Blending Randomness in Closed Queueing Network Models

Giuliano Casale

*Department of Computing, Imperial College London, SW7 2AZ, London, UK,
g.casale@imperial.ac.uk*

Mirco Tribastone

*Electronics and Computer Science, University of Southampton, United Kingdom,
m.tribastone@soton.ac.uk*

Peter G. Harrison

*Department of Computing, Imperial College London, SW7 2AZ, London, UK,
pgh@doc.ic.ac.uk*

Abstract

Random environments are stochastic models used to describe events occurring in the environment a system operates in. The goal is to describe events that affect performance and reliability such as breakdowns, repairs, or temporary degradations of resource capacities due to exogenous factors. Despite having been studied for decades, models that include both random environments and queueing networks remain difficult to analyse. To cope with this problem, we introduce the *blending algorithm*, a novel approximation for closed queueing network models in random environments. The algorithm seeks to obtain the stationary solution of the model by iteratively evaluating the dynamics of the system in between state changes of the environment. To make the approach scalable, the computation relies on a fluid approximation of the queueing network model. A validation study on 1,800 models shows that blending can save a significant amount of time compared to simulation, with an average accuracy that grows with the number of servers in each station. We also give an interpretation of this technique in terms of Laplace transforms and use this approach to determine convergence properties.

Key words: Random environments, fluid models, iterative approximation, transient analysis, Laplace transform

1 Introduction

The recent growth of interest in cloud computing has lead researchers to investigate systems that operate in environments shared by multiple tenants. Modelling these environments is challenging, since performance and availability levels for a user may change over time in a complex manner depending on the activity of the other tenants. For example, cloud virtual machines may experience transient CPU contention periods due to the activity of other virtual machines co-located on the same physical host. Similarly, a web farm deployed on the cloud may experience network bandwidth fluctuations due to traffic generated by other tenants. Accurate models of these systems require the ability to characterise exogenous events that arise in the operational environment of a system.

In performance and availability prediction, the problem of environment modelling may be tackled by associating a state, called a *stage*, to each possible condition of the environment. The active stage of the environment is assumed to evolve in time according to a Markov process. The description of the system then becomes parametric in the currently active stage. A model of this kind is said to describe the system as operating in a *random environment*. Queueing systems operating in random environments have been investigated for decades [30, 33], however models involving queueing networks are often intractable, calling for the development of customised performance evaluation methodologies.

Here we focus on systems that can be modelled as closed queueing network models, for example multi-tier software systems [42]. There is a limited literature on closed systems operating in random environments, often because they do not enjoy closed form expressions of the state probabilities except in particular cases [17, 20, 44]. In comparison, several works have focused on open systems operating in random environments, which are often tractable by means of conditional PASTA arguments [23]. The lack of results for the closed case limit their tractability to small problems where the underlying Markov model can be numerically evaluated. Alternatively, simulation may be used, but its running time makes it impractical for optimisation studies.

A computationally more viable route is offered by analytical approximations of closed networks based on product-form queueing theory. In the *average-environment* approximation [3], the system is studied with exponential state transitions having rates equal to their average value when the environment reaches equilibrium. This captures well environments where transitions happen *frequently* with respect to the representative timescales of the system. Conversely, in the *decomposition* approximation, an isolated model is evaluated for each stage of the random environment. The isolated solutions are

then weighted across stages using the equilibrium distribution of the random environment [14, 43]. This approach describes well an environment where transitions happen *rarely* compared to the representative timescales of the system (e.g., component failures). However, in intermediate situations both types of approximations tend to be loose. While numerical iterative methods exist to address this type of issue [16], they are unable to cope with state spaces of the scale considered in large queueing network models. This calls for developing approximation techniques for queueing network analysis in random environments that are more robust than average-environment and decomposition approximations.

To cope with this challenge, we present the *blending* algorithm, an iterative approximation scheme for closed queueing networks operating in random environments. Blending evaluates at each iteration the dynamics of the system in between two successive stage changes using transient analysis. A sequence of these analyses is used to determine a fixed point for the transient trajectories. The fixed point is then used to compute performance measures of the system for all stages of the random environment. From such embedded averages, it is simple to approximate the long-term equilibrium behaviour of the system and thus estimate its performance measures. We also interpret this value as a Laplace transform evaluated at a specific point and use this interpretation to develop a convergence analysis for the method.

The main challenge of the blending approach is to represent the transient dynamics of the system in between two successive stage changes. Since queueing network models often have an intractably large state space, we cannot use direct numerical methods to evaluate the underlying Markov process. To cope with state space explosion, we use a fluid approximation based on ordinary differential equations (ODEs) defined in the sense of Kurtz [28]. Even though Kurtz’s theory has found a wealth of applications in a broad range of disciplines, including chemistry [29], ecology [35], randomised algorithms [32], and communication networks [10], the application to closed models in random environments seems novel. Using Kurtz’s theory, we analyse the queueing network as a density-dependent process. This process is used to approximate inexpensively the transient evolution of queue-lengths in between successive stage changes. The number of ODEs does not increase with the number of jobs, which is the parameter that most influences the state space size, thus enabling the evaluation of large-scale models.

We evaluate the algorithm by a comprehensive numerical study consisting of 1,800 models. We solve these models by the blending algorithm and compare results with simulation. The study is quite extensive, as it corresponds to a month of computational time on a private compute cluster. The results for blending are in good agreement with simulation, but are far less expensive computationally.

The rest of this paper is organised as follows. Section 2 provides a motivating example. The class of models we focus on is defined in Section 3. The high-level structure of the proposed algorithm is defined in Section 4 and further developed in Section 5, where we introduce a fluid approximation for the class of models considered. In Section 6 we develop the blending algorithm and discuss its convergence in Section 7. Numerical results on random instances are given in Section 8. Related work is discussed in Section 9, followed by conclusions in Section 10. Appendix A provides related material on Markovian random environments.

2 Motivation

To motivate the present work, we present a toy example illustrating the effects of a random environment on the performance measures of a closed queueing network model. The interested reader can find similar examples in previous work, e.g., for open models [23]. The goal of this motivating example is to build intuition on the type of transient behaviours that a closed network can exhibit in presence of a random environment. As we show in Appendix A, the statements in this section hold for a more general class of systems than queueing networks.

Let us consider a tandem closed network composed of a first-come first-served (FCFS) single-server queue and a delay station (infinite-server queue, $-/GI/\infty$). The network topology is cyclic, with a population of $N = 2$ jobs. Service times at the single-server queue are exponentially distributed with rate μ . The queueing network operates in a Markovian random environment that causes server breakdown at the single-server queue with rate α_{21} , and successive repair with rate α_{12} . The random environment is independent of the state of the stations. Before breakdown, the single-server queue processes jobs with rate $\sigma_2 > 0$, but upon a breakdown it switches to a rate $\sigma_1 < \sigma_2$ until repaired. This simple system may be described by a continuous-time Markov chain (CTMC) with infinitesimal generator $\mathbf{Q}$ embedding an environment with generator $\mathbf{E}$, where

$$\mathbf{Q} = \left[\begin{array}{c|c} \mathbf{Q}_1^* - \alpha_{12}\mathbf{I} & \alpha_{12}\mathbf{I} \\ \hline \alpha_{21}\mathbf{I} & \mathbf{Q}_2^* - \alpha_{21}\mathbf{I} \end{array} \right], \quad \mathbf{E} = \begin{bmatrix} -\alpha_{12} & \alpha_{12} \\ \alpha_{21} & -\alpha_{21} \end{bmatrix},$$

$$\mathbf{Q}_1^* = \begin{bmatrix} -2\mu & 2\mu & 0 \\ \sigma_1 & -(\sigma_1 + \mu) & \mu \\ 0 & \sigma_1 & -\sigma_1 \end{bmatrix}, \quad \mathbf{Q}_2^* = \begin{bmatrix} -2\mu & 2\mu & 0 \\ \sigma_2 & -(\sigma_2 + \mu) & \mu \\ 0 & \sigma_2 & -\sigma_2 \end{bmatrix}.$$

Each state in $\mathbf{Q}$ represents a reachable state of the queueing network. The

matrices $\mathbf{Q}_1^*$ and $\mathbf{Q}_2^*$ are the infinitesimal generators for the system conditional on the environment being in stage 1 or 2, respectively. All Markov processes in this paper are ergodic and we also assume them to be irreducible.

To illustrate the challenges in analysing random environments, we recall that the *average-environment approximation* for the model under study can be obtained by solving the Markov process with generator $\mathbf{Q}^{AVG} = \pi_1^{env} \mathbf{Q}_1^* + \pi_2^{env} \mathbf{Q}_2^*$, where $\pi_1^{env} = \frac{\alpha_{21}}{\alpha_{12} + \alpha_{21}}$ is the probability of stage 1 in the random environment. Conversely, the *decomposition approximation* weights by π_1^{env} and $\pi_2^{env} = 1 - \pi_1^{env}$ the equilibrium solutions of the generators $\mathbf{Q}_1^*$ and $\mathbf{Q}_2^*$ considered in isolation, respectively.

We first consider this model with service rate $\mu = 0.1$ at the delay station, $\sigma_1 = 1.0$, and $\sigma_2 = 10.0$ for the single-server queue. The job population size is set to $N = 40$. Figure 1 illustrates queue-length sample paths for different values of $\alpha = \alpha_{12} + \alpha_{21}$, which is called the environment event rate.

Figure 1(a) shows the case of a high event rate with $\alpha_{12} = 2$ and $\alpha_{21} = 1$. In this case, the average-environment approximation is very accurate, see Appendix A for a mathematical justification.

Figure 1(b) shows the opposite case where changes in the environment occur rarely. Here we set $\alpha_{12} = 0.0002$ and $\alpha_{21} = 0.0001$. Note that the time scale represented on the horizontal axis is much longer than that in the other plots since stage transitions happen infrequently. Upon arrival of an event that moves the environment into stage 1, the network enters a long transient period where jobs migrate from the single-server queue to the delay station. Instead, upon a jump to stage 2, the jobs quickly accumulate in the single-server queue. On models of this kind, decomposition is known to be accurate because the length of the transient period is quite short compared to the overall holding time in a stage [7].

Figure 1(c) considers $\alpha_{12} = 0.02$ and $\alpha_{21} = 0.01$, i.e., comparable time scales between environmental events and service rates. Similarly to the decomposable model in Figure 1(b), the network cyclically enters transient periods where jobs migrate from one station to the other. However, the length of such transients is often comparable to the sojourn time of the random environment in each stage. Thus, the queue-length value observed at the instant of a stage change is a random variable that is tightly coupled with the sojourn time at each stage. In this situation, it is easy to show examples where neither average-environment nor decomposition are accurate. We show in Appendix A that linear combinations of the results of the two approximations do not provide a solution to this problem. This motivates our interest for developing novel analysis methods for the problem under study.

Fig. 1. Queue-length sample paths in a random environment

3 Model

We now introduce relevant notation and assumptions. The random environments considered in this paper evolve according to a CTMC jumping between stages. The environment is specified by the following parameters:

- E : number of stages;
- $e, h = 1, \dots, E$: stage indices;
- α_{eh} : rate of jump from stage e to stage h ;
- $\alpha_e = \sum_{h \neq e} \alpha_{eh}$: total outgoing rate from stage e ;
- π_h^{env} : probability of the environment being in stage h at equilibrium

Using the above definitions, the rates of the environment can be organised in an infinitesimal generator

$$\mathbf{E} = \begin{bmatrix} -\alpha_1 & \alpha_{12} & \dots & \alpha_{1E} \\ \alpha_{21} & -\alpha_2 & \dots & \alpha_{2E} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_{E1} & \alpha_{E2} & \dots & -\alpha_E \end{bmatrix} \quad (1)$$

In this paper, we consider only queueing networks processing a constant number of jobs, i.e., closed queueing networks. We let N be the number of circulating jobs in the network, and denote by M the number of stations. Both N and M are assumed to be independent of the currently active stage of the environment. All stations have unbounded buffer capacity and FCFS service discipline. Service times are independent and identically distributed random variables following a Coxian distribution, which includes exponential, hyper-exponential, and Erlang distributions as special cases [7]. We shall refer to the states of the Coxian service processes as *phases*.

When the random environment is in stage e , the queueing network is parameterised as follows:

- $i, j = 1, \dots, M$: station indices;
- $r_{i,j}^e$: routing probability from station i to station j ;
- S_i^e : number of servers at station i ;
- K_i^e : number of phases in the Coxian service process at station i ;
- $k \equiv k(i) = 1, \dots, K_i^e$: phase index for the Coxian service process at station i ;
- $\mu_{i,k}^e$: service rate in phase $k \equiv k(i)$ at station i ;
- $\phi_{i,k}^e$: probability for jobs at station i to complete after receiving service in phase $k \equiv k(i)$.

We shall omit the dependence of k upon the station index i whenever clear from the context.

Finally, we assume that upon a change of stage in the environment, the state of each station is instantaneously modified by moving all active jobs to the first phase of service. Note that exponential servers are insensitive to this rule, since they have a single phase of service. We refer to this assumption as the *reset rule*. The reset rule is a technical assumption used for convenience, but it describes correctly some systems, for example those where a repair after a breakdown leads to restarting the execution of all jobs in a system. Extensions of the blending methodology to other rules appear possible, but we leave this research for future work¹.

Running Case. *Throughout the paper, we use the following running case to illustrate our methodology. We consider again the tandem network in Section 2 and introduce a few modifications. As before, the delay server has rate $\mu_{1,1}^1 = \mu_{1,1}^2 = \mu$ in all stages. Instead, the FCFS queue has now $S_2^1 = S_2^2 = S_2$ servers in all stages: in stage 1 its service times are exponential with rate $\mu_{2,1}^1 = \sigma_1$; in stage 2 they are Erlang-2 with rates $\mu_{2,1}^2 = \mu_{2,2}^2 = \sigma_2$. The model is parametric in S_1^2, S_2^2 , and the job population N . Throughout the paper, we assume $\mu = 1$; $\sigma_1 = 2$, $\sigma_2 = 3$, $\alpha_{12} = 4$, $\alpha_{21} = 5$.*

¹ For example, if the service process at a station is independent of the active stage of the environment, one may simply want to avoid altering job phases when the environment changes. This feature is already supported for servers with exponential service time, but the extension to Coxian service processes require a minor modification of the blending algorithm by changing the definition of the $\mathbf{R}_{he}$ matrices. Early experiments suggest that blending would continue to work under this rule.

3.1 Continuous-Time Markov Chain

From the above definitions, the queueing network model may be readily described by a CTMC having state vector $\mathbf{Y} = (\mathbf{n}, e)$, where $e = 1, \dots, E$ denotes the current stage of the random environment. The distribution of jobs across stations is described by a vector $\mathbf{n} = (\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_M)$, $\mathbf{n}_i$ being the state vector of station i , further specified as $\mathbf{n}_i = (x_i, s_{i,1}, \dots, s_{i,K_i^e})$, where:

- x_i is the sum of two terms, the number of jobs waiting in the queue buffer and the number of jobs that are currently receiving service in phase 1. This representation simplifies the fluid analysis equations in Section 5.
- $s_{i,k}$ is the number of servers that are currently processing jobs in phase k , thus $\sum_{k=1}^{K_i^e} s_{i,k} = S_i^e$. If a server is idle, it is counted within $s_{i,1}$ together with the servers processing jobs in phase 1. If i is a delay station, we set $S_i^e = +\infty$.²

The above definitions imply that the total number of jobs at station i in any state is

$$n_i = x_i + \sum_{k=2}^{K_i^e} s_{i,k} \quad (2)$$

for all stages $e = 1, \dots, E$.

The CTMC defined on the proposed state space may be specified using a collection of *transition rate functions*. These are functions that provide the transition rate between any pair of states $(\mathbf{n}, e)$ and $(\mathbf{n}', h)$ that are connected by a transition in the CTMC. For the class of models we consider, only five transition rate functions are needed to specify the CTMC, they are shown in Table 1. The left column provides the service rate definition and its dependence on the current state $(\mathbf{n}, e)$ and other parameters specific to the transition. The right column describes the destination state $\mathbf{Y}' = (\mathbf{n}', h)$ after the transition occurs.

A description of the transition rate functions shown in Table 1 is as follows:

- $f_1^{dep}(\mathbf{n}, e, i, j)$ describes the rate at which jobs in station i complete after the first phase of service and are then routed to station j , conditional on the environment being in stage e . The case $i = j$ is excluded since the associated transition does not change the state of any station in the network.
- $f_k^{dep}(\mathbf{n}, e, i, j)$ is similar to $f_1^{dep}(\mathbf{n}, e, i, j)$, but provides the rate for completions at the k th phase of service. The case $i = j$ is here valid since completions returning at the same station restart from a different phase, phase 1, thus changing the state of the station.

