 2$ system can be written into one with $K = 4$ and $D = 2$. The arrival rates are still modulated, but in an on-off way. The service rates $\mu_{k,d}$ are now non-modulated.

3 Workload

In this section we consider the workload in a modulated queue and extend the results for an M/G/1-type queue from Falin and Falin [15] and Dimitrov [12] to include modulated service capacity. We derive the mean of the workload and then its distribution in the heavy-traffic regime.

Let $p_{0,d} = P(W = 0, Z = d)$ and let $\mathbf{a} = (a_1, \dots, a_D)^T$ be a vector solving

$$[Q \cdot \mathbf{a}]_d = c_d - \lambda_d h_{d1} - c_\infty(1 - \rho_\infty), \quad (3)$$

for $d = 1, \dots, D$. Note that such a solution always exists since the right-hand side vector of Eq. (3) is orthogonal to $\boldsymbol{\pi}$. We obtain the following result for the mean workload.

Proposition 3.1 For any service requirement distribution $H_d(\cdot)$ and any service discipline that does not depend on the state of the environment, the mean of the workload satisfies

$$\mathbb{E}W = \frac{\sum_d [\pi_d \lambda_d h_{d2}/2 + a_d \pi_d (\lambda_d h_{d1} - c_d) + p_{0,d} c_d a_d]}{c_\infty (1 - \rho_\infty)}, \quad (4)$$

where \mathbf{a} is a solution of Eq. (3). Furthermore, $1 - \rho_\infty = \sum_d p_{0,d} c_d / c_\infty$.

Remark 2 Although the solution vector \mathbf{a} is not unique, the term

$$\sum_{d=1}^D a_d [\pi_d (\lambda_d h_{d1} - c_d) + p_{0,d} c_d],$$

as appearing in Eq. (4), is. This is due to the following argument: Suppose \mathbf{a} and \mathbf{a}^* are two solutions to Eq. (3). Then $0 = Q\mathbf{a} - Q\mathbf{a}^* = Q(\mathbf{a} - \mathbf{a}^*)$, so $(\mathbf{a} - \mathbf{a}^*)$ is in the nullspace of Q . But Q is a generator matrix so it can easily be seen that $Q\mathbf{r} = 0$ for any vector $\mathbf{r} \cdot \mathbf{1}^T = (r, r, \dots, r)^T, r \in \mathbb{R}$. Also, since the environment is an irreducible Markov chain, the nullspace of Q has dimension 1, and therefore, $(\mathbf{a} - \mathbf{a}^*) = r_a \cdot \mathbf{1}^T$, for some $r_a \in \mathbb{R}$. Thus,

$$\sum_{d=1}^D (a_d - a_d^*) [\pi_d (\lambda_d h_{d1} - c_d) + p_{0,d} c_d] = r_a c_\infty (\rho_\infty - 1) + r_a c_\infty (1 - \rho_\infty) = 0,$$

where the first term follows by the definition of ρ_∞ and c_∞ and the second term comes from Eq. (7) in the proof below.

Proof of Proposition 3.1 Define $F_d(x, t) = P(W(t) < x, Z(t) = d)$ for some time $t > 0$. In an infinitesimal time dt , a new arrival requiring service x changes the workload with probability $\lambda_d H_d(x) dt$. The service capacity is scaled by c_d when the environment is in state d , meaning that in dt time, the workload is reduced by $c_d dt$, yielding by a classic birth-and-death argument for the M/G/1 queue,

$$F_d(x, t + dt) = (1 - \lambda_d dt + q_{dd} dt) F_d(x + c_d dt, t) + \sum_{\ell \neq d} q_{\ell d} F_\ell(x + c_\ell dt, t) dt + \lambda_d dt \int_0^x F_d(x + c_d dt - y, t) dH_d(y).$$

We let $t \rightarrow \infty$ to go to steady state, and since

$$\frac{F_d(x + c_d dt) - F_d(x)}{dt} = c_d \frac{F_d(x + c_d dt) - F_d(x)}{c_d dt} \xrightarrow{dt \downarrow 0} c_d F_d'(x),$$

we obtain

$$c_d F_d'(x) = (\lambda_d - q_{dd}) F_d(x) - \sum_{\ell \neq d} q_{\ell d} F_\ell(x) - \lambda_d \int_0^x F_d(x - y) dH_d(y).$$

Denote the LST of $F_d(\cdot)$ by $\varphi_d(s) = \mathbb{E}[e^{-sW(t)}, Z(t) = d] = p_{0,d} + \int_{0+}^{\infty} e^{-sx} dF_d(x)$. The corresponding transform equation becomes

$$\sum_{\ell=1}^D q_{\ell d} \varphi_{\ell}(s) = [\lambda_d(1 - h_d(s)) - sc_d] \varphi_d(s) + sp_{0,d} c_d. \quad (5)$$

It is now convenient to sum over d and divide through Eq. (5) with s to get zero on the left-hand side, leading to

$$\sum_d p_{0,d} c_d = \sum_d \varphi_d(s) \left[c_d - \lambda_d \frac{(1 - h_d(s))}{s} \right]. \quad (6)$$

Using that $\varphi_d(0) = \pi_d$ and by l'Hôpital's rule

$$\lim_{s \downarrow 0} \frac{(1 - h_d(s))}{s} = - \lim_{s \downarrow 0} h'_d(s) = h_{d1}.$$

We get by taking the limit $s \rightarrow 0$ of Eq. (6) that

$$\sum_d \frac{p_{0,d} c_d}{c_{\infty}} = 1 - \rho_{\infty}. \quad (7)$$

We differentiate Eq. (6) w.r.t. s :

$$\sum_d \varphi_d(s) \lambda_d \left[\frac{h'_d(s)}{s} + \frac{1 - h_d(s)}{s^2} \right] = \sum_d \varphi'_d(s) \left[c_d - \lambda_d \frac{1 - h_d(s)}{s} \right],$$

which in the limit $s \rightarrow 0$ results in

$$\sum_d \frac{\pi_d \lambda_d h_{d2}}{2} = \sum_d W_d [c_d - \lambda_d h_{d1}], \quad (8)$$

with the first moment of the workload while in state d being

$$- \lim_{s \downarrow 0} \varphi'_d(s) = \mathbb{E}[W, Z = d] =: W_d.$$

Now multiply Eq. (5) with a_d , sum over d , take the derivative w.r.t. s and let $s \rightarrow 0$ to obtain

$$- \sum_d W_d [c_{\infty}(\rho_{\infty} - 1) + c_d - \lambda_d h_{d1}] = \sum_d [a_d \pi_d (\lambda_d h_{d1} - c_d) + p_{0,d} c_d a_d],$$

by using Eq. (3). Adding this equation to Eq. (8) yields

$$c_\infty(1 - \rho_\infty) \sum_d W_d = \sum_d \left[\frac{\pi_d \lambda_d h_{d2}}{2} + a_d \pi_d (\lambda_d h_{d1} - c_d) + p_{0,d} c_d a_d \right], \quad (9)$$

which gives the desired expression for the mean of the workload, $\mathbb{E}W = \sum_d W_d$. \square

Eq. (7) makes it clear that in heavy traffic, that is, when $\rho_\infty^{(N)} \rightarrow 1$, all the probabilities $p_{0,d}^{(N)}$ go to zero. Also, in heavy traffic, the right-hand side of Eq. (3) reduces to $c_d - \hat{\lambda}_d h_{d1}$. Recalling the parametrization $1 - \rho_\infty^{(N)} = 1/N$ in Sect. 2, we obtain the following result.

Proposition 3.2 *For any service distribution $H_d(\cdot)$ and any service discipline that does not depend on the state of the environment, the mean of the workload in heavy traffic satisfies*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}W^{(N)} = \frac{1}{c_\infty} \sum_d \pi_d \left[\hat{\lambda}_d h_{d2}/2 + a_d (\hat{\lambda}_d h_{d1} - c_d) \right], \quad (10)$$

where \mathbf{a} is a solution of $[Q \cdot \mathbf{a}]_d = c_d - \hat{\lambda}_d h_{d1}$.

Proof Under the heavy-traffic scaling, the empty probabilities $p_{0,d}^{(N)}$ go to zero for $d = 1, \dots, D$. The result then follows immediately from Eq. (4) and $1 - \rho_\infty^{(N)} = 1/N$. \square

This leads to the main result of this section.

Theorem 3.3 *In heavy traffic, the scaled workload $\frac{1}{N}W^{(N)}$ converges in distribution to \hat{W} , where \hat{W} is exponentially distributed with mean given in Eq. (10).*

Proof This follows from combining Proposition 3.2 with Theorem 4 in [12]. The full proof is in the Appendix. \square

The above results yields that \hat{W} is relatively compact, which together with the metric space being separable and complete implies that the scaled workload $\frac{1}{N}W^{(N)}$ is tight, by Prohorov's theorem [6].

4 Queue-length vector under DPS

In the remainder of the paper we focus on the multi-class model under DPS, where we assume that the service requirements of class- k customers are exponentially distributed with rate μ_k . In this section we establish some properties of the joint queue-length distribution. We start with the flow equations, followed by a rate conservation law and an equation for the moments of the queue lengths conditioned on the environment.

