Global attraction of ODE-based mean field models with hyperexponential job sizes

Benny Van Houdt

Dept. Mathematics and Computer Science
University of Antwerp, Belgium

ABSTRACT

Mean field modeling is a popular approach to assess the performance of large scale computer systems. The evolution of many mean field models is characterized by a set of ordinary differential equations that have a unique fixed point. In order to prove that this unique fixed point corresponds to the limit of the stationary measures of the finite systems, the unique fixed point must be a global attractor. While global attraction was established for various systems in case of exponential job sizes, it is often unclear whether these proof techniques can be generalized to non-exponential job sizes.

In this paper we show how simple monotonicity arguments can be used to prove global attraction for a broad class of ordinary differential equations that capture the evolution of mean field models with hyperexponential job sizes. This class includes both existing as well as previously unstudied load balancing schemes and can be used for systems with either finite or infinite buffers.

The main novelty of the approach exists in using a Coxian representation for the hyperexponential job sizes and a partial order that is stronger than the componentwise partial order used in the exponential case.

1 INTRODUCTION

Mean field models are a popular technique to assess the performance of large scale (computer) systems. They have been applied in various areas such as load balancing [1, 5, 19, 22, 34, 38], work stealing [14, 21], caching [15], garbage collection [30, 31], CSMA networks [7], bin packing [36], file swarming systems [18], coupon collector problems [20], etc. In many cases the evolution of the mean field model is described by a simple set of ordinary differential equations (ODEs) and one can show that this set of ODEs

has a unique fixed point (that may even have a closed form). The main idea behind a mean field approximation is that the stationary distribution of a single component in the network should (weakly) converge to (the Dirac measure of) the fixed point of the ODEs as the number of components N tends to infinity. Therefore the fixed point approximates the stationary behavior of any component in a large finite system.

Different approaches exist to prove the convergence of the stationary distributions to the fixed point of the mean field limit as N tends to infinity. The traditional *indirect* method exists in first proving convergence of the stochastic processes over finite time scales, that is, for any fixed T, one shows that the sample paths of the stochastic processes on $[0, T]$ converge towards the solution of the ODEs on $[0, T]$ (with the appropriate initial condition). For this step, one can often rely on Kurtz's theorem [9, 22] or the convergence of transition semigroups of Markov processes [1, 34]. The second step exists in showing that the stochastic systems with finite N each have a stationary measure and that this sequence of stationary measures has a limit point (which follows from the tightness of the stationary measures). The final step then exists in showing that the fixed point is a global attractor and that the limit point of the stationary measures must be the Dirac measure associated with the fixed point. It is fair to state that, given existing mean field theory, proving global attraction of the fixed point is often the most demanding step (especially if the state space is a subset of $\mathbb{R}^n$, see Corollary 1 in Section 6).

A recent *direct* method to prove convergence is to rely on Theorem 1 of [37] or Theorem 3.2 of [13] that were both obtained using Stein's method [6]. This approach does not require proving convergence over finite time scales $[0, T]$. Instead it makes use of the solution of the so-called Poisson equation. The solution of this equation is expressed as an integral that is only properly defined if the fixed point is a global attractor (that is locally exponentially stable). Thus, Stein's method, when applied to ODE-based mean field models, also requires global attraction of the fixed point. In fact the main challenge in verifying the conditions needed to apply Theorem 1 of [37] or Theorem 3.2 of [13] exists in showing that the fixed point is a global attractor.

One approach to prove global attraction of a set of ODEs towards its fixed point relies on defining a Lyapunov function as in [22]. However in general coming up with a suitable Lyapunov function, even in case of exponential job sizes, is highly challenging. A somewhat more flexible approach, that was applied in [1, 19, 34] for systems with exponential job sizes, relies on monotonicity. It is composed of the following three steps. First, one defines the state space Ω in such a way that the set of ODEs maintains the componentwise partial order $\leq$ over time. In other words, if $h \leq \tilde{h}$

in the componentwise ordering, then $h(t) \leq \tilde{h}(t)$ where $h(t)$ and $\tilde{h}(t)$ are the unique solutions to the set of ODEs with $h(0) = h$ and $\tilde{h}(0) = \tilde{h}$. Next, one shows that for any fixed point π and $h \in \Omega$ there exists an $h^{(l)}, h^{(u)} \in \Omega$ such that $h^{(l)} \leq h, \pi \leq h^{(u)}$ in the componentwise ordering. Finally, global attraction on Ω follows by proving attraction for any initial point $h \in \Omega$ for which either $h \leq \pi$ or $\pi \leq h$ in the componentwise ordering.

Although it is easy to generalize ODE-based mean field models with exponential job sizes to hyperexponential job sizes (or even phase-type distributed job sizes), generalizing this monotonicity approach to establish global attraction appears problematic. In this paper we nevertheless demonstrate that for a broad class of ODE-based mean field models with hyperexponential jobs sizes, one can still rely on such a monotonicity argument. In order to do so, we introduce two novel key ideas. First, we set up the ODE-based mean field model using a Coxian representation of the hyperexponential distribution. By using this Coxian representation all jobs necessarily start service in phase one, the service phase can only increase by one at a time and the service completion rate decreases as the phase increases (see Section 2). These three features are essential to find a partial ordering on Ω that is preserved by the set of ODEs over time. Second, we rely on a partial ordering that is stronger than the componentwise ordering used in the exponential case as the set of ODEs does *not* preserve the usual componentwise order over time (see Section 5).

Hyperexponential distributions are often used to model highly variable workloads [10]. Efficient algorithms to fit a hyperexponential distribution to heavy tailed distributions can be found in [11, 16, 25, 27]. The class of hyperexponential distributions is also dense in the set of all distributions with a completely monotone probability density function (pdf) [11, Theorem 3.2], such as the Pareto and Weibull distribution. A pdf f is completely monotone if all its derivatives exist and $(-1)^n f^{(n)}(t) \geq 0$ for all $t > 0$ and $n \geq 1$.

Although various mean field models with non-exponential job sizes have been introduced, e.g., [33], most of these papers only focus on the convergence over finite time scales and the uniqueness of the fixed point. One notable exception is [5] which establishes the convergence of the stationary regime for the classic power-of-d load balancing scheme under FIFO service and any job size distribution with decreasing hazard rate. Their proof is highly technical, while the approach taken in this paper is much more elementary.

Instead of focusing on a single mean field model, we identify a set of sufficient conditions such that our result applies to any mean field model satisfying these conditions. We demonstrate that these conditions are satisfied by various mean field models, such as the classic power-of-d load balancing [22, 34], the pull/push strategies studied in [21] and the power-of-d choices load balancing with batch sampling. Further, we introduce a class of probability distributions C_0, show that the set of hyperexponential distributions is a strict subclass of C_0 and establish global attraction under these sufficient conditions for any job size distribution belonging to the class C_0. In other words, the main result also holds for some job size distributions that are not hyperexponential distributions. We

also theoretically characterize the first three moments that can be matched with a distribution belonging to C_0.

The paper is structured as follows. In Section 2 we derive a Coxian representation of a hyperexponential distribution, define the class of distributions C_0, prove that all hyperexponential distributions belong to C_0 and characterize the first three moments. In Section 3 we introduce the general form of the set of ODEs characterizing the mean field model. Examples are presented in Section 4. The state space and partial order that enable us to use monotonicity arguments are outlined in Section 5. The set of sufficient conditions and the global attraction theorem are discussed in Section 6, where we also show that convergence of the stationary measures then follows from existing results for systems with finite buffers. These conditions are verified in Section 7 for the examples presented in Section 4. The proof of the global attraction theorem is detailed in Section 8. Conclusions are drawn in Section 9.

2 COXIAN REPRESENTATIONS

A cumulative distribution function (cdf) F is a hyperexponential distribution if there exists a set of probabilities $\tilde{p}_1, \ldots, \tilde{p}_n$ such that $\sum_{k=1}^{n} \tilde{p}_k = 1$ and real numbers $\mu_1, \ldots, \mu_n > 0$ such that $F(t) = 1 - \sum_{k=1}^{n} \tilde{p}_k e^{-\mu_k t}$. Further, a cdf F is a phase-type distribution if there exists a non-negative vector $\alpha = (\alpha_1, \ldots, \alpha_n)$ with $\sum_{i=1}^{n} \alpha_i = 1$ and a matrix S with negative diagonal entries, non-negative off-diagonal entries and non-positive row sums such that $F(t) = 1 - \alpha e^{St} \underline{1}$, where $\underline{1}$ is a column vector of ones. In which case (α, S) is called a phase-type representation of F. It is well known that the representation (α, S) of a phase-type distribution is not unique [23].

The most natural phase-type representation of a hyperexponential distributions is clearly given by setting $\alpha = (\tilde{p}_1, \ldots, \tilde{p}_n)$ and

$$S = \begin{bmatrix} -\mu_1 & & & \\ & -\mu_2 & & \\ & & \ddots & \\ & & & -\mu_n \end{bmatrix}.$$

Thus, it is very natural to use this phase-type representation to define an ODE-based mean field model for systems with hyperexponential job sizes. However, by doing so it appears hard (if not impossible) to introduce a partial ordering on the state space that is preserved by the set of ODEs over time. We therefore propose to use a different phase-type representation, being the Coxian representation introduced below. Note that the choice of the phase-type representation does *not* affect the main performance measures of the system, such as the queue length or response time distribution. It obviously does affect measures such as the joint distribution of the queue length and service phase as different representations of the same distributions do not even require to have the same number of phases n.

A cdf F is a Coxian distribution if and only if it has a phase-type representation with $\alpha = (1, 0, \ldots, 0)$ and a matrix S of the

following form

$$S = \begin{bmatrix} -\mu_1 & p_1\mu_1 & & & \\ & -\mu_2 & p_2\mu_2 & & \\ & & \ddots & \ddots & \\ & & & -\mu_{n-1} & p_{n-1}\mu_{n-1} \\ & & & & -\mu_n \end{bmatrix}, \qquad (1)$$

with $\mu_i > 0$ and $0 < p_i < 1$. Thus, $F(t) = 1 - (1, 0, \ldots, 0)e^{St}\underline{1}$. For ease of presentation we define $p_n = 0$.

Cumani [8] showed that any distribution that has a phase type representation (α, S) with S triangular is a Coxian distribution (of at most the same order n). Further, as mixtures of Erlang distributions with common scale parameter are triangular, the class of Coxian distributions is dense on the space of distributions on $\mathbb{R}^+$ [29, p. 163-164]. We now introduce a subclass of the set of all Coxian distributions.

DEFINITION 1. *The class C_0 of distributions on $\mathbb{R}^+$ is defined as the class of distributions with a Coxian representation such that $\mu_i(1-p_i)$ is decreasing in i.*

In this paper we prove global attraction for a class of ODE-based mean field models where the service time distribution belongs to C_0. Let C_{Cox} be the set of all Coxian distributions and C_{hexp} the set of all hyperexponential distributions, then clearly $C_0 \subset C_{Cox}$ and $C_{hexp} \subset C_{Cox}$ (the latter due to Cumani). We now prove that C_{hexp} is a strict subclass of C_0. We first derive a simple explicit expression for the parameters of a Coxian representation of a hyperexponential distribution. To do so we start with a technical lemma.