² In implementations, this value may be set equal to the total job population N .

Transition rate	Destination state $\mathbf{Y}'$
$f_1^{dep}(\mathbf{n}, e, i, j) = \phi_{i1}^e r_{ij}^e \mu_{i1}^e \min(x_i, s_{i1})$ for $i, j = 1, \dots, M; i \neq j, e = 1, \dots, E$	$\mathbf{Y}' = (\mathbf{n}', e) = (\mathbf{n}_1, \dots, \mathbf{n}'_i, \dots, \mathbf{n}'_j, \dots, \mathbf{n}_M, e)$ $\mathbf{n}'_i = (x_i - 1, s_{i1}, s_{i2}, \dots, s_{iK_i^e})$ $\mathbf{n}'_j = (x_j + 1, s_{j1}, s_{j2}, \dots, s_{jK_j^e})$
$f_k^{dep}(\mathbf{n}, e, i, j) = \phi_{ik}^e r_{ij}^e \mu_{ik}^e s_{ik}$ for $i, j = 1, \dots, M, e = 1, \dots, E$ $k = 2, \dots, K_i^e$	$\mathbf{Y}' = (\mathbf{n}', e) = (\mathbf{n}_1, \dots, \mathbf{n}'_i, \dots, \mathbf{n}'_j, \dots, \mathbf{n}_M, e)$ $\mathbf{n}'_i = (x_i, s_{i1} + 1, \dots, s_{ik} - 1, \dots, s_{iK_i^e})$ $\mathbf{n}'_j = (x_j + 1, s_{j1}, s_{j2}, \dots, s_{jK_j^e})$
$f_1^{ph}(\mathbf{n}, e, i) = (1 - \phi_{i1}^e) \mu_{i1}^e \min(x_i, s_{i1})$ for $i = 1, \dots, M, e = 1, \dots, E$	$\mathbf{Y}' = (\mathbf{n}', e) = (\mathbf{n}_1, \dots, \mathbf{n}'_i, \dots, \mathbf{n}_M, e)$ $\mathbf{n}'_i = (x_i - 1, s_{i1} - 1, s_{i2} + 1, \dots, s_{iK_i^e})$
$f_k^{ph}(\mathbf{n}, e, i) = (1 - \phi_{ik}^e) \mu_{ik}^e s_{ik}$ for $i = 1, \dots, M, e = 1, \dots, E,$ $k = 2, \dots, K_i^e$	$\mathbf{Y}' = (\mathbf{n}', e) = (\mathbf{n}_1, \dots, \mathbf{n}'_i, \dots, \mathbf{n}_M, e)$ $\mathbf{n}'_i = (x_i, s_{i1}, \dots, s_{ik} - 1, s_{i(k+1)} + 1, \dots, s_{iK_i^e})$
$f^{env}(\mathbf{n}, e, h) = \alpha_{eh}$ for $e, h = 1, \dots, E, e \neq h$	$\mathbf{Y}' = (\mathbf{n}', h) = (\mathbf{n}'_1, \dots, \mathbf{n}'_M, h)$ $\mathbf{n}'_i = \left(x_i + \sum_{k=2}^{K_i^e} s_{ik}, s_{i1} + \sum_{k=2}^{K_i^e} s_{ik}, \underbrace{0, \dots, 0}_{K_i^e - 1} \right)$

Table 1

Transition rate functions. Source state is $(\mathbf{n}, e) = (\mathbf{n}_1, \dots, \mathbf{n}_i, \dots, \mathbf{n}_M, e)$.

- $f_1^{ph}(\mathbf{n}, e, i)$ describes the rate at which servers in station i advance from the first to the second Coxian phase, conditional on the environment being in stage e .
- $f_k^{ph}(\mathbf{n}, e, i)$ is similar to $f_1^{ph}(\mathbf{n}, e, i)$, but refers to the rate at which jobs in service advance from the k th phase to $(k + 1)$ th phase.
- $f^{env}(\mathbf{n}, e, h)$ describes the rate at which the random environment stage changes from stage e to stage $h \neq e$.

Using the transition rate functions in Table 1, the infinitesimal generator of the model as a whole may be written as

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_{11} & \cdots & \mathbf{Q}_{1e} & \cdots & \mathbf{Q}_{1E} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathbf{Q}_{e1} & \cdots & \mathbf{Q}_{ee} & \cdots & \mathbf{Q}_{eE} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \mathbf{Q}_{E1} & \cdots & \mathbf{Q}_{Ee} & \cdots & \mathbf{Q}_{EE} \end{bmatrix}, \quad (3)$$

where $\mathbf{Q}_{ee} = \mathbf{Q}_e^* - \alpha_e \mathbf{I}$, $\mathbf{Q}_e^*$ is the infinitesimal generator of the queueing network model considered in isolation with the parameterisation for stage e , $\mathbf{I}$ is the identity matrix, and for $e \neq h$, $\mathbf{Q}_{eh} = \alpha_{eh} \mathbf{R}_{eh}$, where $\mathbf{R}_{eh}$ is a

matrix encoding the reset rule. Matrix $\mathbf{R}_{eh}$ has a 1 at each row-index pair (i, j) corresponding to pair of states $\mathbf{Y} = (\mathbf{n}, e)$ and $\mathbf{Y}' = (\mathbf{n}', h)$ mapped by $f^{env}(\mathbf{n}, e, h)$ in Table 1; all the other entries are 0.

Running Case. According to the previous definitions, the state descriptor for the running case is

$$\mathbf{Y} = \begin{cases} \mathbf{n} = ((x_1, s_{1,1}), (x_2, s_{2,1})), & x_1 + x_2 = N, s_{1,1} = +\infty, s_{2,1} = S_2, & \text{if } e = 1 \\ \mathbf{n} = ((x_1, s_{1,1}), (x_2, s_{2,1}, s_{2,2})) & x_1 + x_2 = N, s_{1,1} = +\infty, s_{2,1} + s_{2,2} = S_2, & \text{if } e = 2 \end{cases}$$

For readability, we omit the constant terms in the above expressions, such that we may write $\mathbf{n} = (x_1, x_2)$ in stage $e = 1$ and $\mathbf{n} = (x_1, x_2, s_{2,1}, s_{2,2})$ in stage $e = 2$. Assume for example $N = 2$ and $S_2 = 1$ such that in stage 1 it is

$$\mathbf{n} \in \{(2, 0), (1, 1), (0, 2)\}$$

and in stage 2 it is

$$\mathbf{n} \in \{(2, 0, 1, 0), (1, 1, 1, 0), (1, 1, 0, 1), (0, 2, 1, 0), (0, 2, 0, 1)\}.$$

Given this state ordering, the CTMC of the model has generator

$$\mathbf{Q} = \left[\begin{array}{c|c} \mathbf{Q}_1^* - \alpha_{12}\mathbf{I} & \alpha_{12}\mathbf{R}_{12} \\ \hline \alpha_{21}\mathbf{R}_{21} & \mathbf{Q}_2^* - \alpha_{21}\mathbf{I} \end{array} \right],$$

where

$$\begin{aligned} \mathbf{Q}_1^* &= \begin{bmatrix} -2\mu & 2\mu & 0 \\ \sigma_1 & -(\sigma_1 + \mu) & \mu \\ 0 & \sigma_1 & -\sigma_1 \end{bmatrix}, & \mathbf{R}_{12} &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \\ \mathbf{R}_{21} &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}, & \mathbf{Q}_2^* &= \begin{bmatrix} -2\mu & 2\mu & 0 & 0 & 0 \\ 0 & -(\sigma_2 + \mu) & \sigma_2 & \mu & 0 \\ \sigma_2 & 0 & -(\sigma_2 + \mu) & 0 & \mu \\ 0 & 0 & 0 & -\sigma_2 & \sigma_2 \\ 0 & \sigma_2 & 0 & 0 & -\sigma_2 \end{bmatrix}. \end{aligned}$$

Notice in particular that $\mathbf{R}_{12}$ and $\mathbf{R}_{21}$ reset all FCFS queue servers to phase 1, according to the reset rule.

We now illustrate the transition rates for the system in stage 2 using the entries of $\mathbf{Q}_2^*$. Let $Q_2^*(i, j)$ be the entry of $\mathbf{Q}_2^*$ on row i and column j . The transition rates are then:

- $f_1^{dep}(\mathbf{n}, 2, 1, 2) = \mu_{1,1}^2 \min(x_1, +\infty) = \mu x_1$, e.g., $Q_2^*(1, 2)$, $Q_2^*(2, 4)$, $Q_2^*(3, 5)$.
- $f_1^{dep}(\mathbf{n}, 2, 2, 1) = 0$ since an Erlang server cannot complete in phase 1 ($\phi_{2,1}^2 = 0$), e.g., $Q_2^*(2, 1)$, $Q_2^*(4, 3)$.
- $f_2^{dep}(\mathbf{n}, 2, 2, 1) = \mu_{2,2} s_{2,2} = \sigma_2 s_{2,2}$, since an Erlang server always completes in phase 2, e.g., $Q_2^*(3, 1)$, $Q_2^*(5, 2)$.
- $f_1^{ph}(\mathbf{n}, 2, 2) = \mu_{2,1} \min(x_2, s_{2,1}) = \sigma_2 \min(x_2, s_{2,1})$ is the rate of servers moving from the first to the second Erlang stage, e.g., $Q_2^*(2, 3)$, $Q_2^*(4, 5)$.

4 Iterative Analysis at Stage Transition Times

The CTMC described in the previous section suffers state explosion issues when the number of queues, jobs, and service phases in the model increases. To cope with the state space explosion issue, we propose to characterise the equilibrium behaviour of the model focusing on the network state observed immediately after a stage transition occurs in the random environment. Specifically, we focus on the average queue-lengths seen at instants of stage change in the random environment.

We begin with some definitions. We consider the system at equilibrium and assume to observe it at an instant immediately *after* the environment enters stage h . We call *entry probability vector* for stage h the vector of equilibrium state probabilities for the system seen at this instant and denote it by $\boldsymbol{\eta}^h$. After spending an exponentially distributed time with rate α_h in stage h , the environment moves to a new stage $e \neq h$. We call *exit probability vector* for stage h , the vector of equilibrium state probabilities seen immediately *before* this instant and denote it by $\boldsymbol{\xi}^h$. Then, by construction

$$\boldsymbol{\xi}^h = \int_0^{+\infty} \boldsymbol{\eta}^h e^{\mathbf{Q}_h^* t} \alpha_h e^{-\alpha_h t} dt, \quad (4)$$

where $\alpha_h e^{-\alpha_h t}$ is the probability that the environment leaves stage h after t time units and the vector $\boldsymbol{\eta}^h e^{\mathbf{Q}_h^* t}$ gives the system state probabilities after t time units since entering stage h , as determined by the Kolmogorov forward equations of the underlying Markov process, conditional on the system staying in stage h along the whole sample path.

We now further investigate the relationship between entry and exit probability vectors. Consider the exit probability vector $\boldsymbol{\xi}^h$ for a stage h and assume that the environment next evolves to a stage $e \neq h$ at some instant. Then,

to determine the entry probability vector for stage e conditional on the destination stage e , ξ^h needs to be multiplied by $\mathbf{R}_{he}$ that applies the reset rule for this stage transition. Removing the conditioning on the destination stage, the entry probability vector η^e will be determined as an average over all the vectors $\xi^h \mathbf{R}_{he}$ for all jumps originating from stages $h \neq e$. That is, we can write

$$\eta^e = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \xi^h \mathbf{R}_{he},$$

where p_{he}^{src} is the probability of stage h being the source of the jump to e . This probability term can be readily obtained by standard arguments as

$$p_{he}^{src} = \frac{\pi_h^{env} \alpha_{he}}{\sum_{j \neq e} \pi_j^{env} \alpha_{je}}, \quad \sum_{h \neq e} p_{he}^{src} = 1, \quad (5)$$

Equations (4)-(5) immediately imply the following proposition.

Proposition 1 *In the CTMC with generator (3), at equilibrium, the probability vectors embedded at stage entry instants satisfy the balance equation*

$$\eta^e = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \left(\int_0^\infty \eta^h e^{\mathbf{Q}_h^* t} \alpha_h e^{-\alpha_h t} dt \right) \mathbf{R}_{he}, \quad (6)$$

for all stages $e = 1, \dots, E$, where p_{he}^{src} , given in (5), is the probability of a jump to stage e having originated from stage $h \neq e$.

Running Case. Assume $N = 2$, $S_2 = 2$. Let $\pi_1 = [\pi(x_1, x_2)]$ and $\pi_2 = [\pi(x_1, x_2, s_{21}, s_{22})]$ be the vectors of equilibrium probabilities for stages 1 and 2, respectively, specified according to the state ordering of the running case. Then for the given parameters

$$\pi_1 = [0.1780, 0.2289, 0.1486],$$

$$\pi_2 = [0.1221, 0.1423, 0.0474, 0.0921, 0.0405].$$

By basic Markov chain theory, the entry probability vectors are given by

$$\eta_1 = \frac{\pi_2 \mathbf{R}_{21}}{\pi_2 \mathbf{R}_{21} \mathbf{1}} = [0.2746, 0.4270, 0.2983]$$

$$\eta_2 = \frac{\pi_1 \mathbf{R}_{12}}{\pi_1 \mathbf{R}_{12} \mathbf{1}} = [0.3204, 0.4120, 0, 0.2676, 0]$$

where $\mathbf{1}$ is a column vector of ones. It can be readily verified that η_1 and η_2 satisfy (6), where $p_{12}^{src} = 1$ and $p_{21}^{src} = 1$. For example, using the Laplace transform of the matrix exponential in (6) it can be readily seen that $\eta_1 =$

$\alpha_{21}\boldsymbol{\eta}_2(\alpha_{21}\mathbf{I}-\mathbf{Q}_2^*)^{-1}\mathbf{R}_{21}$ and $\boldsymbol{\eta}_2 = \alpha_{12}\boldsymbol{\eta}_1(\alpha_{12}\mathbf{I}-\mathbf{Q}_1^*)^{-1}\mathbf{R}_{12}$, which can be verified without explicit integration.

We are now ready to characterise the mean queue-lengths embedded at stage transition instants. Let $\bar{n}_i^h(t)$ be a deterministic function representing the mean number of jobs in station i at time t since entering stage h , conditional on the absence of stage changes occurring during these t time units. Similarly, let $\bar{x}_i^h(t)$ and $\bar{s}_k^h(t)$ be the time averaged values of the state descriptors x_i and s_k in stage h , computed at time t after entering that stage. From (2) it follows that

$$\bar{n}_i^h(t) = \bar{x}_i^h(t) + \sum_{k=2}^{K_{max}} \bar{s}_k^h(t)$$

Denote by $\bar{\mathbf{n}}^h(t) = (\bar{n}_1^h(t), \dots, \bar{n}_M^h(t))$ the corresponding mean queue-length vector. Also, let $\mathcal{L}(\mathbf{X}(t); \alpha) = \int_0^\infty \mathbf{X}(t)e^{-\alpha t}dt$ be the Laplace transform of vector $\mathbf{X}(t)$, $\alpha > 0$.

Corollary 1

$$\bar{\mathbf{n}}^e(0) = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \alpha_h \mathbf{N}^h(\alpha_h), \quad (7)$$

where $\mathbf{N}^h(\alpha) = \mathcal{L}(\bar{\mathbf{n}}^h(t); \alpha)$.