Equating the flow in and out of state $(\mathbf{M}, Z) = (\mathbf{m}, d)$ yields (noting that $-q_{dd} = \sum_{\ell \neq d} q_{d\ell}$)

$$\begin{aligned} & \left(\sum_{k=1}^K \left(\lambda_{k,d} + \frac{g_k m_k}{\sum_i g_i m_i} \mu_{k,d} \cdot \mathbf{1}_{\{m_k > 0\}} \right) - q_{dd} \right) p_{\mathbf{m},d} \\ &= \sum_{k=1}^K \left(\lambda_{k,d} p_{\mathbf{m}-\mathbf{e}_k,d} \cdot \mathbf{1}_{\{m_k > 0\}} + \frac{g_k(m_k + 1)}{\sum_i g_i m_i + g_k} \mu_{k,d} p_{\mathbf{m}+\mathbf{e}_k,d} \right) + \sum_{\ell \neq d} q_{d\ell} p_{\mathbf{m},\ell}, \end{aligned} \quad (11)$$

where $p_{\mathbf{m},d} := \mathbb{P}((\mathbf{M}, Z) = (\mathbf{m}, d))$ and \mathbf{e}_k is the vector with 1 in the k -th place and zeros elsewhere. We now define the partial probability generating function (PGF) for when the background process is in state d :

$$\begin{aligned} P_d(\mathbf{z}) &:= \mathbb{E}[z_1^{M_1} \cdots z_K^{M_K} \cdot \mathbf{1}_{\{Z=d\}}] \\ &:= \sum_{m_1=0}^{\infty} \cdots \sum_{m_K=0}^{\infty} \mathbb{P}(M_1 = m_1, \dots, M_K = m_K, Z = d) \cdot z_1^{m_1} \cdots z_K^{m_K} \\ &= \sum_{\mathbf{m} \geq \mathbf{0}} p_{\mathbf{m},d} \mathbf{z}^{\mathbf{m}}, \end{aligned}$$

where $z_1^{m_1} \cdots z_K^{m_K} =: \mathbf{z}^{\mathbf{m}}$ and $(m_1, \dots, m_K) \geq (0, \dots, 0)$, i.e. $\mathbf{m} \geq \mathbf{0}$. Then the overall generating function for the queue length is $P(\mathbf{z}) := \mathbb{E}[z_1^{M_1} \cdots z_K^{M_K}] = \sum_{d=1}^D P_d(\mathbf{z})$. We also define

$$\begin{aligned} R_d(\mathbf{z}) &:= \sum_{\mathbf{m} \geq \mathbf{0}} \frac{p_{\mathbf{m},d} \mathbf{z}^{\mathbf{m}}}{\sum_j g_j m_j} \cdot \mathbf{1}_{\{\sum_{j=1}^K m_j > 0\}}, \quad \text{hence} \\ \frac{\partial R_d(\mathbf{z})}{\partial z_k} &= z_k^{-1} \sum_{\mathbf{m} \geq \mathbf{e}_k} \frac{m_k}{\sum_j g_j m_j} p_{\mathbf{m},d} \mathbf{z}^{\mathbf{m}} \cdot \mathbf{1}_{\{\sum_{j=1}^K m_j > 0\}}. \end{aligned}$$

By multiplying the flow equation (11) with $\mathbf{z}^{\mathbf{m}}$, summing over all vectors $\mathbf{m} \geq \mathbf{0}$ and rearranging terms, we can eventually write it in terms of the PGF $P_d(\mathbf{z})$ and the partial derivative $\partial R_d(\mathbf{z})/\partial z_k$, that is,

$$\sum_{k=1}^K \left[\lambda_{k,d} (1 - z_k) P_d(\mathbf{z}) + \mu_{k,d} g_k (z_k - 1) \frac{\partial R_d(\mathbf{z})}{\partial z_k} \right] = \sum_{\ell=1}^D P_\ell(\mathbf{z}) q_{\ell d}. \quad (12)$$

It will be convenient to write the equation fully in terms of $\partial R_d/\partial z_k$, so we note the relation

$$P_d(\mathbf{z}) = \sum_{k=1}^K g_k z_k \frac{\partial R_d(\mathbf{z})}{\partial z_k} + p_{\mathbf{0},d}, \quad (13)$$

where $p_{\mathbf{0},d} = P((\mathbf{M}, \mathbf{Z}) = (\mathbf{0}, d))$ is the probability of an empty queue in state d , which is equivalent to the probability of no workload defined in Sect. 3. We incorporate this into Eq. (12) to obtain

$$\begin{aligned} & \sum_{k=1}^K \lambda_{k,d}(1-z_k) \left[\sum_{j=1}^K g_j z_j \frac{\partial R_d(\mathbf{z})}{\partial z_j} + p_{\mathbf{0},d} \right] \\ & + \sum_{k=1}^K \mu_{k,d} g_k (z_k - 1) \frac{\partial R_d(\mathbf{z})}{\partial z_k} = \sum_{\ell=1}^D \left[\sum_{k=1}^K g_k z_k \frac{\partial R_\ell(\mathbf{z})}{\partial z_k} + p_{\mathbf{0},\ell} \right] q_{\ell d}. \end{aligned} \quad (14)$$

This equation we will use later when deriving the heavy-traffic limit.

For the M/M/1 queue with modulated arrivals and service times, moments of the queue length and the sojourn times can be found for the FCFS service discipline in [34] and [25], respectively. In [29] the authors establish a recursive formula to calculate moments of the queue length in a non-modulated DPS system. In a similar fashion, we obtain an expression for the sum of the state-dependent moments weighted with the capacity of the server. We also derive a *rate conservation law*, which shows how the average arrival rates per class are proportional to the resources allocated to that same class, and the service they receive. Both results can be found in the following proposition.

Proposition 4.1 *When the queue is stable, the average number of class- k arrivals is proportional to the service resources allocated to class- k customers, i.e.*

$$\lambda_{k,\infty} = \sum_d \mu_{k,d} \mathbb{E} \left[\frac{g_k M_k}{\sum_j g_j M_j} \cdot \mathbf{1}_{\{\sum_j M_j > 0\}} \cdot \mathbf{1}_{\{Z=d\}} \right],$$

for $k = 1, \dots, K$. Furthermore, the state-dependent expectations of M_k satisfy

$$\sum_d c_d \mathbb{E}[M_k \cdot \mathbf{1}_{\{Z=d\}}] = \frac{\lambda_{k,\infty}}{\mu_k} + \sum_{d,j} g_j \frac{\lambda_{k,d} \mathbb{E}[M_j \cdot \mathbf{1}_{\{Z=d\}}] + \lambda_{j,d} \mathbb{E}[M_k \cdot \mathbf{1}_{\{Z=d\}}]}{\mu_k g_k + \mu_j g_j}.$$

Proof We sum Eq. (12) over d and take the derivative w.r.t. z_i to obtain

$$\begin{aligned} & \sum_d \left[-\lambda_{i,d} P_d(\mathbf{z}) + \sum_j \lambda_{j,d} (1-z_j) \frac{\partial P_d(\mathbf{z})}{\partial z_j} + \mu_{i,d} g_i \frac{\partial R_d(\mathbf{z})}{\partial z_i} \right. \\ & \quad \left. + \sum_j \mu_{j,d} g_j (z_j - 1) \frac{\partial^2 R_d(\mathbf{z})}{\partial z_i \partial z_j} \right] = 0. \end{aligned}$$

Letting $\mathbf{z} \rightarrow \mathbf{1}$ yields

$$\sum_d \left[\mu_{i,d} g_i \frac{\partial R_d(\mathbf{z})}{\partial z_i} \Big|_{\mathbf{z} \rightarrow \mathbf{1}} - \lambda_{i,d} P_d(\mathbf{z}) \Big|_{\mathbf{z} \rightarrow \mathbf{1}} \right] = 0. \quad (15)$$

Since

$$\lim_{z \rightarrow 1} P_d(z) = \pi_d,$$

the following conservation law results from Eq. (15), for $k = 1, \dots, K$:

$$\begin{aligned} \lambda_{k,\infty} &= \sum_d \mu_{k,d} g_k \left. \frac{\partial R_d(z)}{\partial z_k} \right|_{z \rightarrow 1} \\ &= \sum_d \mu_{k,d} \mathbb{E} \left[\frac{g_k M_k}{\sum_j g_j M_j} \cdot \mathbf{1}_{\{\sum_j M_j > 0\}} \cdot \mathbf{1}_{\{Z=d\}} \right]. \end{aligned} \quad (16)$$

By taking partial derivatives of Eq. (13), we obtain after standard calculations the recursive relation

$$\left. \frac{\partial^j P_d(z)}{\partial z_{i_1} \cdots \partial z_{i_j}} \right|_{z \rightarrow 1} = \sum_{k=1}^K g_k \frac{\partial^{j+1} R_d(z)}{\partial z_{i_1} \cdots \partial z_{i_j} \partial z_k} + \sum_{\ell=1}^j g_{i_\ell} \frac{\partial^j R_d(z)}{\partial z_{i_1} \cdots \partial z_{i_j}}. \quad (17)$$