LEMMA 1. *For any $\ell > k \geq 1$ and $\mu_j \neq \mu_k$ for $j > k$, we have*

$$\sum_{i=k+1}^{\ell} \frac{\prod_{v=k}^{i-1}(\mu_v - \mu_\ell)}{\prod_{j=k+1}^{i}(\mu_j - \mu_k)} = -1.$$

PROOF. The sum can be written as

$$\frac{\mu_k - \mu_\ell}{\mu_{k+1} - \mu_k}\left(1 + \frac{\mu_{k+1} - \mu_\ell}{\mu_{k+2} - \mu_k}\left(1 + \ldots \frac{\mu_{\ell-2} - \mu_\ell}{\mu_{\ell-1} - \mu_k}\left(1 + \frac{\mu_{\ell-1} - \mu_\ell}{\mu_\ell - \mu_k}\right)\right)\right).$$

As $(\mu_{i-1} - \mu_\ell)/(\mu_i - \mu_k)[1 + (\mu_i - \mu_\ell)/(\mu_\ell - \mu_k)] = (\mu_{i-1} - \mu_\ell)/(\mu_\ell - \mu_k)$ this expression collapses to -1. □

PROPOSITION 1. *Let $F(t) = 1 - \sum_{k=1}^{n} \tilde{p}_k e^{-\mu_k t}$ be a hyperexponential distribution and assume without loss of generality that $\mu_1 > \mu_2 > \ldots > \mu_n > 0$. Then, $F(t)$ has a Coxian representation with parameters μ_i and*

$$p_i = \frac{\sum_{k=i+1}^{n} \tilde{p}_k \prod_{j=1}^{i}(1 - \frac{\mu_k}{\mu_j})}{\sum_{k=i}^{n} \tilde{p}_k \prod_{j=1}^{i-1}(1 - \frac{\mu_k}{\mu_j})}. \qquad (2)$$

PROOF. We show that both the hyperexponential and Coxian representations are equivalent by showing that both distributions have the same Laplace Stieltjes transform $\tilde{F}(s)$ (LST). These transforms are given by

$$\tilde{F}_{hexp}(s) = \sum_{k=1}^{n} \tilde{p}_k \frac{\mu_k}{s + \mu_k},$$

and

$$\tilde{F}_{Cox}(s) = \sum_{i=1}^{n}(1 - p_i)\left(\prod_{j=1}^{i-1} p_j\right)\prod_{j=1}^{i} \frac{\mu_j}{s + \mu_j},$$

as with probability $(1 - p_i)\prod_{j=1}^{i-1} p_j$ we visit the first i phases for the Coxian representation. Using a partial fraction expansion for $1/\prod_{j=1}^{i}(s + \mu_j)$, we get

$$\tilde{F}_{Cox}(s) = \sum_{k=1}^{n}\left(\sum_{i=k}^{n}(1 - p_i)\left(\prod_{j=1}^{i-1} p_j\right)\prod_{j=1, j\neq k}^{i} \frac{\mu_j}{\mu_j - \mu_k}\right)\frac{\mu_k}{s + \mu_k}.$$

Hence, $\tilde{F}_{hexp}(s) = \tilde{F}_{Cox}(s)$ if

$$\tilde{p}_k = \sum_{i=k}^{n}(1 - p_i)\left(\prod_{j=1}^{i-1} p_j\right)\prod_{j=1, j\neq k}^{i} \frac{\mu_j}{\mu_j - \mu_k}. \qquad (3)$$

This is a linear system in the unknowns $(1 - p_i)\prod_{j=1}^{i-1} p_j$ and we now show that its solution can be expressed as

$$(1 - p_i)\prod_{j=1}^{i-1} p_j = \sum_{\ell=i}^{n} \tilde{p}_\ell \frac{\mu_\ell}{\mu_i}\prod_{v=1}^{i-1}(1 - \frac{\mu_\ell}{\mu_v}). \qquad (4)$$

For $i = n$ this is immediate from (3) with $k = n$ (as $p_n = 0$). We now apply backward induction on i. Assume the result holds for $i = k + 1, \ldots, n$. From (3) we find

$$\tilde{p}_k \prod_{j=1}^{k-1}(1 - \frac{\mu_k}{\mu_j}) = (1 - p_k)\prod_{j=1}^{k-1} p_j$$

$$+ \underbrace{\sum_{i=k+1}^{n}(1 - p_i)\left(\prod_{j=1}^{i-1} p_j\right)\prod_{j=k+1}^{i} \frac{\mu_j}{\mu_j - \mu_k}}_{(a)}. \qquad (5)$$

Applying induction and switching sums yields that (a) equals

$$\sum_{\ell=k+1}^{n} \tilde{p}_\ell \frac{\mu_\ell}{\mu_k}\sum_{i=k+1}^{\ell}\left(\prod_{v=1}^{i-1}(1 - \frac{\mu_\ell}{\mu_v})\right)\left(\prod_{j=k+1}^{i} \frac{\mu_j}{\mu_j - \mu_k}\right)\frac{\mu_k}{\mu_i}$$

$$= \sum_{\ell=k+1}^{n} \tilde{p}_\ell \frac{\mu_\ell}{\mu_k}\prod_{v=1}^{k-1}(1 - \frac{\mu_v}{\mu_j})$$

$$\cdot \underbrace{\left(\sum_{i=k+1}^{\ell}\left(\prod_{v=k}^{i-1} \frac{\mu_v - \mu_\ell}{\mu_v}\right)\left(\prod_{j=k+1}^{i} \frac{\mu_j}{\mu_j - \mu_k}\right)\frac{\mu_k}{\mu_i}\right)}_{(b)}.$$

The expression in (b) is equal to -1 due to Lemma 1, which allows us to conclude that (4) holds due to (5). Using (4), $\prod_{j=1}^{i-1} p_j = \prod_{j=1}^{i} p_j + (1 - p_i)\prod_{j=1}^{i-1} p_j$ and backward induction on i, we may conclude that $\tilde{F}_{hexp}(s) = \tilde{F}_{Cox}(s)$ if

$$\prod_{j=1}^{i} p_j = \sum_{k=i}^{n} \tilde{p}_k \prod_{j=1}^{i}(1 - \frac{\mu_k}{\mu_j}) = \sum_{k=i+1}^{n} \tilde{p}_k \prod_{j=1}^{i}(1 - \frac{\mu_k}{\mu_j}). \qquad (6)$$

The expression in (2) is now immediate, where we note that $p_i \in (0, 1)$ as $\tilde{p}_k > 0$ and $0 < (1 - \mu_k/\mu_j) < 1$ for $j < k$. □

The above result may be of separate interest. We now use it to establish the following theorem:

THEOREM 1. *The class of hyperexponential distributions C_{hexp} is a subclass of C_0.*

PROOF. Given Proposition 1, it suffices to show that $(1-p_i)\mu_i$ is decreasing in i. As

$$\mu_i \sum_{k=i+1}^{n} \tilde{p}_k \prod_{j=1}^{i}(1-\frac{\mu_k}{\mu_j}) = \sum_{k=i+1}^{n} \tilde{p}_k \mu_i \prod_{j=1}^{i-1}(1-\frac{\mu_k}{\mu_j})$$
$$- \sum_{k=i+1}^{n} \tilde{p}_k \mu_k \prod_{j=1}^{i-1}(1-\frac{\mu_k}{\mu_j}),$$

one readily obtains from (2) that

$$(1-p_i)\mu_i = \frac{\sum_{k=i}^{n} \tilde{p}_k \mu_k \prod_{j=1}^{i-1}(1-\frac{\mu_k}{\mu_j})}{\sum_{k=i}^{n} \tilde{p}_k \prod_{j=1}^{i-1}(1-\frac{\mu_k}{\mu_j})}. \quad (7)$$

As $(1-\frac{\mu_k}{\mu_i}) = 0$ for $k=i$, we can start both sums in the expression for $(1-p_{i+1})\mu_{i+1}$ in $k=i$. Further, as $(1-\frac{\mu_k}{\mu_j}) > 0$ for $k > j$, we can rewrite $(1-p_i)\mu_i > (1-p_{i+1})\mu_{i+1}$ as

$$\left(\sum_{k=i}^{n} \tilde{p}_k \mu_k \xi_{k,i-1}\right)\left(\sum_{k=i}^{n} \tilde{p}_k \xi_{k,i-1} - \sum_{k=i}^{n} \tilde{p}_k \frac{\mu_k}{\mu_i} \xi_{k,i-1}\right) >$$
$$\left(\sum_{k=i}^{n} \tilde{p}_k \xi_{k,i-1}\right)\left(\sum_{k=i}^{n} \tilde{p}_k \mu_k \xi_{k,i-1} - \sum_{k=i}^{n} \tilde{p}_k \frac{\mu_k^2}{\mu_i} \xi_{k,i-1}\right),$$

where we denoted $\prod_{j=1}^{i-1}(1-\frac{\mu_k}{\mu_j})$ as $\xi_{k,i-1}$. This can be restated as

$$\left(\sum_{k=i}^{n} \tilde{p}_k \mu_k \xi_{k,i-1}\right)^2 < \left(\sum_{k=i}^{n} \tilde{p}_k \xi_{k,i-1}\right)\left(\sum_{k=i}^{n} \tilde{p}_k \mu_k^2 \xi_{k,i-1}\right),$$

which is equivalent to

$$\left(\sum_{k=i}^{n} \mu_k \frac{\tilde{p}_k \xi_{k,i-1}}{\sum_{k=i}^{n} \tilde{p}_k \xi_{k,i-1}}\right)^2 < \sum_{k=i}^{n} \mu_k^2 \frac{\tilde{p}_k \xi_{k,i-1}}{\sum_{k=i}^{n} \tilde{p}_k \xi_{k,i-1}}.$$

By defining X_i such that $P[X_i = \mu_k] = \tilde{p}_k \xi_{k,i-1}/\sum_{k=i}^{n} \tilde{p}_k \xi_{k,i-1}$, the above inequality holds as $E[X]^2 < E[X^2]$ for any random variable X (that is not deterministic). □

When $n=2$ one can show that all Coxian distributions with $(1-p_1)\mu_1 > (1-p_2)\mu_2 = \mu_2$ are also hyperexponential distributions. However for $n > 2$ the example below shows that this is not the case, so the set of hyperexponential distributions C_{hexp} is a *strict* subclass of the class C_0. Consider the Coxian distribution with parameters $\mu_1 = 1, \mu_2 = 2, \mu_3 = 0.1, p_1 = 0.1$ and $p_2 = 0.8$. This distribution belongs to the class C_0. However using (3), we see that its LST is given by

$$\tilde{F}_{Cox}(s) = \frac{83}{90}\frac{1}{s+1} - \frac{3}{190}\frac{2}{s+2} + \frac{16}{171}\frac{0.1}{s+0.1}.$$

This distribution is not a hyperexponential as $\tilde{p}_2$ is negative.

Let R_i be the expected remaining service time of a job in phase i. Clearly, $R_n = 1/\mu_n$ and $R_{i-1} = 1/\mu_{i-1} + p_{i-1}R_i$ for $i = 2,\ldots,n$. Without loss of generality we assume that the mean job size equals

one, which implies that $R_1 = 1$ (as all jobs start in phase 1 and stay there for an exponential amount of time). For later use, we rewrite this as

$$R_i p_{i-1}\mu_{i-1} = \mu_{i-1}R_{i-1} - 1. \quad (8)$$

LEMMA 2. *If $\mu_i(1-p_i)$ is decreasing in i, we have $R_i > R_{i-1}$, for $i = 2,\ldots,n$.*

PROOF. The proof is presented in Appendix A. □

Remark: Coxian distributions are sometimes defined using an alternate (α, S) representation given by $\alpha = (\alpha_1,\ldots,\alpha_n)$ and

$$S = \begin{bmatrix} -\lambda_1 & \lambda_1 & & & \\ & -\lambda_2 & \lambda_2 & & \\ & & \ddots & \ddots & \\ & & & -\lambda_{n-1} & \lambda_{n-1} \\ & & & & -\lambda_n \end{bmatrix}, \quad (9)$$

with $\alpha_{n-i+1} = (1-p_i)\prod_{j=1}^{i-1} p_j$ and $\lambda_{n-i+1} = \mu_i$.

2.1 Moment matching

In this section we study the range of the first three moments that can be matched with a distribution belonging to class C_0. We first establish that any distribution in C_0 has a decreasing hazard rate.