Proof 1 Using the Kolmogorov forward equations, the state of the system after t time units since the environment enters stage h can be written as $\boldsymbol{\eta}^h(t) = \boldsymbol{\eta}^h e^{\mathbf{Q}_h^* t}$. Since the matrix exponential is a continuous mapping, $\boldsymbol{\eta}^h(t)$ is a continuous vector function by definition. Let $\mathbf{v}_i^h$ be the vector of coefficients such that $\bar{n}_i^h(t) = \boldsymbol{\eta}^h(t) \mathbf{v}_i^h = \boldsymbol{\eta}^h e^{\mathbf{Q}_h^* t} \mathbf{v}_i^h$. This vector always exists as it multiplies the probability of a state $(\mathbf{n}, h)$ in $\boldsymbol{\eta}^h(t)$ by $n_i^h = x_i^h + \sum_{k=2}^{K_i^e} s_{i,k}$. Using (6), it is then

$$\bar{n}_i^e(0) = \boldsymbol{\eta}^e(0) \mathbf{v}_i^e = \boldsymbol{\eta}^e \mathbf{v}_i^e = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \left(\int_0^{+\infty} \boldsymbol{\eta}^h e^{\mathbf{Q}_h^* t} \alpha_h e^{-\alpha_h t} dt \right) \mathbf{R}_{he} \mathbf{v}_i^e,$$

where we used that $\boldsymbol{\eta}^e(0) = \boldsymbol{\eta}^e e^{\mathbf{Q}_e^* 0} = \boldsymbol{\eta}^e$. We now observe that we can replace $\mathbf{R}_{he} \mathbf{v}_i^e$ with $\mathbf{v}_i^h$, since the reset rule does not change the total populations at each queue. This implies that

$$\bar{n}_i^e(0) = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \alpha_h \left(\int_0^{+\infty} \bar{n}_i^h(t) e^{-\alpha_h t} dt \right).$$

Finally, the Laplace transform exists since $\bar{n}_i^h(t)$ is continuous, which follows by the continuity of $\boldsymbol{\eta}^h(t)$, and always defined for $\alpha_h > 0$, since $\int_0^\infty \boldsymbol{\eta}^h e^{(-\alpha_h \mathbf{I} + \mathbf{Q}_h^*)t} \mathbf{v}_h^e dt =$

$\boldsymbol{\eta}^h(\alpha_h \mathbf{I} - \mathbf{Q}_h^*)^{-1} \mathbf{v}_h^e$ and $(\alpha_h \mathbf{I} - \mathbf{Q}_h^*)^{-1}$ is always invertible being $\mathbf{Q}_h^*$ a generator. $\square$

Running Case. From $\boldsymbol{\eta}_1$, we can readily compute the mean queue-length at the FCFS queue seen upon stage change as $\bar{n}_2^1(0) = \boldsymbol{\eta}_1[0, 1, 2]^T = 1.0237$ jobs and $\bar{n}_2^2(0) = \boldsymbol{\eta}_2[0, 1, 1, 2, 2]^T = 0.9471$ jobs. To illustrate (7), we plot the trajectory of $\bar{\mathbf{n}}^2(t)$:

The quantity $\alpha_{21} \mathbf{N}^2(\alpha_{21})$ (7) is simply the average of $\bar{\mathbf{n}}^2(t)$ computed with respect to the stage transition time density $\alpha_{21} e^{-\alpha_{21} t}$ and equals $\bar{\mathbf{n}}^1(0) = \alpha_{21} \mathbf{N}^2(\alpha_{21}) = 1.0237$.

Expression (7) cannot be exploited directly to solve the model as it requires the unknown mean queue-length transient trajectories $\bar{n}_i(t)$ in order to compute the Laplace transform terms. The blending algorithm attempts to overcome this issue by developing a fixed point iteration with the goal of approximating these trajectories. A key advantage of this approach compared to direct numerical evaluation of the CTMC underlying the queueing network is that the fixed point iteration we propose uses a system of ODEs that grows in size only as $O(MK_{max})$, being M the number of queues and $K_{max} = \max_{i=1..E} K_i$ the maximum number of phases across all service processes, and independently of the number of jobs in the system.

We begin with the observation that, if such trajectories satisfy an initial value problem defined by an ODE system

$$\frac{d\bar{\mathbf{n}}^h(t)}{dt} = \mathbf{U}^h(\bar{\mathbf{n}}^h(t); \bar{\mathbf{n}}^h(0)), \quad (8)$$

for some piece-wise differentiable vector function $\mathbf{U}^h(\cdot)$ and initial vector $\bar{\mathbf{n}}^h(0)$, then the Laplace transforms are readily determined as

$$\alpha \mathbf{N}^h(\alpha) = \bar{\mathbf{n}}^h(0) + \mathcal{L}(\mathbf{U}^h(\bar{\mathbf{n}}^h(t)); \alpha),$$

such that (7) becomes

$$\bar{\mathbf{n}}^e(0) = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \left(\bar{\mathbf{n}}^h(0) + \mathcal{L}(\mathbf{U}^h(\bar{\mathbf{n}}^h(t)); \alpha_h) \right), \quad e = 1, \dots, E. \quad (9)$$

where the Laplace transform is now evaluated at α_h . The last equation may be used to develop a fixed-point iteration as follows. Let $\bar{\mathbf{n}}_c^e(t)$ be the estimate of the $\bar{\mathbf{n}}^e(t)$ vector at iteration $c \geq 1$. Then we can replace (9) with the iteration

$$\bar{\mathbf{n}}_{c+1}^e(0) = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \left(\bar{\mathbf{n}}_c^h(0) + \mathcal{L}(\mathbf{U}^h(\bar{\mathbf{n}}_c^h(t)); \alpha_h) \right), \quad (10)$$

for all stages $e = 1, \dots, E$. By iteratively evaluating (10), the blending algorithm aims at finding a fixed point solution $\bar{\mathbf{n}}_\infty^e(t)$ satisfying

$$\bar{\mathbf{n}}_\infty^e(0) = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \left(\bar{\mathbf{n}}_\infty^h(0) + \mathcal{L}(\mathbf{U}^h(\bar{\mathbf{n}}_\infty^h(t)); \alpha_h) \right), \quad (11)$$

such that we can operate the approximation $\bar{\mathbf{n}}^e(t) \approx \bar{\mathbf{n}}_\infty^e(t)$, for all stages $e = 1, \dots, E$. As we show later, we can directly use the estimates $\bar{\mathbf{n}}_\infty^e(t)$ to calculate performance measures such as mean throughputs, utilisations, and mean response times at equilibrium for the queueing network.

In the next sections, we investigate the fixed point iteration (10), first defining in Section 5 a suitable vector function $\mathbf{U}^h(\cdot)$ to approximate the $\bar{n}_i(t)$ trajectories in closed queueing networks, then introducing in Section 6 the blending algorithm. The Laplace transform analysis will be exploited in particular in Section 7 for convergence analysis.

5 Fluid Limits

In this section, we use the theory of fluid limits developed by Kurtz [28] to define an approximate ODE analysis method for queueing network models. This will provide the form of the ODE system (8) to be used in the fixed point iteration (10). The parameters and the formulas are assumed to refer to an arbitrary stage h , however we shall omit in this section this index for ease of readability. Under this assumption, the state descriptor of the CTMC simplifies to $\mathbf{Y} = \mathbf{n}$.

For queueing networks, the fluid limit we consider describes the behaviour of the Markov process underlying the queueing network as the population of jobs N and the number of servers S_i at each queue i grow simultaneously

to infinity. Such growth is assumed to be proportional to a scale parameter V , which increases both the number of servers and jobs in the system, but keeps their ratio constant. For example, when $V = 10$, a model with single server queues and $N = 10$ jobs is scaled to a model with $N' = 100$ jobs and multi-server queues having each $S'_i = 10$ servers. Since capacity grows at the same factor of the population growth, the network avoids saturation in the asymptotic regime. In this limit, the network also maintains the same topology, routing matrix, and service rates of the unscaled model.

For the class of models we consider, and recalling our assumption of conditioning upon a stage, a fluid limit with the above properties is obtained by taking a family of CTMCs $\{\mathbf{Y}_V(t), V \in \mathbb{N}\}$ indexed by the scale parameter V . Each vector $\mathbf{Y}_V(t)$ is the state vector for the system at time t , obtained by observing a sample path of the CTMC underlying the queueing network model instantiated with scale parameter V . Let $\mathbf{Y}_0$ be the initial state of the CTMC at time $t = 0$, given by the initial placements of jobs across the stations and the initial phase of all servers. Then $\mathbf{Y}_1(t)$, where $V = 1$, is the state vector reachable from $\mathbf{Y}_0$ through the rates in Table 1, conditional on the environment remaining on the same stage. Note that the case $V = 1$ corresponds to the original model being investigated, thus $\mathbf{Y}_1(t) = \mathbf{n}(t)$. The vector $\mathbf{Y}_2(t)$, where $V = 2$, is the state vector reachable from $2\mathbf{Y}_0$ through the same rates. Generalising, $\mathbf{Y}_V(t)$ is the state vector reachable from $V\mathbf{Y}_0$, for all $V \in \mathbb{N}$. Recalling that each state $\mathbf{Y}$ includes both a description of jobs and servers, it follows that using $V\mathbf{Y}_0$ as initial state scales both the number of jobs and the number of servers in the queueing network.

The fluid approximation of interest here consists in the limit behaviour of the *normalised* CTMC $\mathbf{Y}_V(t)/V$ when V grows asymptotically large, with initial state $\mathbf{Y}_0$. A necessary condition to apply the limit theorem is that all transition rates enjoy the so-called *density-dependent* form [28], which requires that all CTMC transition rates need to be proportional to V once rewritten in terms of the normalised state descriptor $\mathbf{Y}/V = \mathbf{n}/V$. It can be readily verified that all the transition rates in Table 1 can be expressed in such a form, except for f^{env} that however we do not consider here having conditioned the analysis on a given stage. For instance, making explicit the normalised state descriptor $\mathbf{Y}/V = \mathbf{n}/V$ in the f_1^{dep} rates yields

$$\begin{aligned} f_1^{dep}(\mathbf{n}, i, j) &= \phi_{i1} r_{ij} \mu_{i1} \min(x_i, s_{i1}) \\ &= V \phi_{i1} r_{ij} \mu_{i1} \min(x_i/V, s_{i1}/V) \\ &= V f_1^{dep}(\mathbf{n}/V, i, j), \end{aligned} \tag{12}$$

which is exactly the density-dependent form. In Kurtz's theory, this form implies that, in the limit, a sample path of the normalised CTMC takes increasingly small steps, of order $O(1/V)$, at increasingly large rates, of order $O(V)$. As $V \rightarrow +\infty$, a sample path can therefore be related to a continuous trajec-

tory. Kurtz shows that this trajectory is the solution of a system of ordinary differential equations (ODEs), defined in terms of the density-dependent transition rates of the CTMC. The structure of this system of ODEs is discussed in the next subsection.

5.1 ODE Approximation, Convergence, and Accuracy

Using the methodology developed by Kurtz, the fluid approximation defined above is evaluated as follows. We consider the limit $V \rightarrow +\infty$ and approximate the state descriptor $\mathbf{Y}_V(t)$ leveraging on a deterministic trajectory $\mathbf{y}(t)$ defined by the nonlinear ODE system

$$\frac{d\mathbf{y}(t)}{dt} = \sum_{i,j=1}^M \sum_{k=1}^{K_i} f_k^{dep}(\mathbf{y}(t), i, j) \boldsymbol{\delta}_k^{dep}(i, j) + \sum_{i=1}^M \sum_{k=1}^{K_i} f_k^{ph}(\mathbf{y}(t), i) \boldsymbol{\delta}_k^{ph}(i), \quad (13)$$

where $\delta_k^{dep}(i, j)$ and $\boldsymbol{\delta}_k^{ph}(i)$ are *jump vectors* that record the net change of probability mass due to a transition of type $f_k^{dep}(\mathbf{y}(t), i, j)$ and $f_k^{ph}(\mathbf{y}(t), i)$, respectively. The jump vector for each transition rate function is constructed by taking $\mathbf{n}' - \mathbf{n}$ in the corresponding row in Table 1. Note that (13) is mass-conserving by the definitions in Table 1.

Let $\mathbf{y}(0)$ be the initial state of (13). Using Kurtz's theorem [28, Theorem 3.1] and the fact that the transition rates are Lipschitz continuous everywhere, which implies global existence and uniqueness of the initial value problem (13), it can be shown that, for any finite T , it holds that, for any $\varepsilon > 0$,

$$\lim_{V \rightarrow \infty} \mathbb{P} \left\{ \sup_{t \leq T} |\mathbf{Y}_V(t)/V - \mathbf{y}(t)| > \varepsilon \right\} = 0. \quad (14)$$

This implies that a sample path (of finite length) is asymptotically indistinguishable, in probability, from the unique solution to the ODE system (13). Furthermore, this fluid limit can be shown to imply *convergence in mean* [6]:

$$\lim_{V \rightarrow \infty} \mathbb{E} \{ |\mathbf{Y}_V(t)/V - \mathbf{y}(t)| \} = 0, \quad t \geq 0, \quad (15)$$

which provides the motivation for considering the approximation $\mathbb{E}\{\mathbf{Y}_V(t)\} \approx V\mathbf{y}(t)$ that relates the expected queue length to the re-scaled ODE trajectory. The ODE solution $\mathbf{y}(t)$ will be used to approximate the mean performance measures of the original queueing network in which $V = 1$.

Running Case. We assume $N = V$, $S_2 = V$ and determine the standard deviation of the mean number of jobs at the FCFS queue at time $t = 100$ for increasing values of V , computed over 30 sample paths and conditional on the system being in stage 2. The resulting plot is as follows:

Consistently with (15), the figure shows that as V grows the system becomes increasingly deterministic, thus easier to approximate with the trajectory of an ODE system. Notice in particular that even small values of V can produce a sharp decrease of variability.

5.2 Building the ODE System

Despite its apparently complex form, (13) is in practice simple to specify. Let us define $\mathbf{R} = [r_{ij}]_{M \times M}$ to be the routing probability matrix of the queueing network model. For example, in a queueing network with exponential servers $\bar{\mathbf{n}}(t) = \bar{\mathbf{x}}(t)$, where $\bar{\mathbf{n}}(t) = (\bar{n}_1(t), \dots, \bar{n}_M(t))$ and $\bar{\mathbf{x}}(t) = (\bar{x}_1(t), \dots, \bar{x}_M(t))$ being $\bar{x}_i(t)$ the mean value of $x_i(t)$ at time t . Since all servers are in phase 1, the ODE system admits the simplified form

$$\frac{d\bar{\mathbf{n}}(t)}{dt} = (-\mathbf{I} + \mathbf{R}^T)\boldsymbol{\mu}_1 \min(\bar{\mathbf{n}}(t), \mathbf{S}), \quad (16)$$

where the minimum of two vectors is taken element-wise, the vector $\mathbf{S} = (S_1, \dots, S_M)$ collects the number of servers at each station, $\boldsymbol{\mu}_1 = \text{diag}(\mu_{11}, \dots, \mu_{M1})$ is the vector of service rates at the exponential phase, and the ODE coefficient matrix is thus

$$(-\mathbf{I} + \mathbf{R}^T)\boldsymbol{\mu}_1 = \begin{bmatrix} -(1 - r_{1,1})\mu_{1,1} & r_{2,1}\mu_{2,1} & \dots & r_{M,1}\mu_{M,1} \\ r_{1,2}\mu_{1,1} & -(1 - r_{2,2})\mu_{2,1} & \dots & r_{M,2}\mu_{M,1} \\ \vdots & \vdots & \ddots & \vdots \\ r_{1,M}\mu_{1,1} & r_{2,M}\mu_{2,1} & \dots & -(1 - r_{M,M})\mu_{M,1} \end{bmatrix}. \quad (17)$$

Notice that the coefficient matrix includes only rates pertaining to the f_1^{dep} transitions in Table 1 since PHs are exponential and that its transpose is an

infinitesimal generator since $\sum_{j=1}^M r_{i,j} = 1$, for all $i = 1, \dots, M$.³

Running Case. The ODE system (17) can describe stage 1, where service times are exponentially distributed. Let $\mathbf{S} = (S_1, S_2) = (\infty, 2)$, $r_{1,2} = r_{2,1} = 1$, $r_{1,1} = r_{2,2} = 0$, $\mu_{1,1} = \mu$, $\mu_{2,1} = \sigma_1$. Thus

$$(-\mathbf{I} + \mathbf{R}^T)\boldsymbol{\mu}_1 = \begin{bmatrix} -\mu & \sigma_1 \\ \mu & -\sigma_1 \end{bmatrix}$$

and $(-\mathbf{I} + \mathbf{R}^T)\boldsymbol{\mu}_1^T$ is an infinitesimal generator. Recall that $\bar{n}_i^e(t)$ is the mean queue-length at station i at time t , when the environment is in stage e . Then the ODE system may be written as

$$\begin{aligned} \frac{d\bar{n}_1^1(t)}{dt} &= -\mu\bar{n}_1^1(t) + \sigma_1 \min(\bar{n}_2^1(t), S_2) \\ \frac{d\bar{n}_2^1(t)}{dt} &= \mu\bar{n}_1^1(t) - \sigma_1 \min(\bar{n}_2^1(t), S_2) \end{aligned}$$

for arbitrary initial conditions $\bar{n}_1^1(0)$, $\bar{n}_2^1(0)$.