In particular, this yields the explicit form

$$\mathbb{E}[M_k \cdot \mathbf{1}_{\{Z=d\}}] = \left. \frac{\partial P_d}{\partial z_k} \right|_{z \rightarrow 1} = \sum_j g_j \left. \frac{\partial^2 R_d}{\partial z_k \partial z_j} \right|_{z \rightarrow 1} + g_k \left. \frac{\partial R_d}{\partial z_k} \right|_{z \rightarrow 1}.$$

Proceeding from the rate conservation law, Eq. (16), and by using $\mu_{k,d} = \mu_k c_d$, we have

$$\sum_d c_d g_k \left. \frac{\partial R_d}{\partial z_k} \right|_{z \rightarrow 1} = \frac{\lambda_{k,\infty}}{\mu_k}.$$

By taking two partial derivatives of the balance equation Eq. (12), we can solve for a second mixed derivative of R_d , namely

$$\left. \frac{\partial^2 R_d}{\partial z_k \partial z_j} \right|_{z \rightarrow 1} = \frac{\sum_\ell \mathbb{E}[M_k M_j \cdot \mathbf{1}_{\{Z=\ell\}}] q_{\ell d} + \lambda_{k,d} \mathbb{E}[M_j \cdot \mathbf{1}_{\{Z=d\}}] + \lambda_{j,d} \mathbb{E}[M_k \cdot \mathbf{1}_{\{Z=d\}}]}{\mu_{k,d} g_k + \mu_{j,d} g_j},$$

thus also yielding a mixed moment. Summing over the weighted moments of the number of class- k customers while in state d , we obtain a linear equation resembling Eq. (16) in [29] for the non-modulated DPS queue:

$$\begin{aligned} \sum_d c_d \mathbb{E}[M_k \cdot \mathbf{1}_{\{Z=d\}}] &= \frac{\lambda_{k,\infty}}{\mu_k} + \sum_{d,j} c_d g_j \left. \frac{\partial^2 R_d}{\partial z_k \partial z_j} \right|_{z \rightarrow 1} \quad (\text{by Eq. (17)}) \\ &= \frac{\lambda_{k,\infty}}{\mu_k} + \sum_{d,j} g_j \frac{\sum_\ell \mathbb{E}[M_k M_j \cdot \mathbf{1}_{\{Z=\ell\}}] q_{\ell d} + \lambda_{k,d} \mathbb{E}[M_j \cdot \mathbf{1}_{\{Z=d\}}] + \lambda_{j,d} \mathbb{E}[M_k \cdot \mathbf{1}_{\{Z=d\}}]}{\mu_k g_k + \mu_j g_j} \end{aligned}$$

$$= \frac{\lambda_{k,\infty}}{\mu_k} + \sum_{d,j} g_j \frac{\lambda_{k,d} \mathbb{E}[M_j \cdot \mathbf{1}_{\{Z=d\}}] + \lambda_{j,d} \mathbb{E}[M_k \cdot \mathbf{1}_{\{Z=d\}}]}{\mu_k g_k + \mu_j g_j},$$

where the last equality comes from $\sum_d q_{\ell d} = 0$. \square

5 Preliminary results for the queue length in heavy traffic

We proceed to show that in heavy traffic the distribution of the environment and the joint queue length become independent. This result, along with two technical lemmas that we derive in this section, will later help establish the main result about the limiting queue length under DPS, presented in Sect. 6. Here we consider the queue-length vector (M_1, \dots, M_K) scaled with $1/N$ and evaluate the PGF in $z^{1/N}$. The objective is to determine the distribution of $\frac{1}{N}(M_1^{(N)}, \dots, M_K^{(N)}) \cdot \mathbf{1}_{\{Z=d\}}$ as N goes to infinity. We will state the existence of the limiting vector, and thus also the limit of the generating function $P_d^{(N)}(z^{1/N})$, as an assumption. This assumption will be proven in Sect. 6.2. The superscript N denotes dependency on the prelimit parameter $\lambda_d^{(N)}$.

We make use of the change of variables $e^{-s_k} = z_k$, for $s_k > 0$, and denote $e^{-s/N} := (e^{-s_1/N}, \dots, e^{-s_K/N})$. Assuming that the limit exists, we use the new variables to define the heavy-traffic quantities: We let $\hat{p}_{\mathbf{0},d} := \lim_{N \rightarrow \infty} P_{\mathbf{0},d}^{(N)}$, $\hat{P}_d(s) := \lim_{N \rightarrow \infty} P_d^{(N)}(e^{-s/N}) = \lim_{N \rightarrow \infty} \mathbb{E}[e^{-\sum_j s_j M_j/N} \cdot \mathbf{1}_{\{Z=d\}}]$ and $\hat{P}(s) := \sum_d \hat{P}_d(s) = \lim_{N \rightarrow \infty} \mathbb{E}[e^{-\sum_j s_j M_j/N}]$. We denote by $(\hat{M}_1, \dots, \hat{M}_K)$ the random vector corresponding to the LST $\hat{P}(s)$. Finally, let

$$\hat{R}_d(s) := \mathbb{E} \left[\frac{1 - e^{-\sum_j s_j \hat{M}_j}}{\sum_j g_j \hat{M}_j} \cdot \mathbf{1}_{\{\sum_j \hat{M}_j > 0\}} \cdot \mathbf{1}_{\{Z=d\}} \right],$$

where the 1 in the numerator is to ensure that the bracketed expression remains bounded when the queue length quantities \hat{M}_k are all near zero. We can now proceed to the following lemma.

Lemma 5.1 *If $\lim_{N \rightarrow \infty} P_d^{(N)}(e^{-s/N})$ exists, then it satisfies*

$$\hat{P}_d(s) = \sum_{k=1}^K g_k \frac{\partial \hat{R}_d(s)}{\partial s_k} + \hat{p}_{\mathbf{0},d}. \quad (18)$$

Proof From Eq. (13),

$$\lim_{N \rightarrow \infty} P_d^{(N)}(e^{-s/N}) = \lim_{N \rightarrow \infty} \sum_{k=1}^K g_k z_k \frac{\partial R_d^{(N)}(z)}{\partial z_k} \Big|_{z=e^{-s/N}} + \hat{p}_{\mathbf{0},d}. \quad (19)$$

Note that

$$\begin{aligned}
& \lim_{N \rightarrow \infty} z_k \frac{\partial R_d^{(N)}(\mathbf{z})}{\partial z_k} \Big|_{z=e^{-s/N}} \\
&= \lim_{N \rightarrow \infty} \sum_{m \geq \mathbf{e}_k} \frac{m_k}{\sum_j g_j m_j} P_{m,d}^{(N)} z^m \cdot \mathbf{1}_{\{\sum_j m_j > 0\}} \Big|_{z=e^{-s/N}} \\
&= \lim_{N \rightarrow \infty} \sum_{m \geq \mathbf{e}_k} \frac{m_k}{\sum_j g_j m_j} P_{m,d}^{(N)} e^{-s_1 m_1 / N} \dots e^{-s_K m_K / N} \cdot \mathbf{1}_{\{\sum_j m_j > 0\}} \\
&= \lim_{N \rightarrow \infty} \mathbb{E} \left[\frac{M_k^{(N)}}{\sum_j g_j M_j^{(N)}} e^{-\sum_j s_j M_j^{(N)} / N} \cdot \mathbf{1}_{\{\sum_j M_j^{(N)} / N > 0\}} \cdot \mathbf{1}_{\{Z=d\}} \right] \\
&= \mathbb{E} \left[\frac{\hat{M}_k}{\sum_j g_j \hat{M}_j} e^{-\sum_j s_j \hat{M}_j} \cdot \mathbf{1}_{\{\sum_j \hat{M}_j > 0\}} \cdot \mathbf{1}_{\{Z=d\}} \right] \\
&= \frac{\partial \hat{R}_d(s)}{\partial s_k}. \tag{20}
\end{aligned}$$

The second-to-last step follows from the fact that $\frac{M_k^{(N)}}{\sum_j g_j M_j^{(N)}} \cdot e^{-\sum_j s_j M_j^{(N)} / N}$ is upper bounded by $1 / \min_j(g_j)$. By the continuous mapping theorem (see Billingsley [6]), it converges in distribution to $\frac{\hat{M}_k}{\sum_j g_j \hat{M}_j} \cdot e^{-\sum_j s_j \hat{M}_j} \cdot \mathbf{1}_{\{\sum_j \hat{M}_j > 0\}}$. The environment is not affected by the heavy-traffic scaling. Eqs. (19) and (20) now conclude the proof. \square

With the help of Lemma 5.1 we obtain:

Proposition 5.2 *If $\lim_{N \rightarrow \infty} P_d^{(N)}(e^{-s/N})$ exists, the joint queue-length distribution is independent of the environment in heavy traffic, that is,*

$$\hat{P}_d(s) = \pi_d \hat{P}(s).$$

Proof We use the change of variables $z_k = e^{-s_k}$. Since

$$z_k \frac{\partial R_d}{\partial z_k}(\mathbf{z}) \Big|_{z=e^{-s}} = -\frac{\partial R_d}{\partial s_k}(e^{-s}),$$

we obtain, by applying the heavy-traffic scaling to Eq. (14),

$$\begin{aligned}
& \sum_k \left[\lambda_{k,d}^{(N)} \left(1 - e^{-s_k/N} \right) \left[\sum_{j=1}^K g_j \frac{-\partial R_d^{(N)}}{\partial s_j} \left(e^{-s/N} \right) + p_{\mathbf{0},d}^{(N)} \right] \right. \\
& \quad \left. - \mu_{k,d} g_k \left(e^{-s_k/N} - 1 \right) e^{s_k/N} \frac{\partial R_d^{(N)}}{\partial s_k} \left(e^{-s/N} \right) \right]
\end{aligned}$$

$$= \sum_{\ell=1}^D \left[\sum_{k=1}^K g_k \frac{-\partial R_{\ell}^{(N)}}{\partial s_k} \left(e^{-s/N} \right) + p_{\mathbf{0},\ell}^{(N)} \right] q_{\ell d}.$$

With Taylor expansion we obtain

$$\begin{aligned} & \sum_k \left[\lambda_{k,d}^{(N)} \left(\frac{s_k}{N} - \frac{s_k^2}{N^2} \right) \left[\sum_{j=1}^K g_j \frac{-\partial R_d^{(N)}}{\partial s_j} \left(e^{-s/N} \right) + p_{\mathbf{0},d}^{(N)} \right] \right. \\ & \quad \left. - \mu_{k,d} g_k \left(\frac{s_k}{N} + \frac{s_k^2}{N^2} \right) \frac{-\partial R_d^{(N)}}{\partial s_k} \left(e^{-s/N} \right) \right] \\ & = \sum_{\ell=1}^D \left[\sum_{k=1}^K g_k \frac{-\partial R_{\ell}^{(N)}}{\partial s_k} \left(e^{-s/N} \right) + p_{\mathbf{0},\ell}^{(N)} \right] q_{\ell d} + \mathcal{O}(N^{-3}). \end{aligned} \quad (21)$$

Since $\frac{-\partial R_d^{(N)}}{\partial s_j}$ is bounded (see the proof of Lemma 5.1) and converges to $\frac{\partial \hat{R}_d}{\partial s_j}$, we obtain, as $N \rightarrow \infty$,

$$\mathbf{v} \cdot [Q]_d = \sum_{\ell=1}^D \left[\sum_{k=1}^K g_k \frac{\partial \hat{R}_{\ell}(s)}{\partial s_k} + \hat{p}_{\mathbf{0},\ell} \right] q_{\ell d} = 0, \quad s \geq \mathbf{0}, \quad \forall d, \quad (22)$$

where $[Q]_d$ is the d^{th} column of Q and \mathbf{v} is a row vector with $v_{\ell} = \sum_{k=1}^K g_k \frac{\partial \hat{R}_{\ell}(s)}{\partial s_k} + \hat{p}_{\mathbf{0},\ell}$. This implies that $\mathbf{v}Q = 0$, and since Q is a generator we conclude that

$$v_d = \sum_{k=1}^K g_k \frac{\partial \hat{R}_d(s)}{\partial s_k} + \hat{p}_{\mathbf{0},d} = \pi_d x,$$

where x does not depend on d . Observe now that by Lemma 5.1 we have

$$\begin{aligned} \hat{P}_d(s) - \hat{p}_{\mathbf{0},d} &= \sum_{k=1}^K g_k \frac{\partial \hat{R}_d(s)}{\partial s_k} \\ &= \pi_d x - \hat{p}_{\mathbf{0},d} \\ &= \mathbb{E}[\mathbf{1}_{\{Z=d\}}]x - \hat{p}_{\mathbf{0},d}. \end{aligned}$$

Since $\hat{P}_d(s) = \mathbb{E} \left[e^{-\sum_j s_j \hat{M}_j} \cdot \mathbf{1}_{\{Z=d\}} \right]$, this implies that $x = \mathbb{E} \left[e^{-\sum_j s_j \hat{M}_j} \right] = \hat{P}(s)$. This shows that the environment becomes independent from the joint queue-length process in the heavy-traffic limit. \square

The flow equation, Eq. (14), simplifies considerably in heavy traffic, as shown in the following lemma.

Lemma 5.3 If $\lim_{N \rightarrow \infty} P_d^{(N)}(e^{-s/N})$ exists, then $\hat{R}_d(s)$ satisfies the following equation:

$$0 = \sum_{k,d} F_{k,d}(s) \frac{\partial \hat{R}_d(s)}{\partial s_k}, \quad \forall s \geq \mathbf{0},$$

with $F_{k,d}(s)$ defined as

$$F_{k,d}(s) := g_k \left(\sum_j \hat{\lambda}_{j,d} s_j - \mu_{k,d} s_k \right).$$

Proof We start by multiplying through Eq. (21) with N , followed by summing over d . Due to Q being a generator, this eliminates the right-hand side with the transition rates $q_{\ell d}$:

$$\begin{aligned} & \sum_{k,d} \left[\lambda_{k,d}^{(N)} \left(s_k - \frac{s_k^2}{N} \right) \left[\sum_{j=1}^K g_j \frac{-\partial R_d^{(N)}}{\partial s_j} (e^{-s/N}) + p_{\mathbf{0},d}^{(N)} \right] \right] \\ & - \sum_{k,d} \left[\mu_{k,d} g_k \left(s_k + \frac{s_k^2}{N} \right) \frac{-\partial R_d^{(N)}}{\partial s_k} (e^{-s/N}) \right] + \mathcal{O}(N^{-2}) = 0. \end{aligned}$$

Taking the limit $N \rightarrow \infty$ yields

$$\begin{aligned} 0 &= \sum_{k,d} \hat{\lambda}_{k,d} s_k \sum_j g_j \frac{\partial \hat{R}_d(s)}{\partial s_j} - \sum_{k,d} \mu_{k,d} g_k s_k \frac{\partial \hat{R}_d(s)}{\partial s_k} \\ &= \sum_{k,d} g_k \frac{\partial \hat{R}_d(s)}{\partial s_k} \sum_j \hat{\lambda}_{j,d} s_j - \sum_{k,d} \mu_{k,d} g_k s_k \frac{\partial \hat{R}_d(s)}{\partial s_k} \\ &= \sum_{k,d} g_k \left(\sum_j \hat{\lambda}_{j,d} s_j - \mu_{k,d} s_k \right) \frac{\partial \hat{R}_d(s)}{\partial s_k} \\ &= \sum_{k,d} F_{k,d}(s) \frac{\partial \hat{R}_d(s)}{\partial s_k}. \end{aligned} \tag{23}$$

□

In what follows we focus on

$$F_{k,\infty}(s) := \sum_d F_{k,d}(s) \pi_d = g_k \left(\sum_j \hat{\lambda}_{j,\infty} s_j - \mu_{k,\infty} s_k \right),$$

and denote its vector counterpart by $\mathbf{F}_\infty(s) = (F_{1,\infty}(s), \dots, F_{K,\infty}(s))$.

6 Queue-length distribution in heavy traffic

We now state and subsequently prove our main result about the queue-length distribution.

Theorem 6.1 *When scaled by $1/N = (1 - \rho_\infty^{(N)})$, the queue-length vector converges in distribution as $(\rho_{1,\infty}^{(N)}, \dots, \rho_{K,\infty}^{(N)}) \rightarrow (\hat{\rho}_{1,\infty}, \dots, \hat{\rho}_{K,\infty})$, i.e. $\rho_\infty^{(N)} \rightarrow 1$, namely*

$$\begin{aligned} \frac{1}{N} \left(M_1^{(N)}, \dots, M_K^{(N)} \right) \cdot \mathbf{1}_{\{Z=d\}} &\stackrel{d}{\rightarrow} \left(\hat{M}_1, \dots, \hat{M}_K \right) \cdot \mathbf{1}_{\{Z=d\}} \\ &\stackrel{d}{=} \pi_d \left(\frac{\hat{\rho}_{1,\infty}}{g_1}, \dots, \frac{\hat{\rho}_{K,\infty}}{g_K} \right) \cdot X, \end{aligned} \quad (24)$$

where $\stackrel{d}{\rightarrow}$ denotes convergence in distribution and X is exponentially distributed with mean

$$\mathbb{E}X = \frac{\sum_k \hat{\rho}_{k,\infty} / \mu_k - \sum_d c_d \pi_d a_d (1 - \hat{\rho}_d)}{c_\infty \sum_k \hat{\rho}_{k,\infty} / (g_k \mu_k)}, \quad (25)$$

with $\hat{\rho}_d := c_d^{-1} \sum_k \hat{\lambda}_{k,d} / \mu_k$ and $\mathbf{a} = (a_1, \dots, a_D)^T$ being a solution of

$$[Q \cdot \mathbf{a}]_d = c_d (1 - \hat{\rho}_d).$$

We will prove this theorem in the two following subsections, first showing in Sect. 6.1 the state-space collapse observed in Eq. (24) and then in Sect. 6.2 we will show that X is exponentially distributed with the mean given by Eq. (25).

6.1 State-space collapse

The first part of the proof of Theorem 6.1 is the state-space collapse. In this section, we assume $\lim_{N \rightarrow \infty} P_d^{(N)}(e^{-s/N})$ exists.