PROPOSITION 2. *Any distribution belonging to C_0 has a decreasing hazard rate.*

PROOF. Let τ_t be the service phase of a job at time t (given that it started service at time 0) and Y the job size, then the hazard rate $h(t)$ at time t can be written as

$$h(t) = \sum_{i=1}^{n} P[\tau_t = i|Y > t]\mu_i(1-p_i). \quad (10)$$

We need to show that $h(s) \geq h(t)$ for $0 \leq s < t$. As the hazard rate $h(t)$ defined in (10) can be rewritten as

$$h(t) = \underbrace{P[\tau_t \geq 1|Y > t]}_{=1}\mu_1(1-p_1)$$
$$- \sum_{i=2}^{n} P[\tau_t \geq i|Y > t]\underbrace{(\mu_{i-1}(1-p_{i-1}) - \mu_i(1-p_i))}_{>0},$$

we find that $h(s) \geq h(t)$ if

$$P[\tau_s \geq i|Y > s] \leq P[\tau_t \geq i|Y > t],$$

for $s < t$. This inequality is immediate after noting that:

$$P[\tau_t \geq i, Y > t] \geq P[\tau_s \geq i, Y > t]$$
$$= P[\tau_s \geq i, Y > s]P[\tau_s \geq i, Y > t|\tau_s \geq i, Y > s]$$
$$= P[\tau_s \geq i, Y > s]P[Y > t|\tau_s \geq i, Y > s]$$
$$\geq P[\tau_s \geq i, Y > s]P[Y > t|Y > s]$$
$$= P[\tau_s \geq i, Y > s]P[Y > t]/P[Y > s],$$

where the second inequality is due to the fact that the rate at which a service completion can occur decreases as the phase increases. □

Let $m_i = E[Y^i]$ be the i-th moment of the job size distribution Y. For any phase type distribution with representation (α, S) we have $m_i = i!\alpha(-S)^{-i}\underline{1}$. In order to characterize the set of the first three moments that can be matched by the distributions belonging to C_0, we focus on the second and third normalized moments:

$$n_2 = \frac{m_2}{m_1^2}, \quad n_3 = \frac{m_3}{m_1 m_2},$$

where $n_2, n_3 \geq 1$ for any positive valued distribution [24]. The advantage of using the normalized moments is that we no longer need to care about the first moment. Indeed, if (α, S) matches n_2 and n_3 and has mean 1, then $(\alpha, S/m_1)$ still matches n_2 and n_3 and has mean m_1 (as dividing S by m_1 changes the i-th moment by a factor m_1^i, which implies that n_2 and n_3 are not affected by dividing S by m_1). Thus, if we found a distribution with mean 1 in C_0 that matches n_2 and n_3, we can simply multiply the rates μ_i by $1/m_1$ to get any desired mean m_1.

Let $\mathcal{A}_{n_2, n_3}^{(n)}$ be the set of normalized second and third moments that can be matched with a distribution belonging to C_0 with at most n phases.

PROPOSITION 3. *The set* $\mathcal{A}_{n_2, n_3}^{(2)} = \{(n_2, n_3) | n_2 > 2, n_3 > \frac{3}{2}n_2\} \cup \{(2, 3)\}$

PROOF. By Proposition 2 any distribution part of C_0 has a decreasing hazard rate and therefore its squared coefficient of variation $C_X \geq 1$ [28, p. 16-19]. As $n_2 = 1 + C_X$, we have $n_2 \geq 2$. Further if $n_2 = 2$, the distribution is the exponential distribution and n_3 therefore equals 3. For $n_2 > 2$ the value of n_3 must exceed $3n_2/2$ as Theorem 3.1 in [4] indicates that this is the case for any order 2 Coxian distribution with $n_2 > 2$ (see also Theorem 1 in [24]). Thus it remains to show that C_0 contains a distribution that matches n_2 and n_3 for any $n_2 > 2$ and $n_3 > 3n_2/2$.

The proposition in Section 3.1 of [35] shows that the set of normalized moments n_2 and n_3 that can be matched by a hyperexponential distribution is exactly the set $\mathcal{A}_{n_2, n_3}^{(2)}$ and in such case the matching can be achieved with just 2 phases. In fact, the parameters of a two phase hyperexponential distribution that matches n_2 and n_3 are given by (3.5) and (3.6) in [35]. As all hyperexponential distributions belong to C_0, this completes the proof.
□

We note that the proposition in Section 3.1 in [35] indicates that we cannot match a larger range of (n_2, n_3) values by using more than two phases in case we restrict ourselves to hyperexponential distributions. If we consider Coxian distributions with n phases and $n_2 > 2$, then Theorem 3.1 in [4] indicates that we can match any $n_3 > (n+1)n_2/n$. Thus, the larger n, the lower n_3 can become, contrary to the class of hyperexponential distributions.

The next proposition shows that while class C_0 lies somewhere between the class of hyperexponential and Coxian distributions with $n_2 \geq 2$, increasing the number of phases does *not* allow us to match a larger range of n_3 values. Thus, as far as matching the first three moments is concerned, the class C_0 does not provide more flexibility than the set of hyperexponential distributions.

PROPOSITION 4. *The set* $\mathcal{A}_{n_2, n_3}^{(n)} = \mathcal{A}_{n_2, n_3}^{(2)}$ *for any* $n \geq 2$.

PROOF. We use induction on n and note that the result clearly holds for $n = 2$. The proof follows the same line of reasoning as the proof of Theorem 3.1 in [4]. Let $\hat{n}_2$ and $\hat{n}_3$ be the normalized moments of a distribution in C_0 represented by (α, S). Denote the matrix S as

$$S = \begin{bmatrix} -\mu_1 & p_1\mu_1\alpha_{n-1} \\ 0 & A \end{bmatrix},$$

where α_{n-1} is the first row of the size $n - 1$ identity matrix. Note that (α_{n-1}, A) is a phase type representation of a distribution with $n - 1$ phases in C_0. Let n_2 and n_3 be the normalized moments of (α_{n-1}, A) and m_1 be its mean. By induction we know $n_3 > 3n_2/2$ for $n_2 > 2$. Using exactly the same arguments as in the proof of Theorem 3.1 in [4], we find

$$\hat{n}_3 = \frac{3}{g} + \frac{(\hat{n}_2 g - 2)^2}{(g-1)g\hat{n}_2}\frac{n_3}{n_2},$$

with $g = 1 + \mu_1 m_1 p_1 \geq 1$. By induction we know $n_3/n_2 \geq 3/2$, meaning

$$\hat{n}_3 \geq \frac{3}{g} + \frac{(\hat{n}_2 g - 2)^2}{(g-1)g\hat{n}_2}\frac{3}{2}. \tag{11}$$

Further in the proof of Theorem 3.1 in [4] it is shown that $(\hat{n}_2 g - 2)^2/((g-1)g\hat{n}_2)$ is decreasing in g on $(1, \infty)$ whenever $\hat{n}_2 \geq (n + 4)/(n + 1)$. Since $\hat{n}_2 \geq 2$ as any distribution in C_0 has a decreasing hazard rate, we obtain a lower bound for $\hat{n}_3$ by taking the limit of g to infinity in (11). This limit clearly equals $3\hat{n}_2/2$, which completes the proof.
□

3 THE FORM OF THE ODE

ODE-based mean field models of systems with exponential job sizes (with mean 1) are often of the following form (see Section 4 for examples):

$$\frac{d}{dt}h_{\ell,1}(t) = f_{\ell,1}(h(t)) - (h_{\ell,1}(t) - h_{\ell+1,1}(t)), \tag{12}$$

where $h_{\ell,1}(t)$ represents the fraction of the servers with at least ℓ jobs and $f_{\ell,1}(h(t))$ captures events such as job arrivals and job transfers (see Section 4.2). The term $-(h_{\ell,1}(t) - h_{\ell+1,1}(t))$ reflects the drift due to the exponential service completions. The assumption that the mean job size equals 1 is made throughout the paper (without loss of generality).

We now generalize this set of ODEs to the case where the job sizes belong to class C_0 given that the service discipline is first-come-first-served (FCFS). Define $h_{\ell,i}(t)$, for $\ell > 0$ and $i = 1, \ldots, n$, as the fraction of the queues at time t with a queue length of at least ℓ in service phase $j \geq i$. Thus, $(h_{\ell,i}(t) - h_{\ell,i+1}(t))$ is the fraction of queues at time t with ℓ or more jobs that are in service phase i. For ease of notation let $h_{\ell,n+1}(t) = 0$ and $h_{0,1}(t) = 1$. Note that a service completion in a queue with a length of at least ℓ always decreases $h_{\ell,i}(t)$ for $i \geq 2$ as the next customer starts service in phase 1, whereas $h_{\ell,1}(t)$ only decreases if the queue length is exactly ℓ.

Hence, the set of ODEs given by (12) then generalizes to:

$$\frac{d}{dt}h_{\ell,1}(t) = f_{\ell,1}(h(t))$$
$$- \sum_{j=1}^{n} \left[(h_{\ell,j}(t) - h_{\ell,j+1}(t)) - (h_{\ell+1,j}(t) - h_{\ell+1,j+1}(t)) \right] \mu_j (1 - p_j) \tag{13}$$

$$\frac{d}{dt}h_{\ell,i}(t) = 1[\ell > 1]f_{\ell,i}(h(t)) + (h_{\ell,i-1}(t) - h_{\ell,i}(t))p_{i-1}\mu_{i-1}$$
$$- \sum_{j=i}^{n} (h_{\ell,j}(t) - h_{\ell,j+1}(t))\mu_j (1 - p_j) \tag{14}$$

for $\ell \geq 1$ and $i = 2, \ldots, n$, where the sums are due to service completions and the second term in the drift of $h_{\ell,i}(t)$ corresponds to phase changes.

We remark that we can also model systems with a finite buffer of size B by setting $f_{\ell,i}(h) = 0$, for $i = 1, \ldots, n$ and $\ell > B$, as this implies that $h_{\ell,i}(t) = 0$ for $i = 1, \ldots, n$ and $\ell > B$.

4 EXAMPLES

4.1 JSQ(d): Join-the-Shortest-Queue among d randomly selected servers

Let us first consider the classic power-of-d choices load balancer [22, 34], where jobs arrive at rate λN to a dispatcher who immediately assigns incoming jobs among the N servers by routing the job to the server with the least number of jobs among d randomly selected servers. In this case the function f reflects the changes due to arrivals and one finds for $\ell \geq 1$

$$f_{\ell,1}(h(t)) = \lambda(h_{\ell-1,1}(t)^d - h_{\ell,1}(t)^d),$$

as $h_{\ell-1,1}(t)^d - h_{\ell,1}(t)^d$ is the probability that the server with the least number of jobs among d randomly selected servers has queue length $\ell - 1$. Further, since the dispatcher does not take the service phase into account when dispatching jobs and $(h_{\ell-1,i}(t) - h_{\ell,i}(t))/(h_{\ell-1,1}(t) - h_{\ell,1}(t))$ is the probability that a server of length $\ell - 1$ is in service phase $j \geq i$, we have

$$f_{\ell,i}(h(t)) = f_{\ell,1}(h(t))\frac{h_{\ell-1,i}(t) - h_{\ell,i}(t)}{h_{\ell-1,1}(t) - h_{\ell,1}(t)}$$
$$= \lambda \left(\sum_{j=0}^{d-1} h_{\ell-1,1}(t)^j h_{\ell,1}(t)^{d-1-j} \right)(h_{\ell-1,i}(t) - h_{\ell,i}(t)),$$

for $\ell > 1$ and $i = 2, \ldots, n$, as $(a^d - b^d)/(a - b) = \sum_{j=0}^{d-1} a^j b^{d-1-j}$. For convenience we also define $f_{1,i}(h(t)) = 0$ for $i = 2, \ldots, n$.

4.2 Pull and push strategies

In this example we consider the system analyzed in [21]. It consists of N servers that each have local job arrivals with rate λ. Servers that are idle generate probe messages at rate r. A probe message is sent to a random server and if this server has pending jobs, a job is transferred to the idle server. The function f now captures the changes due to arrivals as well as job transfers, hence

$$f_{\ell,1}(h(t)) = \lambda(h_{\ell-1,1}(t) - h_{\ell,1}(t))$$
$$+ r(1 - h_{1,1}(t))[1[\ell = 1]h_{2,1}(t) - 1[\ell > 1](h_{\ell,1}(t) - h_{\ell+1,1}(t))],$$

for $\ell \geq 1$ and

$$f_{\ell,i}(h(t)) = \lambda(h_{\ell-1,i}(t) - h_{\ell,i}(t))$$
$$- r(1 - h_{1,1}(t))(h_{\ell,i}(t) - h_{\ell+1,i}(t)),$$

for $\ell > 1$ and $i = 2, \ldots, n$. Note that $r(1 - h_{1,1}(t))(h_{\ell,i}(t) - h_{\ell+1,i}(t))$ is the rate at which jobs are transferred from a server with length ℓ in phase $j \geq i$ to an idle server. Therefore $r(1 - h_{1,1}(t))h_{2,1}(t)$ is the rate at which idle servers become busy due to the probe messages.

4.3 JSQ(K,d): Join-the-Shortest-K-Queues among d randomly selected servers

This example is a generalization of the first example. Jobs now arrive in batches of size K and the dispatcher assigns the K jobs (with independent sizes) belonging to the same batch to the K servers with the least number of jobs among d randomly selected servers (with $K \leq d$). This load balancing scheme is called *batch sampling* in [38]. The mean field model in [38] is however different than the one presented here, as we assume that both K and d are fixed, i.e., do not grow as a function of N.