In the general case where service processes are Coxian instead of exponential, the ODE system has the following structure. Let us partition the vector of unknowns as $\mathbf{y}(t) = (\bar{\mathbf{x}}(t), \bar{\mathbf{s}}_1(t), \dots, \bar{\mathbf{s}}_{K_{\max}}(t))$, where $K_{\max} = \max_{i=1, \dots, E} K_i$ and $\bar{\mathbf{s}}_k(t) = (s_{1k}(t), \dots, s_{Mk}(t))$ for phase k , where we set $s_{ik}(t) \equiv 0$ if $k > K_i$. Compared to the previous case, we have now that servers may be in different phases, but their number does not change within a stage, thus

$$\sum_{k=1}^{K_{\max}} \bar{\mathbf{s}}_k(t) = \mathbf{S}. \quad (18)$$

Recalling the definitions in Section 3, we now define for phase $k = 1, \dots, K_{\max}$ the diagonal matrix $\boldsymbol{\mu}_k = \text{diag}(\mu_{1k}, \dots, \mu_{Mk})$ to be the service rate matrix and $\boldsymbol{\Phi}_k = \text{diag}(\phi_{1k}, \dots, \phi_{Mk})$ to be the service completion probability matrix, where we set $\phi_{iK_i} = 1$ and $\mu_{ik} = 0$ if $k > K_i$. The ODE system then becomes

$$\frac{d\bar{\mathbf{x}}(t)}{dt} = (-\mathbf{I} + \mathbf{R}^T \boldsymbol{\Phi}_1)\boldsymbol{\mu}_1 \min(\bar{\mathbf{x}}(t), \bar{\mathbf{s}}_1(t)) + \sum_{k=2}^{K_{\max}} \mathbf{R}^T \boldsymbol{\Phi}_k \boldsymbol{\mu}_k \bar{\mathbf{s}}_k(t), \quad (19)$$

$$\frac{d\bar{\mathbf{s}}_2(t)}{dt} = -\boldsymbol{\mu}_2 \bar{\mathbf{s}}_2(t) + (\mathbf{I} - \boldsymbol{\Phi}_1)\boldsymbol{\mu}_1 \min(\bar{\mathbf{x}}(t), \bar{\mathbf{s}}_1(t)), \quad (20)$$

$$\frac{d\bar{\mathbf{s}}_k(t)}{dt} = -\boldsymbol{\mu}_k \bar{\mathbf{s}}_k(t) + (\mathbf{I} - \boldsymbol{\Phi}_{k-1})\boldsymbol{\mu}_{k-1} \bar{\mathbf{s}}_{k-1}(t), \quad k = 3, \dots, K_{\max} \quad (21)$$

³ Throughout the next sections, we use the ODE system convention of storing state variables in column vectors. By taking the transpose of our equations it is possible to retrieve the row-oriented notation used in Markov process analysis.

where by (18) it is $\bar{\mathbf{s}}_1(t) = \mathbf{S} - \sum_{k=2}^{K_{\max}} \bar{\mathbf{s}}_k(t)$. The mean queue-lengths are computed by the usual relationship $\bar{\mathbf{n}}(t) = \bar{\mathbf{x}}(t) + \sum_{k=2}^{K_{\max}} \bar{\mathbf{s}}_k(t)$.

Running Case. The ODE system (19)-(21) can be used to describe the queueing network while the environment is in stage 2. Here $\mathbf{S} = (S_1, S_2) = (\infty, 2)$, $r_{1,2} = r_{2,1} = 1$, $r_{1,1} = r_{2,2} = 0$, $\mu_{1,1} = \mu$, $\mu_{2,1} = \sigma_2$. Furthermore,

$$\mathbf{R} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \mathbf{\Phi}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \mathbf{\Phi}_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \boldsymbol{\mu}_1 = \begin{bmatrix} \mu & 0 \\ 0 & \sigma_2 \end{bmatrix}, \boldsymbol{\mu}_2 = \begin{bmatrix} 0 & 0 \\ 0 & \sigma_2 \end{bmatrix},$$

thus (19)-(21) become

$$\frac{d\bar{x}_1(t)}{dt} = -\mu\bar{x}_1(t) + \sigma_2\bar{s}_{2,2}(t) \quad (22)$$

$$\frac{d\bar{x}_2(t)}{dt} = \mu\bar{x}_1(t) - \sigma_2 \min(\bar{x}_2(t), \bar{s}_{2,1}(t)) \quad (23)$$

$$\frac{d\bar{s}_{2,2}(t)}{dt} = -\sigma_2\bar{s}_{2,2}(t) + \sigma_2 \min(\bar{x}_2(t), \bar{s}_{2,1}(t)) \quad (24)$$

where $\bar{s}_{2,1}(t) = S_2 - \bar{s}_{2,2}(t)$.

The ODE system (19)-(21) can be directly evaluated by a nonlinear ODE solver. The initial values are $\bar{\mathbf{s}}_2(0) = \dots = \bar{\mathbf{s}}_{K_{\max}}(0) = \mathbf{0}$, to impose that jobs start service from phase 1, and a queue-length vector $\bar{\mathbf{x}}(0) \geq \mathbf{0}$, $\sum_j \bar{x}_j(0) = N$, such that the total population of jobs is N . Since the ODE system size is $M + \sum_{i=1}^M (K_i^e - 1)$ equations and variables, the fluid approximation provides major computational savings compared to a CTMC analysis where the number of global balance equations and variables grows exponentially as $O(EN^M \prod_{i=1}^{K_i} S_i^{K_i})$. As explained before, despite its small size, the fluid approximation provides accuracy and convergence guarantees as the system size increases.

5.3 Laplace Transform of Queue-Length Trajectories

We now seek the expression of the Laplace transform used in Corollary 7. From this expressions, we will develop in Section 6 the blending algorithm using the iterative approach outlined in Section 4.

Thanks to Lipschitz continuity, the ODE system (19)-(21) has a unique solution with piece-wise continuous derivatives. Since the network is closed, both vectors $\bar{\mathbf{x}}(t)$ and $\bar{\mathbf{s}}_k(t)$, $k \geq 2$, have to remain bounded, since $\bar{s}_k^h(t) \leq \min(N, S_k)$ and $\bar{x}_k(t) \leq N$ for all queues k and stages h , thus their Laplace

transform exists together with their derivatives. This makes the fluid approach suitable for use within the iterative analysis outlined in Section 4. Let

$$\mathbf{X}(\alpha) = \mathcal{L}(\bar{\mathbf{x}}(t); \alpha), \quad \mathbf{M}(\alpha) = \mathcal{L}(\min(\bar{\mathbf{x}}(t), \bar{\mathbf{s}}_1(t)); \alpha), \quad \mathbf{S}_k(\alpha) = \mathcal{L}(\bar{\mathbf{s}}_k(t); \alpha).$$

The Laplace transform of the ODE system is

$$\alpha \mathbf{X}(\alpha) = \bar{\mathbf{x}}(0) + (-\mathbf{I} + \mathbf{R}^T \Phi_1) \boldsymbol{\mu}_1 \mathbf{M}(\alpha) + \sum_{k=2}^{K_{\max}} \mathbf{R}^T \Phi_k \boldsymbol{\mu}_k \mathbf{S}_k(\alpha), \quad (25)$$

$$\alpha \mathbf{S}_2(\alpha) = -\boldsymbol{\mu}_2 \mathbf{S}_2(\alpha) + (\mathbf{I} - \Phi_1) \boldsymbol{\mu}_1 \mathbf{M}(\alpha), \quad (26)$$

$$\alpha \mathbf{S}_k(\alpha) = -\boldsymbol{\mu}_k \mathbf{S}_k(\alpha) + (\mathbf{I} - \Phi_{k-1}) \boldsymbol{\mu}_{k-1} \mathbf{S}_{k-1}(\alpha), \quad k = 3, \dots, K_{\max} \quad (27)$$

where we used that due to the reset rule $\bar{\mathbf{s}}_k(0) = \mathbf{0}$, $k \geq 2$. Bringing $\mathbf{S}_2(\alpha)$ and $\mathbf{S}_k(\alpha)$ to the left hand side and noting that $\alpha \mathbf{I} + \boldsymbol{\mu}_k$, $k \geq 2$, is a diagonal matrix with strictly positive entries, thus always invertible, (26)-(27) imply

$$\mathbf{S}_k(\alpha) = \prod_{j=k \dots 2} ((\alpha \mathbf{I} + \boldsymbol{\mu}_j)^{-1} (\mathbf{I} - \Phi_{j-1}) \boldsymbol{\mu}_{j-1}) \mathbf{M}(\alpha), \quad k = 2, \dots, K_{\max}$$

Since the Laplace transform of the queue-length vector is $\mathbf{N}(\alpha) = \mathcal{L}(\bar{\mathbf{n}}(t); \alpha) = \mathbf{X}(\alpha) + \sum_{k=2}^{K_{\max}} \mathbf{S}_k(\alpha)$, being by our definitions $n_j = x_j + \sum_{k=2}^{K_j} s_{j,k}$ and having assumed the convention $\mathbf{S}_k = \mathbf{0}$ for $k > K_i$, we conclude that

$$\alpha \mathbf{N}(\alpha) = \bar{\mathbf{n}}(0) + \mathbf{B}(\alpha) \mathbf{M}(\alpha) = \bar{\mathbf{x}}(0) + \mathbf{B}(\alpha) \mathbf{M}(\alpha), \quad (28)$$

where the matrix

$$\mathbf{B}(\alpha) = (-\mathbf{I} + \mathbf{R}^T \Phi_1) \boldsymbol{\mu}_1 + \sum_{k=2}^{K_{\max}} (\mathbf{R}^T \Phi_k \boldsymbol{\mu}_k + \alpha \mathbf{I}) \prod_{j=k \dots 2} (\boldsymbol{\mu}_j + \alpha \mathbf{I})^{-1} (\mathbf{I} - \Phi_{j-1}) \boldsymbol{\mu}_{j-1} \quad (29)$$

only depends on the input parameters of the model and satisfies the following statement.

Theorem 1 *The matrix $(\mathbf{B}(\alpha))^T$ is an infinitesimal generator.*

Proof 2 Let $\mathbf{1}^T$ be a row vector of ones. We begin by proving that $(\mathbf{R}^T \Phi_{K_{\max}} \boldsymbol{\mu}_{K_{\max}} + \alpha \mathbf{I})(\alpha \mathbf{I} + \boldsymbol{\mu}_{K_{\max}})^{-1}$ is a column-stochastic matrix. Note that the entries of this matrix are non-negative being all matrices non-negative and the inverse applied to a diagonal matrix. Thus, we just need to show that columns sum to unity. Recall that $\Phi_{K_{\max}}$ is a binary matrix, since all Coxian service process must either complete at an earlier phase, in which case $\phi_{i,K_{\max}} = 0$ and $\mu_{i,k} = 0$, or complete at phase $K_{\max}$ implying $\phi_{i,K_{\max}} = 1$ and $\mu_{i,k} > 0$. This implies that $\Phi_{K_{\max}} \boldsymbol{\mu}_{K_{\max}} = \boldsymbol{\mu}_{K_{\max}}$. Therefore

$$\mathbf{1}^T (\mathbf{R}^T \Phi_{K_{\max}} \boldsymbol{\mu}_{K_{\max}} + \alpha \mathbf{I})(\boldsymbol{\mu}_{K_{\max}} + \alpha \mathbf{I})^{-1} = (\mathbf{1}^T \boldsymbol{\mu}_{K_{\max}} + \alpha \mathbf{1}^T)(\boldsymbol{\mu}_{K_{\max}} + \alpha \mathbf{I})^{-1} = \mathbf{1}^T$$

where we used that $\mathbf{1}^T \mathbf{R}^T = \mathbf{1}^T$ and the last passage holds since $\boldsymbol{\mu}_{K_{\max}}$ is diagonal.

We now focus on the summation in (29). Using the matrix recurrence

$$\mathbf{F}(j) = (\mathbf{F}(j+1) + (\mathbf{R}^T \Phi_j \boldsymbol{\mu}_j + \alpha \mathbf{I}))(\boldsymbol{\mu}_j + \alpha \mathbf{I})^{-1}(\mathbf{I} - \Phi_{j-1})\boldsymbol{\mu}_{j-1}$$

with $\mathbf{F}(K_{max}+1) = \mathbf{0}$, the summation in (29) can be computed as $\mathbf{F}(1)$. From the above definition, we have then

$$\begin{aligned} \mathbf{1}^T \mathbf{F}(K_{max}) &= \mathbf{1}^T (\mathbf{R}^T \Phi_{K_{max}} \boldsymbol{\mu}_{K_{max}} + \alpha \mathbf{I})(\boldsymbol{\mu}_{K_{max}} + \alpha \mathbf{I})^{-1}(\mathbf{I} - \Phi_{K_{max}-1})\boldsymbol{\mu}_{K_{max}-1} \\ &= \mathbf{1}^T (\mathbf{I} - \Phi_{K_{max}-1})\boldsymbol{\mu}_{K_{max}-1} \end{aligned}$$

that allows us to determine the next value in the recurrence as

$$\begin{aligned} \mathbf{1}^T \mathbf{F}(K_{max} - 1) &= \mathbf{1}^T ((\mathbf{I} - \Phi_{K_{max}-1})\boldsymbol{\mu}_{K_{max}-1} + (\mathbf{R}^T \Phi_{K_{max}-1} \boldsymbol{\mu}_{K_{max}-1} + \alpha \mathbf{I})) \times \\ &\quad \times (\boldsymbol{\mu}_{K_{max}-1} + \alpha \mathbf{I})^{-1}(\mathbf{I} - \Phi_{K_{max}-2})\boldsymbol{\mu}_{K_{max}-2} \\ &= \mathbf{1}^T (\mathbf{I} - \Phi_{K_{max}-2})\boldsymbol{\mu}_{K_{max}-2} \end{aligned}$$

and thus recursively we obtain $\mathbf{1}^T \mathbf{F}(1) = \mathbf{1}^T (\mathbf{I} - \Phi_1)\boldsymbol{\mu}_1$. This implies that

$$\begin{aligned} \mathbf{1}^T \mathbf{B}(\alpha) &= \mathbf{1}^T (-\mathbf{I} + \mathbf{R}^T \Phi_1)\boldsymbol{\mu}_1 + \mathbf{1}^T (\mathbf{I} - \Phi_1)\boldsymbol{\mu}_1 \\ &= -\mathbf{1}^T \boldsymbol{\mu}_1 + (\mathbf{1}^T \mathbf{R}^T \Phi_1 + (\mathbf{I} - \Phi_1))\boldsymbol{\mu}_1 \\ &= -\mathbf{1}^T \boldsymbol{\mu}_1 + \mathbf{1}^T \boldsymbol{\mu}_1 = \mathbf{0}. \end{aligned}$$

thus $\mathbf{B}(\alpha)$ has columns that sum to zero. The proof follows by noting that the only negative term in the definition of $\mathbf{B}(\alpha)$ is the diagonal matrix $-\boldsymbol{\mu}_1$ whereas all off-diagonal entries are non-negative by the structure of (29), where off-diagonal elements results from products of non-negative matrices. $\square$

Note that we cannot further simplify (28), since $\mathbf{M}(\alpha)$ is the Laplace transform of $\min(\bar{\mathbf{x}}(t), \bar{\mathbf{s}}_1(t))$, which cannot be derived analytically unless an explicit expression for the trajectories $\bar{\mathbf{x}}(t)$ and $\bar{\mathbf{s}}_1(t)$ is known. However, we cannot obtain an explicit solution for ODE systems including minimum functions except in special cases. This issue is common in other problem areas as well; for example, process algebras that use Kurtz's fluid all rely on numerical ODE solutions due to presence of minimum functions, e.g. PEPA [41], GPEPA [39].

Running Case. We wish to illustrate (28) under increasing scaling factors V . Assume first $V = 1$, $N = 2V = 2$, $S_2 = V = 1$. Using the results of the Running Case in Section 4, we have $\bar{\mathbf{n}}^2(0) = [1.0529, 0.9471]$. Let $\mathbf{M}^2(\alpha)$ and $\mathbf{B}^2(\alpha)$ be $\mathbf{M}(\alpha)$ and $\mathbf{B}(\alpha)$ conditional on the system being in stage 2. Then we can determine $\mathbf{M}^2(\alpha_2)$ by the definition after evaluating (22)-(24) with initial condition $\bar{\mathbf{n}}^2(0)$ and $s_{2,2}(0) = 0$. We evaluate the ODE system using MATLAB's `ode15s` with numerical tolerance set to 10^{-12} and find

$$\mathbf{M}^2(\alpha_{21}) = \begin{bmatrix} 0.2023 \\ 0.1430 \end{bmatrix}, \mathbf{B}^2(\alpha_{21}) = \begin{bmatrix} -1.0000 & 1.1250 \\ 1.0000 & -1.1250 \end{bmatrix}, \mathbf{B}^2(\alpha_{21})\mathbf{M}^2(\alpha_{21}) = \begin{bmatrix} -0.0414 \\ 0.0414 \end{bmatrix}$$

that implies $\bar{\mathbf{n}}^2(0) + \mathbf{B}^2(\alpha_2)\mathbf{M}^2(\alpha_2) = [1.0115, 0.9885]$. Using the ODE system solution we also find $\alpha_2\mathbf{N}^2(\alpha_2) = [1.0113, 0.9887]$ which equals $\bar{\mathbf{n}}^2(0) + \mathbf{B}^2(\alpha_2)\mathbf{M}^2(\alpha_2)$ except for minor numerical errors. This illustrates that (28) holds, but also that the ODE system is approximate, since $\alpha_2\mathbf{N}^2(\alpha_2) \neq \bar{\mathbf{n}}^1(0) = [0.9763, 1.0237]$, where the latter is the exact value computed from the CTMC.