Observe that due to Proposition 5.2,

$$\begin{aligned} \hat{R}_d(s) &= \mathbb{E} \left[\frac{1 - e^{-\sum_k s_k \hat{M}_k}}{\sum_k g_k \hat{M}_k} \cdot \mathbf{1}_{\{\sum_k \hat{M}_k > 0\}} \cdot \mathbf{1}_{\{Z=d\}} \right] \\ &= \mathbb{E} \left[\frac{1 - e^{-\sum_k s_k \hat{M}_k}}{\sum_k g_k \hat{M}_k} \cdot \mathbf{1}_{\{\sum_k \hat{M}_k > 0\}} \right] \cdot \pi_d \\ &=: \hat{R}(s) \pi_d, \end{aligned} \quad (26)$$

where the last equation defines $\hat{R}(s)$, which is independent of d . We now derive some properties of $\hat{R}(s)$.

Lemma 6.2 $\hat{R}(s)$ is constant on a $(K - 1)$ -dimensional hyperplane \mathcal{H}_c , where

$$\mathcal{H}_c := \left\{ s \geq \mathbf{0} : \sum_k \frac{\hat{\rho}_{k,\infty}}{g_k} s_k = c \right\}, \quad c > 0.$$

Proof We follow closely the steps of the proof of Lemma three in [32]. The proof has three steps: (i) Show that $F_{k,\infty}(s)$ is parallel to the hyperplane. Hence, any flow corresponding to $F_{k,\infty}$ that starts in the plane, stays in the plane. (ii) Show that $\hat{R}(s)$ is constant along each flow in the hyperplane and (iii) show that each flow in the hyperplane converges to a unique point. This implies that $\hat{R}(s)$ is constant on the hyperplane.

(i) $F_{\infty}(s)$ is parallel to \mathcal{H}_c

Observe that with $1 = \sum_k \hat{\rho}_{k,\infty}$ and $\hat{\rho}_{k,\infty} = \hat{\lambda}_{k,\infty} / \mu_{k,\infty}$,

$$\begin{aligned} \sum_k \frac{\hat{\rho}_{k,\infty}}{g_k} F_{k,\infty}(s) &= \sum_k \hat{\rho}_{k,\infty} \left(\sum_j \hat{\lambda}_{j,\infty} s_j - \mu_{k,\infty} s_k \right) \\ &= \sum_j \hat{\lambda}_{j,\infty} s_j - \sum_k \hat{\rho}_{k,\infty} \mu_{k,\infty} s_k \\ &= \sum_k \hat{\lambda}_{k,\infty} s_k - \sum_k \hat{\lambda}_{k,\infty} s_k \\ &= 0. \end{aligned}$$

This indicates that the K -dimensional vector $F_{\infty}(s)$ is parallel to the hyperplane.

(ii) $\hat{R}(s)$ is constant along flows in \mathcal{H}_c

For each state $s \geq \mathbf{0}$, there exists a unique flow $f(u) = (f_1(u), \dots, f_K(u))^T$ parametrized by $u \geq 0$, such that

$$f(0) = s \quad \text{and} \quad \frac{df_k(u)}{du} = F_{k,\infty}(f(u)). \quad (27)$$

Due to (i), any flow that starts in \mathcal{H}_c , stays in \mathcal{H}_c . Now,

$$\begin{aligned} \frac{d\hat{R}(f(u))}{du} &= \sum_{k=1}^K \frac{df_k(u)}{du} \cdot \frac{\partial \hat{R}(s)}{\partial s_k} \Big|_{s=f(u)} \\ &= \sum_{k=1}^K F_{k,\infty}(f(u)) \cdot \frac{\partial \hat{R}(s)}{\partial s_k} \Big|_{s=f(u)} \\ &= \sum_{k=1}^K \sum_{d=1}^D F_{k,d}(f(u)) \pi_d \cdot \frac{\partial \hat{R}(s)}{\partial s_k} \Big|_{s=f(u)} \end{aligned}$$

$$\begin{aligned}
&= \sum_{k=1}^K \sum_{d=1}^D F_{k,d}(\mathbf{f}(u)) \cdot \left. \frac{\partial \hat{R}_d(\mathbf{s})}{\partial s_k} \right|_{\mathbf{s}=\mathbf{f}(u)} \\
&= 0, \quad \text{by Eq. (23),}
\end{aligned}$$

implying that $\hat{R}(\mathbf{f}(u))$ is constant along each flow $\mathbf{f}(u)$ which lies in \mathcal{H}_c .

(iii) Each flow in \mathcal{H}_c converges to a unique point

Here we first write the flow specifications in a vector-matrix form, then show that one eigenvalue of that matrix is zero with eigenvector $\mathbf{s}^* \in \mathcal{H}_1$, and the other eigenvalues are negative, and thus we can write $\mathbf{f}(u) = c \cdot \mathbf{s}^* + \mathbf{g}(u)$ where $\lim_{u \rightarrow \infty} \mathbf{g}(u) = \mathbf{0}$.

Eq. (27) can be written in matrix-vector form as

$$\mathbf{f}'(u) = A\mathbf{f}(u),$$

with

$$A = \begin{pmatrix} g_1(\hat{\lambda}_{1,\infty} - \mu_{1,\infty}) & g_1\hat{\lambda}_{2,\infty} & \cdots & g_1\hat{\lambda}_{K,\infty} \\ g_2\hat{\lambda}_{1,\infty} & g_2(\hat{\lambda}_{2,\infty} - \mu_{2,\infty}) & \cdots & g_2\hat{\lambda}_{K,\infty} \\ \vdots & & \ddots & \vdots \\ g_K\hat{\lambda}_{1,\infty} & \cdots & & g_K(\hat{\lambda}_{K,\infty} - \mu_{K,\infty}) \end{pmatrix}.$$

Let D be the diagonal matrix with $d_i = \hat{\rho}_{i,\infty}/g_i$ on the diagonal. Then with

$$S := DAD^{-1} = \begin{pmatrix} g_1(\hat{\lambda}_{1,\infty} - \mu_{1,\infty}) & g_2 \frac{\hat{\rho}_{1,\infty}}{\hat{\rho}_{2,\infty}} \hat{\lambda}_{2,\infty} & \cdots & g_K \frac{\hat{\rho}_{1,\infty}}{\hat{\rho}_{K,\infty}} \hat{\lambda}_{K,\infty} \\ g_1 \frac{\hat{\rho}_{2,\infty}}{\hat{\rho}_{1,\infty}} \hat{\lambda}_{1,\infty} & g_2(\hat{\lambda}_{2,\infty} - \mu_{2,\infty}) & \cdots & g_K \frac{\hat{\rho}_{2,\infty}}{\hat{\rho}_{K,\infty}} \hat{\lambda}_{K,\infty} \\ \vdots & & \ddots & \vdots \\ g_1 \frac{\hat{\rho}_{K,\infty}}{\hat{\rho}_{1,\infty}} \hat{\lambda}_{1,\infty} & \cdots & & g_K(\hat{\lambda}_{K,\infty} - \mu_{K,\infty}) \end{pmatrix}.$$

S^T is a generator corresponding to a finite-state Markov chain.

From the proof of Lemma 4 in [32], it is easily seen that the Markov chain corresponding to S^T is irreducible (since we assume that all $\lambda_{k,d} > 0$ for all k and at least one d). Retracing the arguments stated there, for completeness, it follows that this Markov chain has a unique equilibrium distribution (column) vector, $\boldsymbol{\eta}$, such that $\boldsymbol{\eta}^T S^T = \mathbf{0}$. In particular, 0 is an eigenvalue with multiplicity one and all other eigenvalues have a strictly negative real part; see [4]. Since the eigenvalues of S^T and A are the same, 0 is also an eigenvalue of A with corresponding right eigenvector $\mathbf{s}^* = D^{-1}\boldsymbol{\eta}$, $\mathbf{s}^* \geq \mathbf{0}$, $\mathbf{s}^* \in \mathcal{H}_1$. The solution of the linear system $\mathbf{f}'(u) = A\mathbf{f}(u)$, $\mathbf{f}(0) \in \mathcal{H}_c$ can now be written as the sum of the homogeneous and the particular solution, i.e. $\mathbf{f}(u) = c \cdot \mathbf{s}^* + \mathbf{g}(u)$, where $\lim_{u \rightarrow \infty} \mathbf{g}(u) = \mathbf{0}$. This implies that all the flows in \mathcal{H}_c converge to one common point $c \cdot \mathbf{s}^*$.

Combining (i), (ii) and (iii), we conclude that the function $\hat{R}(\mathbf{s})$ is constant on \mathcal{H}_c .