In this case $\lambda < 1/K$ in order to have a stable system (as the mean service time of a job equals 1) and the function f once more reflects the changes due to arrivals. Note that

$$p_{k,\ell}(h(t)) = \sum_{s=0}^{k-1} \binom{d}{s}(1 - h_{\ell,1}(t))^s h_{\ell,1}(t)^{d-s},$$

is the probability that the k-th shortest queue has a length of at least ℓ. As such

$$f_{\ell,1}(h(t)) = \lambda \sum_{k=1}^{K} (p_{k,\ell-1}(h(t)) - p_{k,\ell}(h(t))) = \lambda \sum_{s=0}^{K-1} (K - s)\binom{d}{s}$$
$$\cdot \left((1 - h_{\ell-1,1}(t))^s h_{\ell-1,1}(t)^{d-s} - (1 - h_{\ell,1}(t))^s h_{\ell,1}(t)^{d-s} \right),$$

for $\ell \geq 1$ and

$$f_{\ell,i}(h(t)) = f_{\ell,1}(h(t))\frac{h_{\ell-1,i}(t) - h_{\ell,i}(t)}{h_{\ell-1,1}(t) - h_{\ell,1}(t)}.$$

for $\ell > 1$. In addition we define $f_{1,i}(h(t)) = 0$ for $i = 2, \ldots, n$.

Note that in the special case where $K = d$, one finds that $f_{\ell,1}(h(t))$ simplifies to $\lambda K(h_{\ell-1,1}(t) - h_{\ell,1}(t))$. Thus, when $K = d$ the set of ODEs describes the transient evolution of an $M/Cox/1$ queue with arrival rate λK.

5 STATE SPACE AND PARTIAL ORDER

In the case of exponential job sizes the state space is typically defined as

$$\Omega_{expo} = \{(h_{\ell,1})_{\ell>0} | 0 \leq h_{\ell,1} \leq 1, h_{\ell+1,1} \leq h_{\ell,1}, \sum_{\ell} h_{\ell,1} < \infty\},$$

where $h_{\ell,1}$ represents the fraction of queues with length ℓ or more. The partial order used to prove global attraction on Ω_{expo} in case of exponential job sizes is the componentwise order. In this section we introduce the state space and partial order needed in case of a job size distribution belonging to class C_0.

We define the state space Ω of the mean field model in terms of the variables $h_{\ell,i}$ as follows

$$\Omega = \{(h_{\ell,i})_{\ell>0,\,i\in\{1,\ldots,n\}} \mid 0 \le h_{\ell,i} \le 1, h_{\ell,i+1} \le h_{\ell,i}, h_{\ell+1,i} \le h_{\ell,i},$$
$$h_{\ell,i} + h_{\ell+1,i+1} \ge h_{\ell+1,i} + h_{\ell,i+1}, \sum_{\ell} h_{\ell,1} < \infty\}.$$

The conditions $h_{\ell,i+1} \le h_{\ell,i}$ and $h_{\ell+1,i} \le h_{\ell,i}$ are obvious as $h_{\ell,i}$ is the fraction of servers with at least ℓ jobs in service phase $j \ge i$. The inequality $h_{\ell,i} + h_{\ell+1,i+1} \ge h_{\ell+1,i} + h_{\ell,i+1}$ may seem a bit unexpected. This inequality can be understood by noting that Ω corresponds to

$$\bar{\Omega} = \{(x_0, (x_{\ell,i})_{\ell>0,\,i\in\{1,\ldots,n\}}) \mid x_0 \ge 0, x_{\ell,i} \ge 0,$$
$$x_0 + \sum_{\ell,i} x_{\ell,i} = 1, \sum_{\ell,i} \ell x_{\ell,i} < \infty\}.$$

after a change of variables (i.e., $h_{\ell,i} = \sum_{\ell' \ge \ell} \sum_{i' \ge i} x_{\ell',i'}$), where $x_{\ell,i}$ is the fraction of servers with exactly ℓ jobs in service phase i. Therefore the inequality $h_{\ell,i} + h_{\ell+1,i+1} \ge h_{\ell+1,i} + h_{\ell,i+1}$ follows from the fact that $(h_{\ell,i} - h_{\ell,i+1}) - (h_{\ell+1,i} - h_{\ell+1,i+1}) = x_{\ell,i} \ge 0$.

In the case of a system with a finite buffer of size B the state space reduces to

$$\Omega_B = \{(h_{\ell,i})_{\ell\in\{1,\ldots,B\},\,i\in\{1,\ldots,n\}} \mid 0 \le h_{\ell,i} \le 1, h_{\ell,i+1} \le h_{\ell,i},$$
$$h_{\ell+1,i} \le h_{\ell,i}, h_{\ell,i} + h_{\ell+1,i+1} \ge h_{\ell+1,i} + h_{\ell,i+1}\}.$$

Whenever the buffer size is finite, we can replace Ω in all subsequent statements by Ω_B.

PROPOSITION 5. *For any fixed point $\pi \in \Omega$ of the set of ODEs given by (13-14), we have $\pi_{1,i} = \pi_{1,1} \sum_{j=i}^{n} \frac{1}{\mu_j} \prod_{s=1}^{j-1} p_s$, for $i = 1,\ldots,n$, where μ_i and p_s are the parameters of the Coxian representation.*

PROOF. See Appendix B. □

We introduce the following partial order on Ω which reduces to the usual componentwise order in case of exponential job sizes (i.e., when $n = 1$).

DEFINITION 2 (PARTIAL ORDER $\le_C$). *Let $h, \tilde{h} \in \Omega$. We state that $h \le_C \tilde{h}$ if and only if*

$$h_{\ell,i} \le \tilde{h}_{\ell,i}, \tag{15}$$

for all ℓ, i, and

$$h_{\ell_1,1} + \sum_{i=2}^{n}(h_{\ell_i,i} - h_{\ell_{i-1},i}) \le \tilde{h}_{\ell_1,1} + \sum_{i=2}^{n}(\tilde{h}_{\ell_i,i} - \tilde{h}_{\ell_{i-1},i}), \tag{16}$$

for any set of integers $\ell_1 \ge \ell_2 \ge \ldots \ge \ell_n \ge 1$ with $\ell_1 > \ell_n$.

It is useful to note that $h_{\ell_1,1} + \sum_{i=2}^{n}(h_{\ell_i,i} - h_{\ell_{i-1},i})$ is the fraction of the servers for which the queue length is at least ℓ_i and the service phase equals i for some $i \in \{1,\ldots,n\}$.

Without condition (16) the order would correspond to the usual componentwise partial order. To illustrate the need for a stronger partial order, let $n = 2$ and consider $h, \tilde{h} \in \Omega$ with $h_{1,1} = \tilde{h}_{1,1} = 1$, $h_{1,2} = \tilde{h}_{1,2} = 1/2$, $h_{2,1} = \tilde{h}_{2,1} = 1/2$, $h_{2,2} = 0$, $\tilde{h}_{2,2} = 1/2$ and $h_{3,1} = \tilde{h}_{3,1} = 0$. Thus, in both states half of the servers have queue length one and the other half has queue length 2. In state h the servers with length 1 are in service phase 2 and the servers with

length 2 are in phase 1, while in state $\tilde{h}$ the phases are reversed (queues with length i are in phase i, for $i = 1, 2$). Note that h is smaller than $\tilde{h}$ in the componentwise order, but condition (16) is violated with $\ell_1 = 2$ and $\ell_2 = 1$, meaning $h \not\le_C \tilde{h}$. If we now look at the drift of the number of busy servers due to service completions, we see that it equals $-\mu_2/2$ in state h and $-\mu_1(1-p_1)/2$ in state $\tilde{h}$. Hence, $h_{1,1} = \tilde{h}_{1,1} = 1$, but $h_{1,1}$ decreases more slowly than $\tilde{h}_{1,1}$ (when $\mu_2 < \mu_1(1-p_1)$). This example therefore shows that the componentwise partial order used for the set of ODEs with exponential job sizes, is *not* preserved over time by the set of ODEs with a job size distribution in C_0 and we need to replace it by a stronger partial order, which turns out to be the order $\le_C$ defined above.

We end by noting that due to the condition $h_{\ell,i} - h_{\ell+1,i} \ge h_{\ell,i+1} - h_{\ell+1,i+1}$ in Ω, we have $h_{\ell,1} - h_{\ell+1,1} \ge h_{\ell,i} - h_{\ell+1,i}$ for any $i = 2,\ldots,n$ and therefore

$$h_{\ell_1,1} + \sum_{i=2}^{n}(h_{\ell_i,i} - h_{\ell_{i-1},i}) = h_{\ell_1,1} + \sum_{i=2}^{n}\sum_{\ell=\ell_i}^{\ell_{i-1}-1}(h_{\ell,i} - h_{\ell+1,i})$$

$$\le h_{\ell_1,1} + \sum_{i=2}^{n}\sum_{\ell=\ell_i}^{\ell_{i-1}-1}(h_{\ell,1} - h_{\ell+1,1})$$

$$= h_{\ell_1,1} + \sum_{i=2}^{n}(h_{\ell_i,1} - h_{\ell_{i-1},1}) = h_{\ell_n,1}, \tag{17}$$

for any $h \in \Omega$.

6 GLOBAL ATTRACTION

We now list the assumptions needed to establish the main result. Note that some of the intermediate results do not require all of the assumptions.

ASSUMPTION A0. *The functions $f_{\ell,i}(h) : \Omega \to \mathbb{R}$ are such that for any $h \in \Omega$, the set of ODEs given by (13-14) has a unique solution $h(t) : [0,\infty) \to \mathbb{R}$ with $h(0) = h$ and there exists a fixed point π in Ω.*

The existence of a unique (continuously differentiable) solution $h(t)$ is guaranteed by defining a norm on $\mathbb{R}^{\mathbb{N}}$ such that the drift is locally Lipschitz and bounded on Ω. When the buffer size $B < \infty$, the existence of a fixed point follows almost immediately from Brouwer's fixed point theorem as Ω_B is a convex and compact subset of $\mathbb{R}^{Bn}$ and Ω_B is clearly a forward invariant set [3].

The next two assumptions are needed to establish that the partial order $\le_C$ is preserved over time by the set of ODEs.

ASSUMPTION A1. *The functions $f_{\ell,i}(h) : \Omega \to \mathbb{R}$ are non-decreasing in $h_{\ell',i'}$ for any $(\ell',i') \ne (\ell,i)$.*

For any set of integers $\ell_1 \ge \ell_2 \ge \ldots \ge \ell_n \ge 1$ with $\ell_1 > \ell_n$, define $g_{(\ell_1,\ldots,\ell_n)}$ and $f_{(\ell_1,\ldots,\ell_n)}$ as a function from Ω to $\mathbb{R}$ such that

$$g_{(\ell_1,\ldots,\ell_n)}(h) = h_{\ell_1,1} + \sum_{i=2}^{n}(h_{\ell_i,i} - h_{\ell_{i-1},i}), \tag{18}$$

and

$$f_{(\ell_1,\ldots,\ell_n)}(h) = f_{\ell_1,1}(h) + \sum_{i=2}^{n}(f_{\ell_i,i}(h) - f_{\ell_{i-1},i}(h)). \quad (19)$$

Due to (16), $h \leq_C \tilde{h}$ implies that $g_{(\ell_1,\ldots,\ell_n)}(h) \leq g_{(\ell_1,\ldots,\ell_n)}(\tilde{h})$.

ASSUMPTION A2. *The functions* $f_{(\ell_1,\ldots,\ell_n)}(h) : \Omega \to \mathbb{R}$ *are such that*

$$f_{(\ell_1,\ldots,\ell_n)}(h) \leq f_{(\ell_1,\ldots,\ell_n)}(\tilde{h})$$

for any $h, \tilde{h} \in \Omega$ *such that* $h \leq_C \tilde{h}$ *and* $g_{(\ell_1,\ldots,\ell_n)}(\tilde{h}) = g_{(\ell_1,\ldots,\ell_n)}(h)$.

The next assumption is used to prove that for any $h \leq_C \pi$ or $\pi \leq_C h$ the trajectory starting in h of the set of ODEs converges to the fixed point π.