Assume now $V = 10$, $N = 2V = 20$, $S_2 = V = 10$. Under this parameterisation the CTMC solution gives $\bar{\mathbf{n}}^2(0)/V = [1.1978, 0.8022]$. Solving the ODE system with this initial point and $s_{2,2}(0) = 0$ we find $\alpha_2\mathbf{N}^2(\alpha_2)/V = [1.1189, 0.8811]$, $\bar{\mathbf{n}}^2(0)/V + \mathbf{B}^2(\alpha_2)\mathbf{M}^2(\alpha_2)/V = [1.1185, 0.8815]$. The exact CTMC value is $\bar{\mathbf{n}}^1(0)/V = [1.1120, 0.8880]$ showing, compared to previous case, a progressive convergence of $\alpha_2\mathbf{N}^2(\alpha_2)/V$ to $\bar{\mathbf{n}}^2(0)/V$ the solutions as V grows.

6 Blending Algorithm

6.1 Blending Iteration

Recall that (28) has been obtained subject to the ODE evolving in a give stage h . This dependence can now be reinstated in the notation as

$$\alpha\mathbf{N}^h(\alpha) = \bar{\mathbf{n}}^h(0) + \mathbf{B}^h(\alpha)\mathbf{M}^h(\alpha), \quad (30)$$

where all the matrices defining $\mathbf{B}^h(\alpha)$ are instantiated with the stage h model parameters and indicated with $\mathbf{R}^h$, Ψ_k^h , μ_k^h , $k = 1, \dots, K_{max}$, $h = 1, \dots, E$. The above expression holds under the assumption that the time variable t in the Laplace transforms now represents the time since entering stage h . Under this assumption, (30) uses the fact that at time $t = 0$ all jobs are in phase 1 due to the reset rule, thus $\bar{\mathbf{s}}_{k,c}^h = \mathbf{0}$, $k \geq 2$, and we can use in the expression $\bar{\mathbf{n}}_c^h(0)$ since $\bar{\mathbf{n}}_c^h(0) = \bar{\mathbf{x}}_c^h(0)$.

Following the line of reasoning used in Section 4, we can formulate two equivalent expressions for the fixed point iteration (10). The first expression follows by combining (30) with Corollary 1 and gives the blending algorithm iteration needed to update the mean queue-lengths in each stage, as observed immediately after a stage change:

$$\bar{\mathbf{n}}_{c+1}^e(0) = \sum_{\substack{h=1 \\ h \neq e}}^E p_{he}^{src} \left(\bar{\mathbf{n}}_c^h(0) + \mathbf{B}^h(\alpha_h)\mathbf{M}_c^h(\alpha_h; \bar{\mathbf{n}}_c^E(0)) \right), \quad (31)$$

for $e = 1, \dots, E$, where $\mathbf{M}_c^h(\alpha_h; \bar{\mathbf{n}}_c^h(0)) = \mathcal{L} \left(\min(\bar{\mathbf{x}}_c^h(t), \mathbf{S} - \sum_{k=2}^{K_{max}^h} \bar{\mathbf{s}}_{k,c}^h(t)), \alpha_h \right)$,

in which we make explicit the dependence of the transform on the initial vector $\bar{\mathbf{n}}_c^h(0)$ used in the ODE system for stage h at iteration c .

Next, we develop an expression that consolidates all equations (31) into a single matrix formula. This second expression is useful for the convergence analysis developed in Section 7. Let $\bar{\mathbf{n}}_c(0) = (\bar{\mathbf{n}}_c^1(0), \dots, \bar{\mathbf{n}}_c^E(0))$, for all $c \geq 1$, be the vector of mean queue-lengths conditional on the current stage, such that $\|\bar{\mathbf{n}}_c(0)\|_1 = NE$. Also, let $\mathbf{P} = [p_{he}^{src} \mathbf{I}_{M \times M}]_{e,h=1,\dots,E}$ that is a row-stochastic matrix of order $E^2 M^2$ with zero diagonal.⁴ Expression (31) can now be rewritten as

$$\bar{\mathbf{n}}_{c+1}(0) = \mathbf{P}(\bar{\mathbf{n}}_c(0) + \mathbf{B}\mathbf{M}_c(\bar{\mathbf{n}}_c(0))), \quad (32)$$

for all iterations $c \geq 1$, being

$$\mathbf{B} = \text{diag}(\mathbf{B}^1(\alpha_1), \dots, \mathbf{B}^E(\alpha_E)) \quad (33)$$

$$\mathbf{M}_c(\bar{\mathbf{n}}_c(0)) = [\mathbf{M}_c^1(\alpha_1; \bar{\mathbf{n}}_c^1(0)); \dots; \mathbf{M}_c^E(\alpha_E; \bar{\mathbf{n}}_c^E(0))] \quad (34)$$

where $\mathbf{B}$ is block diagonal and $\mathbf{M}_c(\bar{\mathbf{n}}_c(0))$ is a column vector obtained by concatenation. Equations (31) and (32) are equivalent and define the fundamental equations of the fixed point iteration used by the blending algorithm to approximate queueing network models in random environments. An algorithmic description together with details on initial and termination conditions and implementation guidelines are provided in the next section.

Running Case. Assume $V = 1$, $N = 2V = 2$, $S_2 = V = 1$. We initialise the iteration using the exact values determined by the CTMC, i.e., $\bar{\mathbf{n}}_1(0) = [0.9763, 1.0237, 1.0529, 0.9471]^T$. Using the definitions we have

$$\mathbf{P} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} -1 & 2 & 0 & 0 \\ 1 & -2 & 0 & 0 \\ 0 & 0 & -1 & 1.1250 \\ 0 & 0 & 1 & -1.1250 \end{bmatrix}, \mathbf{M}_1(\bar{\mathbf{n}}_1(0)) = \begin{bmatrix} 0.2904 \\ 0.2096 \\ 0.2238 \\ 0.1285 \end{bmatrix}$$

thus $\bar{\mathbf{n}}_2(0) = \mathbf{P}(\bar{\mathbf{n}}_1(0) + \mathbf{B}\mathbf{M}_1(\bar{\mathbf{n}}_1(0))) = [1.0114, 0.9886, 1.1503, 0.8497]$. Reapplying the blending iteration (32) until $c = 20$ we find

$$\bar{\mathbf{n}}_{20}(0) = [1.1354, 0.8646, 1.2208, 0.7792].$$

If we chose a different initial condition, e.g., initialising all jobs at the delay server such that $\bar{\mathbf{n}}_1(0) = [2, 0, 2, 0]^T$, then at iteration 20 it is $\bar{\mathbf{n}}_{20}(0) =$

⁴ Recall that the probabilities p_{eh}^{src} are the probability of jumping to stage h after leaving stage e . Since we assume that the stage changes, it must be $p_{ee}^{src} = 0, \forall h$.

$[1.1357, 0.8643, 1.2210, 0.7790]$ and at iteration 30 it is

$$\bar{\mathbf{n}}_{30}(0) = [1.1354, 0.8646, 1.2208, 0.7792].$$

This illustrates that the iteration reaches the same fixed point regardless of the initial condition, but the latter affects convergence speed. The maximum absolute error of the final estimate compared to the exact value is 17.72%.

6.2 Algorithm and pseudo-code

We now combine the results in the previous sections to formalise the blending algorithm and provide implementation guidelines. A pseudo-code to guide the implementation of the method is given in Algorithm 1.

The blending algorithm consists in iteratively evaluating (32) until a user-specified number of iterations C is reached or until convergence of the estimated queue-lengths vectors $\bar{\mathbf{n}}_c(0)$ is achieved within a tolerance, e.g., $\|\bar{\mathbf{n}}_{c+1}(0) - \bar{\mathbf{n}}_c(0)\|_1 \leq \delta_{max}$ for some $c \leq C$. Iteration (32) requires to compute the $\mathbf{M}_c^h$ matrices by approximating the state trajectories $\bar{\mathbf{x}}^h(t)$ and $\bar{\mathbf{s}}_k^h(t)$, $k = 1, \dots, K_{max}$, $h = 1, \dots, E$, for all instants $t \in [0, T]$, where T is a user specified maximum time. Such time should be chosen such that the probability of a stage jump is negligible in any stage, e.g., $T = \max\{t : \max_{h,e} e^{\alpha_{he}t} \leq \epsilon\}$, where ϵ is a numerical tolerance (e.g., $\epsilon = 10^{-8}$). At each iteration, the knowledge of the $\bar{\mathbf{x}}^h(t)$ and $\bar{\mathbf{s}}_k^h(t)$ trajectories gives the ability of computing the mean queue-length trajectories by the usual expression $\mathbf{n}_c^h(t) = \bar{\mathbf{x}}_c^h(t) + \sum_{k=2}^{K_{max}} \bar{\mathbf{s}}_{k,c}^h(t)$, $0 \leq t \leq T$, for all iterations $c \geq 1$.

The blending algorithm aims at finding a fixed point for the trajectories $\bar{\mathbf{n}}_c^h(t)$, $\bar{\mathbf{x}}_{c,k}^h(t)$ and $\bar{\mathbf{s}}_{\infty,k}^h(t)$ over a set of $c \leq C$ iteration, where C is the maximum number of iterations allowed by the user. We have empirically noted that a fixed point exists on most instances, but a formal convergence analysis is carried out in Section 7 to better characterise the algorithm. From this fixed point, which we indicate with the notation $(\bar{\mathbf{n}}_\infty^h(t), \bar{\mathbf{x}}_\infty^h(t), \bar{\mathbf{s}}_\infty^h(t))$, performance measures can be easily estimated using the expressions we develop in Section 6.3.

An important implementation decision for the blending algorithm is the choice of the initial queue-lengths estimates for the M queues in each stage $h = 1, \dots, E$. Due to the reset rule, it is $\mathbf{s}_{j,1}(0) = \mathbf{S}^h$, $\mathbf{s}_{j,k}(0) = \mathbf{0}$, $k = 2, \dots, K_{max}$, and thus $\mathbf{n}^h(0) = \mathbf{x}^h(0)$ implies that it is sufficient to decide the values for the $\mathbf{x}^h(0)$ variables for initialisation of the ODEs. Let $\mathbf{x}^h(0) = \mathbf{x}_0^h$, then $\mathbf{x}_0^h$ may be chosen according to an heuristic, for example by placing all jobs at a certain station or by equally dividing them across all queues, but preference should be given to distributions that are expected to be close to the equilibrium distribution. For example, decomposition or average-environment based

on approximate mean value analysis results may be alternatively used as alternatives for initialisation, which however come at the price of complicating the implementation.

6.3 Mean Performance Measures

Assume that an exact solution for the model is available, then the mean queue-length at station j at equilibrium, denoted by Q_j , is readily computed as

$$Q_j = \sum_{h=1}^E \pi_h \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \bar{n}_j^h(t) dt = \sum_{h=1}^E \pi_h \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \left(\bar{x}_j^h(t) + \sum_{k=2}^{K_j^h} \bar{s}_{jk}^h(t) \right) dt, \quad (35)$$

which is simply the time average of $\bar{n}_j^h(t)$ in each stage, scaled by the stationary probabilities π_h of that stage. The mean throughput at station j , denoted by X_j , is obtained as

$$X_j = \sum_{h=1}^E \pi_h \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \left(\phi_{j,1}^h \mu_{j,1}^h \min(\bar{x}_j^h(t), \bar{s}_{j1}^h(t)) + \sum_{k=2}^{K_j^h} \phi_{j,k}^h \mu_{j,k}^h \bar{s}_{j,k}^h(t) \right) dt, \quad (36)$$

which weights by the probabilities π_h the average departure rates from each Coxian server. Finally, the mean response time W_j per visit at station j may be readily obtained by Little's law as

$$W_j = \frac{Q_j}{X_j} \quad (37)$$

In order to compute the above measures, the blending algorithm can either obtain a fixed point $(\bar{n}_\infty^h(t), \bar{x}_\infty^h(t), \bar{s}_{\infty,k}^h(t))$ or return a vector $(\bar{n}_C^h(t), \bar{x}_C^h(t), \bar{s}_{C,k}^h(t))$ in case the maximum number of iterations is reached before convergence. The entries of these vectors can be used to directly estimate (35) and (36) by replacing $\bar{x}_j^h(t)$ and $\bar{s}_{j,k}^h(t)$ with the corresponding estimates returned by the algorithm. In particular, the limits in the formulas can be estimated by time-averaging the occupancy measures in the station state vector up to a large enough T , i.e., the T_{max} parameter provided in input to Algorithm 1. Other measures of interest can be easily computed if they can be expressed as a function of the state trajectory.

An issue that may arise is the case where the solution of the model does not converge to a fixed point, but oscillations with some period p are observed in the mean queue-length vectors $\bar{n}_c^h(0)$ throughout the iterations. These cases may arise as a result of the fluid approximation, for example we describe in Section 7 a degenerate case of balanced network showing this behaviour

in heavy-load. For such models, the last p vectors $\bar{\mathbf{n}}_c^h(0)$ should be averaged to obtain the stationary value of the queue lengths seen at stage transition instants. Such pathological cases may be detected by heuristic tests, e.g., the MSER-5 rule [26].

7 Convergence Analysis

Evaluating the convergence of (32) is challenging due to the intractability of the Laplace transform of the minimum function that appears in the definition of $\mathbf{M}_c(\bar{\mathbf{n}}_c(0))$. Bounds on the transforms in the expressions are possible, e.g., $\mathcal{L}(\min(\bar{\mathbf{x}}(t), \mathbf{S}); \alpha) \leq \min(\mathbf{X}(\alpha), \alpha^{-1}\mathbf{S})$, but in our experience they yield after a few iterations trivial values because they violate the conservation of the fluid mass of the ODE system.

To cope with this intractability, we here focus on two special cases where (32) can be evaluated analytically, namely light-load and heavy-load regimes. The following subsections provide detailed convergence analysis results. We here anticipate the main qualitative findings. The light-load and heavy-load analyses suggest that the blending algorithm provides stronger convergence guarantees in light-load, although the heavy-load case is only partly tractable. Our analysis suggests that the algorithm features in both regimes a structure similar to the power method iteration. The power method is an eigenvalue algorithm applied in various domains, including Markov chain analysis, and it is known to converge on most instances except for a few pathological cases, e.g., cyclic matrices. It is possible to construct toy examples to show that such cases can arise also in the application of the blending algorithm, but these appear contrived examples. Indeed, the numerical experiments reported in Section 8 suggest that the algorithm is quite accurate on average. For pathological cases where the solution oscillates, averaging performance measures over the last few iterations provides a simple way to cope with the lack of a fixed point.