□

As a consequence of Lemma 6.2, the function $\hat{R}(s)$ depends on s only through the sum $\sum_{k=1}^K (\hat{\rho}_{k,\infty}/g_k)s_k$. Therefore, there exists a function $\hat{R}^* : \mathbb{R} \rightarrow \mathbb{R}$ such that $\hat{R}(s) = \hat{R}^*(\sum_{k=1}^K (\hat{\rho}_{k,\infty}/g_k)s_k)$. Then

$$\frac{\partial}{\partial s_k} \hat{R}(s) = \frac{\hat{\rho}_{k,\infty}}{g_k} \frac{d\hat{R}^*(v)}{dv} \Big|_{v=\sum_{k=1}^K (\hat{\rho}_{k,\infty}/g_k)s_k},$$

so we obtain

$$\begin{aligned} \mathbb{E}[e^{-\sum_{k=1}^K s_k \hat{M}_k}] &= \lim_{N \rightarrow \infty} \sum_{d=1}^D P_d^{(N)}(e^{-s/N}) \\ &= \sum_{d=1}^D \sum_{k=1}^K g_k \frac{\partial \hat{R}_d(s)}{\partial s_k} \\ &= \sum_{d=1}^D \sum_{k=1}^K g_k \frac{\partial \hat{R}(s)}{\partial s_k} \pi_d \\ &= \sum_{k=1}^K \hat{\rho}_{k,\infty} \frac{d\hat{R}^*(v)}{dv} \Big|_{v=\sum_{k=1}^K (\hat{\rho}_{k,\infty}/g_k)s_k} \\ &= \frac{d\hat{R}^*(v)}{dv} \Big|_{v=\sum_{k=1}^K (\hat{\rho}_{k,\infty}/g_k)s_k}, \end{aligned} \tag{28}$$

which only depends on $v = \sum_{k=1}^K (\hat{\rho}_{k,\infty}/g_k)s_k$. Since we also have

$$\begin{aligned} \mathbb{E}[e^{-\sum_{k=1}^K s_k \hat{M}_k}] &= \mathbb{E} \left[e^{-\frac{g_1}{\hat{\rho}_{1,\infty}} v \hat{M}_1} \cdot e^{-\frac{\hat{\rho}_{2,\infty}}{g_2} s_2 \left(\frac{g_2}{\hat{\rho}_{2,\infty}} \hat{M}_2 - \frac{g_1}{\hat{\rho}_{1,\infty}} \hat{M}_1 \right)} \right. \\ &\quad \left. \dots \cdot e^{-\frac{\hat{\rho}_{K,\infty}}{g_K} s_K \left(\frac{g_K}{\hat{\rho}_{K,\infty}} \hat{M}_K - \frac{g_1}{\hat{\rho}_{1,\infty}} \hat{M}_1 \right)} \right], \end{aligned}$$

this, together with Eq. (28), implies that $\left(\frac{g_j}{\hat{\rho}_{j,\infty}} \hat{M}_j - \frac{g_1}{\hat{\rho}_{1,\infty}} \hat{M}_1 \right) = 0$, for all $j = 1, \dots, K$. Thus $(g_k/\hat{\rho}_{k,\infty})\hat{M}_k \stackrel{d}{=} (g_j/\hat{\rho}_{j,\infty})\hat{M}_j$, for all k, j , almost surely. Combining this finding with that of Eq. (22), we obtain Eq. (24) with X distributed as $(g_1/\hat{\rho}_{1,\infty})\hat{M}_1$.

Remark 3 (Continuously modulated service requirements) In Sect. 3 we saw that the critical load is indeed reached when $\rho_\infty \rightarrow 1$, since then $p_{0,d} \rightarrow 0$. This indicates that $(1 - \rho_\infty)$ is the right heavy-traffic scaling when $\mu_{k,d} = \mu_k c_d$. This is less clear for a general $\mu_{k,d}$, that is, for continuously modulated service requirements, where the environment can influence the departure rate of customers present in the system. For that setting, the workload process is no longer independent of the employed

scheduling discipline, since the decision on which class to serve impacts the rate at which customers leave. We are not aware of any results on workload and waiting time distributions where the service distribution is a general function of both class and environment.

The majority of the preceding queue-length results in this paper can, however, be proven without the restriction of the product form, i.e. for continuously modulated service requirements. The traffic intensity for this variant is defined as for the multi-class model above, only this time one cannot split the average class- k service rate into $\mu_{k,\infty} = \mu_k c_\infty$. The traffic intensity per class k , $\rho_{k,\infty} = \lambda_{k,\infty} / \mu_{k,\infty}$, is in line with the Markov-modulated single-server queues. Assuming there exists a scaling $f(N)$ such that $f(N)(M_1, \dots, M_K)\mathbf{1}_{\{Z=d\}}$ converges in distribution, it can be shown that the empty probabilities $p_{0,d}$ vanish in heavy traffic as $N \rightarrow \infty$, for a general $\mu_{k,d}$. This property then follows from Proposition 5.2 without relying on the workload results from Sect. 3 and the product form assumed there. Furthermore, under this assumption, all results in Sect. 6.1 hold, implying that a state-space collapse will appear. In other words, we can prove the first half of Theorem 6.1. However, we do not know what the distribution of the common factor X will be.

6.2 Distribution of the common factor

In order to prove that the limiting queue-length distribution exists and to find the common factor of the queue-length distribution in heavy traffic, the random variable X , we make use of the results on the workload of the total system. From [32] and Eq. (24) we have that

$$\hat{W} \stackrel{d}{=} \sum_{k=1}^K \frac{\hat{M}_k}{\mu_k} = X \cdot \sum_{k=1}^K \frac{\hat{\rho}_{k,\infty}}{g_k \mu_k}. \quad (29)$$

In order to apply the workload result of Sect. 3, we first derive the service requirement of an arbitrary customer while in state d . If $H_k(\cdot)$ is the distribution function of a class- k customer's service requirement, then the probability of a class- k customer arriving and requiring service not exceeding x is $\alpha_{k,d} H_k(x)$. Summing over k now yields the desired distribution,

$$H_d(x) := \sum_{k=1}^K \alpha_{k,d} H_k(x).$$

The overall service requirement distribution thus depends on the state of the environment at its arrival. With exponential service requirements, the corresponding LST is given by

$$h_d(s) = \sum_{k=1}^K \frac{\alpha_{k,d} \mu_k}{\mu_k + s}, \quad s \geq 0, \quad (30)$$

and the first and second moment are given by

$$h_{d1} = \sum_k \frac{\alpha_{k,d}}{\mu_k}, \quad h_{d2} = 2 \sum_k \frac{\alpha_{k,d}}{\mu_k^2}. \quad (31)$$

We can now apply the result of Theorem 3.3 with moments as given in Eq. (31). Hence, we have that \hat{W} is exponentially distributed with mean

$$\mathbb{E}\hat{W} = c_\infty^{-1} \left(\sum_k \hat{\rho}_{k,\infty} / \mu_k - \sum_d c_d \pi_d a_d (1 - \hat{\rho}_d) \right),$$

where \mathbf{a} is a solution of $[Q \cdot \mathbf{a}]_d = c_d - \hat{\lambda}_d \sum_k \frac{\alpha_{k,d}}{\mu_k} = c_d (1 - \hat{\rho}_d)$. Along with Eq. (29), this yields the mean of the exponential random variable X :

$$\begin{aligned} \mathbb{E}X &= \frac{\mathbb{E}\hat{W}}{\sum_k \hat{\rho}_{k,\infty} / (g_k \mu_k)} \\ &= \frac{\sum_k \hat{\rho}_{k,\infty} / \mu_k - \sum_d c_d \pi_d a_d (1 - \hat{\rho}_d)}{c_\infty \sum_k \hat{\rho}_{k,\infty} / (g_k \mu_k)}. \end{aligned} \quad (32)$$

The first term of the numerator is in accordance with the results of [29] and [32], and the second term is a result of the random environment.

The results in Sects. 5 and 6.1 are based on the assumption that $\lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{M} \cdot \mathbf{1}_{\{Z=d\}}$ exists. Since the scaled workload is tight, see Sect. 3, so is the scaled queue length. Then, by Prohorov's theorem ([6]) there exists a subsequence of N such that $\frac{1}{N} M_k$ converges in distribution, and hence for this subsequence $\lim_{N \rightarrow \infty} P_d^{(N)}(e^{-s/N})$ exists. Since each converging subsequence yields the same limit, the limit itself exists (see corollary page 59 in [6]), i.e. $\frac{1}{N} (\mathbf{M}, Z = d) \xrightarrow{d} \hat{\mathbf{M}} \cdot \mathbf{1}_{\{Z=d\}}$, as $N \rightarrow \infty$, with the limiting vector as in Eq. (24).

This concludes the proof of Theorem 6.1.

7 Conclusion and future work

We first studied the workload for a queue with modulated arrivals, service requirements and service capacity, and derived that the scaled workload converges to an exponentially distributed random variable in heavy traffic. The workload results obtained are valid for any service distribution and for any service discipline which does not depend on the environment. We then focussed on the special setting of a multi-class queue under the DPS policy and showed that the joint queue-length distribution for such a system undergoes a state-space collapse in heavy traffic. Under the scaling of $(1 - \rho_\infty)$, the vector-valued limiting distribution is independent of the modulating environment and converges in distribution to a one-dimensional random variable times a deterministic vector. In this derivation, the distribution of the scaled workload is a key quantity. With this we extend known results about the DPS queue to a Markov-modulated setting.