ASSUMPTION A3. *The functions* $f_{\ell,1}(h)$ *are such that for any fixed point* π *and* $L \geq 1$ *we have*

$$\sum_{\ell \geq L}(f_{\ell,1}(h) - f_{\ell,1}(\pi)) = \sum_{\ell=1}^{L-1}\sum_{j=1}^{n} b_{L,\ell,j}(h)(h_{\ell,j} - \pi_{\ell,j}) - a_{L,\pi}(h),$$

for some bounded functions $b_{L,\ell,j}(h)$ *on* Ω *and functions* $a_{L,\pi}(h)$ *for which* $a_{L,\pi}(h) \geq 0$ *if* $\pi \leq_C h$ *and* $a_{L,\pi}(h) \leq 0$ *if* $h \leq_C \pi$.

The main theorem is stated below.

THEOREM 2 (GLOBAL ATTRACTION). *Consider the set of ODEs given by (13-14). Assume that (A0-A3) hold and* $\mu_i(1 - p_i)$ *is decreasing in* i. *Then, for any* $h(0) \in \Omega$, $h(t)$ *converges pointwise to the unique fixed point* $\pi \in \Omega$ *as t tends to infinity.*

PROOF. In Section 8 we show that any fixed point $\pi \in \Omega$ is a global attractor, which implies that the fixed point is unique. □

COROLLARY 1. *Consider a density dependent population process as defined by Kurtz [17] on* Ω_B *such that its drift given by the right hand side of (13-14) is Lipschitz continuous. Let* $X_\infty^{(N)}$ *be the stationary measure of the N-th population process and assume (A0-A3) hold, then* $X_\infty^{(N)}$ *converges weakly to the Dirac measure on* π.

PROOF. As the N-th population process has a finite number of states it has a unique stationary measure $X_\infty^{(N)}$. This sequence of measures is tight as Ω_B is compact, thus any subsequence has a further subsequence that converges to some limit point. Due to Theorem 3.5 and Corollary 3.9 in [26] any such limit point has support on the Birkhoff center of the set of ODEs[1]. As π is a global attractor, the Birkhoff center is the singleton $\{\pi\}$ and the only possible limit point is therefore the Dirac measure on π. Thus, every subsequence of $X_\infty^{(N)}$ has a further subsequence that converges to the same limit, which implies that the entire sequence converges to this limit. □

Note that the above corollary is very general. In order to apply it, we do need to truncate the buffer to some finite size B (as in [12, 14]). This is not a real restriction from a practical point of view as there is virtually no difference between having an infinite buffer or a huge finite buffer, say of size $B = 10^{15}$ (provided that

---

[1]Note that this result holds in a more general setting that the one considered here, where the drift is not necessarily Lipschitz continuous and is characterized by a differential inclusion.

the system is stable in case of an infinite buffer). Establishing a similar result for infinite buffers is technically more demanding as one needs to establish the existence of $X_\infty^{(N)}$ and prove that this sequence converges.

An issue regarding the convergence is that Ω is not compact due to the condition $\sum_\ell h_{\ell,1} < \infty$ (note that as Ω is not a finite dimensional Euclidean space, compactness of a set depends on the norm used). For systems with exponential job sizes the following approach is often used, see [1, 19, 34]. One first drops the condition that prevents Ω from being compact, thus in our case we consider Ω_∞ which equals Ω without the requirement $\sum_\ell h_{\ell,1} < \infty$. Then one picks a suitable norm, for instance $||h - \tilde{h}|| = \sum_{i=1}^{n}\sum_\ell \frac{|h_{\ell,i} - \tilde{h}_{\ell,i}|^2}{2^\ell}$ in our case[2], such that Ω_∞ is compact. Thus, by Prokhorov's theorem any subsequence $X_\infty^{(N_k)}$ of $X_\infty^{(N)}$ has a further subsequence $X_\infty^{(N'_k)}$ that converges to some measure on Ω_∞. Next one argues that any such limit point π^* necessarily concentrates on Ω. Note that while $X_\infty^{(N)}(\Omega) = 1$ for all N, weak convergence does not immediately imply that $\pi^*(\Omega) = 1$ as Ω is an open set.

To show that $\pi^*(\Omega) = 1$, it suffices that $E_{\pi^*}[\sum_\ell h_{\ell,1}] < \infty$. As $E_{\pi^*}[\sum_\ell h_{\ell,1}] \leq \liminf E_{X_\infty^{(N'_k)}}[\sum_\ell h_{\ell,1}]$ (due to Portmanteau's theorem as $\sum_\ell h_{\ell,1}$ is continuous and bounded from below), $\pi^*(\Omega) = 1$ if $E_{X_\infty^{(N)}}[\sum_\ell h_{\ell,1}]$ is bounded by some constant c for all N. Finally, this constant c is shown to be the mean queue length in some finite stable queueing system (an M/M/1 queue in [34], a set of J queues with random routing in [1] and a classic Jackson network in [19]).

Having established that $\pi^*(\Omega) = 1$ for any limit point π^*, one can use Theorem 1 of [2] and the global attraction to show that π^* is the Dirac measure δ_π. To apply this theorem weak convergence over finite time scales suffices.

In case of phase-type service exactly the same line of reasoning can be applied. The main step that requires extra care is to show that the mean queue length of a queue in the N-th system is bounded by some constant c, for instance by letting c be the mean queue length of a queue in a set of N independent $M/PH/1$ queues (which should hold for any load balancing strategy that performs better than random). Recall that an $M/PH/1$ queue (with a load below one) has a finite mean queue length as the second moment of a phase-type distribution is finite.

7 EXAMPLES REVISITED

In this section we discuss assumptions A0 to A3 for the examples listed in Section 4. With respect to assumption A0, we only briefly discuss the existence of a fixed point as the existence of a unique solution $h(t)$ with $h(0) = h$ for $h \in \Omega$ can be easily verified by checking the Lipschitz continuity of the drift on Ω.

7.1 JSQ(d): Join-the-Shortest-Queue among d randomly selected servers

The existence of a fixed point when the buffer size B is finite is easy to establish (see Section 7.3 with $K = 1$). For an infinite buffer size

---

[2]Note that pointwise convergence of $h(t)$ to π also implies convergence under this norm.

B, the existence of a fixed point in Ω follows from [5, Section 8] as the distributions belonging to C_0 have a decreasing hazard rate.

Assumption A1 is trivial to verify. To check whether Assumption A2 holds, we can write $g_{(\ell_1,\dots,\ell_n)}(h)$ as $\sum_{i=1}^{n}(h_{\ell_i,i} - h_{\ell_i,i+1})$ and similarly $f_{(\ell_1,\dots,\ell_n)}(h)$ equals $\sum_{i=1}^{n}(f_{\ell_i,i}(h) - f_{\ell_i,i+1}(h))$. Therefore,

$$
\begin{aligned}
f_{(\ell_1,\dots,\ell_n)}(h) &= \lambda \sum_{i=1}^{n} 1[\ell_i > 1]\left(\sum_{j=0}^{d-1} h_{\ell_i-1,1}^{j} h_{\ell_i,1}^{d-1-j}\right) \\
&\quad \cdot [(h_{\ell_i-1,i} - h_{\ell_i,i}) - (h_{\ell_i-1,i+1} - h_{\ell_i,i+1})] \\
&= \lambda \sum_{i=1}^{n} 1[\ell_i > 1]\left(\sum_{j=0}^{d-1} h_{\ell_i-1,1}^{j} h_{\ell_i,1}^{d-1-j}\right) \\
&\quad \cdot [(h_{\ell_i-1,i} - h_{\ell_i-1,i+1}) - (h_{\ell_i,i} - h_{\ell_i,i+1})] \quad (20)
\end{aligned}
$$

If $\ell_1 > \ell_2 > \dots > \ell_n$, this can be written as

$$
\begin{aligned}
f_{(\ell_1,\dots,\ell_n)}(h) &= \lambda \sum_{i=1}^{n}\left(\sum_{j=0}^{d-1} h_{\ell_i-1,1}^{j} h_{\ell_i,1}^{d-1-j}\right) \\
&\quad \cdot [g_{(\ell_1,\dots,\ell_{i-1},\ell_i-1[\ell_i>1],\ell_{i+1},\dots,\ell_n)}(h) - g_{(\ell_1,\dots,\ell_n)}(h)].
\end{aligned}
$$

and Assumption A2 holds as $h \leq_C \tilde{h}$ implies that $h_{\ell,i} \leq \tilde{h}_{\ell,i}$ and $g_{(\ell'_1,\dots,\ell'_n)}(h) \leq g_{(\ell'_1,\dots,\ell'_n)}(\tilde{h})$ for $(\ell'_1,\dots,\ell'_n) \neq (\ell_1,\dots,\ell_n)$.

If $\ell_i = \ell_{i+1} > 1$ for some i, the above expression cannot be directly used as $\ell_i - 1[\ell_i > 1] \geq \ell_{i+1}$ does not hold. In general assume $\tilde{\ell}_1 > \tilde{\ell}_2 > \dots > \tilde{\ell}_k$ are the unique values appearing in the sequence $\ell_1 \geq \ell_2 \geq \dots \geq \ell_n$ and let ℓ_{j_i} be the first element in this sequence equal to $\tilde{\ell}_i$, then (20) becomes

$$
\begin{aligned}
f_{(\ell_1,\dots,\ell_n)}(h) &= \lambda \sum_{i=1}^{k} 1[\tilde{\ell}_i > 1]\left(\sum_{j=0}^{d-1} h_{\tilde{\ell}_i-1,1}^{j} h_{\tilde{\ell}_i,1}^{d-1-j}\right) \\
&\quad \cdot [(h_{\tilde{\ell}_i-1,j_i} - h_{\tilde{\ell}_i-1,j_{i+1}}) - (h_{\tilde{\ell}_i,j_i} - h_{\tilde{\ell}_i,j_{i+1}})].
\end{aligned}
$$

Assumption A2 now follows as $(h_{\tilde{\ell}_i-1,j_i} - h_{\tilde{\ell}_i-1,j_{i+1}}) - (h_{\tilde{\ell}_i,j_i} - h_{\tilde{\ell}_i,j_{i+1}})$ can be written as $g_{(\ell'_1,\dots,\ell'_n)}(h) - g_{(\ell_1,\dots,\ell_n)}(h)$ with $\ell'_s = \ell_s - 1$ for $j_i \leq s < j_{i+1}$ and $\ell'_s = \ell_s$ otherwise. Note that $\ell'_1 \geq \ell'_2 \geq \dots \geq \ell'_n$ as required because $\tilde{\ell}_i > \tilde{\ell}_{i+1}$.

Assumption A3 with $L = 1$ is immediate as $\sum_{\ell \geq 1} f_{\ell,1}(h) = \lambda$ for any $h \in \Omega$, meaning we can pick $a_{1,\pi}(h) = 0$. Finally, as $\sum_{\ell \geq L}(f_{\ell,1}(h) - f_{\ell,1}(\pi)) = \lambda(h_{L-1,1}^{d} - \pi_{L-1,1}^{d})$, setting $a_{L,\pi}(h) = 0$, $b_{L,L-1,1}(h) = \lambda \sum_{j=0}^{d-1}(h_{L-1,1})^{j}(\pi_{L-1,1})^{d-1-j} \leq \lambda d$ and $b_{L,\ell,j}(h) = 0$ for $(\ell,j) \neq (L-1,1)$ verifies Assumption A3 with $L > 1$.

When the buffer is finite of size B, the above discussion remains valid, except that we need to set $a_{L,\pi}(h) = \lambda(h_{B,1}^{d} - \pi_{B,1}^{d})$, for $L \geq 1$, such that Assumption A3 holds.