7.1 Convergence analysis in light-load

The light-load convergence analysis assumes that in each stage h the number of jobs at each queues does not exceed its available servers at any point in time. This is a good approximation for a lightly-loaded system where all stations have no significant backlog of waiting jobs and therefore there is a negligible probability that all servers in the queue will become busy. Formally, the assumption is given by

$$\mathcal{L}(\min(\bar{\mathbf{n}}_c^h(t), \mathbf{S}); \alpha_h) = \mathcal{L}(\bar{\mathbf{n}}_c^h(t); \alpha_h) \quad (38)$$

where it should be noted that the minimum in the Laplace transform differs from the one used in the ODE systems discussed in the previous sections, which consider $\bar{\mathbf{x}}_c^h(t)$ instead of $\bar{\mathbf{n}}_c^h(t)$. However, observing that

$$\mathcal{L}(\bar{\mathbf{n}}_c^h(t); \alpha_h) = \mathbf{X}_{c,k}^h(\alpha_h) + \sum_{k=2}^{K_{max}} \mathbf{S}_{c,k}^h(\alpha_h)$$

where $\mathbf{X}_c^h(\alpha) = \mathcal{L}(\bar{\mathbf{x}}_c^h(t); \alpha)$, and that

$$\begin{aligned} \mathcal{L}(\min(\bar{\mathbf{n}}_c^h(t), \mathbf{S}); \alpha_h) &= \mathcal{L}(\min(\bar{\mathbf{n}}_c^h(t) - \sum_{k=2}^{K_{max}} \mathbf{s}_{c,k}^h(t), \mathbf{S} - \sum_{k=2}^{K_{max}} \mathbf{s}_{c,k}^h(t)); \alpha_h) + \sum_{k=2}^{K_{max}} \mathbf{S}_{c,k}^h(\alpha_h) \\ &= \mathcal{L}(\min(\bar{\mathbf{x}}_c^h(t), \mathbf{S} - \sum_{k=2}^{K_{max}} \mathbf{s}_{c,k}^h(t)); \alpha_h) + \sum_{k=2}^{K_{max}} \mathbf{S}_{c,k}^h(\alpha_h) \\ &= \mathbf{M}_{c,k}^h(\alpha_h) + \sum_{k=2}^{K_{max}} \mathbf{S}_{c,k}^h(\alpha_h) \end{aligned}$$

we see that the light-load assumption implies $\mathbf{M}_{c,k}^h(\alpha_h) = \mathbf{X}_{c,k}^h(\alpha_h)$. Inserting the expression in (32) and simplifying, the blending iteration becomes

$$\bar{\mathbf{n}}_{c+1}(0) = \mathbf{P}(\bar{\mathbf{n}}_c(0) + \mathbf{B}\mathbf{X}_c), \quad (39)$$

where $\mathbf{X}_c = \text{diag}(\mathbf{X}_c^1(\alpha_1), \dots, \mathbf{X}_c^E(\alpha_E))$. We can further simplify the expressions as follows. Using $\mathbf{M}_{c,k}^h(\alpha_h) = \mathbf{X}_{c,k}^h(\alpha_h)$, it follows that (25)-(27) becomes a system of linear ODEs with transforms satisfying

$$\alpha_h \begin{bmatrix} \mathbf{X}^h(\alpha_h) \\ \mathbf{S}_2^h(\alpha_h) \\ \vdots \\ \mathbf{S}_{K_{max}}^h(\alpha_h) \end{bmatrix} = \mathbf{V}^h \begin{bmatrix} \mathbf{X}^h(\alpha_h) \\ \mathbf{S}_2^h(\alpha_h) \\ \vdots \\ \mathbf{S}_{K_{max}}^h(\alpha_h) \end{bmatrix} + \begin{bmatrix} \bar{\mathbf{x}}^h(0) \\ \bar{\mathbf{s}}_2^h(0) \\ \vdots \\ \bar{\mathbf{s}}_{K_{max}}^h(0) \end{bmatrix}.$$

where

$$\mathbf{V}^h = \begin{bmatrix} (-\mathbf{I} - \mathbf{R}^T \Phi_1) \boldsymbol{\mu}_1^h & \mathbf{R}_h^T \Phi_2^h \boldsymbol{\mu}_2^h & \mathbf{R}_h^T \Phi_3^h \boldsymbol{\mu}_3^h & \cdots & \mathbf{R}_h^T \boldsymbol{\mu}_{K_{max}}^h \\ (\mathbf{I} - \Phi_1^h) \boldsymbol{\mu}_1^h & -\boldsymbol{\mu}_2^h & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & (\mathbf{I} - \Phi_2^h) \boldsymbol{\mu}_2^h & -\boldsymbol{\mu}_3^h & \ddots & \mathbf{0} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & (\mathbf{I} - \Phi_{K_{max}-1}^h) \boldsymbol{\mu}_{K_{max}-1}^h & -\boldsymbol{\mu}_{K_{max}}^h \end{bmatrix}.$$

and its solution may be expressed in terms of Laplace transforms as

$$\begin{bmatrix} \mathbf{X}^h(\alpha_h) \\ \mathbf{S}_2^h(\alpha_h) \\ \vdots \\ \mathbf{S}_{K_{max}}^h(\alpha_h) \end{bmatrix} = (\alpha_h \mathbf{I} - \mathbf{V}^h)^{-1} \begin{bmatrix} \bar{\mathbf{x}}^h(0) \\ \bar{\mathbf{s}}_2^h(0) \\ \vdots \\ \bar{\mathbf{s}}_{K_{max}}^h(0) \end{bmatrix}.$$

Recalling that due the reset rule $\bar{\mathbf{s}}_2^h(0) = \dots = \bar{\mathbf{s}}_{K_{\max}}^h(0) = \mathbf{0}$, the above expression implies that $\mathbf{X}^h(\alpha_h) = \mathbf{C}^h \bar{\mathbf{x}}^h(0)$ and $\mathbf{C}^h = [(\alpha_h \mathbf{I} - \mathbf{V}^h)^{-1}]_{M \times M}$, where $[\cdot]_{I \times J}$ is the sub-matrix defined by the first I rows and J columns of its argument. Using the last expression and using again the reset rule $\bar{\mathbf{x}}_c(0) = \bar{\mathbf{n}}_c(0)$, (39) may be rewritten as

$$\bar{\mathbf{n}}_{c+1}(0) = \mathbf{P}(\mathbf{I} + \mathbf{BC}) \bar{\mathbf{n}}_c(0), \quad (40)$$

where $\mathbf{C} = \text{diag}(\mathbf{C}^1, \dots, \mathbf{C}^E)$.

Lemma 1 $\mathbf{P}(\mathbf{I} + \mathbf{BC})$ has the following block structure

$$\mathbf{P}(\mathbf{I} + \mathbf{BC}) = \begin{pmatrix} \mathbf{0} & p_{12}(\mathbf{I} + \mathbf{B}^2 \mathbf{C}^2) & \dots & p_{1E}(\mathbf{I} + \mathbf{B}^E \mathbf{C}^E) \\ p_{21}(\mathbf{I} + \mathbf{B}^1 \mathbf{C}^1) & \mathbf{0} & \dots & p_{2E}(\mathbf{I} + \mathbf{B}^E \mathbf{C}^E) \\ \vdots & \vdots & \ddots & \vdots \\ p_{E1}(\mathbf{I} + \mathbf{B}^1 \mathbf{C}^1) & p_{E2}(\mathbf{I} + \mathbf{B}^2 \mathbf{C}^2) & \dots & \mathbf{0} \end{pmatrix},$$

where the matrices $(\mathbf{I} + \mathbf{B}^h \mathbf{C}^h)$ are column stochastic.

Proof 3 The sparsity and block structure follow immediately from the definition of $\mathbf{P}$, since $p_{hh} = 0$, for all stages h . To prove that $(\mathbf{I} + \mathbf{B}^h \mathbf{C}^h)$ is column stochastic, let us observe that by the given definitions

$$\begin{aligned} \mathbf{B}^h &= (-\mathbf{I} + (\mathbf{R}^h)^T \Phi_1) \boldsymbol{\mu}_1^h + \sum_{k=2}^{K_{\max}} ((\mathbf{R}^h)^T \Phi_k \boldsymbol{\mu}_k^h + \alpha_h \mathbf{I}) \prod_{j=k \dots 2} (\alpha_h \mathbf{I} + \boldsymbol{\mu}_j^h)^{-1} (\mathbf{I} - \Phi_{j-1}^h) \boldsymbol{\mu}_{j-1}^h \\ \mathbf{C}^h &= (\alpha_h \mathbf{I} + \mathbf{D}^h - \mathbf{B}^h)^{-1} \end{aligned}$$

where $\mathbf{D}^h = \sum_{k=2}^{K_{\max}} \alpha_h \prod_{j=k \dots 2} (\alpha_h \mathbf{I} + \boldsymbol{\mu}_j^h)^{-1} (\mathbf{I} - \Phi_{j-1}^h) \boldsymbol{\mu}_{j-1}^h$ is a diagonal matrix.

Carrying out the products yields after some manipulations

$$\begin{aligned} (\mathbf{I} + \mathbf{B}^h \mathbf{C}^h) &= ((\mathbf{C}^h)^{-1} + \mathbf{B}^h) \mathbf{C}^h = (\alpha_h \mathbf{I} + \mathbf{D}^h) \mathbf{C}^h \\ &= (\alpha_h \mathbf{I} + \mathbf{D}^h) (\alpha_h \mathbf{I} + \mathbf{D}^h - \mathbf{B}^h)^{-1} = \mathbf{I} + \mathbf{B}^h (\alpha_h \mathbf{I} + \mathbf{D}^h - \mathbf{B}^h)^{-1} \end{aligned}$$

Let $\mathbf{1}^T$ be a row vector of ones, then we get

$$\mathbf{1}^T (\mathbf{I} + \mathbf{B}^h \mathbf{C}^h) = \mathbf{1}^T + \mathbf{1}^T \mathbf{B}^h (\alpha_h \mathbf{I} + \mathbf{D}^h - \mathbf{B}^h)^{-1} = \mathbf{1}^T$$

since $(\mathbf{B}^h)^T$ is an infinitesimal generator and thus $\mathbf{1}^T \mathbf{B}^h = \mathbf{0}^T$.

We are thus left to prove that all entries of $(\mathbf{I} + \mathbf{B}^h \mathbf{C}^h)$ are non-negative. We consider again the expression $(\mathbf{I} + \mathbf{B}^h \mathbf{C}^h) = (\alpha_h \mathbf{I} + \mathbf{D}^h) (\alpha_h \mathbf{I} + \mathbf{D}^h - \mathbf{B}^h)^{-1}$. Now observe that $\mathbf{D}^h$ is a diagonal non-negative matrix, thus $\alpha_h \mathbf{I} + \mathbf{D}^h - \mathbf{B}^h$ is

a M -matrix [34]. Since the inverse of a non-singular M -matrix is non-negative and $(\alpha_h \mathbf{I} + \mathbf{D}^h)$ is also non-negative, we conclude that all entries of $(\mathbf{I} + \mathbf{B}^h \mathbf{C}^h)$ are non-negative. $\square$

Theorem 2 *In light load, the blending algorithm has a unique fixed point equal to the dominant right eigenvector of $\mathbf{P}(\mathbf{I} + \mathbf{BC})$.*

Proof 4 Using Lemma 1, we see that $\mathbf{P}(\mathbf{I} + \mathbf{BC})$ is a non-negative matrix, being the product of a row-stochastic matrix with a column-stochastic matrix. We now wish to determine the spectral radius of $\mathbf{P}(\mathbf{I} + \mathbf{BC})$ that characterises the convergence of (40) and then characterise uniqueness using the Perros-Frobenius theorem. First observe that since (40) is an exact reformulation of the ODE system (19)-(21), it preserves the fluid mass of the $\bar{\mathbf{n}}_c(0)$ vectors for all iterations $c \geq 1$, and so their 1-norm. This implies that $\rho(\mathbf{P}(\mathbf{I} + \mathbf{BC})) \geq 1$ since a spectral radius less than unity would lead to $\lim_{c \rightarrow +\infty} \bar{\mathbf{n}}_c(0) = \mathbf{0}$, against the conservation of mass. Suppose now that $\rho(\mathbf{P}(\mathbf{I} + \mathbf{BC})) > 1$, then $\|(\mathbf{P}(\mathbf{I} + \mathbf{BC}))^k\|_p$ would not be bounded for increasing k values under any p -norm. However this is not possible due to the structure of the involved matrices. This is because, by Lemma 1, powers of $\mathbf{P}(\mathbf{I} + \mathbf{BC})$ are composed by blocks that are convex combinations, through the coefficients $p_{h1}, \dots, p_{hE}$ of products of column-stochastic matrices. Since products of column-stochastic matrices also yield column stochastic matrices, we conclude that column entries in any power of $\mathbf{P}(\mathbf{I} + \mathbf{BC})$ will be upper bounded by the numbers of blocks in each column, i.e., E . Thus, $\|(\mathbf{P}(\mathbf{I} + \mathbf{BC}))^k\|_1 \leq E, \forall k$, against the assumption that $\rho(\mathbf{P}(\mathbf{I} + \mathbf{BC})) > 1$. This proves that $\rho(\mathbf{P}(\mathbf{I} + \mathbf{BC})) = 1$.

Assuming that $\mathbf{P}$ and the routing probability matrices $\mathbf{R}$ are irreducible, then $\mathbf{P}(\mathbf{I} + \mathbf{BC})$ is also irreducible and by the Perron-Frobenius theorem its dominant eigenvalue is 1, being equal to the spectral radius, it is unique, and the associated unique right eigenvector has strictly positive entries. $\square$

Expression (40) can be recognised as the power method iteration used for stationary analysis of Markov chains, but applied to the matrix $\mathbf{P}(\mathbf{I} + \mathbf{BC})$. Note that the power method is guaranteed to converge, although its convergence rate can be slower than other eigenvalue algorithms [40].

7.2 Convergence analysis in heavy-load

In the heavy-load regime we assume that

$$\mathcal{L}(\min(\bar{\mathbf{n}}_c^h(t), \mathbf{S}^h); \alpha) = \mathcal{L}(\mathbf{S}^h; \alpha) = \alpha^{-1} \mathbf{S}^h.$$

This assumption models an heavy-loaded regime where all servers are busy at all stations. Using the heavy-load assumption in (32) and considering the transform of (18), it is found that

$$\bar{\mathbf{n}}_{c+1}(0) = \mathbf{P}\bar{\mathbf{n}}_c(0) + \mathbf{P}\mathbf{B}\mathbf{S}_{1,c}, \quad (41)$$

where $\mathbf{S}_{1,c} = (\mathbf{S}_{1,c}^1(\alpha_1), \dots, \mathbf{S}_{1,c}^E(\alpha_E))$ is the Laplace transform vector for the $\bar{\mathbf{s}}_{1,c}^h(t)$ trajectories. We now make the observation that, under the reset rule, since all servers are continuously utilised, the trajectories $s_{1,c}(t)$ are independent of $\bar{\mathbf{n}}_c(0)$ and identical at each iteration. Thus

$$\bar{\mathbf{n}}_{c+1}(0) = \mathbf{P}\bar{\mathbf{n}}_c(0) + \mathbf{P}\mathbf{B}\mathbf{S}_{1,1}, \quad (42)$$

This is a first-order matrix difference equation with solution

$$\bar{\mathbf{n}}_c(0) = \mathbf{P}^c \bar{\mathbf{n}}_1(0) + \sum_{j=1}^c \mathbf{P}^j \mathbf{B}\mathbf{S}_{1,1}, \quad (43)$$

Since $\mathbf{P}$ is a stochastic matrix, different cases are possible depending on $\mathbf{P}$ and the choice of the initial point.

As an example of instance where blending does not converge in heavy-load, consider a balanced network, for example in a cyclic network of exponential stations having the same rates in a stage, but which have different rates in each stage. We assume an identical number of servers S at each stations, thus $\mathbf{S}_{1,1} = S\mathbf{1}$. In these models, it is possible to verify from the definitions that $\mathbf{B}$ is doubly-stochastic and thus $\mathbf{B}\mathbf{S}_{1,1} = S\mathbf{B}\mathbf{1} = \mathbf{0}$. This implies that the heavy-load iteration has a simplified structure

$$\bar{\mathbf{n}}_{c+1}(0) = \mathbf{P}\bar{\mathbf{n}}_c(0)$$

with initial vector $\bar{\mathbf{n}}_1(0)$. It is well-known from the Markov chain literature that an iteration of this kind may not convergence, for example in cases where $\mathbf{P}$ is cyclic. We tested a cyclic $\mathbf{P}$ and $\bar{\mathbf{n}}_1^1(0) \neq \bar{\mathbf{n}}_1^2(0) \neq \dots \neq \bar{\mathbf{n}}_1^E(0)$ confirming that the solution oscillates. Conversely, if the same model is initialised with the invariant eigenvector, i.e., $\bar{\mathbf{n}}_1^1(0) = \bar{\mathbf{n}}_1^2(0) = \dots = \bar{\mathbf{n}}_1^E(0)$, the algorithm converges remaining at this fixed point. This suggests that on some instances the initial points may affect the convergence of the blending algorithm. Still, the procedure we have outlined in Section 6.3 for calculation of performance measures is robust, as it guarantees that blending returns an estimate also on instances where the iteration does not converge.

$M = 2, E = 2, S = 1$				$M = 2, E = 2, S = 5$		
<i>Dist.</i>	<i>Error</i>	<i>Runtime [s]</i>		<i>Error</i>	<i>Runtime [s]</i>	
		<i>SIM</i>	<i>BLE</i>		<i>SIM</i>	<i>BLE</i>
Erl	0.062	806	75	0.062	366	47
Exp	0.078	99	6	0.020	27	6
Cox	0.075	398	11	0.024	204	10

Table 2

Validation results: networks with 2 stations and 2 stages.