Clearly an interesting question for future consideration is whether the state-space collapse for the DPS policy carries over to continuously modulated service requirements, as discussed in Remark 3. Another open question concerns the characterization of the moments of the queue lengths for the modulated DPS queue, outside of heavy

traffic. Last but not least, modulating the weights of the DPS would open the possibility of dynamical scheduling based on the environment. The latter would be a study on its own, as already the stability conditions will no longer be independent of the weights of the DPS policy.

Acknowledgments The authors would like to thank Urtzi Ayesta, Joke Blom and Michel Mandjes for helpful discussions. This research was partially supported by the SMI Program of INP Toulouse.

Appendix: Proof of Theorem 3.3

The proof Theorem 3.3 is based on Theorem 4 in [12], which can be adapted to our model as follows:

We start with notation and some preliminaries. Let $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_D)$, $\bar{H}(s) = \text{diag}(1 - h_1(s), \dots, 1 - h_D(s))$, $C = \text{diag}(c_1, \dots, c_D)$ and $\mathbf{p}_0 = (p_{0,1}, \dots, p_{0,D})$. Furthermore \bar{H}_1 and \bar{H}_2 are the diagonal matrices corresponding to the moments h_{d1} and h_{d2} , respectively, for $d = 1, \dots, D$. Recall Eq. (3), $[\mathcal{Q} \cdot \mathbf{a}]_d = c_d - \lambda_d h_{d1} - c_\infty(1 - \rho_\infty)$. We will now construct a partial inverse of \mathcal{Q} to make it easier to find a vector \mathbf{a} which solves this equation. Let Q_1 and R be matrices such that

$$Q_1 = \begin{pmatrix} q_{22} & q_{23} & \cdots & q_{2D} \\ \vdots & & & \\ q_{D2} & q_{D3} & \cdots & q_{DD} \end{pmatrix}, \quad R = \begin{pmatrix} 0 & 0 & \cdots & 0 \\ 0 & & & \\ \vdots & Q_1^{-1} & & \\ 0 & & & \end{pmatrix}.$$

Then $\det Q_1 \neq 0$ and, due to Q being a generator (for more details, see [12]), we have

$$QR = \begin{pmatrix} 0 & \frac{-\pi_2}{\pi_1} & \frac{-\pi_3}{\pi_1} & \cdots & \frac{-\pi_D}{\pi_1} \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots & \ddots \\ 0 & 0 & \cdots & \cdots & 1 \end{pmatrix}.$$

It follows that for any vector \mathbf{x} , it holds that

$$\mathbf{x}QR = \mathbf{x} - \frac{x_1}{\pi_1} \boldsymbol{\pi}. \quad (33)$$

Then it can be verified with straightforward calculations that

$$\begin{aligned} \mathbf{a} &= (a_1, \dots, a_D) = R(c_1 - \lambda_1 h_{11} - c_\infty(1 - \rho_\infty), \dots, c_D - \lambda_D h_{D1} - c_\infty(1 - \rho_\infty))^T \\ &= R[C - \Lambda \bar{H}_1] \mathbf{e} - c_\infty(1 - \rho_\infty) \mathbf{R} \mathbf{e} \\ &= R[C - \Lambda \bar{H}_1] \mathbf{e} - \mathbf{r} \end{aligned} \quad (34)$$

is a possible solution vector, with $\mathbf{r} := c_\infty(1 - \rho_\infty) \mathbf{R} \mathbf{e}$.

Define the vector $\boldsymbol{\varphi} = (\varphi_1(s), \dots, \varphi_D(s))$ and write Eq. (5) in matrix-vector terms:

$$\boldsymbol{\varphi}(s)Q = \boldsymbol{\varphi}(s)[\Lambda\bar{H}(s) - sC] + s\mathbf{p}_0C. \quad (35)$$

Observe that, according to Eq. (7),

$$\mathbf{p}_0C\mathbf{e} = c_\infty(1 - \rho_\infty).$$

Now multiply from the right both sides of the new vector-matrix equation, Eq. (35), with a D -dimensional vector of 1's, \mathbf{e} , to obtain

$$\boldsymbol{\varphi}(s)[\Lambda\bar{H}(s) - sC]\mathbf{e} + sc_\infty(1 - \rho_\infty) = 0. \quad (36)$$

Multiply from the right both sides of Eq. (35) with the matrix R to get

$$\boldsymbol{\varphi}(s)QR = \boldsymbol{\varphi}(s)[\Lambda\bar{H}(s) - sC]R + s\mathbf{p}_0CR.$$

Rewrite this equation by using the property of Eq. (33) to obtain

$$\boldsymbol{\varphi}(s) = \frac{\varphi_1(s)}{\pi_1}\boldsymbol{\pi} + \boldsymbol{\varphi}(s)[\Lambda\bar{H}(s) - sC]R + s\mathbf{p}_0CR. \quad (37)$$

Iterate Eq. (37) with itself by inserting $\boldsymbol{\varphi}(s)$ into the right-hand side of the equation to obtain, after some algebraic transformations,

$$\boldsymbol{\varphi}(s) = \frac{\varphi_1(s)}{\pi_1}\boldsymbol{\pi}[I + G(s)R] + \mathbf{y}(s), \quad (38)$$

with $G(s) := \Lambda\bar{H}(s) - sC$ and

$$\mathbf{y}(s) := \boldsymbol{\varphi}(s)[G(s)R]^2 + s\mathbf{p}_0CR[G(s)R + I]. \quad (39)$$

Substitute Eq. (38) into Eq. (36) to obtain an expression for $\varphi_1(s)$:

$$\begin{aligned} 0 &= \left[\frac{\varphi_1(s)}{\pi_1}\boldsymbol{\pi}[I + G(s)R] + \mathbf{y}(s) \right] \cdot G(s)\mathbf{e} + sc_\infty(1 - \rho_\infty) \\ &= \frac{\varphi_1(s)}{\pi_1} [\boldsymbol{\pi}G(s)\mathbf{e} + \boldsymbol{\pi}G(s)RG(s)\mathbf{e}] + \mathbf{y}(s)G(s)\mathbf{e} + sc_\infty(1 - \rho_\infty) \\ &= \frac{\varphi_1(s)}{\pi_1} [B_2(s) + B_3(s)] + B_1(s), \end{aligned} \quad (40)$$

with $B_1(s) = \mathbf{y}(s)G(s)\mathbf{e} + sc_\infty(1 - \rho_\infty)$, $B_2(s) = \boldsymbol{\pi}G(s)\mathbf{e}$ and $B_3(s) = \boldsymbol{\pi}G(s)RG(s)\mathbf{e}$.

The next step is to insert the scaling $s \mapsto s/N$ for each term. Recall that using the heavy-traffic parametrization introduced in Sect. 2, we have $(1 - \rho_\infty) = 1/N$. Now observe that, as $N \rightarrow \infty$,

$$\frac{\bar{H}(s/N)}{s/N} \rightarrow \bar{H}_1, \quad \frac{\bar{H}_1 s/N - \bar{H}(s/N)}{(s/N)^2} \rightarrow \frac{\bar{H}_2}{2}.$$

Therefore the limit

$$\frac{G(s/N)}{s/N} \rightarrow \Lambda \bar{H}_1 - C,$$

is a constant, and since $|\varphi(s/N)| \leq 1$ and $\mathbf{p}_0^{(N)} \rightarrow 0$ (see Eq. (7)), we have

$$\begin{aligned} \frac{\mathbf{y}(s/N)}{s/N} &= \varphi(s/N) \frac{[G(s/N)R]^2}{s/N} + \mathbf{p}_0^{(N)} CR[G(s/N)R + I] \\ &= \varphi(s/N) \left(\frac{s}{N}\right) \left[\frac{G(s/N)}{s/N} R\right]^2 + \mathbf{p}_0^{(N)} CR[G(s/N)R + I] \\ &\rightarrow 0, \end{aligned}$$

as $N \rightarrow \infty$. Combining the above we obtain

$$\begin{aligned} B_1(s/N) &= \mathbf{y}(s/N)G(s/N)\mathbf{e} + sc_\infty(1 - \rho_\infty)/N \\ &= sc_\infty(1 - \rho_\infty)/N + o(N^{-2}) = sc_\infty/N^2 + o(N^{-2}). \end{aligned}$$