7.2 Pull and push strategies

In this example it is possible to show that for $\lambda < 1$ the set of ODEs has a unique fixed point that can be computed by determining the invariant distribution of an ergodic Quasi-Birth-Death Markov chain [32]. Note that the issue of global attraction is not addressed in [32]. Assumption A1 is readily verified. To verify Assumption

A2 it is not hard to show that $f_{(\ell_1,\dots,\ell_n)}(h)$ can be written as

$$
\begin{aligned}
f_{(\ell_1,\dots,\ell_n)}(h) &= \lambda g_{(\ell_1-1,\ell_2-1[\ell_2>1],\dots,\ell_n-1[\ell_n>1])}(h) \\
&\quad - \lambda g_{(\ell_1,\dots,\ell_n)}(h) - r(1-h_{1,1})g_{(\ell_1,\dots,\ell_n)}(h) \\
&\quad + r(1-h_{1,1})g_{(\ell_1+1,\ell_2+1[\ell_2>1],\dots,\ell_n+1[\ell_n>1])}(h).
\end{aligned}
$$

Thus if $g_{(\ell_1,\dots,\ell_n)}(h) = g_{(\ell_1,\dots,\ell_n)}(\tilde{h})$, then

$$
\begin{aligned}
f_{(\ell_1,\dots,\ell_n)}(\tilde{h}) - f_{(\ell_1,\dots,\ell_n)}(h) &= \lambda g_{(\ell_1-1,\ell_2-1[\ell_2>1],\dots,\ell_n-1[\ell_n>1])}(\tilde{h}) \\
&\quad - \lambda g_{(\ell_1-1,\ell_2-1[\ell_2>1],\dots,\ell_n-1[\ell_n>1])}(h) \\
&\quad + r(1-\tilde{h}_{1,1})g_{(\ell_1+1,\ell_2+1[\ell_2>1],\dots,\ell_n+1[\ell_n>1])}(\tilde{h}) \\
&\quad - r(1-h_{1,1})g_{(\ell_1+1,\ell_2+1[\ell_2>1],\dots,\ell_n+1[\ell_n>1])}(h) \\
&\quad + r(\tilde{h}_{1,1} - h_{1,1})g_{(\ell_1,\dots,\ell_n)}(\tilde{h}) \geq 0
\end{aligned}
$$

if $h \leq_C \tilde{h}$ as $g_{(\ell_1,\dots,\ell_n)}(\tilde{h}) \geq g_{(\ell_1+1,\ell_2+1[\ell_2>1],\dots,\ell_n+1[\ell_n>1])}(\tilde{h})$.

Assumption A3 with $L = 1$ follows by noting that $\sum_{\ell \geq 1} f_{\ell,1}(h) = \lambda = \sum_{\ell \geq 1} f_{\ell,1}(\pi)$. Finally, Assumption A3 with $L > 1$ can be verified as follows. Note that

$$
\begin{aligned}
\sum_{\ell \geq L}(f_{\ell,1}(h) - f_{\ell,1}(\pi)) &= \lambda(h_{L-1,1} - \pi_{L-1,1}) \\
&\quad - r((1-h_{1,1})h_{L,1} - (1-\pi_{1,1})\pi_{L,1}) = \lambda(h_{L-1,1} - \pi_{L-1,1}) \\
&\quad + r(h_{1,1} - \pi_{1,1})h_{L,1} - r(1-\pi_{1,1})(h_{L,1} - \pi_{L,1}),
\end{aligned}
$$

meaning Assumption A3 with $L > 1$ holds with $a_{L,\pi}(h) = r(1-\pi_{1,1})(h_{L,1} - \pi_{L,1})$, $b_{L,1,1}(h) = rh_{L,1} \leq r$, $b_{L,L-1,1}(h) = \lambda$ and all other $b_{L,\ell,j}(h)$ equal to zero.

7.3 JSQ(K,d): Join-the-Shortest-K-Queues among d randomly selected servers

With respect to assumption A0, we limit ourselves to the case where B is finite. The existence of a fixed point in Ω_B follows from the fact that a convex compact forward invariant set $\mathcal{K} \subset \mathbb{R}^{Bn}$ of a dynamical system has a fixed point in $\mathcal{K}$ [3], which is not hard to prove using Brouwer's fixed point theorem.

Contrary to the previous two examples, verifying Assumption A1 requires some work. First note that $f_{\ell,1}(h)$ only depends on $h_{\ell-1,1}$ and $h_{\ell,1}$ and therefore the functions $f_{\ell,1}(h)$ are non-decreasing in $h_{\ell',i}$ for $(\ell',i) \neq (\ell,1)$ if

$$
\phi_K(x) = \sum_{s=0}^{K-1}(K-s)\binom{d}{s}x^{d-s}(1-x)^{s},
$$

is non-decreasing on $[0, 1]$. We now prove that

$$
\phi'_K(x) = \sum_{s=0}^{K-1} d\binom{d-1}{s}x^{d-s-1}(1-x)^{s}, \quad (21)
$$

which is clearly positive on $[0, 1]$. By definition of $\phi_K(x)$ we have

$$
\phi'_K(x) = \sum_{s=0}^{K-1}(K-s)\binom{d}{s}x^{d-s-1}(1-x)^{s-1}((1-x)d - s).
$$

By induction on K we find

$$\phi'_K(x) = \sum_{s=0}^{K-1} \binom{d}{s} x^{d-s-1}(1-x)^{s-1}((1-x)d - s)$$
$$+ \sum_{s=0}^{K-2} d\binom{d-1}{s} x^{d-s-1}(1-x)^s.$$

Hence, (21) is equivalent to showing that

$$\sum_{s=0}^{K-1} d\binom{d}{s} x^{d-s-1}(1-x)^s =$$
$$\sum_{s=1}^{K-1} d\binom{d-1}{s-1} x^{d-s-1}(1-x)^{s-1} + d\binom{d-1}{K-1} x^{d-K}(1-x)^{K-1}.$$

which is easy to establish (using induction on K once more).

Let us now focus on the functions $f_{\ell,i}(h)$ with $i > 1$. Clearly, these functions are increasing in $h_{\ell-1,i}$. It remains to show that they are also increasing in $h_{\ell-1,1}$ and $h_{\ell,1}$, which holds if

$$\xi_K(x_1, x_2) = \frac{\phi_K(x_2) - \phi_K(x_1)}{x_2 - x_1},$$

is increasing in both components for $0 \le x_1 \le x_2 \le 1$. As $\xi_K(x_1, x_2)$ is symmetric, it suffices to argue that $\xi_K(x_1, x_2)$ is increasing in x_1. Further, demanding that the derivative of $\xi_K(x_1, x_2)$ with respect to x_1 is non-negative is equivalent to

$$\phi_K(x_2) \ge \phi_K(x_1) + \phi'_K(x_1)(x_2 - x_1),$$

which holds if and only if $\phi_K(x)$ is convex. Using (21) we have for $K < d$

$$\phi''_K(x) = \sum_{s=0}^{K-1} d\binom{d-1}{s} x^{d-s-2}(1-x)^s((d-1)(1-x) - s).$$

Using induction on K this can be rewritten as

$$\phi''_K(x) = d(d-1)\binom{d-2}{K-1} x^{d-K-1}(1-x)^{K-1},$$

which is clearly positive on $[0, 1]$. For $K = d$, we have $\phi''_K(x) = 0$ as $\phi_K(x) = dx$.

We now proceed with Assumption A2. As $f_{(\ell_1,...,\ell_n)}(h)$ equals $\sum_{i=1}^n (f_{\ell_i,i}(h) - f_{\ell_i,i+1}(h))$ and $\ell_1 > 1$, we note that

$$f_{(\ell_1,...,\ell_n)}(h) = \lambda \sum_{i=1}^n 1[\ell_i > 1]\xi_K(h_{\ell_i,1}, h_{\ell_i-1,1})$$
$$\cdot [(h_{\ell_i-1,i} - h_{\ell_i,i}) - (h_{\ell_i-1,i+1} - h_{\ell_i,i+1})]$$

If $\ell_1 > \ell_2 > \ldots > \ell_n$, this can be written as

$$f_{(\ell_1,...,\ell_n)}(h) = \lambda \sum_{i=1}^n \xi_K(h_{\ell_i,1}, h_{\ell_i-1,1})$$
$$\cdot [g_{(\ell_1,...,\ell_{i-1},\ell_i-1[\ell_i>1],\ell_{i+1},...,\ell_n)}(h) - g_{(\ell_1,...,\ell_n)}(h)],$$

where $\xi_K(x_1, x_2)$ was increasing in x_1 and x_2. The case where $\ell_i = \ell_{i+1}$ for some i can be dealt with in the same manner as in Example 7.1.

Assumption A3 with $L = 1$ holds as $\sum_{\ell \ge 1} f_{\ell,1}(h) = \lambda K = \sum_{\ell \ge 1} f_{\ell,1}(\pi)$. Finally, with respect to Assumption A3 with $L > 1$ we have

$$\sum_{\ell \ge L}(f_{\ell,1}(h) - f_{\ell,1}(\pi)) = \lambda(\phi_K(h_{L-1,1}) - \phi_K(\pi_{L-1,1}))$$
$$= \lambda\xi_K(\pi_{L-1,1}, h_{L-1,1})(h_{L-1,1} - \pi_{L-1,1}),$$

and $\xi_K(x_1, x_2) \le \phi'_K(1) = d$ due to the convexity of $\phi_K(x)$.

8 PROOF OF THEOREM 2

In this section we define $v_i = \mu_i(1 - p_i)$ to ease the notation. We start by showing that the order $\le_C$ is preserved over time.

PROPOSITION 6. *Assume that (A0-A2) hold and let $h, \tilde{h} \in \Omega$. Let $h(t)$ and $\tilde{h}(t)$ be the unique solution of (13-14) with $h(0) = h$ and $\tilde{h}(0) = \tilde{h}$, respectively. If $\mu_i(1 - p_i)$ is decreasing in i and $h \le_C \tilde{h}$, then $h(t) \le_C \tilde{h}(t)$ for any $t > 0$.*

PROOF. Assume that at some time t we have $h_{\ell,i}(t) = \tilde{h}_{\ell,i}(t)$ for some ℓ and i, while $h(t) \le_C \tilde{h}(t)$. We need to argue that

$$d\tilde{h}_{\ell,i}(t)/dt \ge dh_{\ell,i}(t)/dt,$$

as the order is otherwise violated at time $t+$.

As $h(t) \le_C \tilde{h}(t)$ implies that $h_{\ell',i'}(t) \le \tilde{h}_{\ell',i'}(t)$ for all ℓ' and i', it would be sufficient that $dh_{\ell,i}(t)/dt$ is non-decreasing in all $h_{\ell',i'}(t)$ with $(\ell', i') \ne (\ell, i)$. Looking at (13-14) and due to Assumption A1, we see that this is clearly the case, except perhaps for the sums over j (that are due to the service completions).

For $i > 1$, we have

$$-\sum_{j=i}^n (h_{\ell,j}(t) - h_{\ell,j+1}(t))v_j = \sum_{j=i+1}^n h_{\ell,j}(t)(v_{j-1} - v_j) - h_{\ell,i}(t)v_i,$$

meaning $dh_{\ell,i}(t)/dt$ is non-decreasing in any $h_{\ell',i'}(t)$ with $(\ell', i') \ne (\ell, i)$ when $i > 1$, as $v_i = \mu_i(1 - p_i)$ is decreasing in i (and positive).

For $i = 1$, we find

$$-\sum_{j=1}^n \left[(h_{\ell,j}(t) - h_{\ell,j+1}(t)) - (h_{\ell+1,j}(t) - h_{\ell+1,j+1}(t))\right] v_j$$
$$= \sum_{j=2}^n (h_{\ell,j}(t) - h_{\ell+1,j}(t))(v_{j-1} - v_j)$$
$$- (h_{\ell,1}(t) - h_{\ell+1,1}(t))v_1$$
$$= \sum_{j=2}^n (h_{\ell+1,1}(t) + h_{\ell,j}(t) - h_{\ell+1,j}(t))(v_{j-1} - v_j)$$
$$+ h_{\ell+1,1}(t)v_n - h_{\ell,1}(t)v_1.$$

This expression is decreasing in $h_{\ell+1,j}(t)$, for $j > 1$, which may appear as a problem. However, as $h(t) \le_C \tilde{h}(t)$, (16) with $\ell_1 = \ldots = \ell_{j-1} = \ell + 1$ and $\ell_j = \ldots = \ell_n = \ell$ implies that $h_{\ell,j}(t) + h_{\ell+1,1}(t) - h_{\ell+1,j}(t) \le \tilde{h}_{\ell,j}(t) + \tilde{h}(t)_{\ell+1,1} - \tilde{h}_{\ell+1,j}(t)$. As a result $dh_{\ell,1}(t)/dt$ does not exceed $d\tilde{h}_{\ell,1}(t)/dt$ as required.

We also need to verify that (16) remains valid, which corresponds to verifying that $g_{(\ell_1,...,\ell_n)}(h(t)) \le g_{(\ell_1,...,\ell_n)}(\tilde{h}(t))$ remains valid.