$M = 2, E = 10, S = 1$				$M = 2, E = 10, S = 5$		
<i>Dist.</i>	<i>Error</i>	<i>Runtime [s]</i>		<i>Error</i>	<i>Runtime [s]</i>	
		<i>SIM</i>	<i>BLE</i>		<i>SIM</i>	<i>BLE</i>
Erl	0.208	6188	15	0.127	5322	18
Exp	0.152	609	4	0.058	166	5
Cox	0.139	1150	19	0.039	1284	23

Table 3

Validation results: networks with 2 stations and 10 stages.

$M = 10, E = 2, S = 1$				$M = 10, E = 2, S = 5$		
<i>Dist.</i>	<i>Error</i>	<i>Runtime [s]</i>		<i>Error</i>	<i>Runtime [s]</i>	
		<i>SIM</i>	<i>BLE</i>		<i>SIM</i>	<i>BLE</i>
Erl	0.105	64923	17359	0.024	29284	8673
Exp	0.116	5923	1555	0.025	1873	434
Cox	0.158	18238	4127	0.028	14504	1141

Table 4

Validation results: networks with 10 stations and 2 stages.

8 Numerical Validation

To study the error behaviour of the blending algorithm, we consider 6 groups of networks with randomly generated parameters. The model parameters are drawn from uniform distributions with rates in $[0.0, 50.0]$ for the random environment. Within a group, we keep fixed the number of queues (M), the number of stages in the random environment (E), and the number of servers

in each queue (S). The population is equal to $N = 50$ jobs. Service time distributions are varied in a controlled form. For each such network, we consider 3 possible service distributions: exponential, Erlang-3 (squared coefficient of variation $c^2 = 1/3$), and two-stage Coxian distribution with balanced phase probabilities ($c^2 = 10$). The service distribution has mean chosen at random uniformly in $[0.001, 10.0]$.

For each group we have analysed 300 models, thus 1800 models were considered overall. Each model was analysed by simulation, stopped after the 95% confidence intervals were within 5% of the mean. Both simulation and the blending algorithm were implemented in MATLAB, in order to use the same environment to compare their runtimes. The implementation of blending alternates execution of a non-stiff solver (`ode45`) and a stiff solver (`ode15s`) until both methods cannot further improve the estimate of the fixed point. We use an iteration tolerance $\delta_{max} = 0.01$. The maximum number of iterations is set to $C = 100$. To speed-up convergence, we use a decomposition approximation to determine the initial vectors $\bar{\mathbf{n}}_1^h(0)$. Such decomposition considers each stage h in isolation and returns the value of $\bar{\mathbf{n}}^h(t)$ for large t obtained by the fluid ODEs. This allows the initial point to accumulate mass on the bottleneck station for each stage, often saving some iterations to the blending algorithm. The time to generate these estimates is included in the figures shown in the tables.

The results are shown in Tables 2-4. For a model with a population of N jobs, the accuracy error is expressed as in [12], which normalises the absolute error with respect to the total job population as follows:

$$Error = \max_{1 \leq i \leq N} \frac{|Q_i^{\text{BLE}} - Q_i^{\text{SIM}}|}{N},$$

Here Q_i^{BLE} and Q_i^{SIM} are the estimates of the mean queue length at station i computed by blending and by simulation, respectively. Each table shows the average errors and execution times for a set of 100 randomly generated networks, with controlled service time distributions and server multiplicities. The computational cost of the study is approximately equal to 1 month on two 16-core Xeon E5540 2.53Ghz, most of which spent for the simulations.

The following observations arise from the results:

- The mean accuracy of blending is adequate in all cases, with the error decreasing with bigger network topologies and more stages of the random environment.
- Blending is generally faster than simulation. In the networks with two stations, the difference in runtimes is two order of magnitudes on average, but it tends to decrease with larger network topologies. In the groups of networks with 10 stations (see Table 4), an analysis of the behaviour of

blending reveals that increased time of the analysis is mainly caused by a larger runtime per single iteration; this is due to the fact that the system of ODEs increases linearly with the number of stations and that larger problems may lead to stiffness. The number of iterations, not reported in the table, are found to be quite similar in all groups.

- By comparing the errors across the same row in every table, we find that increasing the number of servers at each station, while keeping all the other parameters unchanged, leads to an appreciable increase in the quality of the approximation. This is important for practical purposes, since the modelling of modern multi-core servers often requires to consider queues with multiple servers.
- The service-time distributions have an impact on the accuracy error, but this is model-dependent and difficult to interpret. For instance, the instances with Erlang-distributed service times incur the largest errors in Table 3, but they display the smallest errors in the groups of networks of Table 4. This may indicate that, overall, the choice of service distribution does not determine approximation accuracy as other parameters such as the number of servers.

Summarising, our study reveals that the blending algorithm is reasonably accurate in the majority of instances. Quite often, the largest errors are associated to stiff models or early termination ($c > C = 100$). The algorithm is faster than simulations run at 95% confidence intervals, an useful property for optimisation studies which may require the evaluation of thousands of models.

9 Related Work

With the exception of average-environment and decomposition, we are not aware of general analytical approximation techniques for closed queueing networks in random environments. The authors of [37] study the application of transient analysis for studying the behaviour of a queueing system in a Markov-modulated environment. Based on an analytical expression for the transient probability distribution of an $M/M/1$ queue, the authors derive numerical approximations for the $MMPP/M/1$ queue at equilibrium. Results are found to be accurate. Our approach is different in three ways. First, we focus on random environments independent of the queueing model state. Next, we provide a general method for a class of closed queueing networks. Finally, our method is technically different because the blending algorithm uses a fluid limit approximation.

Related to this paper are also work on queueing under breakdown and repair [11,31] and recent work on queues with fluctuating loads [23]. Such works feature random environments similar to the ones illustrated in Section 2 and

show commonalities with our approach. For example, [23] uses Laplace transforms to study the state trajectory. [18, 19] develop semi-numerical methods to evaluate the Laplace transform in fluid queues with semi-Markovian inputs. [24] applies fluid analysis to open queues and outlines a moment-matching methodology for tandem networks. Compared to the present work, these papers focus on open models, whereas the focus of this paper is on closed queueing networks.

Fluid limits have been extensively studied for both discrete-time and continuous-time Markov chains (e.g., [5] and [28], respectively). A wealth of applications to specific case studies can be found in the literature; for instance to MAC protocols [8, 38], TCP protocols [1, 2], peer-to-peer networks [36], caching algorithms [27] and load balancing strategies [21, 22].

Important requirements for the existence of a deterministic limit are conditions on the nature of the transition rates. In some specific cases of stochastic processes, such as the queueing networks considered in this paper, it is possible to define a general framework of convergence which holds for a large class of models; for instance, this is due to the fact that the transition rates give rise to a vector field for the ODE model which is Lipschitz continuous globally. This result has been exploited for other classes of stochastic models, for instance those that are induced by high-level modelling languages such as process algebra [25, 41].

In all these cases, all the transition rates enjoy properties that are analogous to the density-dependent form discussed in Section 5. In our model, instead, the transitions due to the random environment do not scale with increasing population levels, but are fixed. Although approaches for multi-scale stochastic processes are available [4], they cannot be applied to our model because they require a clear separation between *fast* and *slow* processes. In general, however our model of random environment is quite general and allows transitions between stages to be in the same time scale as the network's service rates.

A suitable mathematical framework for defining our model could be that of Piece-wise Deterministic Markov Processes (PDMPs) [15], where continuous flows defined as systems of ODEs are *reset* by stochastic jumps distributed according to exponential distributions. Therefore, one would naturally associate changes in the random environment with such jumps, and the queueing network dynamics within one stage as a continuous flow by the fluid model. It would be in principle possible to show that, with such association, a suitable family of queueing networks with random environment converges to a PDMP [9]. The dynamics of a PDMP is characterised by a partial differential equation defining the joint probability distribution of the process being in one (discrete) stage and the continuous flow being less than a certain level. Un-

like our blending algorithm, such a model is very difficult to solve analytically in general, but can be simulated. An exception is [13], where the authors were able to provide an elegant solution by exploiting the decoupling between the stochastic and the deterministic dynamics.

10 Conclusion

Blending is a new technique for the analysis of closed queueing networks in Markovian random environments. The approach iteratively evaluates the mean number of jobs at each queue observed at the instants when the active stage changes. It does so by applying a fluid approximation to the queueing network model, conditioning its evolution in a given environment stage. This information is used to approximate the equilibrium behaviour of the system at steady state. We have used a large evaluation study to prove that the algorithm has good scalability and accuracy.

Several possible extensions of this work may be investigated, in particular multi-class closed networks and open models. While an extension to multi-class models appears possible, possibly under some more restrictive conditions for the scheduling disciplines, it is unclear at present if the blending algorithm could also be applied to open models. This is because the convergence analysis reveals a dependence on the property of conservation of the fluid mass, which exists because of the closed nature of the system.

Another direction for future work consists in extending the approach to systems without the reset rule introduced in Section 3. For example, it would be valuable to investigate models where the number of phases in a service process is independent of the current stage or reset rules that instantaneously move jobs to other stations.

Acknowledgement

The work of Giuliano Casale and Peter Harrison is supported by the European project MODAClouds (FP7-318484). Mirco Tribastone is supported by the European project ASCENS. Peter Harrison is also supported by the EPSRC project APROPOS (EP/I030921/1). The authors thank Juan Fernando Pérez for helpful comments during the revision of this work.

References

- [1] M. Alizadeh, A. Javanmard, and B. Prabhakar. Analysis of DCTCP: stability, convergence, and fairness. In *Proceedings of ACM SIGMETRICS 2011*, pages 73–84, 2011.
- [2] M. Alizadeh, A. Kabbani, B. Atikoglu, and B. Prabhakar. Stability analysis of qcn: the averaging principle. In *Proceedings of ACM SIGMETRICS 2011*, pages 49–60, 2011.
- [3] V. V. Anisimov. *Switching Processes in Queueing Models*. Wiley, 2008.
- [4] K. Ball, T. Kurtz, L. Popovic, and G. Rempala. Asymptotic analysis of multiscale approximations to reaction networks. *Ann. Appl. Probab.*, 14(4):1925–1961, 2006.
- [5] M. Benaïm and J.-Y. Le Boudec. A class of mean field interaction models for computer and communication systems. *Performance Evaluation*, 65(11-12):823–838, 2008.
- [6] P. Billingsley. *Probability and Measure*. Wiley, 3rd edition, 1995.
- [7] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi. *Queueing Networks and Markov Chains*. 2nd ed., John Wiley and Sons, 2006.
- [8] C. Bordenave, D. McDonald, and A. Proutière. A particle system in interaction with a rapidly varying environment: Mean field limits and applications. *NHM*, 5(1):31–62, 2010.
- [9] L. Bortolussi. Limit behavior of the hybrid approximation of stochastic process algebras. In *Proceedings of 17th International Conference on Analytical and Stochastic Modeling Techniques and Applications, ASMTA 2010*, volume 6148 of *Lecture Notes in Computer Science*, pages 367–381. Springer, 2010.
- [10] A. Chaintreau, J.-Y. L. Boudec, and N. Ristanovic. The age of gossip: spatial mean field regime. In *SIGMETRICS/Performance*, pages 109–120, 2009.
- [11] R. Chakka and I. Mitrani. Heterogeneous multiprocessor systems with breakdowns: performance and optimal repair strategies. *Theoretical Computer Science*, 125(1):91–109, 14 Mar. 1994.
- [12] K. M. Chandy and D. Neuse. Linearizer: A heuristic algorithm for queueing network models of computing systems. *Commun. ACM*, 25(2):126–134, 1982.
- [13] F. Clévenot and P. Nain. A simple model for the analysis of squirrel. In *Proceedings of INFOCOM 2004*, 2004.
- [14] P. Courtois. Decomposability, instabilities, and saturation in multiprogramming systems. *Comm. of the ACM*, 18(7):371–377, 1975.
- [15] M. Davis. *Markov Models and Optimization*. Chapman & Hall, 1993.

- [16] H. De Sterck, T. A. Manteuffel, S. F. McCormick, K. Miller, J. Pearson, J. Ruge, and G. Sanders. Smoothed aggregation multigrid for Markov chains. *SIAM Journal on Scientific Computing*, 32(1):40–61, 2010.
- [17] A. Economou. Generalized product-form stationary distributions for markov chains in random environments with queueing applications. *Advances in Applied Probability*, 37(1):pp. 185–211, 2005.
- [18] A. J. Field and P. G. Harrison. An approximate compositional approach to the analysis of fluid queue networks. *Perform. Eval.*, 64(9-12):1137–1152, 2007.
- [19] A. J. Field and P. G. Harrison. Busy periods in fluid queues with multiple emptying input states. *Journal of Applied Probability*, 47(2):474–497, 06 2010.
- [20] J. M. Fourneau, B. Plateau, and W. Stewart. Product form for stochastic automata networks. In *Proc. of ValueTools*, pages 1–10, 2007.
- [21] A. Ganesh, S. Lilienthal, D. Manjunath, A. Proutiere, and F. Simatos. Load balancing via random local search in closed and open systems. In *Proceedings of ACM SIGMETRICS 2010*, pages 287–298, 2010.
- [22] N. Gast and B. Gaujal. A Mean Field Model of Work Stealing in Large-Scale Systems. *ACM SIGMETRICS*, 2010.
- [23] V. Gupta, M. Harchol-Balter, A. Scheller-Wolf, and U. Yechiali. Fundamental characteristics of queues with fluctuating load. In *SIGMETRICS/Performance*, pages 203–215, 2006.
- [24] V. Gupta and P. G. Harrison. Fluid level in a reservoir with an on-off source. *SIGMETRICS Performance Evaluation Review*, 36(2):128–130, 2008.
- [25] R. Hayden and J. T. Bradley. A fluid analysis framework for a Markovian process algebra. *Theoretical Computer Science*, 2010.
- [26] P. Heidelberger and P. D. Welch. A spectral method for confidence interval generation and run length control in simulations. *Comm. of the ACM*, 24(4):233–245, 1981.
- [27] S. Ioannidis, L. Massoulié, and A. Chaintreau. Distributed caching over heterogeneous mobile networks. *Queueing Syst.*, 72(3-4):279–309, 2012.
- [28] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure Markov processes. *J. Appl. Prob.*, 7(1):49–58, April 1970.
- [29] T. G. Kurtz. The relationship between stochastic and deterministic models for chemical reactions. *The Journal of Chemical Physics*, 57(7):2976–2978, 1972.
- [30] I. Mitrani. *Probabilistic Modelling*. Cambridge University Press, 1998.
- [31] I. Mitrani and A. A. Puhalskii. Limiting results for multiprocessor systems with breakdowns and repairs. *Queueing Syst*, 14(3-4):293–311, 1993.
- [32] M. Mitzenmacher. *The power of two choices in randomized load balancing*. PhD thesis, University of California at Berkeley, Department of Computer Science, Berkeley, CA, 1996.

- [33] M. F. Neuts. *Matrix-Geometric Solutions in Stochastic Models, an Algorithmic Approach*. The Johns Hopkins University Press, Baltimore, MD, 1981.
- [34] R. J. Plemmons. M -matrix characterizations. I—nonsingular M -matrices. *Linear Algebra and Its Applications*, 18:175–188, 1977.
- [35] P. Pollet. On a model for interference between searching insect parasites. *J. Austral. Math. Soc. Ser. B*, 32(2):133–150, 1990.
- [36] D. Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *Proceedings of ACM SIGCOMM 2004*, pages 367–378, 2004.
- [37] P. Romano, B. Ciciani, A. Santoro, and F. Quaglia. Fast computation of hyper-exponential approximations of the response time distribution of MMPP/M/1 queues. In *Annual Simulation Symposium*, pages 113–120. IEEE Computer Society, 2008.
- [38] G. Sharma, A. Ganesh, and P. Key. Performance analysis of contention based medium access control protocols. *IEEE Transactions on Information Theory*, 55(4):1665–1682, 2009.
- [39] A. Stefanek, R. A. Hayden, and J. T. Bradley. GPA - A tool for fluid scalability analysis of massively parallel systems. In *QEST*, pages 147–148. IEEE Computer Society, 2011.
- [40] W. J. Stewart. *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, Princeton, NJ, USA, 1994.
- [41] M. Tribastone, S. Gilmore, and J. Hillston. Scalable differential analysis of process algebra models. *IEEE Trans. Software Eng.*, 38(1):205–219, 2012.
- [42] B. Urgaonkar, G. Pacifici, P. J. Shenoy, M. Spreitzer, and A. N. Tantawi. An analytical model for multi-tier internet services and its applications. In *Proc. of ACM SIGMETRICS*, pages 291–302. ACM Press, 2005.
- [43] J. Zahorjan, E. D. Lazowska, and R. L. Garner. A decomposition approach to modelling high service time variability. *Performance Evaluation*, 3:35–54, 1983.
- [44] Y. Zhu. Markovian queueing networks in a random environment. *Operations Research Letters*, (1):11 – 17.