Then

$$\begin{aligned} B_2(s/N) &= \boldsymbol{\pi} [\Lambda \bar{H}(s/N) - sC/N] \mathbf{e} \\ &= \boldsymbol{\pi} \Lambda [\bar{H}(s/N) - \bar{H}_1 s/N] \mathbf{e} + \boldsymbol{\pi} [\Lambda \bar{H}_1 s/N - sC/N] \mathbf{e} \\ &= \boldsymbol{\pi} \Lambda \frac{\bar{H}(s/N) - \bar{H}_1 s/N}{(s/N)^2} (s/N)^2 \mathbf{e} - sc_\infty(1 - \rho_\infty)/N \\ &= -\boldsymbol{\pi} \Lambda \frac{\bar{H}_2}{2} (s/N)^2 \mathbf{e} + o(N^{-2}) - sc_\infty/N^2 \\ &= -(s/N)^2 \sum_{d=1}^D \pi_d \lambda_d h_{d2}/2 - sc_\infty/N^2 + o(N^{-2}), \end{aligned}$$

since $\boldsymbol{\pi} \Lambda \bar{H}_1 \mathbf{e} = \sum_{d=1}^D \pi_d \lambda_d h_{d1} = \rho_\infty c_\infty$. Furthermore,

$$\begin{aligned} B_3(s/N) &= \boldsymbol{\pi} [\Lambda \bar{H}(s/N) - Cs/N] R [\Lambda \bar{H}(s/N) - Cs/N] \mathbf{e} \\ &= \boldsymbol{\pi} (s/N)^2 \left[\Lambda \frac{\bar{H}(s/N)}{(s/N)} - C \right] R \cdot \left[\Lambda \frac{\bar{H}(s/N)}{(s/N)} - C \right] \mathbf{e} \\ &= \boldsymbol{\pi} (s/N)^2 [\Lambda \bar{H}_1 - C] R [\Lambda \bar{H}_1 - C] \mathbf{e} + o(N^{-2}) \\ &= -\boldsymbol{\pi} (s/N)^2 [\Lambda \bar{H}_1 - C] (\mathbf{a} + \mathbf{r}) + o(N^{-2}) \\ &= -\sum_d \pi_d [(a_d + r_d)(\lambda_d h_{d1} - c_d)] (s/N)^2 + o(N^{-2}), \end{aligned}$$

due to $R[C - \Lambda \bar{H}_1]e = a + r$; see Eq. (34). Under the heavy-traffic scaling,

$$r = c_\infty(1 - \rho_\infty)Re = c_\infty N^{-1}Re$$

is an $o(1)$ term. Observe that

$$\begin{aligned} & - (B_2(s/N) + B_3(s/N)) \\ &= (s/N)^2 \sum_{d=1}^D \pi_d [\lambda_d h_{d2}/2 + (a_d + o(1))(\lambda_d h_{d1} - c_d)] + s c_\infty / N^2 + o(N^{-2}). \end{aligned}$$

Rearranging Eq. (40) yields

$$\begin{aligned} \varphi_1(s/N) &= \pi_1 \frac{B_1(s/N)}{-(B_2(s/N) + B_3(s/N))} \\ &= \pi_1 \frac{c_\infty s / N^2 + o(N^{-2})}{(s/N)^2 \sum_{d=1}^D \pi_d [\lambda_d h_{d2}/2 + a_d(\lambda_d h_{d1} - c_d)] + c_\infty s / N^2 + o(N^{-2})} \\ &= \pi_1 \frac{1 + o(1)}{1 + c_\infty^{-1} \sum_d \pi_d [\lambda_d h_{d2}/2 + a_d(\lambda_d h_{d1} - c_d)] s + o(1)}. \end{aligned}$$

Let M be the desired mean stated in Theorem 3.3, that is,

$$M := c_\infty^{-1} \sum_d \pi_d \left[\hat{\lambda}_d h_{d2}/2 + a_d(\hat{\lambda}_d h_{d1} - c_d) \right].$$

Then, taking the heavy-traffic limit,

$$\lim_{N \rightarrow \infty} \varphi(s/N) = \lim_{N \rightarrow \infty} \frac{\varphi_1(s/N)}{\pi_1} \pi = \frac{\pi}{1 + Ms},$$

i.e. the LST $\varphi(s)$ converges in distribution to the LST of an exponentially distributed random variable with mean M .

References

1. Altman, E., Jimenez, T., Kofman, D.: DPS queues with stationary ergodic service times and the performance of TCP in overload. Proc. IEEE Infocom **2**, 975–983 (2004)
2. Altman, E., Avrachenkov, K., Ayesta, U.: A survey on discriminatory processor sharing. Queueing Syst. **53**, 53–63 (2006)
3. Asmussen, S.: The heavy traffic limit of a class of Markovian queueing models. Oper. Res. Lett. **6**(6), 301–306 (1987)
4. Asmussen, S.: Applied Probability and Queues, 2nd edn. Springer, New York (2003)
5. Avrachenkov, K., Ayesta, U., Brown, P., Núñez-Queija, R.: Discriminatory processor sharing revisited. Proc. IEEE Infocom **2**, 784–795 (2005)
6. Billingsley, P.: Convergence of Probability Measures. Wiley, New York (1999)
7. Blom, J., Kella, O., Mandjes, M., Thorsdottir, H.: Markov-modulated infinite-server queues with general service times. Queueing Syst. **76**(4), 403–424 (2014)

8. Borst, S.C., van Ooteghem, D., Zwart, A.: Tail asymptotics for discriminatory processor sharing queues with heavy-tailed service requirements. *Perform. Eval.* **61**, 281–298 (2005)
9. Borst, S.C., Núñez-Queija, R., Zwart, A.: Sojourn time asymptotics in processor sharing queues. *Queueing Syst.* **53**, 31–51 (2006)
10. Boxma, O.J., Kurkova, I.A.: The M/G/1 queue with two service speeds. *Adv. Appl. Probab.* **33**(2), 520–540 (2001)
11. Budhiraja, A., Ghosh, A., Liu, X.: Scheduling control for Markov-modulated single-server multiclass queueing systems in heavy traffic. *Queueing Syst.* **78**(1), 57–97 (2014)
12. Dimitrov, M.: Single-server queueing system with Markov-modulated arrivals and service times. *Pliska Stud. Math. Bulg.* **20**, 53–62 (2011)
13. Dorsman, J.L., Vlasiou, M., Zwart, B.: Heavy-traffic asymptotics for networks of parallel queues with Markov-modulated service speeds. *Queueing Syst.* **79**(3–4), 293–319 (2015)
14. Grishechkin, S.: On a relationship between processor sharing queues and Crump–Mode–Jagers branching processes. *Adv. Appl. Probab.* **24**(3), 653–698 (1992)
15. Falin, G., Falin, A.: Heavy traffic analysis of M/G/1 type queueing systems with Markov-modulated arrivals. *Soc. Estad. Invest. Oper.* **7**(2), 279–291 (1999)
16. Fayolle, G., Mitrani, I., Iasnogorodski, R.: Sharing a processor among many job classes. *J. ACM* **27**(3), 519–532 (1980)
17. Fiems, D., Altman, E.: Markov-modulated stochastic recursive equations with applications to delay-tolerant networks. *Perform. Eval.* **70**(11), 965–980 (2013)
18. Hassin, R., Haviv, M.: To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems. Kluwer Academic Publishers, Boston (2003)
19. Izagirre, A., Ayesta, U., Verloop, I.M.: Interpolation approximations for the steady-state distribution in multi-class resource-sharing systems. *Perform. Eval.* **91**, 56–79 (2015)
20. Izagirre, A., Ayesta, U., Verloop, I.M.: Heavy-traffic analysis of a multi-class queue with relative priorities. *Probab. Eng. Inform. Sc.* **29**(2), 153–180 (2015)
21. Katsuda, T.: State-space collapse in stationarity and its application to a multiclass single-server queue in heavy traffic. *Queueing Syst.* **65**(3), 237–273 (2010)
22. Kim, B., Kim, J.: A single server queue with Markov modulated service rates and impatient customers. *Perform. Eval.* **83–84**, 1–15 (2015)
23. Kingman, J.F.C.: The single server queue in heavy traffic. *Proc. Camb. Philos. Soc.* **57**(4), 902–904 (1961)
24. Kleinrock, L.: Time-shared systems: a theoretical treatment. *J. ACM* **14**(2), 242–261 (1967)
25. Mahabhashyam, S., Gautam, N.: On queues with Markov modulated service rates. *Queueing Syst.* **51**(1–2), 89–113 (2005)
26. Neuts, M.F.: Matrix-geometric Solutions in Stochastic Models, an Algorithmic Approach. The Johns Hopkins University Press, Baltimore (1981)
27. Núñez-Queija, R.: Processor-sharing models for integrated-services networks. PhD thesis, Eindhoven University of Technology (2000)
28. Purdue, P.: The M/M/1 queue in a Markovian environment. *Oper. Res.* **22**, 562–569 (1974)
29. Rege, K.M., Sengupta, B.: Queue-length distribution for the discriminatory processor-sharing queue. *Oper. Res.* **44**(4), 653–657 (1996)
30. Regterschot, G.J.K., de Smit, J.H.A.: The queue M/G/1 with Markov-modulated arrivals and service. *Math. Oper. Res.* **11**, 465–483 (1986)
31. Takine, T.: Single-server queues with Markov-modulated arrivals and service speed. *Queueing Syst.* **49**, 7–22 (2005)
32. Verloop, I.M., Ayesta, U., Núñez-Queija, R.: Heavy-traffic analysis of a multiple-phase network with discriminatory processor sharing. *Oper. Res.* **59**(3), 648–660 (2011)
33. Wu, Y., Bui, L., Johari, R.: Heavy traffic approximation of equilibria in resource sharing games. *IEEE J. Sel. Areas Commun.* **30**(11), 2200–2209 (2012)
34. Yechiali, U., Naor, P.: Queueing problems with heterogeneous arrival and service. *Oper. Res.* **19**, 722–734 (1971)
35. Zorine, A.V.: On ergodicity conditions in a polling model with Markov modulated input and state-dependent routing. *Queueing Syst.* **76**, 223–241 (2014)