Assume that $g_{(\ell_1,\ldots,\ell_n)}(h(t)) = g_{(\ell_1,\ldots,\ell_n)}(\tilde{h}(t))$ for some $(\ell_1,\ldots,\ell_n)$, then we need to argue that

$$dg_{(\ell_1,\ldots,\ell_n)}(\tilde{h}(t))/dt \geq dg_{(\ell_1,\ldots,\ell_n)}(h(t))/dt,$$

whenever $h(t) \leq_C \tilde{h}(t)$ to complete the proof. Due to Assumption A2, we can restrict ourselves to showing that the terms of $dg_{(\ell_1,\ldots,\ell_n)}(h(t))/dt$ corresponding to phase changes and service completions are increasing in $g_{(\ell'_1,\ldots,\ell'_n)}(h(t))$ when $(\ell'_1,\ldots,\ell'_n) \neq (\ell_1,\ldots,\ell_n)$.

Phase changes increase $g_{(\ell_1,\ldots,\ell_n)}(h(t))$ if such a change occurs in a queue in phase i with a length in $[\ell_{i+1},\ell_i - 1]$. Therefore phase changes increase $g_{(\ell_1,\ldots,\ell_n)}(t)$ at rate

$$\sum_{i=1}^{n-1} \mu_i p_i (g_{(\ell_1,\ldots,\ell_{i-1},\ell_{i+1},\ell_{i+1},\ldots,\ell_n)}(h(t)) - g_{(\ell_1,\ldots,\ell_n)}(h(t))).$$

Service completions in queues in phase i that have a length in $[l_i, l_1]$ decrease $g_{(\ell_1,\ldots,\ell_n)}(h(t))$ at rate v_i (as the initial service phase of the next job in service is phase 1). The drift is therefore given by

$$-\sum_{i=1}^{n} v_i (g_{(\underbrace{\ell_1+1,\ldots,\ell_1+1}_{i-1 \text{ times}},\ell_i,\ldots,\ell_n)}(h(t)) - g_{(\underbrace{\ell_1+1,\ldots,\ell_1+1}_{i \text{ times}},\ell_{i+1},\ldots,\ell_n)}(h(t)))$$

$$= \sum_{i=2}^{n} g_{(\underbrace{\ell_1+1,\ldots,\ell_1+1}_{i-1 \text{ times}},\ell_i,\ldots,\ell_n)}(h(t))(v_{i-1} - v_i)$$

$$- g_{(\ell_1,\ldots,\ell_n)}(h(t))v_1 + g_{(\ell_1+1,\ldots,\ell_1+1)}(h(t))v_n.$$

$\square$

The next proposition shows that it suffices to prove attraction for points $h \in \Omega$ for which $h \leq_C \pi$ or $\pi \leq_C h$, where π is a fixed point of the ODEs (13-14).

PROPOSITION 7. *Let $h \in \Omega$ and assume $\mu_i(1 - p_i)$ is decreasing in i, then the trajectory $h(t)$ starting in $h(0) \in \Omega$ at time 0 converges pointwise to π provided that for any $h \in \Omega$ with $h \leq_C \pi$ or $\pi \leq_C h$, $h(t)$ with $h(0) = h$ converges pointwise to π.*

PROOF. Due to Proposition 6 it suffices to show that for any $h \in \Omega$ there exists a $h^{(l)}, h^{(u)} \in \Omega$, with $h^{(l)} \leq_C \pi$ and $\pi \leq_C h^{(u)}$, such that $h^{(l)} \leq_C h \leq_C h^{(u)}$. For $h^{(l)}$ we can simply pick the zero vector as $0 \leq_C h$ for any $h \in \Omega$. For $h^{(u)}$ we set $h^{(u)}_{\ell,i} = \max(h_{\ell,1},\pi_{\ell,1})$, for $i = 1,\ldots,n$. Hence,

$$h_{\ell,i} \leq h_{\ell,1} \leq \max(h_{\ell,1},\pi_{\ell,1}) = h^{(u)}_{\ell,i}$$

and for $\ell_1 \geq \ell_2 \geq \ldots \geq \ell_n \geq 1$ with $\ell_1 > \ell_n$

$$h_{\ell_1,1} + \sum_{i=2}^{n}(h_{\ell_i,i} - h_{\ell_{i-1},i}) \leq h_{\ell_n,1} \leq \max(h_{\ell_n,1},\pi_{\ell_n,1}) = h^{(u)}_{\ell_n,1}$$

$$= h^{(u)}_{\ell_1,1} + \sum_{i=2}^{n}(h^{(u)}_{\ell_i,1} - h^{(u)}_{\ell_{i-1},1}) = h^{(u)}_{\ell_1,1} + \sum_{i=2}^{n}(h^{(u)}_{\ell_i,i} - h^{(u)}_{\ell_{i-1},i}),$$

where the first inequality follows from (17). This shows that $h \leq_C h^{(u)}$ and similarly one finds that $\pi \leq_C h^{(u)}$. $\square$

Remark that when B is finite we can simply use $h^{(u)}_{B,n} = 1$ and $h^{(u)}_{\ell,i} = 0$ for $(\ell,i) \neq (B,n)$.

LEMMA 3. *Define $z_{1,L}(h(t)) = \sum_{\ell \geq L} h_{\ell,1}(t)$ and*

$$z_2(h(t)) = \sum_{i=2}^{n} h_{1,i}(t)(R_i - R_{i-1}).$$

Then,

$$\frac{d}{dt} z_{1,L}(h(t)) = \sum_{\ell \geq L} f_{\ell,1}(h(t)) - \sum_{j=1}^{n}(h_{L,j}(t) - h_{L,j+1}(t))v_j, \quad (22)$$

and

$$\frac{d}{dt} z_2(h(t)) = -h_{1,1}(t) + \sum_{j=1}^{n}(h_{1,j}(t) - h_{1,j+1}(t))v_j. \quad (23)$$

PROOF. The expression for $dz_{1,L}(h(t))/dt$ is immediate from (13). For $dz_2(h(t))/dt$ we can make use of (14) and obtain (after exchanging the order of the sums):

$$\frac{d}{dt} z_2(h(t)) = \sum_{i=2}^{n}(h_{1,i-1}(t) - h_{1,i}(t))p_{i-1}\mu_{i-1}R_i$$

$$- \sum_{i=2}^{n}(h_{1,i-1}(t) - h_{1,i}(t))p_{i-1}\mu_{i-1}R_{i-1}$$

$$- \sum_{j=2}^{n}\left(\sum_{i=2}^{j}(R_i - R_{i-1})\right)(h_{1,j}(t) - h_{1,j+1}(t))v_j.$$

Using (8) on the first sum, we can rewrite this as

$$\frac{d}{dt} z_2(h(t)) = \sum_{i=2}^{n}(h_{1,i-1}(t) - h_{1,i}(t))v_{i-1}R_{i-1}$$

$$- \sum_{i=2}^{n}(h_{1,i-1}(t) - h_{1,i}(t)) - \sum_{j=2}^{n}(h_{1,j}(t) - h_{1,j+1}(t))v_jR_j$$

$$+ \sum_{j=2}^{n}(h_{1,j}(t) - h_{1,j+1}(t))v_jR_1$$

$$= (h_{1,1}(t) - h_{1,2}(t))v_1R_1 - h_{1,1}(t) + h_{1,n}(t) - h_{1,n}(t)v_nR_n$$

$$+ \sum_{j=2}^{n}(h_{1,j}(t) - h_{1,j+1}(t))v_jR_1.$$

The result follows by noting that $R_1 = 1$ and $R_n = 1/\mu_n$. $\square$

PROPOSITION 8. *Assume that (A0-A3) hold and that $\mu_i(1 - p_i)$ is decreasing in i. For any $h(0) \in \Omega$ with $h(0) \leq_C \pi$ or $\pi \leq_C h(0)$, $h(t)$ converges pointwise to π.*

PROOF. We assume $\pi \leq_C h(0)$, the proof for $h(0) \leq_C \pi$ proceeds similarly. We first show that $h_{1,1}(t)$ converges to $\pi_{1,1}$. As $\pi \leq_C h(t)$ for $t \geq 0$ due to Proposition 6, it suffices to show $\int_{t=0}^{\infty}(h_{1,1}(t) - \pi_{1,1})dt < \infty$. Let $z(h) = z_{1,1}(h) + z_2(h)$, then by Lemma 3 and Assumption A3 with $L = 1$, we have $dz(h(t))/dt = \sum_{\ell \geq 1} f_{\ell,1}(\pi) - h_{1,1}(t) - a_1(h(t)) \leq 0$. Further $\sum_{\ell \geq 1} f_{\ell,1}(\pi) = \pi_{1,1}$ as $dz(\pi)/dt = 0$. Therefore,

$$\int_{t=0}^{\tau}(h_{1,1}(t) - \pi_{1,1})dt = -\int_{t=0}^{\tau}\frac{dz(t)}{dt}dt - \int_{t=0}^{\tau} a_1(h(t))dt$$

$$\leq z(h(0)) - z(h(\tau)) \leq z(h(0)),$$

as $\pi \leq_C h(t)$ and $z(h(\tau)) \geq 0$ for $\tau \geq 0$. Hence, $\int_{t=0}^{\tau}(h_{1,1}(t)-\pi_{1,1})dt$ is uniformly bounded in $\tau \geq 0$, meaning $\int_{t=0}^{\infty}(h_{1,1}(t)-\pi_{1,1})dt < \infty$.

We now show that $h_{1,j}(t)$ converges to $\pi_{1,j}$, for $j = 2, \ldots, n$, by arguing that $\int_{t=0}^{\tau}(h_{1,j}(t)-\pi_{1,j})dt$ is uniformly bounded in $\tau \geq 0$. As $v_{j-1} - v_j > 0$ it suffices to show that

$$\int_{t=0}^{\tau} \sum_{j=2}^{n}(h_{1,j}(t)-\pi_{1,j})(v_{j-1}-v_j)dt$$

is uniformly bounded in $\tau \geq 0$. As $z_2(h(\tau)) \geq 0$, we have $z_2(h(0)) \geq -\int_{t=0}^{\tau}\frac{dz_2(h(t))}{dt}dt$ and Lemma 3 implies

$$z_2(h(0)) \geq \int_{t=0}^{\tau} h_{1,1}(t)dt - \int_{t=0}^{\tau}\sum_{j=1}^{n}(h_{1,j}(t)-h_{1,j+1}(t))v_jdt$$

$$= \int_{t=0}^{\tau} h_{1,1}(t)(1-v_1)dt + \int_{t=0}^{\tau}\sum_{j=2}^{n}h_{1,j}(t)(v_{j-1}-v_j)dt$$

$$= \int_{t=0}^{\tau}(h_{1,1}(t)-\pi_{1,1})(1-v_1)dt$$

$$+ \int_{t=0}^{\tau}\sum_{j=2}^{n}(h_{1,j}(t)-\pi_{1,j})(v_{j-1}-v_j)dt,$$

where the last inequality is due to the fact that π is a fixed point. This shows the uniform boundedness in τ as $0 \leq \int_{t=0}^{\tau}(h_{1,1}(t)-\pi_{1,1})dt \leq z(h(0))$.