A Markovian Random Environments

In this appendix, we illustrate under what conditions the average-environment and decomposition approximations are appropriate for modelling systems evolving in random environments. The results in this section are not limited to queueing networks and hold for any discrete-state Markov model.

Similarly to Section 2, we focus on an environment composed of two stages with infinitesimal generator

$$\mathbf{E} = \begin{bmatrix} -\alpha_{12} & \alpha_{12} \\ \alpha_{21} & -\alpha_{21} \end{bmatrix},$$

We consider a system evolving in it with infinitesimal generation $\mathbf{Q}_1^*$ in stage 1 and $\mathbf{Q}_2^*$ in stage 2. Since the environment is unaffected by the system state, the Markov process for the entire model has generator

$$\mathbf{Q} = \left[\begin{array}{c|c} \mathbf{Q}_1^* - \alpha_{12}\mathbf{I} & \alpha_{12}\mathbf{I} \\ \hline \alpha_{21}\mathbf{I} & \mathbf{Q}_2^* - \alpha_{21}\mathbf{I} \end{array} \right], ,$$

where the diagonal blocks may have different orders. Let the equilibrium distribution of the network be the probability vector $\boldsymbol{\pi} = [\boldsymbol{\pi}_1, \boldsymbol{\pi}_2]$ such that $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$, $\boldsymbol{\pi}\mathbf{1} = 1$. Here, $\boldsymbol{\pi}_1$ represents the equilibrium distribution vector for the states in $\mathbf{Q}_1^*$ corresponding to environment stage 1. Similarly, $\boldsymbol{\pi}_2$ is the equilibrium distribution for $\mathbf{Q}_2^*$ for stage 2. We have the following necessary condition.

Proposition 2 *At equilibrium*

$$\boldsymbol{\pi}_1 \left(\left(\frac{\alpha_{21}}{\alpha_{12} + \alpha_{21}} \right) \mathbf{Q}_1^* + \left(\frac{\alpha_{12}}{\alpha_{12} + \alpha_{21}} \right) \mathbf{Q}_2^* - \frac{\mathbf{Q}_1^* \mathbf{Q}_2^*}{\alpha_{12} + \alpha_{21}} \right) = \mathbf{0}, \quad (\text{A.1})$$

$$\boldsymbol{\pi}_2 \left(\left(\frac{\alpha_{21}}{\alpha_{12} + \alpha_{21}} \right) \mathbf{Q}_1^* + \left(\frac{\alpha_{12}}{\alpha_{12} + \alpha_{21}} \right) \mathbf{Q}_2^* - \frac{\mathbf{Q}_2^* \mathbf{Q}_1^*}{\alpha_{12} + \alpha_{21}} \right) = \mathbf{0}. \quad (\text{A.2})$$

Proof 5 *Let us first study the evolution of the Markov process $\mathbf{Q}$ embedded at arrival instants of events that change the stage from 1 to 2. Define $\boldsymbol{\eta}_1^{(i)}$ to be the exit probability distribution vector from $\mathbf{Q}_1$ after the i th transition from stage 1 to 2. Then for the Markov random environment under study*

$$\boldsymbol{\eta}_1^{(i)} = \boldsymbol{\eta}_1^{(i-1)} \int_0^{+\infty} \int_0^{+\infty} \alpha_{12} e^{(\mathbf{Q}_2 - \alpha_{21}\mathbf{I})b_i} \alpha_{21} e^{(\mathbf{Q}_1 - \alpha_{12}\mathbf{I})r_i} db_i dr_i, ,$$

for $i \geq 1$, where b_i (resp. r_i) denotes the inter-arrival time of the j th breakdown (resp. repair) in the random environment. Using the fact that for a sub-generator A is

$$\int_0^{+\infty} e^{At} dt = (-A)^{-1}$$

we write

$$\boldsymbol{\eta}_1^{(i+1)} = \alpha_{12}\alpha_{21}\boldsymbol{\eta}_1^{(i)}(\mathbf{Q}_2 - \alpha_{21}\mathbf{I})^{-1}(\mathbf{Q}_1 - \alpha_{12}\mathbf{I})^{-1},$$

Observe now that a limiting vector $\boldsymbol{\eta}_1 = \lim_{i \rightarrow +\infty} \boldsymbol{\eta}_1^{(i)}$ exists if $\mathbf{Q}$ admits an equilibrium distribution $\boldsymbol{\pi}$, since $\boldsymbol{\eta}_1 = \boldsymbol{\pi}_1 \alpha_{12} / \alpha_{12}^{tot}$, being α_{12}^{tot} the aggregate

arrival rate of repair events. (Note that for the random environment here it is $\alpha_{12}^{tot} = \alpha_{12}$). Thus we have

$$\boldsymbol{\eta}_1 = \alpha_{12}\alpha_{21}\boldsymbol{\eta}_1(\mathbf{Q}_2 - \alpha_{21}\mathbf{I})^{-1}(\mathbf{Q}_1 - \alpha_{12}\mathbf{I})^{-1},$$

or equivalently since $\boldsymbol{\eta}_1 \propto \boldsymbol{\pi}_1$

$$\boldsymbol{\pi}_1(\alpha_{12}\alpha_{21}\mathbf{I} - (\mathbf{Q}_1 - \alpha_{12}\mathbf{I})(\mathbf{Q}_2 - \alpha_{21}\mathbf{I})) = 0,$$

Using a similar argument for the exit distribution $\boldsymbol{\eta}_2$ after a breakdown we write

$$\boldsymbol{\pi}_2(\alpha_{12}\alpha_{21}\mathbf{I} - (\mathbf{Q}_2 - \alpha_{21}\mathbf{I})(\mathbf{Q}_1 - \alpha_{12}\mathbf{I})) = 0.$$

The characterization (A.1)-(A.2) is then obtained with simple passages after expanding the products and dividing both sides by $\alpha_{12} + \alpha_{21}$. $\square$

This characterization highlights the relation between the exact solution $\boldsymbol{\pi} = [\boldsymbol{\pi}_1, \boldsymbol{\pi}_2]$ and the average-environment and decomposition approximations. Note first that $\mathbf{Q}_1^*$ and $\mathbf{Q}_2^*$ are the generators used by the decomposition approximation for the model, which has solution $\boldsymbol{\pi}^{DEC} = [\boldsymbol{\pi}_1^{DEC}, \boldsymbol{\pi}_2^{DEC}]$ where

$$\begin{aligned} \boldsymbol{\pi}_1^{DEC}\mathbf{Q}_1^* &= \mathbf{0}, & \boldsymbol{\pi}_1^{DEC}\mathbf{1} &= \pi_1 = \frac{\alpha_{21}}{\alpha_{12} + \alpha_{21}}, \\ \boldsymbol{\pi}_2^{DEC}\mathbf{Q}_2^* &= \mathbf{0}, & \boldsymbol{\pi}_2^{DEC}\mathbf{1} &= \pi_2 = \frac{\alpha_{12}}{\alpha_{12} + \alpha_{21}}. \end{aligned}$$

Observing that the coefficients of $\mathbf{Q}_1^*$ and $\mathbf{Q}_2^*$ in (A.1)-(A.2) are the equilibrium probabilities of the random environment for stages 1 and 2, we get

$$\boldsymbol{\pi}_1 \left(\mathbf{Q}^{AVG} - \frac{\mathbf{Q}_1^*\mathbf{Q}_2^*}{\nu} \right) = \mathbf{0} \quad (\text{A.3})$$

$$\boldsymbol{\pi}_2 \left(\mathbf{Q}^{AVG} - \frac{\mathbf{Q}_2^*\mathbf{Q}_1^*}{\nu} \right) = \mathbf{0} \quad (\text{A.4})$$

where $\nu = \alpha_{12} + \alpha_{21}$ and

$$\mathbf{Q}^{AVG} = \left(\frac{\alpha_{21}}{\alpha_{12} + \alpha_{21}} \right) \mathbf{Q}_1^* + \left(\frac{\alpha_{12}}{\alpha_{12} + \alpha_{21}} \right) \mathbf{Q}_2^*$$

is the AVG approximation for the model under study. Denote by $\boldsymbol{\pi}^{AVG}$, $\boldsymbol{\pi}_1^{AVG}\mathbf{1} = 1$, the equilibrium of $\mathbf{Q}^{AVG}$. Then, as the aggregate event rate ν decreases, the solution of the queueing network approaches that of the decomposition approximation up to scaling factors, i.e.,

$$\boldsymbol{\pi}_1 \lim_{\nu \rightarrow 0^+} \left(\mathbf{Q}^{AVG} - \frac{1}{\nu} \mathbf{Q}_1^*\mathbf{Q}_2^* \right) = \mathbf{0} \Rightarrow \boldsymbol{\pi}_1 \mathbf{Q}_1^*\mathbf{Q}_2^* = \mathbf{0} \Rightarrow \boldsymbol{\pi}_1 \propto \boldsymbol{\pi}_1^{DEC}$$

$$\boldsymbol{\pi}_2 \lim_{\nu \rightarrow 0^+} \left(\mathbf{Q}^{AVG} - \frac{1}{\nu} \mathbf{Q}_2^*\mathbf{Q}_1^* \right) = \mathbf{0} \Rightarrow \boldsymbol{\pi}_2 \mathbf{Q}_2^*\mathbf{Q}_1^* = \mathbf{0} \Rightarrow \boldsymbol{\pi}_2 \propto \boldsymbol{\pi}_2^{DEC}$$

Conversely, if ν becomes asymptotically large one finds that the AVG solution becomes exact:

$$\pi_1 \lim_{\nu \rightarrow 0^+} \left(Q^{AVG} - \frac{1}{\nu} Q_1^* Q_2^* \right) = \mathbf{0} \Rightarrow \pi_1 Q^{AVG} = \mathbf{0} \Rightarrow \pi_1 = \pi^{AVG}$$

$$\pi_2 \lim_{\nu \rightarrow 0^+} \left(Q^{AVG} - \frac{1}{\nu} Q_2^* Q_1^* \right) = \mathbf{0} \Rightarrow \pi_2 Q^{AVG} = \mathbf{0} \Rightarrow \pi_2 = \pi^{AVG}$$

As expected, these results indicate that the average-environment and decomposition approximations provide exact results for models in random environments where the frequency of events is respectively asymptotically large or asymptotically small compared to the rates of the queues. However, in intermediate cases, we have established by (A.1)–(A.2) that the probability distributions π_1 and π_2 depend also on matrices $Q_1^* Q_2^*$ and $Q_2^* Q_1^*$, which are not infinitesimal generators, having negative off-diagonal entries. This suggests that there may *not* exist a simple way to approximate well in all cases π_1 and π_2 based on weighting of the equilibrium solutions for Q^{AVG} , Q_1^* , or Q_2^* , a strategy that is common in previous work [14, 43].

Algorithm 1. BLENDING ALGORITHM

- 1: **input:**
 - $\mathbf{R}^h$: routing matrix for all stages $h = 1, \dots, E$
 - $\mathbf{S}^h$: number servers for each of the M queues in stage h
 - $\mathbf{x}_0^h$: initial queue-lengths estimates for the M queues in stage h
 - $\mathbf{E}$: infinitesimal generator of the random environment
 - $\boldsymbol{\mu}_k^h, k = 1, \dots, K_{max}; h = 1, \dots, E$
 - $\boldsymbol{\Phi}_k^h, k = 1, \dots, K_{max}; h = 1, \dots, E$
 - ϵ : numerical tolerance
 - T : maximum integration time
 - δ_{max} : convergence tolerance
 - C : maximum number of iterations
- 2: **output:** approximate $\bar{\mathbf{n}}^h(t), \bar{\mathbf{x}}^h(t), \bar{\mathbf{s}}_k^h(t), k = 1, \dots, K_{max}, h = 1, \dots, E, 0 \leq t \leq T$
- 3: **algorithm:**
- 4: $c = 1$
- 5: **for** $h = 1, \dots, E$ **do**
- 6: $\bar{\mathbf{n}}_c^h(0) = \bar{\mathbf{x}}_c^h(0) = \mathbf{x}_0^h,$
- 7: $\bar{\mathbf{s}}_{c,k}^h(0) = \mathbf{0}, k = 2, \dots, K_{max}$
- 8: $\mathbf{B}^h = (-\mathbf{I} + (\mathbf{R}^h)^T \boldsymbol{\Phi}_1) \boldsymbol{\mu}_1^h + \sum_{k=2}^{K_{max}} ((\mathbf{R}^h)^T \boldsymbol{\Phi}_k \boldsymbol{\mu}_k^h + \alpha_h \mathbf{I}) \prod_{j=k \dots 2} (\alpha_h \mathbf{I} + \boldsymbol{\mu}_j^h)^{-1} (\mathbf{I} - \boldsymbol{\Phi}_{j-1}^h) \boldsymbol{\mu}_{j-1}^h$
- 9: **end for**
- 10: compute $\boldsymbol{\pi}^{env} = [\pi_h^{env}]_{h=1, \dots, E}$ such that $\boldsymbol{\pi}^{env} \mathbf{E} = \boldsymbol{\pi}^{env}, \sum_{h=1}^E \pi_h^{env} = 1$
- 11: compute $p_{he}^{src} = \pi_h^{env} \alpha_{he} \left(\sum_{j \neq e} \pi_j^{env} \alpha_{je} \right)^{-1},$ for $h, e = 1, \dots, E; h \neq e$
- 12: **repeat**
- 13: $c = c + 1$
- 14: **for** $h = 1, \dots, E$ **do**
- 15: solve for all $t \leq T$, with numerical tolerance ϵ , the initial value problem
$$\begin{aligned} \frac{d\bar{\mathbf{x}}_c^h(t)}{dt} &= (-\mathbf{I} + \mathbf{R}^T \boldsymbol{\Phi}_1) \boldsymbol{\mu}_1 \min(\bar{\mathbf{x}}_c^h(t), \bar{\mathbf{s}}_{c,1}^h(t)) + \sum_{k=2}^{K_{max}} \mathbf{R}^T \boldsymbol{\Phi}_k \boldsymbol{\mu}_k \bar{\mathbf{s}}_{c,k}^h(t), \\ \frac{d\bar{\mathbf{s}}_{c,2}^h(t)}{dt} &= -\boldsymbol{\mu}_2 \bar{\mathbf{s}}_{c,2}^h(t) + (\mathbf{I} - \boldsymbol{\Phi}_1) \boldsymbol{\mu}_1 \min(\bar{\mathbf{x}}_c^h(t), \bar{\mathbf{s}}_{c,1}^h(t)), \\ \frac{d\bar{\mathbf{s}}_{c,k}^h(t)}{dt} &= -\boldsymbol{\mu}_k \bar{\mathbf{s}}_{c,k}^h(t) + (\mathbf{I} - \boldsymbol{\Phi}_{k-1}) \boldsymbol{\mu}_{k-1} \bar{\mathbf{s}}_{c,k-1}^h(t), \quad k = 3, \dots, K_{max} \end{aligned}$$

subject to

$$\begin{aligned} \bar{\mathbf{x}}_c^h(0) &= \bar{\mathbf{n}}_c^h(0) \\ \bar{\mathbf{s}}_{c,1}^h(0) &= \mathbf{S}^h, \\ \bar{\mathbf{s}}_{c,k}^h(0) &= \mathbf{0}, \quad k = 2, \dots, K_{max} \\ \bar{\mathbf{s}}_{c,1}^h(t) &= \mathbf{S}^h - \sum_{k=2}^{K_{max}} \bar{\mathbf{s}}_{c,k}^h(t), \quad t \geq 0 \end{aligned}$$
- 16: compute $\bar{\mathbf{n}}_c^h(t) = \bar{\mathbf{x}}_c^h(t) + \sum_{k=2}^{K_{max}} \bar{\mathbf{s}}_{c,k}^h(t), 0 \leq t \leq T$
- 17: compute $\mathbf{M}_c^h = \int_0^T \bar{\mathbf{n}}_c^h(t) e^{-\alpha_h t} dt$
- 18: compute $\bar{\mathbf{n}}_{c+1}^h(0) = \sum_{h=1}^E p_{he}^{src} \bar{\mathbf{n}}_c^h(0) + \mathbf{B}^h \mathbf{M}_c^h,$
- 19: **end for**
- 20: **until** $\max_{h=1 \dots E} \|\bar{\mathbf{n}}_{c+1}^h(0) - \bar{\mathbf{n}}_c^h(0)\|_1 < \delta_{max}$ **or** $c \leq C$
- 21: **return** $\bar{\mathbf{n}}_c^h(t), \bar{\mathbf{x}}_c^h(t), \bar{\mathbf{s}}_{c,k}^h(t), k = 1, \dots, K_{max}, 0 \leq t \leq T$