We complete the proof by showing that $(h_{L,j}(t)-h_{L,j+1}(t))$ converges to $(\pi_{L,j}-\pi_{L,j+1})$, for $L > 1$ and $j = 1, \ldots, n$. Note that $(h_{L,j}(t)-h_{L,j+1}(t))$ is not necessarily larger than $(\pi_{L,j}-\pi_{L,j+1})$ when $\pi \leq_C h(t)$. We do however have that $h_{L,j}(t) + h_{L-1,j+1}(t) - h_{L,j+1}(t) \geq \pi_{L,j} + \pi_{L-1,j+1} - \pi_{L,j+1}$ when $\pi \leq_C h(t)$ due to (16). Thus, using induction on L it suffices to show that

$$\Psi_{L,\tau} = \int_{t=0}^{\tau}\sum_{j=1}^{n}((h_{L,j}(t)+h_{L-1,j+1}(t)-h_{L,j+1}(t))$$

$$- (\pi_{L,j}+\pi_{L-1,j+1}-\pi_{L,j+1}))v_jdt \quad (24)$$

is uniformly bounded in $\tau \geq 0$. As $z_{1,L}(h(\tau)) \geq 0$ for $\tau \geq 0$, we have $z_{1,L}(h(0)) \geq -\int_{t=0}^{\tau}\frac{dz_{1,L}(h(t))}{dt}dt$ and Lemma 3 for $L > 1$ implies

$$z_{L,1}(h(0)) \geq$$

$$- \int_{t=0}^{\tau}\sum_{\ell \geq L}f_{\ell,1}(h(t))dt + \int_{t=0}^{\tau}\sum_{j=1}^{n}(h_{L,j}(t)-h_{L,j+1}(t))v_jdt$$

$$= - \int_{t=0}^{\tau}\sum_{\ell \geq L}(f_{\ell,1}(h(t))-f_{\ell,1}(\pi))dt$$

$$+ \int_{t=0}^{\tau}\sum_{j=1}^{n}((h_{L,j}(t)-h_{L,j+1}(t))-(\pi_{L,j}-\pi_{L,j+1}))v_jdt$$

where the last equality holds as π is a fixed point. Therefore we find

$$z_{L,1}(h(0)) + \int_{t=0}^{\tau}\sum_{j=1}^{n}(h_{L-1,j+1}(t)-\pi_{L-1,j+1})v_jdt$$

$$\geq - \int_{t=0}^{\tau}\sum_{\ell \geq L}(f_{\ell,1}(h(t))-f_{\ell,1}(\pi))dt + \Psi_{L,\tau}$$

By relying on Assumption A3 with $L > 1$ this can be restated as

$$z_{L,1}(h(0)) + \int_{t=0}^{\tau}\sum_{j=1}^{n}(h_{L-1,j+1}(t)-\pi_{L-1,j+1})v_jdt$$

$$+ \int_{t=0}^{\tau}\sum_{\ell=1}^{L-1}\sum_{j=1}^{n}b_{L,\ell,j}(h(t))(h_{\ell,j}(t)-\pi_{\ell,j})dt$$

$$\geq \Psi_{L,\tau} + \int_{t=0}^{\tau}a_{L,\pi}(h(t))dt.$$

As $b_{L,\ell,j}(h)$ is bounded on Ω, the left hand side is uniformly bounded in τ by induction on L and therefore so are the (positive) integrals on the right hand side. □

Theorem 2 follows by combining Proposition 7 and 8.

9 CONCLUSIONS

In this paper we demonstrated that monotonicity arguments can still be applied to prove global attraction of mean field models with hyperexponential job sizes, which is a widely used class of distributions for systems exhibiting large job size variability. The key ideas to enable the use of such monotonicity arguments existed in formulating the ODE-based mean field model using a Coxian representation and introducing a partial order that is stronger than the usual componentwise order.

We believe that the approach presented in this paper can be extended to heterogeneous systems and systems in which a server can serve multiple jobs simultaneously (i.e., in which each server behaves as an $\cdot/Cox/C$ server). Whether the assumption on the first-come-first-served scheduling discipline can be relaxed is unclear at this moment and is the topic of future work.

REFERENCES

[1] A. Karthik A. Mukhopadhyay and R. R. Mazumdar. Randomized assignment of jobs to servers in heterogeneous clusters of shared servers for low delay. *Stochastic Systems*, 6(1):90 – 131, 2016.
[2] M. Benaïm and J. Le Boudec. On mean field convergence and stationary regime. *CoRR*, abs/1111.5710, Nov 24 2011.
[3] N. P. Bhatia and G. P. Szegö. *Stability theory of dynamical systems*. Springer Science & Business Media, 2002.
[4] A. Bobbio, A. Horvth, and M. Telek. Matching three moments with minimal acyclic phase type distributions. *Stochastic Models*, 21(2-3):303–326, 2005.
[5] M. Bramson, Y. Lu, and B. Prabhakar. Asymptotic independence of queues under randomized load balancing. *Queueing Syst.*, 71(3):247–292, 2012.
[6] A. Braverman, JG Dai, and J. Feng. Stein's method for steady-state diffusion approximations: an introduction through the erlang-a and erlang-c models. *Stochastic Systems*, 6(2):301–366, 2017.
[7] F Cecchi, SC Borst, and JSH van Leeuwaarden. Mean-field analysis of ultra-dense csma networks. *ACM SIGMETRICS Performance Evaluation Review*, 43(2):13–15, 2015.
[8] A. Cumani. On the canonical representation of homogeneous markov processes modelling failure - time distributions. *Microelectronics Reliability*, 22(3):583 – 602, 1982.

[9]   S.N. Ethier and T.C. Kurtz. *Markov processes: characterization and convergence.* Wiley, 1986.

[10] DG Feitelson. *Workload Modeling for Computer Systems Performance Evaluation.* Cambridge University Press, New York, NY, USA, 1st edition, 2015.

[11] A. Feldmann and W. Whitt. Fitting mixtures of exponentials to long-tail distributions to analyze network performance models. *Performance Evaluation*, 31(3):245 – 279, 1998.

[12] A. Ganesh, S. Lilienthal, D. Manjunath, A. Proutiere, and F. Simatos. Load balancing via random local search in closed and open systems. *SIGMETRICS Perform. Eval. Rev.*, 38(1):287–298, June 2010.

[13] N. Gast. Expected values estimated via mean-field approximation are 1/n-accurate. *Proc. ACM Meas. Anal. Comput. Syst.*, 1(1):17:1–17:26, June 2017.

[14] N. Gast and B. Gaujal. A mean field model of work stealing in large-scale systems. *SIGMETRICS Perform. Eval. Rev.*, 38(1):13–24, June 2010.

[15] N. Gast and B. Van Houdt. Transient and steady-state regime of a family of list-based cache replacement algorithms. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 123–136. ACM, 2015.

[16] REA Khayari, R. Sadre, and B. R. Haverkort. Fitting world-wide web request traces with the EM-algorithm. *Performance Evaluation*, 52(2):175–191, 2003.

[17] T. Kurtz. *Approximation of population processes.* Society for Industrial and Applied Mathematics, 1981.

[18] M. Lin, B. Fan, J.C.S. Lui, and D. Chiu. Stochastic analysis of file-swarming systems. *Performance Evaluation*, 64(9):856–875, 2007.

[19] J. B. Martin and Yu. M. Suhov. Fast jackson networks. *Ann. Appl. Probab.*, 9(3):854–870, 08 1999.

[20] L. Massoulié and M. Vojnović. Coupon replication systems. *SIGMETRICS Perform. Eval. Rev.*, 33(1):2–13, June 2005.

[21] W. Minnebo and B. Van Houdt. A fair comparison of pull and push strategies in large distributed networks. *IEEE/ACM Transactions on Networking*, 22:996–1006, 2014.

[22] M. Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.*, 12:1094–1104, October 2001.

[23] C. O'Cinneide. On non-uniqueness of representations of phase-type distributions. *Communications in Statistics. Stochastic Models*, 5(2):247–259, 1989.

[24] Takayuki Osogami and Mor Harchol-Balter. Closed form solutions for mapping general distributions to quasi-minimal ph distributions. *Perform. Eval.*, 63(6):524–552, June 2006.

[25] A. Riska, V. Diev, and E. Smirni. An EM-based technique for approximating long-tailed data sets with ph distributions. *Performance Evaluation*, 55(1):147 – 164, 2004.

[26] Grégory Roth and William H Sandholm. Stochastic approximations with constant step size and differential inclusions. *SIAM Journal on Control and Optimization*, 51(1):525–555, 2013.

[27] D. Starobinski and M. Sidi. Modeling and analysis of power-tail distributions via classical teletraffic methods. *Queueing Systems*, 36(1-3):243–267, 2000.

[28] Dietrich Stoyan and Daryl J Daley. Comparison methods for queues and other stochastic models. *JOHN WILEY & SONS, INC., 605 THIRD AVE., NEW YORK, NY 10158, USA, 1983*, 1983.

[29] H.C. Tijms. *Stochastic models: an algorithmic approach.* Wiley series in probability and mathematical statistics. John Wiley & Sons, 1994.

[30] B. Van Houdt. A mean field model for a class of garbage collection algorithms in flash-based solid state drives. *ACM SIGMETRICS Perform. Eval. Rev.*, 41(1):191–202, 2013.

[31] B. Van Houdt. Performance of garbage collection algorithms for flash-based solid state drives with hot/cold data. *Perform. Eval.*, 70(10):692–703, 2013.

[32] B. Van Houdt. Randomized work stealing versus sharing in large-scale systems with non-exponential job sizes. *arXiv preprint*, 2018. arXiv:1810.13186.

[33] T. Vasantam, A. Mukhopadhyay, and R. R. Mazumdar. Mean-field analysis of loss models with mixed-erlang distributions under power-of-d routing. In *2017 29th International Teletraffic Congress (ITC 29)*, volume 1, pages 250–258, Sept 2017.

[34] N.D. Vvedenskaya, R.L. Dobrushin, and F.I. Karpelevich. Queueing system with selection of the shortest of two queues: an asymptotic approach. *Problemy Peredachi Informatsii*, 32:15–27, 1996.

[35] W. Whitt. Approximating a point process by a renewal process, i: Two basic methods. *Oper. Res.*, 30(1):125–147, February 1982.

[36] Q. Xie, X. Dong, Y. Lu, and R. Srikant. Power of d choices for large-scale bin packing: A loss model. *SIGMETRICS Perform. Eval. Rev.*, 43(1):321–334, June 2015.

[37] L. Ying. On the approximation error of mean-field models. In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*, SIGMETRICS '16, pages 285–297, New York, NY, USA, 2016. ACM.

[38] L. Ying, R. Srikant, and X. Kang. The power of slightly more than one sample in randomized load balancing. In *2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 1131–1139, April 2015.

A PROOF OF LEMMA 2

By definition of R_i we have

$$R_i = \frac{1}{\mu_i} + \sum_{j=i}^{n-1}\left(\prod_{k=i}^{j} p_k\right)\frac{1}{\mu_{j+1}},$$

which implies that

$$R_i - R_{i-1} = \frac{1-p_{i-1}}{\mu_i} + \sum_{j=i}^{n-1}\left(\prod_{k=i}^{j} p_k\right)\frac{1-p_{i-1}}{\mu_{j+1}} - \frac{1}{\mu_{i-1}}.$$

As $\mu_i(1-p_i)$ is decreasing in i, we have $(1-p_{i-1})/\mu_{j+1} > (1-p_{j+1})/\mu_{i-1}$ for $j \geq i-1$. Hence,

$$R_i - R_{i-1} > \frac{1-p_i}{\mu_{i-1}} + \sum_{j=i}^{n-1}\left(\prod_{k=i}^{j} p_k\right)\frac{1-p_{j+1}}{\mu_{i-1}} - \frac{1}{\mu_{i-1}}$$

$$= \frac{1}{\mu_{i-1}}\left(\sum_{j=i}^{n-1}\left(\prod_{k=i}^{j} p_k\right) - \sum_{j=i-1}^{n-1}\left(\prod_{k=i}^{j+1} p_k\right)\right)$$

$$= -\frac{1}{\mu_{i-1}}\prod_{k=i}^{n} p_k = 0,$$

as $p_n = 0$.

B PROOF OF PROPOSITION 5

Let β be the unique invariant vector of $S + (-Se)\alpha$, that is, $\beta(S + (-Se)\alpha) = 0$ and $\beta e = 1$. It is easy to verify that $\beta_i = (\prod_{j=1}^{i-1} p_j)/\mu_i$ (as the mean service time $\sum_i \beta_i \mu_i$ equals one).

Due to (14) with $i = 2$, we immediately have that

$$\sum_{j=1}^{n}(\pi_{1,j} - \pi_{1,j+1})v_j = (\pi_{1,1} - \pi_{1,2})\mu_1. \tag{25}$$

By means of (14) we have for $i \geq 2$:

$$\frac{d}{dt}(h_{1,i}(t) - h_{1,i+1}(t)) = (h_{1,i-1}(t) - h_{1,i}(t))p_{i-1}\mu_{i-1} - (h_{1,i}(t) - h_{1,i+1}(t))\mu_i,$$

meaning

$$(\pi_{1,i-1} - \pi_{1,i})p_{i-1}\mu_{i-1} = (\pi_{1,i} - \pi_{1,i+1})\mu_i. \tag{26}$$

Combining (25) and (26), we see that the vector

$$(\pi_{1,1} - \pi_{1,2}, \pi_{1,2} - \pi_{1,3}, \ldots, \pi_{1,n-1} - \pi_{1,n}, \pi_{1,n}),$$

is an invariant vector of $S + (-Se)\alpha$ and thus a multiple of the vector β. The result then follows as $\beta e = 1$.